

Hot Coffee: Associative Memory with Bump Attractor Cell Assemblies of Spiking Neurons

Christian Robert Huyck · Alberto Arturo Vergani

Received: date / Accepted: date

Abstract Networks of spiking neurons can have persistently firing stable bump attractors, to represent large number spaces (like temperature). This can be done with a topology that has local excitatory synapses, with surround inhibition, but with any particular topology there are critical points in the weight space. Activating large ranges in the bump can lead to streams, that show repeller attractor dynamics, however, these streams can be merged by overcoming these repeller dynamics. A simple associative memory can include these bump attractors, allowing the use of continuous variables in these memories. This is a step toward a more complete cognitive associative memory.

1 Introduction

How does the brain represent concepts that are continuously valued, like height, weight, and temperature? How can these be included in the brain's associative memory. For example, what is the neural basis of the representation of hot coffee? Moreover, how can coffee be considered cold at one temperature, and another drink, say a coke, be considered warm at the same temperature?

The authors firmly believe that concepts, like hot and coffee, are represented in the brain by Cell Assemblies (CAs) (see section 2.4). Associations are represented either by topological connections or by CAs in the brain. In the simulations described in this paper, basic semantic concepts, like coffee are represented by a simple version of CAs called binary CAs.

Other concepts, like temperature, are represented by continuously valued networks that fall into a category called Winner Take All (WTA) networks, or more

C. Huyck
Middlesex University
Tel.: +123-45-678910
E-mail: c.huyck@mdx.ac.uk

A. Vergani
Middlesex University

precisely, bump attractors (see sections 2.1). These bump attractors support another simple version of CAs that differs from binary CAs. In essence, a network of n neurons can support n of these CAs.

While well studied, these bump attractors have an unusual problem when they are activated by a large field; they split into streams of firing neurons (see sections 3.4 and 4.1). This paper proposes a neural topology for a particular associative memory involving hot coffee. The multiple stream problem is resolved, leading to a complete solution for this particular hot coffee associative memory. The memory is generalisable.

Motivation and goals

The authors' overall motivation is to understand the brain and the mind that emerges from it. It is an understatement to say that this is an enormous task, so a smaller version is to understand relatively broad cognitive tasks, in this case, associative memory with components that have continuous values. It should be noted that the authors are not claiming that these simulations are good neuron by neuron representations of the a human brain or even a small part of it; instead these are simplified models. None the less, the point neural models are commonly used, albeit simple, models of biological neurons. Similarly, the topologies are relatively sparse, so that they could reasonably be subsets of the actual biological topology. In this paper, continuously valued features are represented by discrete neural steps. The granularity can be increased arbitrarily.

The goal of this paper is to broaden the reader's and the authors' understanding of CAs, associative memory, and neural processing. A particularly important aspect of this understanding is to show that stable bump behaviour can support persistent firing, and thus CAs. Moreover, this stable bump behaviour that is known to work as an attractor, also works as a repeller. Finally, the paper (and associated code ¹ should provide topologies, and an underlying theory for new topologies, that could be used as improved associative memories in, for example, neural agents.

2 Literature Review

A great deal of literature relates to this paper, but this review concentrates on four bodies. The first is Winner Take All Networks, and bump attractors in particular (section 2.1). There has been a great deal of mathematical analysis of this relatively simple form of neural topology. The main cognitive component of the paper is the associative memory (section 2.2) that allows in grounding concepts representation in the brain (section 2.3). Finally, the bridge between the two is the Cell Assembly (section 2.4).

2.1 Winner Take All Networks

The brain is a part of the central nervous system, which processes multi-modal information. Although there are several sources of stimuli (coming from inside and outside of the body), the brain selectively analyzes this huge amount of information. The information is often ambiguous, so the brain must select one of

¹ Code can be found at <http://www.cwa.mdx.ac.uk/NEAL/wta.html>.

the possible options. In simulated neural and connectionist systems, one model to select between options is a winner take all (WTA) system. One connectionist system that uses this model is the self organizing map (Kohonen, 1982). The map is made up of several nodes, and when an item is presented, the nearest node wins.

The stationary bump, which is another mechanism proposed for feature selectivity in the brain (Somers et al., 1995; Laing et al., 2001), can be used in a spiking neuron network. The bump activity is an example of WTA neural behaviour because a group of co-firing neurons can be considered as winners of competition via inhibitory synapses. Adopting a general point of view, both WTA network and bump activity of multiple neurons are a type of pattern formation process in a population with excitatory and inhibitory neurons (i.e., they work on the patterns of a stable grid (Wilson and Cowan, 1973) and on the process of activity dependent neural group selection (Edelman, 1987)).

In the typical stationary bump model, distance is considered either in one or two dimensions. There are local excitatory synapses, and more broad inhibitory synapses. The recurrent neural network is able to select particular neurons using inhibitory synapses that sustain the competition between neurons (e.g. Chen's surround inhibition (Chen, 2017)).

There is evidence that excitatory cells (i.e., principal neurons) are associated with specialized inhibitory cells (i.e., interneurons or secondary cells) that synapse to principal cells as well as other interneurons. The proper dynamics in the neural network can only be sustained if the excitatory behaviour of principal cells is modulated by the stopping function of secondary cells. If there were only excitatory neurons, their positive spikes could lead to an excitation that produces more excitation (an avalanche effect potentially leading to simulated epilepsy), and therefore, it would be difficult to observe transiently active groups of co-firing neurons such as, for example, CAs (see section 2.4).

This type of stationary bump is widely used in neural simulations. For example, it is used to manage a robot's direction in a path integration task (Kreiser et al., 2018).

2.2 Associative Memory

It is widely agreed that in the brain (and mind) concepts do not exist in isolation, but instead are associated with each other. These associations are part of the underlying concepts, making an associative memory. This is a long standing psychological theory (Quillian, 1967).

This associative memory is the basis of priming effects (Collins and Loftus, 1975). If a concept is activated, it spreads its activation to associated concepts. The associative memory can be thought of as a symbolic Semantic Network, and Semantic Nets are widely used in AI for knowledge representation (Brachman and Schmolze, 1989).

Early versions of associative memories in simulated biological neural nets typically refer to associating vectors of firing neurons (Willshaw et al., 1969). An input vector of neurons, when fired once, causes an associated output vector to fire once. This work however is not particularly well suited for bridging the gap between actual neural behaviour and the emerging psychological behaviour of, in

this case, associative memory, because individual concepts are not represented by a vector of neurons firing once, but by persistently firing cell assemblies.

2.3 Concepts Representation in the brain

Continuous concepts are mental representations able to ground physics dimensions, like time, space, temperature, pressures, force and so on. Those concepts share a common mathematical structure that is processed by specific anatomical regions. Nieder and Dehaene (2009) made a review about a broad range of methodologies in humans and non-humans describing that the numerical information is encoded in the prefrontal and posterior parietal lobes. From a more general perspective, continuous concepts are abstract ideas opposed to the concrete ones, but this dichotomy is too simple to be related with the complex notion of a concept represented in the brain. In fact, Ghio et al. (2013) setted a psycholinguistic experiment by using auditory and visual neuroimaging founding that the abstract-concrete dichotomy is insufficient to account for the entire semantic variability within the two macro-domains.

Moreover, quantity related concepts helps in representing multifaceted ideas shaped with different features, for example animals, tools and objects. For example, Martin (2007) collected the evidences from functional neuroimaging studies about the storage in the cortex of salient properties of an object, like movements, functions, shapes, etc, and they found out that those features are stored in separately sensory and motor systems, suggesting that object concepts emerge from a weighted activity of property-based brain regions. Handjaras et al Handjaras et al. (2016), instead, investigated modality independent and category based organization of semantic knowledge in human brain; they concluded that there are patterns of neural activity spreaded in a large semantic cortical network that comprised parahippocampal, lateral occipital, temporo-parieto-occipital and inferior parietal cortices, that are correlated with domain linguistic production and were both independent from the modality of stimulus presentation.

The concept representation regarding this paper is focused on three main concepts: i) the continuous-like physical scale of the temperature, ii) the discrete representation of beverages (coke and coffee) and iii) the crisp judgments about the drink thermal state (cold, warm and hot).

2.4 Cell Assemblies

In his book, Hebb (1949) developed his famous synaptic learning rule, and used that rule to propose the CA as the neural basis of concepts. That is, semantic concepts, such as coffee are represented by CAs. A CA is a relatively small group of neurons that have high mutual synaptic strength, which is formed by Hebbian learning. When some of the neurons fire, that mutual synaptic weight supports firing in the other neurons in the CA, and this allows a cascade of firing so that the neurons in the CA can fire persistently for a considerable amount of time (seconds). This firing is the neural basis of short term psychological memory.

