Preprints.org

Article

Introduction to the E-Sense Artificial Intelligence System

Kieran Greer *

Article

Introduction to the E-Sense Artificial Intelligence System †

Kieran Greer

Distributed Computing Systems, Belfast, UK; kgreer@distributedcomputingsystems.co.uk

† http://distributedcomputingsystems.co.uk, Version 1.4.

Abstract: This paper describes the E-Sense Artificial Intelligence system. It comprises of a memory model with 2 levels of information and then a more neural layer above that. The lower memory level stores source data in a Markov (n-gram) structure that is unweighted. Then a middle ontology level is created from a further 3 aggregating phases that may be deductive. Each phase re-structures from an ensemble to a tree, where the information transposition is from horizontal set-based sequences into more vertical, typed-based clusters. The base memory is essentially neutral, but bias can be added to any of the levels through associative networks. The success of the ontology typing is open to question, but results suggested related associations more than direct ones. The third level is more functional, where each function can represent a subset of the base data and learn how to transpose across it. The functional structures are shown to be quite orthogonal, or separate and are made from nodes with a progressive type of capability, including unordered to ordered. Comparisons with the columnar structure of the neural cortex can be made and the idea of ordinal learning, or just learning relative positions, is introduced. While this is still a work in progress, it offers a different architecture to the current frontier models and is probably one of the most biologically-inspired designs.

Keywords: brain model; memory model; neural model; cortex; statistical clustering

1. Introduction

Artificial Intelligence is now at the apex of Computer Science. With advancements in pattern recognition and learning [27,32,35,38,45] and recently in prediction [43,44,49,61], the systems can perform many specific tasks as well as humans. Improvements in computing power and automated learning (for example, [16,34]) have also contributed. If the final bastions of reasoning and understanding can be mastered, then AI systems may well challenge humans in a general sense. However, the proponents are quick to point out that the systems are still mostly statistical, even though a new property of emergence has been realised in the very large distributed systems (Large Language Models [44]) that is not statistically predictable. While the path to success seems clear, there are still some hurdles and the problems that autonomous vehicles [13] have could be an example of this. Researchers are always inventing new ways to do things and this paper offers a different architecture to the established theory that may be able to complement the existing systems.

The new system is called E-Sense (Electronic Sense, or Essence) and it is a whole-brain cognitive model. It comprises of a memory model with 2 levels of information and then a more neural layer above that. See Figure 2, Section 4.1, later. The lower memory level stores source data in a Markov (n-gram) [14] structure that is unweighted. In a spatial sense, this would still mean that similar patterns would be clustered together. Then a middle ontology level is created from a further 3 phases of aggregating source information. Each phase re-structures from an ensemble to a tree, where the information transposition is from horizontal set-based sequences into more vertical, typed-based clusters [18]. In this model, 'type' means converting the text from word clusters that make up a sentence to word clusters that are shared between sentences. It is probably a level below symbolic understanding and is based more on context or use. The potential of it is described in Section 4 and the Appendix A examples. The design and implementation of the model is an extension of the

author's earlier publications. For example, exactly how the information transposition for the ontology should be done is an open question, but the philosophy behind it is consistent with earlier work. The base memory is essentially neutral, where any weighted constraints or preferences should be stored in the calling module. This would allow different weight sets to be imposed on the same linked structures, for example, or bias could be added using associative networks (see Section 4.4). The third level is more functional, where each function can represent a subset of the base data and learn how to transpose across it. The functional structures are shown to be quite orthogonal, or separate and are made from nodes with a progressive type of capability, from unordered to ordered. This was a surprising result that is described further in Section 5. Comparisons with the columnar structure of the neural cortex can even be made.

E-Sense is more biologically inspired, which could add flexibility, thus allowing it to be more generic. It also appears to be more sensitive to the data, allowing it to learn from smaller amounts. Current systems rely on large amounts of data to build-up a reliable picture. They can also be sensitive to parameter values and not always return the same result. E-Sense is less random and requires less fine-tuning. However, the proposed system is nowhere near as accurate, and it is not known if adding more data will make it more accurate, or how it will manage conflicts in larger corpuses. New algorithms have resulted from the design that do not appear to be part of current systems. Thus, these types of algorithms are not found in the mostly statistical and formulaic solutions, but maybe require more algorithmic solutions instead. For example, self-organising systems can suffer from a problem of data order, relying heavily on the order of the input. This is the case with the Frequency Grid [23] used in E-Sense and is apparent in some neural networks, or even LLM's. Typing is also typically not considered, maybe requiring a neuro-symbolic [58] approach. Having an unweighted Markov model is also novel. The brain cannot be heavily weighted and so one solution may be to make the design more orthogonal. This architecture could even go back to the on/off firing of the original neuron designs. But ultimately, structure is required for useful content and so bias must be added at some point. This is only a first implementation of the model and any proof that it can improve the performance of current systems is not available. But it has some interesting biological comparisons that may be able to give different views of the data. Direct comparisons with the real human brain are made throughout the paper, where previous work by the author includes [18,19,24].

The rest of the paper is organised as follows: Section 2 briefly introduces the original model again and Section 3 gives some related work. Section 4 describes the new memory model and then Section 5 introduces the new upper neural level. Section 6 introduces the ordinal learning algorithm and Section 7 then gives some preliminary test results for the two main algorithms. Section 8 makes comparisons with purely biological concepts, while Section 9 gives some conclusions on the work.

2. The Original Cognitive Architecture

The original architecture [24], is shown in Figure 1. The design began with the idea that neural networks with their 'single function' are not flexible enough to design complex brain processes. Thus, local links, such as those produced by Stigmergy [17,11], were preferred. Then, to convert the local structures into higher-level knowledge required merging and aggregating them. This resulted in 3 levels of increasing complexity, as shown. The lower-level created structures from local links and tried to optimise that. The middle-level aggregated the lower-level structures, to give knowledge about best values, or if any actually exist. The upper-level aggregated single concepts into more complex ones, thus producing higher-level knowledge. There was also an idea that one complex concept would trigger another one, as part of thinking.
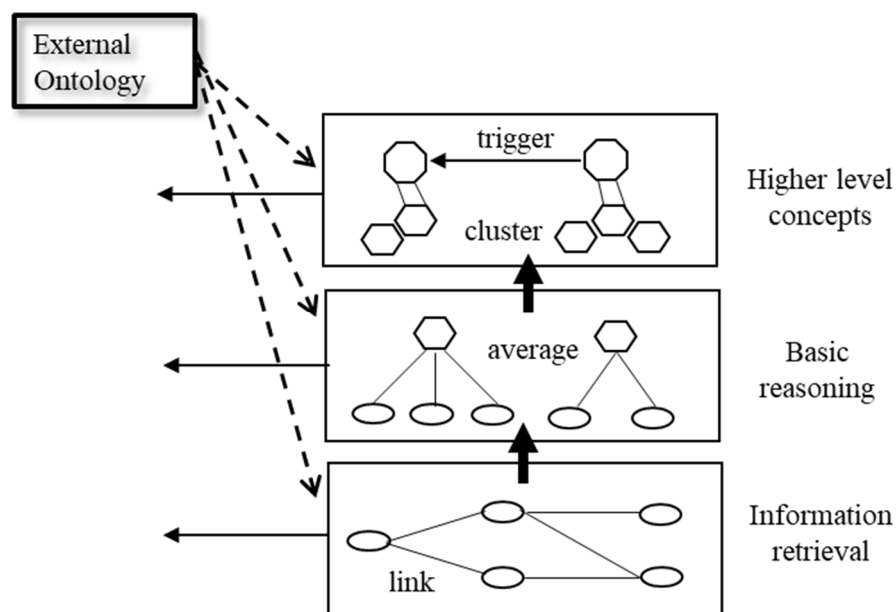
Figure 1. The original 3-Level Cognitive Model [24] with a related ontology. Stigmergic linking at the base passes information to an aggregating layer for best values, that passes information to a clustering layer for more complex concepts.