In the intervening 70 years, there has been significant and growing evidence of the existence of CAs in brains (Singer et al., 1997; Harris, 2005; Buzsaki, 2010),

and there is evidence of CAs in all major cortical areas (Huyck and Passmore, 2013). Neurons in a CA fire persistently, once activated, and fire synchronously.

While the neurobiological evidence is accumulating, there are not very good neural simulations of CAs. While the authors have spent considerable time simulating CAs, and used them in many tasks, recent work has made extensive use of binary CAs (Huyck and Mitchell, 2018). In these simulated CAs, the neurons are either mostly firing, or none are firing, so it is binary, either on or off. This can be implemented with a well connected topology of neurons. Once the neurons start firing, they will fire persistently until some external source shuts them off. This is obviously a poor model of CAs because, among many reasons, in a normal case, CAs would stop firing on their own, like normal short term memories stop on their own.

None the less, binary CAs have been used in simulations of associative memory (Huyck and Ji, 2018). In this case, three concepts are associated, and when two become active, the associated third comes on. That is a $2/3$ associative memory. This is an example of work in attractor networks to form associative memory (Lansner, 2009). Using spiking point neural models starts to bring together more biologically accurate simulations to manage more complex, and psychologically accurate neuropsychological simulations.

The elevated firing of a CA is the neural correlate of a short term or working memory item. In the case of the $2/3$ associative memory items, two CAs, instantiating two concepts, are firing at an elevated rate, and they cause the third to fire at an elevated rate, retrieving the concepts associated with both of the first two. There is evidence that bump attractors instantiate CAs for continuously valued phenomena. For instance, recordings of pre-frontal cortical neurons of monkeys in oculomotor delay response tasks is consist with bump attractors (Wimmer et al., 2014) representing position. Similarly, single cell recordings of grid cells, are consistent with 2-D continuous attractors (Yoon et al., 2013). Bump attractors have been used to model hippocampal place cells (Stringer et al., 2002), and head direction cells (Redish et al., 1996). While these models, and indeed the models used in this paper have no short term plasticity, it seems that there is a sound basis for supporting the use of the importance of this type of plasticity in continuous attractor neural models for short term memory (Seeholzer et al., 2019). Here we propose a relatively simple neural model that represents continuously valued concepts, and uses them as a component of an associative memory.

2.5 Learning

While spiking neural networks possess many benefits, such as parallelism, perhaps the main benefit is their ability to learn. Learning in biological neural nets can be divided into three categories: structural plasticity, long term plasticity, and short term plasticity.

Structural plasticity is the growth or death of neurons and synapses, or synaptic movement from one post-synaptic neuron to another (see (Butz et al., 2009) for review); it is a change in gross topology. Though structural plasticity is more prevalent in the young brain, it does occur in the mature brain. Long term plasticity is the increase (potentiation) or decrease (depression) of synaptic efficiency; this is typically considered in computational models to be permanent, but biologically

seems to be merely long-term. undone Short term plasticity is a short lasting increase (potentiation) or decrease (depression) of synaptic efficiency. undone

Clearly structural plasticity has an effect on long and short term plasticity, since synaptic weight cannot be modified if there is no synapse, and there can be no synapse without a neuron. It seems likely that long and short term plasticity have an effect on structural plasticity and each other. undone citations

The boundary between long and short term plasticity is somewhat arbitrary. undone

Hebbian activity dependant vs homeostatic Plasticity has two functions in the brain, learning and retaining neural firing stability. Maintaining firing stability requires there to be ongoing neural firing, but not too much firing. Plasticity can lead to too much firing, and maintaining this homeostasis can come from, for instance, synaptic scaling (Goel et al., 2019) synaptogenesis (De Paola et al., 2006).

3 Materials and Methods

The overall goal of these simulations is persistently firing WTA networks (see section 3.2). These are based on reasonably biologically accurate point neuron models (see section 3.1).

3.1 Neural and Synaptic Models

The biophysical neuron model used in the simulations described in this paper is a leaky integrate and fire model with a fixed threshold. Synaptic conductance is transmitted at a decaying-exponential rate from the pre to post-synaptic neurons (Gerstner et al., 2014). The simulations are coded in PyNN (Davison et al., 2007) to specify the topology, flow of inputs, and recording. The neurons themselves are simulated using NEST (Gewaltig and Diesmann, 2007).

The model used in this paper follows Fourcaud-Trocmé et al. (2003) (but also see (Richardson and Gerstner, 2003; Gewaltig and Diesmann, 2007)). The activation is the current voltage V_M . Equation 1 describes the change in voltage, V_M is the membrane potential and C_M is the membrane capacity. The four currents are the leak current, the currents from excitatory and inhibitory synapses, and the input current (from some external source). The variable currents are governed by equations 2, 3 and 4. In equations 2 and 3 E_{Ex}^{rev} and E_{In}^{rev} are the reversal potentials; excitation and inhibition respectively change slow as the voltage approaches these reversal potentials. In equation 4, V_{rest} is the resting potential of the neuron, and τ_M is the leak constant.

$$\frac{dV_M}{dt} = \frac{(-I_{Leak} - I_{Ex}^{syn} - I_{In}^{syn} + I_{Ext})}{C_M} \quad (1)$$

$$I_{Exc}^{syn} = G_{Ex} \times (V_M - E_{Ex}^{rev}) \quad (2)$$

$$I_{Inh}^{syn} = G_{In} \times (V_M - E_{In}^{rev}) \quad (3)$$

$$I_{Leak} = \frac{C_M(V_M - V_{rest})}{\tau_M} \quad (4)$$

$$G_{Ex}(t) = k_{Ex} \times t \times e^{-\frac{t}{\tau_{Ex}^{syn}}} \quad (5)$$

$$G_{In}(t) = k_{In} \times t \times e^{-\frac{t}{\tau_{In}^{syn}}} \quad (6)$$

In equations 5 and 6, G_{Ex} and G_{In} the results are the conductance in mS/cm^2 to scale the post-synaptic potential amplitudes used in equation 2, and 3. t is the time step. The constant k_{Ex} and k_{In} are chosen so that $G_{Ex}(\tau_{Ex}^{syn}) = 1$ and $G_{In}(\tau_{In}^{syn}) = 1$. The τ_{Ex}^{syn} and the τ_{In}^{syn} are the decay rate of excitatory and inhibitory synaptic current.

When the voltage reaches the threshold, there is a spike and the voltage is reset. No current is transferred during the refractory period $\tau_{refract}$. In these simulations $v_{thresh} = -48.0mV$, $\tau_{refract} = 2.0$ ms. The time step t is 1ms. $C_M = 1.0nF$, $v_{reset} = -70.0mV$, $v_{rest} = -65.0mV$, $E_{Ex}^{rev} = 0.0mV$, $E_{In}^{rev} = -70mV$, $\tau_{Ex}^{syn} = 5.0ms$, $\tau_{In}^{syn} = 5.0ms$ and $\tau_M = 20.0ms$. These are all the default values. The particular parameters v_{thresh} , $\tau_{refract}$, and t , were selected as the authors have used them for prior simulations; they are the parameters used in the binary CAs for the semantic portion of the associative memory².

3.2 The Winner-Take-All (WTA) Model: Stationary Bumps

This paper describes work on a linear WTA model, instead of a planar model or hyper-planar model. A line of neurons is connected with local excitatory synapses, and a surround of inhibitory synapses. There are many synaptic matrices that lead to persistent behaviour once the initial neurons are stimulated. What is needed is sufficient local synaptic excitatory strength to allow the neurons within the winning group to fire persistently. This needs to be balanced with sufficient inhibitory synaptic strength to prevent spread beyond the initial group.

One example of this is a 2-4 excitatory inhibitory network. The WTA model is implemented in a neural network with 100 spiking neurons. Each neuron has excitatory connections to the nearest two neurons on both sides ($d \leq 2$) and inhibitory connections to the next nearest 4 neurons on both sides ($3 \leq d \leq 6$). This is also called a stationary bump (Laing et al., 2001), or a bump attractor. For the sake of explanation, the WTA spiking neural network approximates the continuous value activation of a temperature scale in which every neuron represents a single degree from 0° to 99° C.

Table 1 describes the behaviour of this network as the inhibitory and excitatory synaptic weights vary in the range 0.01 - 0.1 μS (microsiemens) in steps of .01. Initially, three neurons are forced to spike, representing input from the environment. (Three neurons are chosen as it is typically thought that several neurons are needed to cause another to spike (Churchland and Sejnowski, 1999), so this is the minimum input needed to (ignite) start a CA persistently firing.) After the initial stimulation, each simulation is run for 1000ms. The value in the cells of the table is the number of neurons that are firing at the end of the simulation.

² Note that the model expressed in equation 1 about the exponential integrate-and-fire neuron is a particular case of the AdEx model by removing the adaptation current (w) (Gerstner and Brette, 2009).