Section 4 describes that these levels now form the basis for the new memory and neural models. An external ontology [26] was also part of the original design and that is now integrated as the new middle level. The memory model uses a statistical clustering process, rather than semantics and rules, however. The author supposes that this effect is covered in modern AI programs by using Word Vector models [43], for example. Fully-connected neuron structures are central to some themes in the system, where this idea is quite common (for example, [1,29] and some of the author's earlier papers).

3. Related Work

There are a few notable AI systems that already produce human-like results, where most systems would claim to represent one or more parts of the human brain. In-line with the author's theory, the review paper [50] describes that cognition calls for a mechanistic explanation, where intelligence may depend on specific structural and functional features in the brain. It gives an overview on what types of neural network are used to model which parts of the brain. For example, auto-associative networks have been used to model the cortical regions [58], while feedforward networks have modelled memory or vision. No network type has modelled the whole brain however, probably because of their fixed structures. The paper [41] also describes modular networks for modelling the human brain. For this task, 'heavy-tailed' connectivity becomes important and several papers have discovered this phenomenon when mathematically modelling the neural connectivity (for example, [40,52]). With heavy-tailed connectivity, synaptic connectivity is concentrated among a select few pairs of neurons that are attractors. This results in a sparse network of strong connections dominating over other interactions. The paper [40] has shown that they can occur simply through a mixture of Hebbian and random dynamics, or the preferential attachment model [3]. The paper [52] studied how different rewiring functions would affect the resulting structure that was generated. They observed that random structural connectivity was reshaped by 'ordered' functional connectivity towards a modular topology, which also indicates synchronous firing patterns. One example was that rewiring peripheral nodes inside a module was homogeneous, with high synchrony and low-dimensional chaotic dynamics. On the other hand, central hub nodes connected with other modules, exhibited unsynchronized, high-dimensional stochastic dynamics. To reduce chaotic activity and energy therefore, it makes sense that between-module interaction would occur through a select number of key nodes only.

The paper [55] describes a model of the Hippocampus that uses auto-associative networks with generative learning. It describes how episodic memory is constructive, rather than the retrieval of a copy. But it needs the resource of semantic memory, from the neocortex, which is factual knowledge. They used a modern Hopfield network, where feature units activated by an event were bound together by a memory unit. The two layers here, binding features, is considered in Section 5, for example. The generative networks were implemented as variational autoencoders [34], which are autoencoders with special properties, so that the most compressed layer represents a set of latent variables. These variables can be thought of as hidden factors behind the observed data and can be used to re-create the data again. The paper [58] considered if an auto-associative network can accurately model the cortical region. It considered the different levels in the cortex [46,28] and the different types of connection and function in those levels. The conclusion was that an auto-associative network has sufficient capacity to be used as a memory model, but may require the addition of these other factors. Since then, the creation of modern Hopfield networks [36] has shown that this type of network can have sufficient capacity in general, but needs multi-dimensional functions. Gestalt theory has been mentioned before [23]. With Gestalt psychology, objects are seen independently of their separate pieces. The booklet [5] gives a formal description of the theory and a mathematical proof that links the psychology theories of memory span [42,6] and duality [51]. The relation to Gestalt theory is discussed in Section 8.1.

3.1. State-Of-The-Art in AI

Deep Neural Networks [32,35,38] probably kicked the current revolution off, but other notable successes would include Decision Trees [27], for example. Category Trees [21] might be an interesting alternative. Deep Learning [45] then combined Deep Neural Networks with Reinforcement Learning. It can automatically learn features from the data, which makes it well-suited for tasks like object classification and speech recognition. DeepMind (the company behind deep learning) introduced neural Turing machines (neural networks that can access external memory like a conventional Turing machine), resulting in a computer that loosely resembles short-term memory in the human brain. The model [45] used a convolutional neural network, which is organized similarly to the human visual cortex. The advantage of this kind of network is that the system can pick out particular features from the data, automatically. It is then able to comb through massive amounts of data and identify repeated patterns that can be used to create rules and processes. The general architecture means that DeepMind's algorithms have taught themselves to play Atari games and beat the best humans in Go or Chess. DeepMind has since moved on to tackling more and more real-world problems, such as unravelling the likely structures of Proteins. Then recently, Large Language Models [44], such as OpenAI's ChatGPT and Chat-GPT4 [49] have advanced the state-of-the-art again. Some argue that GPT4 already exhibits a level of Artificial General Intelligence, maybe because of the emergence property. These systems can make use of Word Vector models [43] and Transformers [61], to predict what comes next in a sequence, rather like an n-gram [4,14] for text. They can be trained on a large corpus of data but are then able to create answers for any type of question, even ones not known about. They can use the same process to manage images, mathematical equations and even computer code. The discovery of transformers allowed them to predict to a level not encountered before and together with deep learning, huge distributed models with billions of nodes can be built. But even with all the recent advances, some papers show that there can still be problems, even with benchmark datasets [10,33,48]. The ARC prize [2] has been designed, specifically to highlight where the current systems do not perform well, described as follows:

- *Symbolic interpretation*: AI struggles to assign semantic significance to symbols, instead focusing on shallow comparisons like symmetry checks.
- *Compositional reasoning*: AI falters when it needs to apply multiple interacting rules simultaneously.
- *Contextual rule application*: Systems fail to apply rules differently based on complex contexts, often fixating on surface-level patterns.

E-Sense is not yet at a level where it can address these challenges. New solutions would also want to make the models more economic and reduce their reliability on data.