	-0.01	-0.02	-0.03	-0.04	-0.05	-0.06	-0.07	-0.08	-0.09	-0.10
0.01	0	0	0	0	0	0	0	0	0	0
0.02	0	0	0	0	0	0	0	0	0	0
0.03	0	0	0	0	0	0	0	0	0	0
0.04	0	0	0	0	0	0	0	0	0	0
0.05	0	0	0	0	0	0	0	0	0	0
0.06	100	100	100	9	7	7	7	5	5	5
0.07	100	100	100	100	13	7	7	7	7	5
0.08	100	100	100	100	100	85	9	7	7	7
0.09	100	100	100	100	100	100	100	9	7	7
0.10	100	100	100	100	100	100	100	9	7	7

Table 1 Number of neurons firing at the end of 1000ms. of simulation when three adjacent neurons are forced to spike in a 2-4 network. The horizontal axis is the weight of each inhibitory synapse, and the vertical the weight of each excitatory synapse.

In the first rows of table 1, there is insufficient excitatory synaptic strength to enable the neurons to continue to activate each other. In the excitatory 0.06 row, there is enough spread of activation to enable the neurons to fire persistently. In the first columns (e.g. cell .06 -.02) there is insufficient inhibition to prevent the neural activation spreading, and all of the neurons fire. On the right however (e.g. cell .08 -.09), a small reverberating population fires throughout the simulation (in this case 7 neurons). Note that when more than three neurons are initially fired, lower excitatory weights (e.g. weights .05 -.03) cause persistent firing.

The dynamics of the network, with high excitation and inhibition, is interesting. Table 1 shows that a persistent stable firing occurs after three adjacent neurons are spiked. This persistent spiking is similar to the behaviour of a binary CA. However, when a larger range of neurons are initially spiked (as may be the case in an associative memory), there is further interesting behaviour. Table 2 shows how larger inputs cause streams, or multiple bumps, of neurons to fire.

Table 2 shows how many sets of adjacent neurons persistently fire. Zero means no neurons are persistently firing and D means that all of the neurons fire persistently. One means that all of the neurons that are firing persistently are adjacent to each other; they are a stream. This could be inferred from table 1. The bottom part of the table refers to 75 neurons being initially spiked (25°C - 99°C). As in the top portion, some excitatory inhibitory weight pairs lead to no persistence, and some lead to all of the neurons firing.

In the bottom half of the table, those with one in the cell have more than 75 neurons persistently firing, but not all 100. Most of the table cells show two streams firing. These two streams are always on the edge. The edge neurons inhibit the interior neurons, as do the interior neurons themselves so that they do not fire a second time. The edge neurons have fewer incoming inhibitory connections, so they can persistently fire. After the initial burst, the interior neurons do not fire, the two streams do not influence each other, and they fire persistently as if they were ignited by two individual sets of three inputs. It is also interesting to note that several of these cells have more than two streams; this table shows four, six and seven streams. Again these are all quite thin, with approximately seven persistently firing neurons, and they have a relatively small number of non-firing neurons in between them.

	3 Input							
	-0.03	-0.04	-0.05	-0.06	-0.07	-0.08	-0.09	-0.10
0.05	0	0	0	0	0	0	0	0
0.06	D	1	1	1	1	1	1	1
0.07	D	D	1	1	1	1	1	1
0.08	D	D	D	D	1	1	1	1
0.09	D	D	D	D	D	1	1	1
0.10	D	D	D	D	D	1	1	1
	75 Input							
0.05	1	2	2	2	2	0	0	0
0.06	D	1	4	2	2	2	2	2
0.07	D	D	1	7	2	2	2	2
0.08	D	D	D	D	7	2	2	2
0.09	D	D	D	D	D	6	2	2
0.10	D	D	D	D	D	D	D	4

Table 2 Table of persistently firing streams of neurons for a 2-4 stable bump topology. The top refers to input of three adjacent neurons, and the bottom to an input of 75 contiguous neurons. The value in the cell represents the number of persistent streams; D (divergent) refers to all of the neurons persistently firing.

Note that it is possible to have local excitation with inhibition to all other neurons. When there are a small number of inputs, this performs largely the same as, for instance, a 2-4 stationary bump. However, with a larger number of inputs, say 75, the inhibition from the initial firing prevents all the neurons from firing. In table 2, the 75 input cells would all be 0. Let's call this topology with inhibition to all other neurons a 2-n topology. It is possible to set the weights so that a 2-n topology leads to persistent firing, but the width of the stream would be very large. For instance, a 2-n topology with 0.08 excitation and 0.005 inhibition has a persistent stream 68 neurons wide when 75 neurons initially spiked, and 56 neurons wide when 3 neurons are initially spiked.

3.3 Simple Hot Coffee Network

The basic simulation that has driven the development of this paper is an associative memory of beverages with a semantic value for temperature, and an underlying temperature in celsius. The gross topology of the spiking network that implements this memory simulation is shown in figure 1. There are two beverages, *Coffee* and *Coke*, and three temperature values *Hot*, *Warm* and *Cold*. These are all represented by binary CAs.

The input temperature from the environment is represented by a stable bump network of 100 neurons. The topology used in the remainder of the paper is a 2-4 topology with excitatory connections from a given neuron to the two adjacent neurons on either side, then four inhibitory neurons beyond. So, neuron 20 has excitatory connections to neurons 18, 19, 21 and 22, and inhibitory connections to 14-17 and 23-26. The weights are .08 excitatory and -.08 inhibitory. This network represents the temperature values between 0 and 99.

There are also neurons that represent the temperature of the individual beverages. These networks are used in the associations. The *Inhibition* neurons take input from the beverage temperature neurons, and in return inhibit them. This

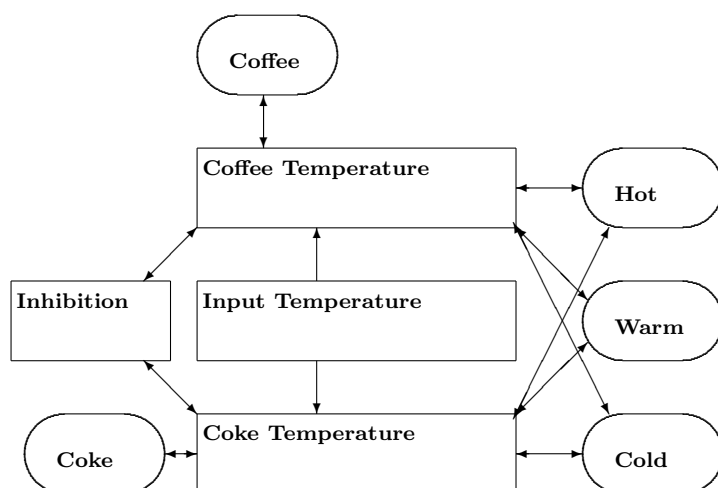


Fig. 1 Hot Coffee Gross Topology.

Beverage	Physical Temperature	Semantic Temperature
Coffee	0-24	Cold
Coffee	25-69	Warm
Coffee	70-99	Hot
Coke	0-9	Cold
Coke	10-24	Warm
Coke	25-99	Hot

Table 3 Beverage to Physical and Semantic Associations

prevents the spread of activation from one beverage to another. For instance, if the input temperature is 75-77, and *Coke* is queried, the associated *Coke Temperature*, will come on, igniting *Hot*; now that the *Input Temperature* and *Hot* are firing, they would ignite *Coffee* without the inhibition.

The arrows represent several synapses from a given set of neurons to another set to support an association. For example, each *Input Temperature* neuron excites its associated *Coffee Temperature* neuron. As these synapses are meant to associate CAs, in themselves, they are insufficient to cause neurons in another CA to fire.

Synapses inside simple CAs are not represented in the figure. Being binary CAs, the neurons in the semantic CAs have internal synapses as does the bump attractor *Input Temperature* net.

The individual CAs can be ignited by external stimulation. When this happens they all persist, and do not cause any other CAs to fire. There are the five semantic CAs and the *Input Temperature* CAs.

The basic temperature associations that were used are described in table 3. The temperatures and associated labels seem about right to the authors, and include a variety of ranges, which is good for expository purposes.

The basic idea is a two of three (2/3) associative memory. That is, if two of the concepts are firing, they should ignite the third. This largely works on the 18 basic

inputs. For instance, an input temperature of 36° and *Coffee* causes the *Warm* CA to ignite and fire persistently. In this case, the *Input Temperature* and *Coffee* CA cause the associated *Coffee Temperature* neurons to ignite, which in turn leads to the ignition of warm.

There are two types of atypical case to mention. The first case is the activation of a low input temperature ($< 10^\circ$) along with *Cold* or a high input temperature ($\geq 70^\circ$) with *Hot*. In this case, both beverages are activated, which is of course the correct result.

The second is the activation of two semantic CAs, for example, *Hot* and *Coffee*. Here the full range of beverage temperature neurons fire, but they fire at a low rate. In this case, the *Coffee Temperature* neurons from 70 - 99 fire. While this topology, using weakly connected beverage temperature neurons, in a sense solves the problem, a better result might be the prototypical beverage temperature neurons firing persistently. For instance, for *Hot* and *Coffee*, the neurons around 85° might fire.

An obvious modification is to replace the internally unconnected beverage temperature nets with bump attractor nets. However, a straight forward switch elicits a flaw. If for example *Hot Coffee* is stimulated, all 30 neurons (70 - 99°) fire, but the stable state that the *Coffee Temperature* network settles into is two streams of neurons, one from 70° and one to 99° , fire persistently with those in between silent. This is similar to the two streams of table 2.

3.4 Full Hot Coffee Network

While it seems reasonable for the full range of temperature neurons to fire due to direct semantic information, it is somewhat inconsistent with the firing behaviour from direct temperature input. Perhaps a better result would be to have a prototypical or average temperatures fire persistently. So, in the case of *Hot Coffee* semantic input, the neurons that represent coffee at 81° to 88° might fire persistently.

It is relatively simple to modify the simple hot coffee topology to get the full range of neurons to fire initially, and then to get persistent activity. However, in large ranges, there are two streams of activity at the ends of the range. For example, in the *Hot Coffee* case, neurons 70° to 74° and 95° to 99° fire persistently. This raises the issue of attractor repeller dynamics of stable bump topologies, which the authors have not seen discussed in the literature. Note that having all neurons inhibit all other neurons does not resolve this problem, as it results in very wide streams of neurons. This may be a property of the biologically unrealistic neuron model used with both excitatory and inhibitory connections (Eccles, 1986). A more sophisticated topology using different types of inhibitory neurons (Hirsch and Gilbert, 1991) is another possible avenue of exploration. The attractor repeller dynamics can be broken into two related issues.

Firstly, can the two streams be brought together? A mechanism, implemented with a separate population, can force the two streams toward each other (see section 4.1.2). The second population is activated by the first in a manner that only the neurons that correspond with the two firing streams fire. These have inhibitory connections to the original nets that increase with distance, so that the firing streams are moved towards each other. However, when the streams approach

each other, there is an enormous amount of inhibition sent to the neurons between the streams.

This is the second issue, when close, the streams repel each other. This can be overcome with a second inhibitory mechanism that is similar to a larger width stable bump (see section 4.1.1).

When the two mechanisms are combined, the initial wide range of temperatures still breaks into two streams. However, the two streams are then forced together, and then merged into one stream.

Finally, the new merge topology is used in the full hot coffee topology. These results are described in section 4.2.

4 Results

In this section, a topology for replacing two adjacent streams of firing neurons with one stream in between is described. Results on ongoing spiking and energy behaviour are shown, pointing to the repeller dynamics and how they are overcome. A separate modification to the initial topology that merges two distant streams is then described. Again, the ongoing spiking and energy behaviour are shown. Finally, the two mechanisms are combined (see section 4.1.3) and put into a modified hot coffee topology (see section 4.2) to complete the associative memory.

4.1 Bump Attractor Repeller Dynamics

Merging the two streams that emerge from presenting a broad input to the stable bump consists of two subproblems: the first is overcoming the repeller effect around the two streams (section 4.1.1). This has been solved using a rather extreme topology with quite heavy excitatory and inhibitory weights. The second and easier problem, at least for these simulations, is moving distant streams toward each other (section 4.1.2). Finally, once these two problems are addressed, they need to be combined into a single network (section 4.1.3).

4.1.1 Overcoming the Repeller

It is important to show that the stream in the bump attractor is both an attractor and a repeller. A rastergram is insufficient to show this, but figure 2 includes the energy and rastergram of a single stream, and of two streams that are near to each other.

With a single stream, note that the neurons in centre of the stream spike more frequently than those on the outside (figure 2A). However, the rastergram does not show the effect on non-spiking neurons. The energy diagrams (figures 2B and 2D) do. In figure 2 it can be seen that the neurons adjacent to the stream are excited but not to the firing threshold, and those beyond are inhibited. Those in the stream are, of course, excited; their voltage changes from the reset voltage $-70mV$ toward the firing threshold $-48mV$. Those further away are unaffected and remain at the base level of activation, $-65mV$ as are all neurons before the initial input.

Figure 2A and 2B represent three neurons being sent initial spikes from outside the system to ignite the stable bump, as in the associative memory simulations. The size of this bump, the number of neurons that fire persistently remains largely the same as more neurons are initially spiked (not shown) until 14 neurons are sent initial spikes. When 14 neurons are initially spiked, two streams are generated. This behaviour is represented by figure 2C and 2D.

Note how the initial neurons are not the only ones in the streams. The streams have repelled each other and moved to include new neurons not initially stimulated. This behaviour is amenable to study by thermodynamics or statistical mechanics.

The inhibition in the neurons between the streams is really quite high even when compared with the inhibition on the neurons on the outside of the stream. So, what is needed is a burst of energy into those central neurons with inhibition to the outside, but only in the case where the two streams are quite close together.

So, an extra set of neurons is used that only fires when two nearby streams are firing. There are the same number of neurons in the new *Overcome* population as in the original stable bump population, and they are aligned. The neurons in the *Overcome* population get excitatory input from the corresponding neurons below in a small window. Where the two windows overlap, the neurons fire. If only one stream is firing in the bump, or if the streams are quite distant, no neurons in the *Overcome* net fire.

When the *Overcome* neurons fire they send excitation to the *Bump* neurons directly below, and inhibition more distantly. In this case, the *Overcome* neurons excite in a window of three about themselves, and inhibit the next eight. This is quite similar to a 3-8 bump topology.

The result of this is that the interior neurons fire, and then remain persistently firing. Figure 3 shows this. The bottom figures (C and D) are the raster and energy plots of the bump attractor. Note that the initial firing behaviour leads to a split into two streams in figure 3C. Energy slowly builds in the appropriate neurons in the *Overcome* network 3B, causing a single set of neural firing that shifts the behaviour in the stable bump attractor.

4.1.2 Merging Streams

When the temperature range is quite large, the two initial streams do not influence each other. For example, in the simple hot coffee associative memory, *Hot Coffee* has a temperature range of 30° , and the streams are from 70° to 74° and 95° to 99° . (Note that there are edge effects at 0° and 99° , and behaviour there is different than in the middle of the bump attractor.) As there is no theoretical limit to the range of temperatures for a specific semantic category, the 3-8 bump attractor approach, or indeed any *X-Y* bump attractor approach will not work. In this case, the largest range is 75° (*Hot Coke*). This can be solved with an inhibitory topology increasing with distance, activated by the two firing streams.

As in the overcome case, there is an extra population of neurons of the same size, which is called the *Merge* net. It is important that these neurons do not fire unless there are two streams firing, so each of the *Bump* neurons has a small excitatory connection to each of the merge neurons. There are also direct one to one excitatory connections from the *Bump* net to the *Merge* net so that the neurons associated with the firing streams fire; so, these nets are also aligned.