3.2. Alternative Models

Other designs are described in [39], where one option, used in SPAUN [12], was to transform an ensemble mass into a vector-style structure, with weighted sets of features. SPAUN is one of the most realistic brain model designs and used spiking neurons, but context is still a problem. This is also clear in one of the original designs called SOAR [37]. That system adhered strictly to Newell and Simon's physical symbol system hypothesis [47], which states that symbolic processing is a necessary and sufficient condition for intelligent behaviour. SOAR exploited symbolic representations of knowledge (called chunks) and used pattern matching to select relevant knowledge elements. Basically, where a production matched the contents of declarative (working) memory the rule fired and then the content from the declarative memory was retrieved. SOAR suffers from problems of memory size and heterogeneity. There is also the problem that production rules are not general knowledge but are specific and so there is still not a sufficient understanding at the symbolic level. IBM's Watson [30] is also declarative, using NLP and relies on the cross-referencing of many heuristic results (hyperheuristics) to obtain intelligent results. Context is a key feature of the Watson system however. A recent paper [62] describes a process for recognising 'relation patterns' between objects. Humans acquire abstract concepts from limited knowledge, not the massive databases that current systems use. Their relational bottleneck principle suggests that by restricting information processing to focus only on relations, it will encourage abstract symbol-like mechanisms to emerge in neural networks and they suggest a neuro-symbolic [53] approach. They argue that an inductive process can be used to learn a relation like 'ABA', where A or B can then be anything. The symbolic and connectionist approaches can be reconciled to focus on relations between objects rather than the attributes of individual objects. They propose to use inner products, which naturally capture a notion of relations. A 'small changes' theory is part of this paper's model, described later in Section 5.2. Also to be noted for the idea of massive overlap is the 'Thousand Brains Theory' [28].

A good organisation ability may be an inherent property of humans, or even the animal kingdom and would be something that can be improved in the current systems. This is discussed further in Section 6. The paper [31] suggests a theoretical framework that would try to convert the fixed neural network architecture into one that can represent images in more abstract part-whole hierarchies. It is something that humans do, but neural networks cannot currently do. This is quoted in [15] that again proposes a vector model along with the image parts. The paper [56] used statistical mechanics to try to explain some of the mechanisms that occur in the biological brain. They then showed how their results suggest a thermodynamic limit to the neural activity, but have no definite explanation of why, and this limit suggests a boundary. They also noted that the brain is a nonequilibrium system and asked the question of how it then obtains equilibrium. Most of these papers consider the brain activity to be more entropic than local. The paper [9] is very mathematical, but it might give a solution to the problem of these looser constructs. It proposes to use sheaves and writes about unary and binary typing's. Rather than global, they argue that time can be constructed, like local events, when it might also be thought of as an ordering.

4. The Memory Model

The original cognitive model was based on a 3-level architecture of increasing complexity, which included an external ontology that would be available to all the levels. Ontologies [26] describe the relations between concepts in a very structured and formal way. They are themselves high-level structures and it is not clear how they could be built simply from statistical processes. The ontology of this model is therefore not at that level, but instead, it provides a looser clustering of concepts.

4.1. Memory Model Levels

Aligned with the cognitive model, the memory part is implemented in two lower levels, with some referencing in the upper neural level. It is thought that memory is stored in all parts of the human brain. The 3 memory levels are therefore as follows, also shown in Figure 2:

(1)   The lowest level is an n-gram structure that is sets of links only, between every source concept that has been stored. The links describe any possible routes through the source concept sequences, but are unweighted.

(2)   The middle level is an ontology that aggregates the source data through 3 phases and this converts it from set-based sequences into type-based clusters.

(3)   The upper level is a combination of the functional properties of the brain, with whatever input and resulting conversions they produce, being stored in the same memory substrate.
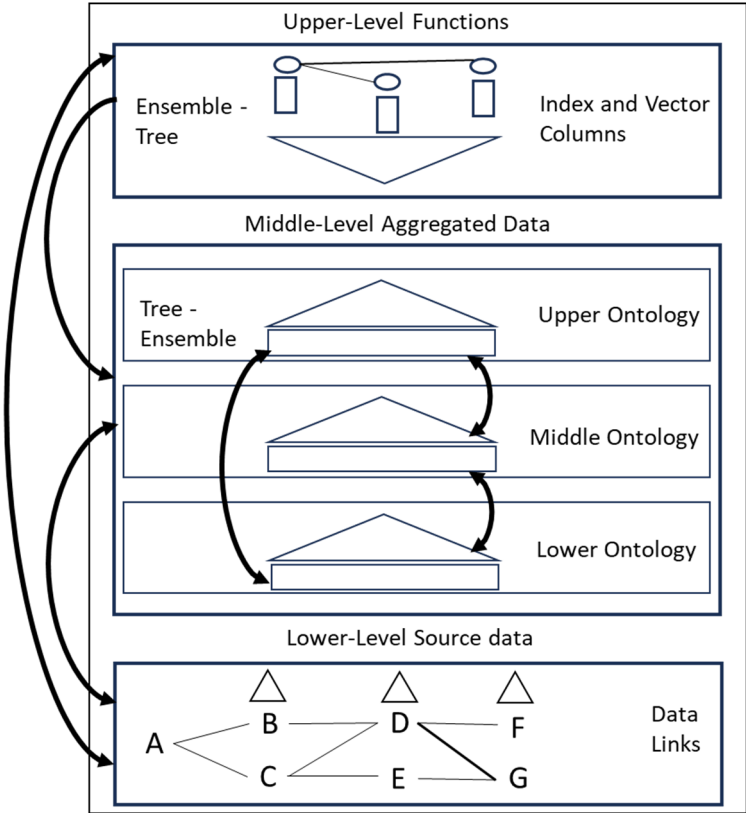

Figure 2. The new 3-Level Cognitive Model. A lower level, links source concepts and optimises storage. A middle level aggregates source sequences into ontology relations and an upper level provides discrete functions on the data.

It may be that the first 2 levels can be made from simpler structures, in the sense that it does not have to be functional. For example, the paper [18] describes that more recently, the perineuronal network [59] has received a lot of attention and may be a sort of simpler memory substrate. An earlier paper [19] described the types of knowledge transitions that may occur. If using a classification of experience-based or knowledge-based information, then experience-based information is dynamic and derived from the use of the system. Knowledge-based information is static and built from the experiences. That paper described that the transitions in the cognitive model may be:

1.   Experience to knowledge.

2.   Knowledge to knowledge.

3.   Knowledge to experience.

If looking at the new Figure 2 design, then from the top level to the bottom level we get these transitions again. The two outer-most layers would be the sensory input (lower level) or the cortical columns (upper level) and would be experience-based. Then between them are transitions into and

out of knowledge. At least 1 time-based layer needs to be added to this region, which will be considered in future work. The paper [50] does not note auto-associative networks as being whole-brain models, but the new cognitive model could be seen as an auto-associative one, where the middle knowledge-based level would represent compressed knowledge variables. The knowledge-to-knowledge transition would be between the middle and upper levels, but currently, that information flow is not implemented. It is also shown that knowledge is stored in tree-like structures, while experience is stored in ensemble-like structures. The upper-level 'ensemble to tree' is reversed in the middle level.

4.2. Lower Memory Level

In the model, the lowest level is an n-gram, where each node contains links to each possible next node. Thus, tracing through these links can return different sequences, but to provide some direction, possibly a 5-gram needs to be used. It is also possible to note start or end nodes in a sequence, to help with defining the sequences better. The database structure is appealing because it is very compact, but while it works well for most of the time, it may not always return every sequence, exactly as it was entered. Unlike neural representations however, this structure does not have to be weighted. A link between nodes in a region is noted only once, no matter how often it occurs during the input. The structure therefore only stores equal possibilities, where preferences or weighted choices would have to be transferred over to a calling module. To make it more accurate however, it can be combined with an associative network that can recognise specific sequences, where that would introduce a certain amount of bias. If the n-gram is sufficient, then the theory would state that a Markov process may be sufficient to describe a Gestalt process. This is discussed further in Section 8.

4.3. Middle Ontology Level

The middle ontology level uses transitions from ensemble to tree, where the final trees look more like single vector lists. While the ontology is still a low-level structure, the type conversion introduces a small amount of knowledge that a search process can make use of. The merging process is novel and produces types of relations. While the lower-level database can be described by an n-gram, the middle ontology level is more complicated. It makes use of the Frequency Grid [23] to generate clusters, but there would be more than 1 way to implement the aggregation process. The author has chosen a version that prunes the most nodes from the result, to save on time and resources. This means that a text document the size of a book may return at the top ontology tree, only a few words clustered together (see Appendix A), but as the search could move back down the structure to the lower levels, it will be able to discover most of the text from linking horizontally there. The transitions may be quite small however, where most changes take place during the early phases.

4.4. Unit of Work

A 'unit of work' structure has been suggested in earlier papers, for example [19] (Section 3), and also the related ensemble-hierarchy [23]. A unit of work is a group of neurons that are inter-linked and fire as a single unit. The unit then links with other units. Thus, a single neuron is converted into a more complex assembly. The intention of the ensemble-hierarchy structure was to make the signal it produces more distinct, where the hierarchy provides structure and would fire with the ensemble, giving it that extra definition, rather like an action potential. In the new model it can take different forms, depending on the functional level. In the middle ontology level, it is an ensemble-tree, for filtering purposes only. In the bottom memory level, the sequences can be combined with an auto-associative network, which would be able to recognise when a sequence represents something. In this design, it is interesting that the base ensemble would store the ordering for the recognised object, while the associative network only recognises it as a whole concept. With a myriad of possibilities, this could still be very useful and is probably what Transformers [61] over the LLMs do in modern systems. There are some indications that base level associative networks could work and the author

is confident that they will be part of the final architecture. A similar type of structure has also been suggested for the upper neural level, as part of another research. This considered contributions from Kolmogorov and Shannon. Shannon based his Information Theory [54] on Entropy and a Markov model. Kolmogorov Complexity theory ([7], chapter 7) states that the shortest sequence is the most likely, and therefore the best. This is also known as Occam's Razor. Combining these gave rise to a related design in [20], where in this case, the hierarchy would be Shannon loops, which represent repeating functional acts that extend from base Kolmogorov sequences, intended to get an entity from A to B in the most economical way.

5. The Neural Level

The neural level is the top level of the design and contains more cognitive or functional units. There are to be 2 different types of neural level in the final model – one is interested in cognitive processes that may be described by a Cognitive Process Language [22], while the other is more interested in logical processes (for example, [19] and the related papers). This paper deals only with the logical neural model, which comprises of functions that operate on the source data of the lower level. It does not follow-on from the middle ontology, but is separate from it, although in theory, it would query the ontology if needed. A typical definition of a function is something that maps a set of input values to a set of output values. It may do this by generating the output, or if the output already exists, then it is a selection process. If it can make use of existing output, then the function reduces to something more like a constraint, with a description like - a function in some cases, may be seen as something that reduces the possibilities for the next step. As the following sections describe, the neural level now looks quite a lot like the cortical regions in the human brain.

5.1. Function Identity

If there are lots of functions in the neural level, then they want to be recognised and made distinct. One option is to consider every node in the function and comparing this between functions would probably produce differences. Another option may be to consider only a set of base 'marker' nodes, when the rest of the function can be built on-top of these. In fact, these key nodes can become an index set that should be found first in the source data and then a larger set of vector values, representing the whole function, can be searched for. This is in fact what modern search engines do [8]. The index set could also help with maintenance. It could be checked for first, to see if some process returns the same index key. If not, then it is likely that something has changed in the base data, linked to by the function, where further processes would need to determine what to do.

5.2. Function Structure

A function is therefore based on an ensemble of these index nodes, each with a set of index values. Each node then links to a larger set of vector values that represent a feature in the function. Each index node matches with 1 or 2 n-gram sequences from the lower-level database and it also stores the relative position(s) of the sequence(s), so that the correct ordering can be re-constructed from values that may be arbitrarily ordered as input. The operation might again be relational, but again in an open sense, where different data can match with the template. Thus, the functions described here are primarily for recognising order and are not even fully logical. Each function part, in fact, resembles quite closely, the Symbolic Neural Network that was written about in [25]. The index values are quite orthogonal, or do not overlap very much. The features, as whole sets, are mostly unique as well, but some index nodes can share the same feature set. There could thus be closures at both the top and bottom of this structure. This type of structure was shown in [25] to be able to filter-out noise quite well, for example. Both the most commonly occurring terms and the least commonly occurring terms are stored in a feature for a sequence. The most common allow potential matches with the sequence to be found in a larger database. Then the least common allow this potential set to be filtered further. A group of the SNNs are stored in an 'Ordinator,' which is really

the whole function for a particular operation. Because there is a lot of overlap in the results returned from each SNN, this produces only small changes in the statistical result, but usually, any changes then need to be included. When building the structure, some weighted components might be used, but when using the structure afterwards, it is mostly an un-weighted process. There might be some frequency counts, but not much else. It is probably the case that the orthogonal nature of the structure reduces the need for weights.

5.3. Index Types

The process of creating the index nodes and related vectors seems to have produced 3 different types. While it was not the case every time, a clear pattern of 3 distinct types emerged. One type may have a longer list of index terms but no related feature set. The other 2 had both index terms and related feature sets, but differ as explained next. These types map quite closely to known neuron types, as follows:

- Unipolar Type: this has a list of index terms that is generally a bit longer and is unordered. It can be matched with any sequence in the input set, but to only 1 sequence.
- Bipolar Type: this has a list of index terms and a related feature set. The index terms should be matched to only 1 sequence and some of the feature values should also match with that sequence. This matching should be in order however, where the order in the feature should be repeated in the sequence. Then the rest of the feature values can match with any other sequence and in any order.
- Pyramidal Type: this has a list of index terms and a related feature set. The index terms however are split over 2 specific sequences. Both the index terms and the related feature set should match with 2 specific sequences and the matching should be ordered in both.