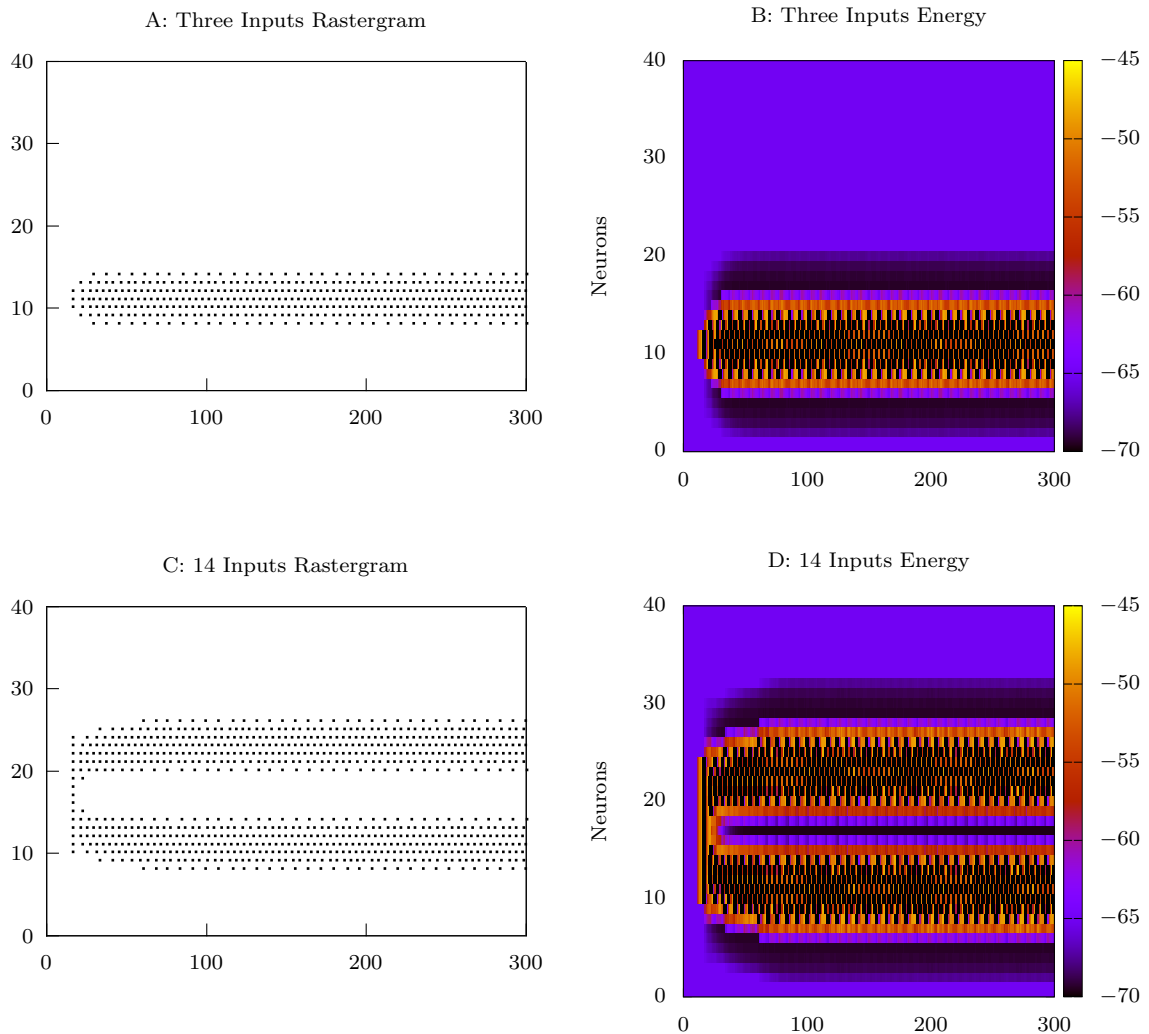


Fig. 2 A rastergram of a single stream (A), energy signature of that stream (B), rastergram of two adjacent streams (C), and the energy signature of adjacent streams(D).

Inhibitory synapses from the *Merge* neurons to the *Bump* neurons are distance biased with more distant neurons being more inhibited. Thus the outside neuron of the neurons in the opposing stream is more inhibited than the inside neuron, and this, metaphorically, pushes the streams together. There are no synapses to nearby neurons.

There is a difficulty that as the distance increases, the inhibitory strength becomes too large and the streams stop each other. The inhibition described in

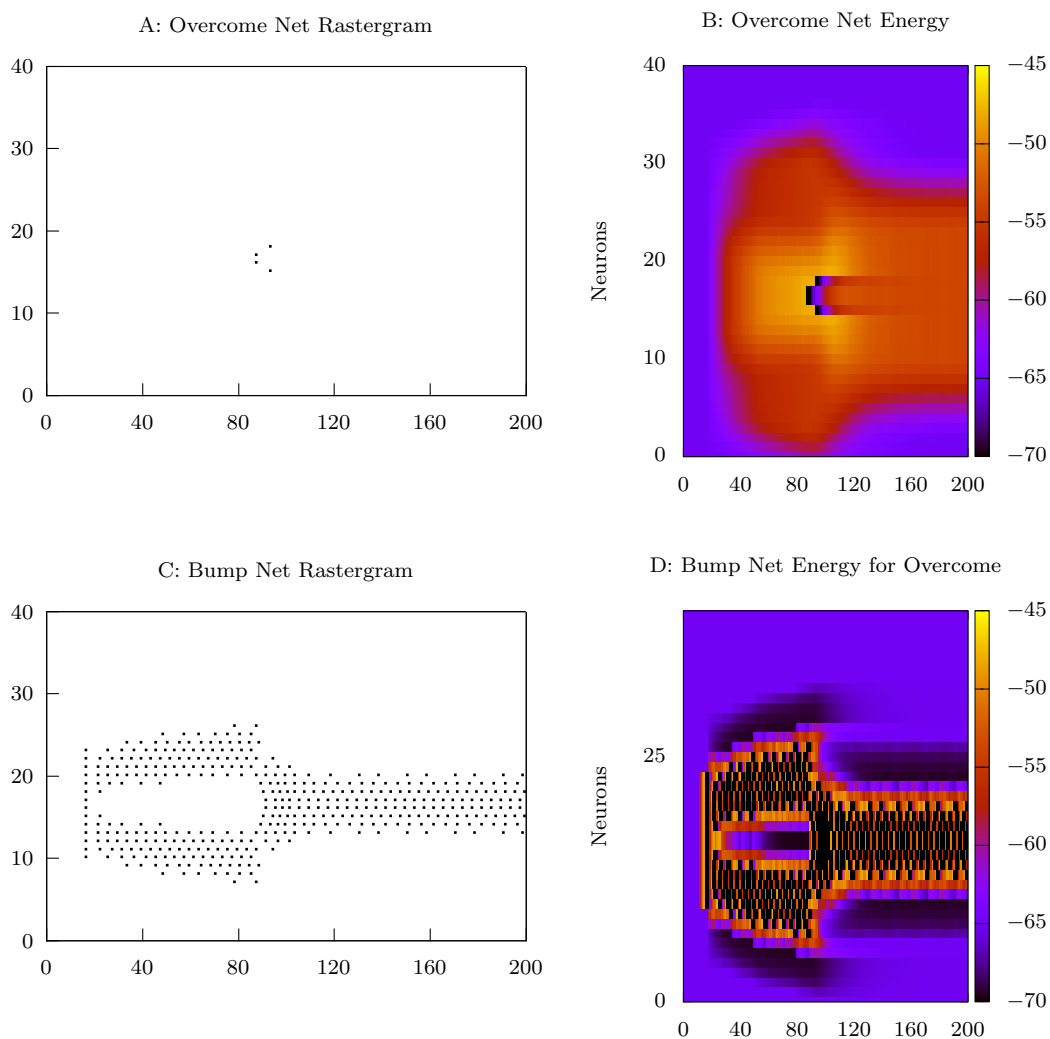


Fig. 3 Results from the Overcome topology initially spiking neurons 10 through 23. A rastergram of the inhibitory network for the Overcome simulation(A), energy signature of the inhibitory network for Overcome (B), rastergram of the bump network for Overcome (C), and energy signature of the bump network for Overcome(D).

this paper, increases exponentially with distance. It is sufficient to cope with 75° of difference, but will not work with much larger differences.

Figure 4 shows the behaviour of this system. Figure 4A shows the spiking behaviour of the inhibitory *Merge* net. Its firing is sparser than the bump attractor 4C, but follows it. Note that once the inside edge of the stream stops being inhibited, because it is near enough to have no more inhibitory synapses, it quickly moves because the inside is not inhibited, but the outside is. Also note that the