There is thus a very interesting progression through the 3 types and suggests a progression in functionality as well. While the index structure maps to these neuron types, it could also be used to create more columnar structures. It would make quite a good basis for the neocortex columns [46,28], for example, with a columnar unit comprising an index node and feature set, and the index nodes would also have horizontal connections. This would occur in Figure 2, in the top neural level.

6. Ordinal Learning

The author would like to introduce the idea of Ordinal Learning. It is being given a specific name, because current methods do not appear to do it. It is not a sorting algorithm and it is not an optimisation algorithm. The ordinal learner can produce order from chaos, for example, or 'sort things out.' It is maybe more algorithmic than functional. Ordinal learning is concerned with re-creating the order of sequences it was trained with. But in this case, it can also interpret for previously unknown input that is statistically close to what it was trained with. Having a sense of order may be deeply inherent in animals, even insects. The papers [40,41], for example, map the neural connectome for some animals and the heavy tails that separate the network into modules are described. The brain is also thought to have a scheduling functionality at the top of the cortex, probably to perform such tasks. Neural networks are able to interpret what a pattern is with some missing information, but do not typically re-order the information that is presented. Although, a second network or module might learn pattern ordering, for example. Large language models also predict across a known sequence but would not intuitively know how to change a faulty sequence order. A change in the sequence order would lead to a change in the question and thus prediction. But this is still a common algorithmic problem with many solutions already. It could probably be solved in a few lines of code in many cases, and so it remains to be seen if the much more complicated method of this paper is more useful.

The following is very hypothetical, but gives the ordinator a biological base. The Ordinator could use something like a heavy-tailed neuron at each ordinal position. Thus, while heavy-tailed neurons are due to a preferential attachment or rich-get-richer mechanism [40], they would now have a

particular function as well, that is to order the surrounding neurons. The ordinator would use these neurons to place the sequence of another neuron into that ordinal position. There would then be a vote from each ordinator node, to claim their position for the neuron, where the most connections would dominate. But this is still a statistical process, where if the train example is missing, the sequence with the closest statistical match will be selected instead. A schematic of the ordinator is given in Figure 3. The system would firstly create and store the feature network for the learned positions. This is the node types described in Section 5.3. Then, new input sets can match with it. There is a hierarchical path leading to each of the ordinator position nodes. It is not clear what exactly should be in the path, but the 2 positions from the neuron types and maybe the context of the query would be possible. Then, query answers that result in sequences being selected for either position 1 or 2 can be added, or in a simpler model, they could be simply linked with the position itself. A majority vote can then be done and would resemble the neurons competing for their position.
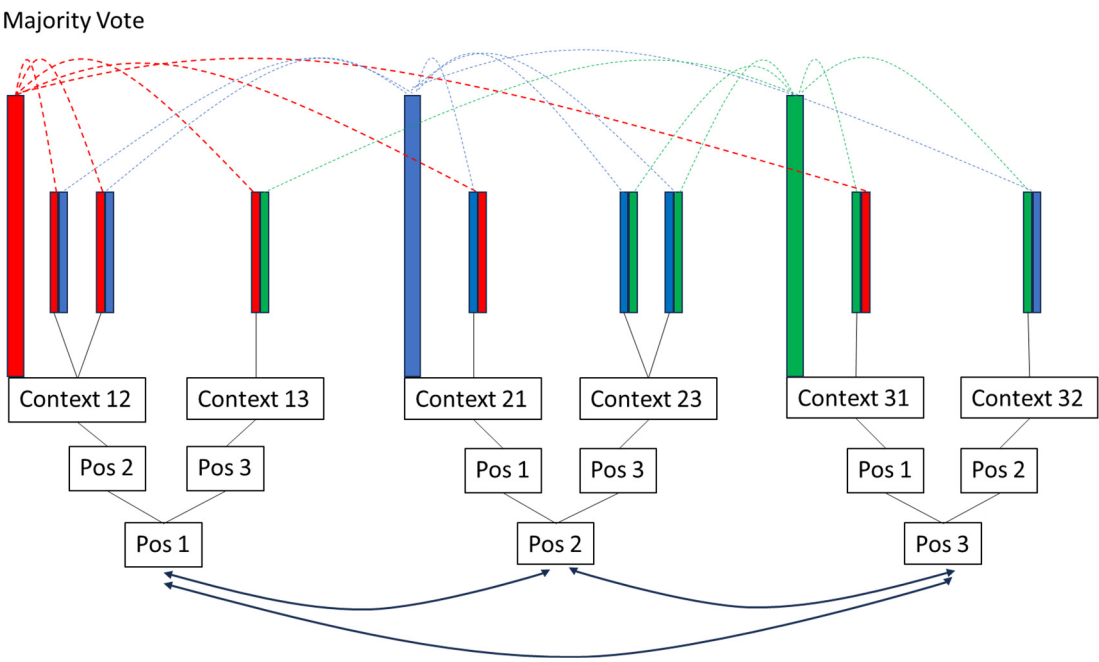

Figure 3. Schematic of the Ordinator function with 3 positions (red, blue, green). The larger structures analogy is heavy-tailed neurons and the smaller ones are (binary) feature sets. If red is also a feature type, then matching with those feature sets will link it with relative position 1 in the ordinator template. The most links with this position will assign that position to the input node.

7. Algorithm Testing

The two main algorithms that have been described have undergone some preliminary testing, where the results are presented here.

7.1. Ontology Tests

A computer program has been written in Java that implements the two memory-model levels, for basic testing. Before the text was processed, a list of over 120 common words were removed, along with any punctuation. Words like 'and', 'it', or 'a' were not included in the text, for example. The success of a test was measured rather arbitrarily, by judging if the words in a cluster had some relation. The author has judged that this is often the case. Each result was from clustering only on a single document however. It is even more difficult to judge how accurate the clusters are when texts are combined, which will be future work. Two tests were carried out. The first was simply to show that the program recognises word groups that occur more often. The second applied the clustering to some classic texts that are not particularly tuned to the problem.

7.1.1. Statistical Counting

A small text file was tested to try to determine more exactly, what the algorithm is clustering. In the file, some animal names occurred more often than others. For example, the file had repeating sequences with shared names, such as:

dog cat horse zebra.
dog cat giraffe lion wolf elephant.
dog cat horse zebra.
dog cat giraffe lion wolf elephant. …

The result was 'dog' and 'cat' being recognised as a cluster and the only cluster. Thus, it appears that the algorithm recognises this statistical variation.

7.1.2. Classic Texts

Appendix A lists some well-known texts [57], together with the final upper ontology sets that the program produced. The author hopes that it is possible to see how some level of real meaning has been derived from the statistical process. Some clusters were more obvious than others and it is very much a work in progress. The middle ontology level simply had more clusters than the upper level and some of them were then removed in the upper level, rather than clusters being merged multiple times. This is probably to be expected if the cluster values occur only once, when there is little opportunity to merge values after that. One problem that has occurred with the frequency grid is when 2 or more smaller clusters are joined together. This can result in a single cluster with apparently 2 or more meanings. This also occurred in some of the final upper ontology clusters. Rather than the program recognising associated antonyms, or something like that, it may have incorrectly combined 2 lower-level clusters somewhere, so this is a minor problem to be solved. It however becomes an interesting feature when constructing the neural level, with binary neurons, which is described in Section 5.

The resulting upper-level structures were vectors rather than trees. One test processed the story 'The Wonderful Wizard of Oz', where 4 of the clusters selected a total of 8 words. The text was also run through a skipgram word vector [63] and then the 5 most common associations with these 8 words was determined. These two cluster sets are listed in Table 1, as follows:

Table 1. E-Sense Typing and Word Vector clusters for 'The Wonderful Wizard of Oz'.

E-Sense Typing	**Type Clusters**	
	dorothy	
	asked, came, see	
	city, emerald	
	great, oz	
Word2Vec	**Type Word**	**Associations**
	dorothy	dorothy, back, over, scarecrow, girl
	asked	asked, cowardly, sorrowfully, promised, courage
	came	came, next, morning, flew, carried
	see	see, lived, away, many, people
	city	city, emerald, brick, streets, led
	emerald	emerald, brick, city, streets, gates
	great	great, room, throne, head, terrible
	oz	oz, throne, room, heart, tell

Trying to interpret the type clusters from E-Sense: the cluster 'asked', 'came', 'see' could be about the journey of discovery, for example. Or one cluster associate's 'Oz' with 'great'. This may have originated from an association with the word 'wizard'. The word vector does not convey the same meaning, but is context free and clusters sequentially related terms. E-Sense has maybe displayed some deduction, but still context free in the LLM sense. The learned part is really the pivot[1], not the deduction and would allow 'any' matching sequences to be compared and make use of the cluster information. Because it is statistically more prominent, it could also form a basis for trees. For example, the wizard is from Oz and great, but the relation describes that anything that links with it, might be from Oz and great. Or if sequences with 'came', 'asked' or 'see' are found, then a journey context can be compared with something else. The same test was carried out on the 'Romeo and Juliet' story, with the results shown in Table 2. The word 'O' was not recognised by Word2Vec and so 'thou' has been displayed instead.

Table 2. E-Sense Typing and Word Vector clusters for 'Romeo and Juliet'.

E-Sense Typing	**Type Clusters**	
	love, o, thy	
	romeo, shall	
	death, eye, hath	
Word2Vec	**Type Word**	**Associations**
	thou	thou, art, wilt, thy, hast
	love	love, art, wit, thou, fortunes
	thy	thy, thou, tybalt, dead, husband
	romeo	romeo, slain, thou, hes, art
	shall	shall, thou, till, art, such
	death	death, romeo, slain, tybalt, thy
	eye	eye, hoar, sometime, dreams, ladies
	hath	hath, husband, slain, tybalt, dead

'Romeo and Juliet' produced less obvious relations, but the Word2Vec clusters are not easy to understand either. With E-Sense, the relation 'love', 'O' and 'thy' could be associated with different characters and may be about emotion. The relation 'romeo' and 'shall' suggests deliberation, so can other deliberators be related with Romeo? The results therefore suggest that some hidden meaning, even if very basic, has been derived and can be applied to any matching sequences. So, this may loosely be called an ontology, but one that can then be used more openly. Also, it would not have to replace a word vector, but could be used with it. The clusters are also local to the document, where a different document might not make these associations. Thus, it may be the case that each document learns its' own small ontology and then these can be linked together horizontally. Or possibly if the word association exists more often, then it may be applicable in a global sense. Global relations would get pushed down vertically, but keep links with their original structures. So, the final architecture could be local, or a global and weighted semantic net, for example, or a combination of both.

7.2. Ordinal Tests

The ordinal learning process has also been implemented in Java code, for basic testing purposes. The code is in no way optimised and it could take minutes to learn the ordering for a 3K document (40 sentences) on a standard laptop. The learned structures were also much larger in size, than the raw data, but the intention is that they now contain knowledge and can be re-used. Also, the amount

of train data is quite small and so it appears to be able to use the data quite efficiently. The process comprises lots of smaller separate algorithms and so it is likely that it could be parallelised. A train text document can be read and stored in a source (lower level) database. It can then be queried for information about the main concepts in it. In fact, an ensemble approach is required, but the process is mostly automatic. A bag-of-words, for example, can determine the main terms in the document and so queries can be run to retrieve the related sequences and build the logic structures. But the process is also recursively repeated for each query, with results from that query. Thus, lots of the results are mostly the same. The index values are also part of the feature vectors and so a structure representing this would probably be quite self-contained. The frequency grid relies on the order of the input data rows, as do some neural networks, but this recursive method seems to solve that issue. After the ordinator has been created, a different test document can be loaded, or the same train document can be used to test it as well. There is no specific ordering in the base database for a whole document, and so the retrieval of sequences from the database will not be in the original document order. The purpose of the function is to recognise this order again. It therefore tries to match with the source data and then sort that into the ordering that it has learnt. The ordering is only relative, where the number of sequences in the train and test documents can be different. For example, ordinal positions of 1 and 2 in the Ordinator simply mean that the sequence in relative position 1 should be before the sequence in relative position 2, not that the positions are exactly 1 and 2.

Appendix B gives the results for some basic tests. Two cooking instruction documents were selected. The first train document described how to cook a hard-boiled egg and the second described how to make Panna Cotta. If the test document was the same, then the data would be returned correctly. For a second test, an ordinator was generated for either train document and then the database was changed to one which contained both - a different description of how to cook a hard-boiled egg and the Panna Cotta description. The egg ordinator was able to select the sequences relating to the second description and also place them in order, as shown. The Panna Cotta ordinator also performed, as shown. So, while it did perform to 100% accuracy for these two small documents, using slightly different words might change the result. The process is more entropic than local connections, however and so there would be a balancing act to adding specific rules about something. It is also slightly stochastic, but the ensemble training method helps to keep the results mostly the same. The distinct concepts in a feature were almost as important. Results also showed however, that for larger documents, of even 40 lines or more, some sequences would typically be missed, or not even retrieved from the source database. It is therefore unlikely that the process can be used to simply rote learn a large corpus of information. But it could be a useful guide and along with search processes, sort through information and make some sense of it. There are still options to be tried, to make it more accurate.