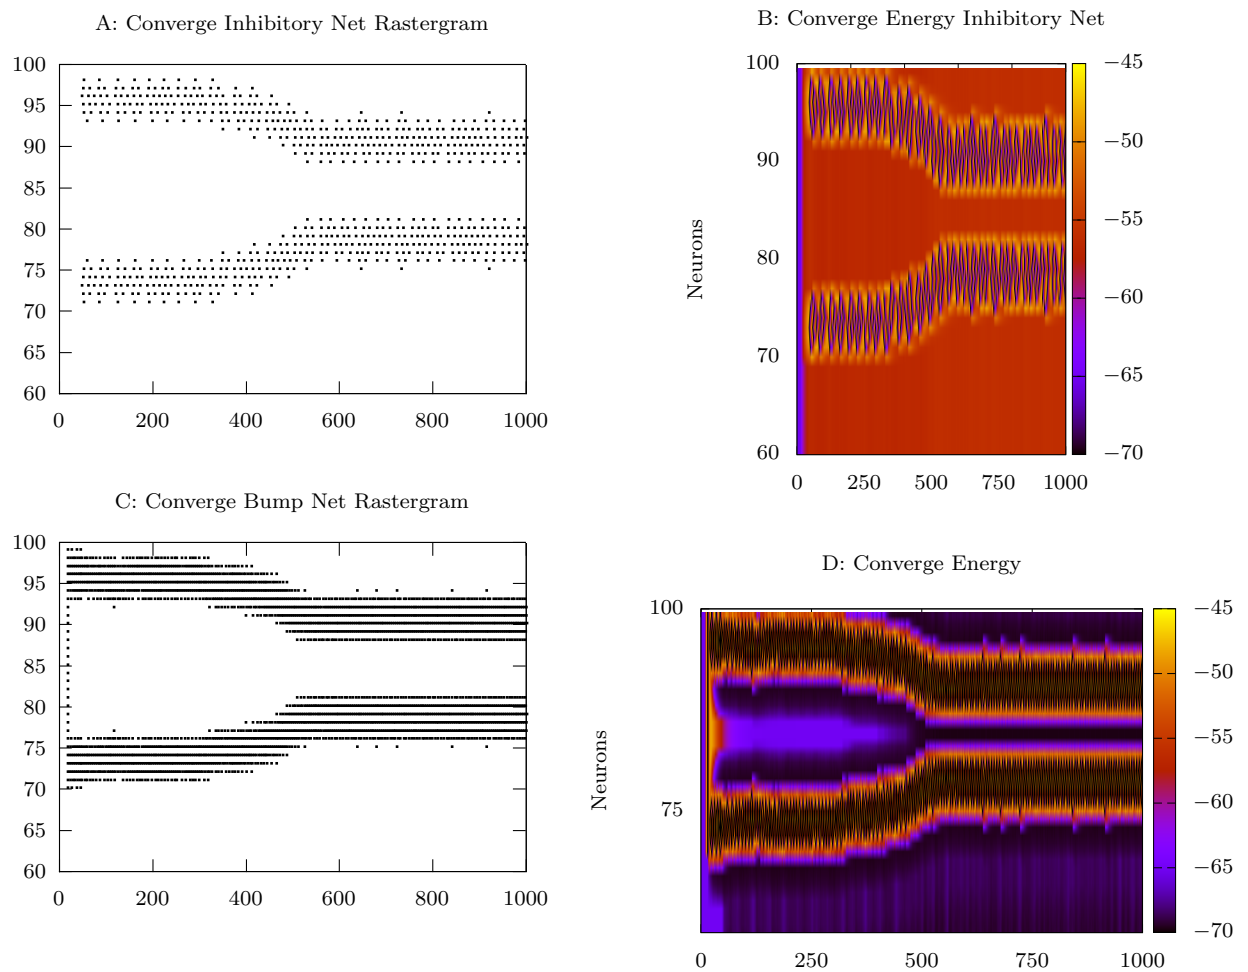


Fig. 4 Results from the Merging topology initially spiking neurons 70 through 99. A rastergram of spiking behaviour for the Inhibitory Network (A), energy signature of the Inhibitory Network (B), rastergram of spiking behaviour for the Bump Attractor Network (C), and energy signature of the Bump Attractor (D).

energy in the *Merge* net is high throughout the run after initial stimulation in figure 4B. This is below firing threshold, but is due to the all to all synapses.

This then moves two streams toward each other. It does not however push them together as they repel each other. Fortunately, this extra merge topology is compatible with the overcome topology.

4.1.3 Combining the Merging and Overcoming

A complete topology to move two streams into one combines the initial bump attractor, the *Merge* net, to bring distant streams close to each other, and the *Overcome* net, to bring the two adjacent, streams into one. The behaviour is that an initial range of neurons fire, based on a range of temperatures from the associated semantic values. If the range is long enough, two streams of neurons on either end of the range fire persistently. The *Merge* net is activated, and causes the two streams to move together. Once together, the merge net cannot overcome the inhibition from the streams repelling each other, but the *Overcome* net fires, causing the two streams to move into one stable state. As both the *Overcome* and *Merge* nets require two streams in the bump attractor to fire, they do not fire, and the attractor is stable.

This is shown in figure 5. Do note how there is no spiking in either the *Merge* or *Overcome* net once the streams are merged. Figure 5D shows the energy of the *Bump* net. Initially, it follows the change of energy from the merge mechanism; once the streams are adjacent, the repelling force of the two streams is overcome by the *Overcome* net, and the resulting single stream is stable.

4.2 Full Hot Coffee Network Results

The full hot coffee network is still described by figure 1, but the *Coffee Temperature* and *Coke Temperature* boxes are now three populations each, a stable bump attractor, a *Merge* net, and an *Overcome* net. Note that the dynamics of these three nets in isolation differ from those in the full topology because of the *Inhibition* network in the full topology; the neurons in the bump attractors fire at a lower rate. Extra excitatory and inhibitory strength are needed in the *Overcome* network as the stable bump is firing at a lower rate due to the effect of the *Inhibition* network.

Now, as in the simple hot coffee topology from section 3.3, the basic one semantic feature and one temperature input work properly, and quickly. For instance, temperature input of 85° to 87° and the semantic value *Coffee*, turns on the semantic value *Hot*. The ambiguous inputs (the semantic value *Cold* with a temperature below 10° , and the semantic value *Hot* with a temperature above 70°) turn on both semantic beverages.

The additional merging topology now cause the double semantic input queries to generate the appropriate temperature outputs. Each of the six pairs (e.g. *Cold Coke*, or *Warm Coffee*) produce a persistently firing output. Unlike the simple topology, this output is self sustaining. That is, it is persistent in its own right. Table 4 shows the association temperature range and the output results. It also shows the time to converge, noting how wider ranges take significantly longer to converge. However, small ranges (like *Cold Coke*) converge almost immediately even when they would break into two streams (*Warm Coke*).

4.3 Associative Memory with Learning

undone using a Hebbian calculation, the network is presented triplets of inputs (semantic beverage, semantic temperature, actual temperature), and the firing

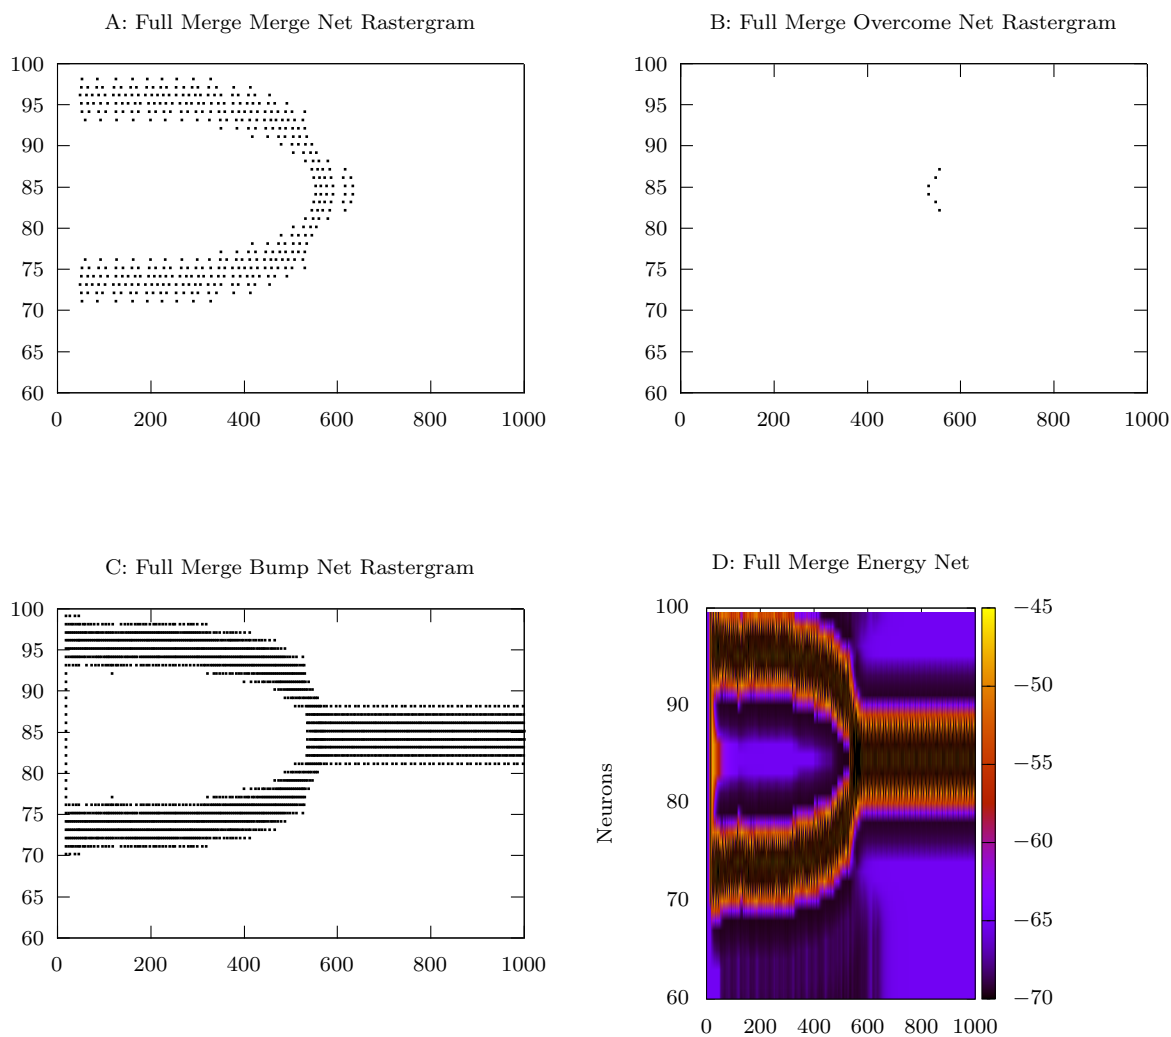


Fig. 5 Results from the full merging topology, initially spiking neurons 70 through 99. All neurons below 60 behave identically. A rastergram of spiking behaviour for the Converge Network (A), rastergram of spiking behaviour for the Overcome Network (B), rastergram of spiking behaviour for the Bump Attractor Network (C), and the energy signature of the Bump Attractor (B).