8. Some Biological Comparisons

This section makes comparisons with some purely biological ideas.

8.1. Gestalt Psychology

With Gestalt psychology, objects have an 'other' interpretation of the sub-features and not just a summed whole of them. Gestalt theory makes use of ideas like similarity and proximity (and good continuation) to group objects and it believes that the brain has an internal order and structure that it places external stimuli into. With the theory, the brain may try to fit the outside world into its internal structure, rather than copy the external world exactly. Gestalt theory could be realised in the lower model level. Because the links are unbiassed, one interesting aspect of the structure is that it may not return exactly what was input, thus satisfying the theory that the whole may be different to the parts. Consider this example where the following 2 sentences are added to the lowest-level memory:

The cat sat on the mat and drank some milk.

The dog barked at the moon and chased its tail.

Start words would now include 'the' and end words would include 'milk' and 'tail.' Each word in the sequence also has a link to any words that immediately follow it. If, for example, the memory system is asked to retrieve a sequence that results in the 'and' word being considered, then there are two possibilities after that – 'drank' or 'chased.' Therefore, a question about a dog could retrieve either of the following two sentences:

The dog barked at the moon and chased its tail, or
The dog barked at the moon and drank some milk.

If the second sentence was returned, then it would not be violating the memory system and the person would probably not have a reason to disbelieve it. If an associative network is to be included however, then that will add some stability and bias, but then the second sentence could be seen as a new option that might be considered if the learned sentences are not suitable. It might also be argued that changing the information in this way is not creative, but simply taking a different route through the linked structure.

8.1.1. Numbers in the Gestalt Process

As part of the computer model, an n-gram depth in the lower level can add accuracy, requiring that several previous concepts are present to establish a sequence, where in fact a 5-gram is currently used. The decision to copy 5 tokens in order was a guess, but the number 5 is also important in Gestalt theory. Buffart [5] developed a mathematical method to formulate Gestalt theory. The outside or environment should contain at least 4 elements. Then, the process requires regularity when comparing elements and one of the theorems states that the relations between the elements in one event can be interpreted at most in two ways if, and only if, the number of elements in that event is at most seven. Moreover, for events with five, six or seven elements the possible interpretations are the same. 'Two interpretations' is also well known in perception as the duality principle and seven minus two is well known in memory and attention research. In his 'seven minus two and structure' section, Buffart presents a proof that with up to seven materials in focus, each complete structure can be described by the combination of at most two complete interpretations. For spans larger than seven, at least three interpretations can be required to represent a complete structure. Thus, the numbers 2 and 5 (to 7) become very important to brain theory. One might consider two to be a stable state, allowing a single comparison. Five to seven then relates to structural stability, where less or more than this can change the dynamics of what elements are linked together. The n-gram length is 5, but could be 7 and the new model process did result in node types, described in Section 5.3, that are either unary or binary. Thus, if this is the stable state, the number 3 has also occurred in the paper for more dynamic states. A value of 3, for example, would allow something to be rooted (1), while two parts are being compared (2 and 3). If the brain likes to synchronise to balanced states, for example, then this asynchrony might encourage the system to explore a step further.

8.2. Cortical Columns

A comparison with these has been made for the upper functional level, primarily because the structure becomes distributed units. This was recognised in the seminal work by Mountcastle [46], where the columns are not because of physical boundaries, but because the neurons connect more vertically and together, than horizontally. It is only from preliminary results, but the idea of different neuron types can be found in Section 5.3. If the vector lists correspond to dendrites, for example, then there was either two or one input dendrites. Two separate dendrites would correspond to a Pyramidal neuron, for example, discovered by Cajal[2]. One other version of the dendrites could be larger (and unordered). Therefore, maybe this larger vector would represent an attractor for the column that would also provide exclusivity of the input signal, against the neighbouring columns.

8.3. Animal Brain Function

This section follows-on from ideas in [18] that describe how cells may have evolved from invertebrates to the human brain. Organisation was a key principle there, and also the conversion to types. The brain is typed, where even insects like ants can recognise types, but the typing might not yet be at the symbolic level. Thus, if traversing through the new cognitive model of Figure 2, the middle ontology layer converts from the lower-level ensembles of set-based values to more type-based vectors. Then, between the middle and upper levels the types can be clustered back again into sequences, based on time and use. What if these new clusters are missing some of the sensory information? If the cortex is mostly about actions, or how to do something, then it does not require additional sensory feedback, which might also help to protect mental states. It would only want to know what the objects are and how to use them. So possibly sensory information can remain deeper in the brain, where it would also be closer to the senses themselves. Again, if thinking about a conscious experience, then it would require the sensory feedback from a whole-brain activation, not just local circuits in the cortex.

8.3.1. Theory of Small Changes

Thus, a constructive process can be seen in the building of brain structures, but the new structure should be added in small amounts each time. If we do not know about a subject for example, then the first structures for it should probably contain the basics or fundamentals for the subject. What this means is not clear, but possibly, what other concepts link with. If we already have some memory about a subject, then we can add to it instead, enriching the information. But the most important blocks and a basic structure is added first. Turing [60] wrote about sub-critical and super-critical processes that a human brain would recognise more than an animal. But he suggested that single events were key and that maybe even the mind had some influence on what gets stored. A mechanical process could simply make small changes until enough of them became critical and forced a larger change (the single event). But then the mechanical process should not try to store every piece of information. Therefore, some type of reinforcement from a more intelligent region that helps to select important input could indeed happen. But again, critical selection could result not only from more intelligence, but also less emotion, as described in the last section. More basic animals make use of memory, recognition and sensory feedback, while more intelligent ones can also manipulate and understand better.

9. Conclusions and Future Work

This paper describes a first implementation of the 3-level cognitive model, now called E-Sense. While it is essentially a computer model, it is based strongly on our understanding of the human brain. The new design is really an improved version of the original one. The stigmergic links have been replaced by ones that still prefer minimal path distances, for example and so it is really only the decay factor that is missing. Averaged values give a sense of what is in the data, which has been replaced by minimal type-based relations. Then, higher-level concepts have been replaced by functions that process subsets of the underlying data. The 2 memory levels are economic in space and can transpose the information into deductive relations. This introduces new knowledge and may in fact be helpful to an upper level that wants to know more about objects and the technical 'how'. The upper level is more neural. Generating it produced a kind of progressive functionality that could be loosely mapped to neuron types, or maybe just standard unary and binary operators. Two new algorithms have been developed that could enhance modern systems, but maybe require changes to the architecture. It might be convenient to look at the whole brain model as auto-associative. The bottom memory substrate / sensory level and the cortical levels map to each other through the middle-level transpositions. The two views are not the same however and so there is still an economy of storage, but the related parts in each map together.