Semantic Pair	Temperature Range	Final Value	Time to Converge
Cold Coffee	0-24	10-14	1080 ms
Warm Coffee	25-69	45-49	7040 ms
Hot Coffee	70-99	82-87	463 ms
Cold Coke	0-9	2-7	90 ms (one stream)
Warm Coke	10-24	14-19	100 ms
Hot Coke	25-99	60-64	13920 ms

Table 4 Semantic Pair Input and Temperature Output

behaviour is stored. From this behaviour, weights are calculated using a compensatory Hebbian learning rule.

4.3.1 One of Each Category

The first training mechanism was to present the system with one of each category triplet. The triplets were: *Hot Coffee* 80° to 82°; *Warm Coffee* 40° to 42°; *Cold Coffee* 15° to 17°; *Hot Coke* 40° to 42°; *Warm Coke* 15° to 17°; and *Cold Coke* 4° to 6°.

These were each presented for 300 ms, and firing was stopped via inhibition between epochs. Firing behaviour was recorded and a compensatory Hebbian learning rule was applied. Weights from the semantic temperatures to the beverage bump attractors was calculated using a post-compensatory rule, and weights from the beverage bump attractors to the semantic temperatures was calculated using a pre-compensatory rule. The total target synaptic weight was 0.03 and 0.1 respectively.

The system was then tested with all 18 pairs (e.g. *Cold* and *Coke*, *Coke* and 4° to 6°, and *Cold* and 4° to 6°). In each case the correct remaining third member was retrieved, and no spurious elements were retrieved. Unlike the original test from section 3.3, a low temperature and *Cold* only retrieves one beverage, because they are associated with different temperature inputs. Indeed, many temperatures not presented are not associated at all.

4.3.2 The Full Range of Inputs

Next the full range of temperatures were presented with their appropriate semantic temperatures and for both beverages. This included roughly 200 runs of the system, with intervening inhibition to stop firing between epochs.

This also leads to a system with correct results. As in the simulations in the prior section, all the base double inputs lead to correct results. The cold and hot temperatures with semantic *Cold* and *Hot* lead to both beverages being retrieved.

undone
 full range of temperature inputs (hot and warm coffee)
 cold and temperature odd
 hot coke result

5 Discussion

This paper has focussed on a particular stable bump attractor, and extending an associative memory topology around it to account for associating two binary concepts and one continuous concept; the particular attractor was a 2-4 attractor, using a particular leaky integrate and fire model (Gerstner et al., 2014) simulated at 1ms time steps. This basic attractor should be readily extendable to different local connectivity, different neural models and different time steps. Moreover, the overall associative memory will be generalisable. For instance, during development of the system while trying to resolve the merge problem, a 1-10 attractor was developed; unlike the six neuron streams that emerged from the 2-4 topology, the

1-10 topology had 11 neuron streams. In general, the larger the basin, the larger the stream.

This beverage association directly extends to more beverages. One could add tea, hot chocolate, and orange juice by simply adding three sets of temperature attractor nets, the three semantic concepts, and the relationships. Perhaps more interesting and neuropsychologically plausible is the use of hierarchy in the concept structure. The authors earlier work on associative memory (Huyck and Ji, 2018) made use of hierarchical relations. In this case, there could be higher level categories, like fruit drinks, that had default values, that could later be refined. It should also be noted that association in these simulations have been done by synaptic connectivity; an independent concept is linked to two others by excitatory synapses; in the brain, particularly salient associations can also be CAs.

Similarly, other continuously valued concepts like time, weight and height could be used. A long time to hold your breath is not a long time to wait for your friend. Another reasonable thing is to combine two continuous attractors with binary semantic categories. For instance, one might combine weight and height to determine when someone is skinny or chubby. Interestingly, this might also apply in vowel recognition in the auditory cortex (Peterson and Barney, 1952) with the primary and secondary formants producing the vowel.

The authors prior work with stable bump attractors (Nadh and Huyck, 2010) shows that they can readily be extended from the one dimensional attractors discussed above to two dimensions, and they can be further extended to higher dimensions. The bump attractors here represent a large number of simple CAs. These CAs can act as short term memories when firing, but CAs should do some sort of calculation, and once stable these do no calculation (Tetzlaff et al., 2015). However, the composite associations are also CAs, so *Hot Coffee* and 82° is a CA composed of three other CAs. When the full hot coffee network is presented with the semantic terms *Hot* and *Coffee*, the merge topology does a calculation to convert the broad temperature range to a single coffee temperature CA. That is a relatively complex calculation, though once in that state the CA does no further calculation.

Therefore, one of the next directions regards using 2D or n-D bump attractors. This could be done analytically, or could follow a particular anatomic topology. For example, there is evidence that biological head direction cells are bump attractor CAs since their activity does not stop when the light is turned off and the bump is stable in the absence of input (Gerstner et al., 2014). The bump, in this biological case, is a localized blob of activity emerging from an initial perturbation of its network by a stimuli or memory recall operation.

The model could be extended from NEST with these particular neurons to other platforms and other neural models. In particular, neuromorphic simulations and emulations (e.g., SpiNNaker (Furber et al., 2013) or other neuromorphic hardware) could allow for much broader use of these persistent bump attractors. Runs on SpiNNaker of the bump attractor largely duplicate the results reported above in NEST.

Another topic to be investigated in more detail is the role of inhibition. Taking a macroscopic point of view, for example considering results from neuroimaging (Biswal et al., 2010), there are functional systemic networks that persist in their activation given a balance between positive and negative relations among different brain regions. For example, during the resting state (that is when subjects do

nothing with eyes open or closed), activation moves from one area to another when measured by fMRI (Van Den Heuvel and Pol, 2010). Each region of interest has CAs or portions of extended CAs, and the authors think that inhibition is used to move from activation in one area to another. This inhibition is similar to that in a bump attractor, and inhibitory effects probably occur at a range of scales.

Learning can be explored. Hebbian learning can be used to learn associations between concepts. When three individual CAs are stimulated, Hebbian learning can be used to associate them. In the case of bump attractors, some generalisation can be used reducing the number of instances that need to be presented; if the triple *Hot*, *Coffee* and $82-85^\circ$ are presented, the adjoining temperature neurons will also gain some association, so that $86-88^\circ$ does not need to be presented.

The literature about bump attractors and WTA functionality also investigates the thermodynamics of the network (see for example (Tkačik et al., 2015; Hahn et al., 2017; Pena et al., 2018)). Using the framework of dynamical system theory and related concepts (Meiss, 2007), this work could be extended by focusing on energy properties of attractors and repellers, stability and instability of the bumps, dynamics of pattern formation, and forking behaviour as bifurcation phenomena.

A note should be made about the terms winner take all networks and bump attractors. It seems the literature often equates the two, but as commonly used, winner take all refers to a single winner amongst multiple competitors. Take for instance the winner in a British Parliamentary election; there may be many competitors, but the one with the most votes wins. This is not what is being modelled in our bump attractor. If several neurons in different areas were given different stimulation, the one with the largest would not typically win.

This paper has described a spiking associative memory with associations between two binary CAs and a bump attractor. Activation of two of the three associated concepts is sufficient to ignite the third. When a large range is presented to the bump attractor, as in the case when the temperature range associated with two semantic values is large, extra neural topology forces the range into its middle, with those neurons firing in a self-sustaining persistent manner. This example can be readily duplicated to account for other similar associative memories.

The simulations in this paper have centred on stable bumps to represent large valued properties, beverage temperatures in particular. These have been included as components in associative memories, leading to an, as far the authors are aware, novel problem of large valued inputs to the bump. While bump attractors are a common and long standing model, these large valued inputs force a reduction in the number of inhibitory connections so that the bump can have a small number of persistently firing neurons. This has the added benefit of using fewer synapses $O(n)$ instead of $O(n^2)$ for well connected inhibitory nets. Large numbers of inputs do lead to streams, and our topology for merging the streams does have $O(n^2)$ synapses. This linear bump is almost the same as a ring attractor, with the ends included, which does seem more likely to emerge biologically for linear phenomena. We reiterate that these models are not good neuron for neuron models of the hot coffee representation. However, by using the combination of simple bump attractors for the linear phenomenon with binary cell assemblies for semantic primitives, the new issue of broad input to the bump has been raised. It is hoped that this will inform future work that moves from these relatively simple neural models and topologies to more biologically and psychologically informative models.

Developing simulations of simple neural circuits furthers understanding of their behaviour, as is the case with persistent bump attractors. These can be combined with other simple circuits to further understanding of more complex behaviour, such as associative memory. Progress in this manner will hopefully lead to a deeper understanding of more complex circuits, such as CAs, and eventually to the full brain.

Conflict of Interest Statement

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Author Contributions

CH and AV have both participated in developing ideas, simulations, and the paper itself. Both authors are accountable for the content of the work.

Funding

This work has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 720270 (the Human Brain Project) and the UK EPSRC grants EP/D059720 and EP/P00542X/1.

References

- Biswal BB, Mennes M, Zuo XN, Gohel S, Kelly C, Smith SM, Beckmann CF, Adelman JS, Buckner RL, Colcombe S, et al. (2010) Toward discovery science of human brain function. *Proceedings of the National Academy of Sciences* 107(10):4734–4739
- Brachman R, Schmolze J (1989) An overview of the kl-one knowledge representation system. In: *Readings in artificial intelligence and databases*, Morgan Kaufmann, pp 207–230
- Butz M, Woergoetter F, van Ooyen A (2009) Activity-dependent structural plasticity. *Brain research reviews* 60:2:287–305
- Buzsaki G (2010) Neural syntax: cell assemblies, synapsembles, and readers. *Neuron* 68:3:362–385
- Chen Y (2017) Mechanisms of winner-take-all and group selection in neuronal spiking networks. *Frontiers in computational neuroscience* 11:20
- Churchland P, Sejnowski T (1999) *The Computational Brain*. MIT Press
- Collins A, Loftus E (1975) A spreading-activation theory of semantic processing. *Psychological Review* 82:6:407–428
- Davison A, Yger P, Kremkow J, Perrinet L, Muller E (2007) Pynn: towards a universal neural simulator api in python. *BMC neuroscience* 8(S2):P2

- De Paola V, Holtmaat A, Knott G, Song S, Wilbrecht L, Caroni P, Svoboda K (2006) Cell type-specific structural plasticity of axonal branches and boutons in the adult neocortex. *Neuron* 49:6:861–875
- Eccles J (1986) Chemical transmission and dale’s principle. *Progress in Brain Research* 68:3–13
- Edelman GM (1987) *Neural Darwinism: The theory of neuronal group selection*. Basic books
- Fourcaud-Trocmé N, Hansel D, Van Vreeswijk C, Brunel N (2003) How spike generation mechanisms determine the neuronal response to fluctuating inputs. *Journal of Neuroscience* 23(37):11628–11640
- Furber S, Lester D, Plana L, Garside J, Painkras E, Temple S, Brown A (2013) Overview of the spinnaker system architecture. *IEEE Transactions on Computers* 62(12):2454–2467
- Gerstner W, Brette R (2009) Adaptive exponential integrate-and-fire model. *Scholarpedia* 4(6):8427, DOI 10.4249/scholarpedia.8427, revision #90944
- Gerstner W, Kistler WM, Naud R, Paninski L (2014) *Neuronal dynamics: From single neurons to networks and models of cognition*. Cambridge University Press
- Gewaltig M, Diesmann M (2007) Nest (neural simulation tool). *Scholarpedia* 2(4):1430
- Ghio M, Vaghi MMS, Tettamanti M (2013) Fine-grained semantic categorization across the abstract and concrete domains. *PloS one* 8(6):e67090
- Goel P, Khan M, Howard S, Kim G, Kiragasi B, Kikuma K, Dickman D (2019) A screen for synaptic growth mutants reveals mechanisms that stabilize synaptic strength. *Journal of Neuroscience* 39:21:4051–4065
- Hahn G, Ponce-Alvarez A, Monier C, Benvenuti G, Kumar A, Chavane F, Deco G, Frégnac Y (2017) Spontaneous cortical activity is transiently poised close to criticality. *PLoS computational biology* 13(5):e1005543
- Handjaras G, Ricciardi E, Leo A, Lenci A, Cecchetti L, Cosottini M, Marotta G, Pietrini P (2016) How concepts are encoded in the human brain: a modality independent, category-based cortical organization of semantic knowledge. *Neuroimage* 135:232–242
- Harris K (2005) Neural signatures of cell assembly organization. *Nature Reviews Neuroscience* 6:399–407
- Hebb DO (1949) *The Organization of Behavior: A Neuropsychological Theory*. J. Wiley & Sons
- Hirsch J, Gilbert C (1991) Synaptic physiology of horizontal connections in the cat’s visual cortex. *Journal of Neuroscience* 11:6:1800–1809
- Huyck C, Ji Y (2018) Two simple neurocognitive associative memory models. In: *Sixteenth International Conference on Cognitive Modelling*, pp 31–36
- Huyck C, Mitchell I (2018) Cabots and other neural agents. *Frontiers in Neuro-robotics* 12:79
- Huyck C, Passmore P (2013) A review of cell assemblies. *Biological Cybernetics* 107:3:263–288
- Kohonen T (1982) Self-organized formation of topologically correct feature maps. *Biological cybernetics* 43(1):59–69
- Kreiser R, Cartiglia M, JMartel, Conradt J, Sandamirskaya Y (2018) A neuromorphic approach to path integration: a head-direction spiking neural network with vision-driven reset. In: *IEEE International Symposium on Circuits and Systems (ISCAS)*, pp 1–5

- Laing C, R, Chow, Carson C (2001) Stationary bumps in networks of spiking neurons. *Neural computation* 13(7):1473–1494
- Lansner A (2009) Associative memory models: from the cell-assembly theory to biophysically detailed cortex simulations. *Trends in Neuroscience* 32:3:178–186
- Martin A (2007) The representation of object concepts in the brain. *Annu Rev Psychol* 58:25–45
- Meiss J (2007) Dynamical systems. *Scholarpedia* 2(2):1629, DOI 10.4249/scholarpedia.1629, revision #137210
- Nadh K, Huyck C (2010) A pong playing agent modelled with massively overlapping cell assemblies. *Neurocomputing* 73:2928–2934
- Nieder A, Dehaene S (2009) Representation of number in the brain. *Annual review of neuroscience* 32:185–208
- Pena RF, Zaks MA, Roque AC (2018) Dynamics of spontaneous activity in random networks with multiple neuron subtypes and synaptic noise. *Journal of computational neuroscience* 45(1):1–28
- Peterson G, Barney H (1952) Control methods used in a study of the vowels. *Journal of the Acoustical Society of America* 24:2:175–184
- Quillian M (1967) Word concepts: A theory of simulation of some basic semantic capabilities. *Behavioral Science* 12:410–30
- Redish A, Elga A, , Touretzky D (1996) A coupled attractor model of the rodent head direction system. *Network: Computation in Neural Systems* 7(4):671–685
- Richardson M, Gerstner W (2003) Conductance versus current-based integrate-and-fire neurons: is there qualitatively new behaviour? *Lausanne lecture*
- Seeholzer A, Deger M, Gerstner W (2019) Stability of working memory in continuous attractor networks under the control of short-term plasticity. *PLoS computational biology* 15:4:e1006928
- Singer W, Engel A, Kreiter A, Munk M, Neuenschwander S, Roelfsema P (1997) Neuronal assemblies: necessity, signature and detectability. *Trends in Cognitive Sciences* 1:7:252–261
- Somers DC, Nelson SB, Sur M (1995) An emergent model of orientation selectivity in cat visual cortical simple cells. *Journal of Neuroscience* 15(8):5448–5465
- Stringer S, Rolls E, Trappenberg T, Araujo ID (2002) Self-organizing continuous attractor networks and path integration: two-dimensional models of place cells. *Network: Computation in Neural Systems* 13:4:429–446
- Tetzlaff C, Dasgupta S, Kulvicius T, Florentin W (2015) The use of hebbian cell assemblies for nonlinear computation. *Scientific reports* 5:12866
- Tkačik G, Mora T, Marre O, Amodei D, Palmer SE, Berry MJ, Bialek W (2015) Thermodynamics and signatures of criticality in a network of neurons. *Proceedings of the National Academy of Sciences* 112(37):11508–11513
- Van Den Heuvel MP, Pol HEH (2010) Exploring the brain network: a review on resting-state fmri functional connectivity. *European neuropsychopharmacology* 20(8):519–534
- Willshaw D, Buneman O, Longuet-Higgins H (1969) Non-holographic associative memory. *Nature* 222:960–962
- Wilson HR, Cowan JD (1973) A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Kybernetik* 13(2):55–80
- Wimmer K, Nykamp D, Constantinidis C, Compte A (2014) Bump attractor dynamics in prefrontal cortex explains behavioral precision in spatial working memory. *Nature neuroscience* 17(3):431

Yoon K, Buice M, Barry C, Hayman R, Burgess N, Fiete I (2013) Specific evidence of low-dimensional continuous attractor dynamics in grid cells. *Nature neuroscience* 16(8):1077