A lot of new ideas have been presented and a recurring question is if it might be useful in practice. The system is more biologically inspired than most AI systems. This should make it more generic, which is a very important point. Condensing information into abstract relational models or having clusters with loose boundaries are part of new research, covered in Section 3.2 on alternative models. So, this is part of a new research direction. Also, as part of this new research, would be compression of memory and symbolic interpretation with the reduced data. The two new algorithms that are described could prove to be useful on their own. Current systems do not produce the same transitions to use-based types, which is a level below symbolic understanding. This is evidenced by the difference between the E-Sense clusters and the word vector clusters. Therefore, any successes are likely to produce results that will enhance the knowledge about AI and help with future research.

There is still a lot of work to be done in all the levels. Some biological analogies have been made and while it looks good on paper, it will probably not be as clear in practice. The results are not too bad and as the name suggests, they may be more about some type of general understanding, than the more direct results that the current systems provide. Some specifics probably need to be added, to increase the intelligence level. Considering future work, the source database does not always return exactly what was entered, so it would need to be determined if this is critical. Should it return everything exactly, or would that require storing the original source as it is? It is a very compact structure. Then, there is not a clear progression yet from the middle to the top level. The top level is currently created from accessing the source database directly, but middle to top could be included in a more dynamic system. Then also, making the system more accurate and useful.

Appendix A. Upper Ontology Trees for Book Texts

This appendix lists the upper-level ontology trees that were created for some well-known books. The clustering relates to the use of the word. Each row is a child node of the row immediately before it, but in fact the row ordering can change.

The Wonderful Wizard of Oz, L. Frank Baum [57].
The upper ontology level reduced the number of clusters from 17 to 9.

Clusters
dorothy
scarecrow
lion, woodman
great, oz
city, emerald
asked, came, see
little, out, tin
again, answered, away, before, down, made, now, shall, toto, up
back, come, girl, go, green, head, heart, man, one, over, upon, very, witch

Romeo and Juliet, William Shakespeare [57].
The upper ontology level reduced the number of clusters from 17 to 8.

Clusters
romeo, shall
thou
love, o, thy

death, eye, hath
come, thee
go, good, here, ill, night, now
day, give, lady, make, one, out, up, well
man, more, tybalt

The Adventures of Sherlock Holmes, Arthur Conan Doyle [57].

Clusters
back, before, came
down, know
more, room, think, well
day, eye, face, found, matter, tell
upon
holmes, very
little, man, now
one
away, case, good, heard, house, much, nothing, quite, street, such, through, two, ye
go, here
come, hand, over, shall, time
asked, never
door, saw
mr, see
out, up
made, way

Computing Machinery and Intelligence, A.M. Turing [60].

Clusters
answer, computer, man, question, think
machine
one
such

Appendix B. Documents and Test Results for the Neural-Level Sorting

This appendix lists the train and test files for testing the ordinal learning. The results of applying the ordering to the test files is also shown.

Train and Test Files

Train File – Hard-Boiled Egg

Place eggs at the bottom of a pot and cover them with cold water.

Bring the water to a boil, then remove the pot from the heat.

Let the eggs sit in the hot water until hard-boiled.

Remove the eggs from the pot and crack them against the counter and peel them with your fingers.

Train File – Panna Cotta
For the panna cotta, soak the gelatine leaves in a little cold water until soft.
Place the milk, cream, vanilla pod and seeds and sugar into a pan and bring to a simmer.
Remove the vanilla pod and discard.
Squeeze the water out of the gelatine leaves, then add to the pan and take off the heat.
Stir until the gelatine has dissolved.
Divide the mixture among four ramekins and leave to cool.
Place into the fridge for at least an hour, until set.
For the sauce, place the sugar, water and cherry liqueur into a pan and bring to the boil.
Reduce the heat and simmer until the sugar has dissolved.
Take the pan off the heat and add half the raspberries.
Using a hand blender, blend the sauce until smooth.
Pass the sauce through a sieve into a bowl and stir in the remaining fruit.
To serve, turn each panna cotta out onto a serving plate.
Spoon over the sauce and garnish with a sprig of mint.
Dust with icing sugar.

Test File – Hard-Boiled Egg and Panna Cotta
Remove the vanilla pod and discard.
For the panna cotta, soak the gelatine leaves in a little cold water until soft.
As soon as they are cooked drain off the hot water, then leave them in cold water until they are cool enough to handle.
Squeeze the water out of the gelatine leaves, then add to the pan and take off the heat.
Spoon over the sauce and garnish with a sprig of mint.
Stir until the gelatine has dissolved.
Place the eggs into a saucepan and add enough cold water to cover them by about 1cm.
Pass the sauce through a sieve into a bowl and stir in the remaining fruit.
Divide the mixture among four ramekins and leave to cool.
Place into the fridge for at least an hour, until set.
To peel them crack the shells all over on a hard surface, then peel the shell off starting at the wide end.
For the sauce, place the sugar, water and cherry liqueur into a pan and bring to the boil.
Place the milk, cream, vanilla pod and seeds and sugar into a pan and bring to a simmer.

Reduce the heat and simmer until the sugar has dissolved.

Take the pan off the heat and add half the raspberries.

Using a hand blender, blend the sauce until smooth.

Bring the water up to boil then turn to a simmer.

To serve, turn each panna cotta out onto a serving plate.

Dust with icing sugar.

Test Results

Selected Sequences from the Hard-Boiled Egg Function
[place, the, eggs, into, a, saucepan, and, add, enough, cold, water, to, cover, them, by, about]
[bring, the, water, up, to, boil, then, turn, to, a, simmer]
[as, soon, as, they, are, cooked, drain, off, the, hot, water, then, leave, them, in, cold, water, until, they, are, cool, enough, to, handle]
[to, peel, them, crack, the, shells, all, over, on, a, hard, surface, then, peel, the, shell, off, starting, at, the, wide, end]

Selected Sequences from the Panna Cotta Function
[for, the, panna, cotta, soak, the, gelatine, leaves, in, a, little, cold, water, until, soft]
[place, the, milk, cream, vanilla, pod, and, seeds, and, sugar, into, a, pan, and, bring, to, the, boil]
[remove, the, vanilla, pod, and, discard]
[squeeze, the, water, out, of, the, gelatine, leaves, then, add, to, the, pan, and, take, off, the, heat]
[stir, until, the, gelatine, has, dissolved]
[divide, the, mixture, among, four, ramekins, and, leave, to, cool]
[place, into, the, fridge, for, at, least, an, hour, until, set]
[for, the, sauce, place, the, sugar, water, and, cherry, liqueur, into, a, pan, and, bring, to, the, boil]
[reduce, the, heat, and, simmer, until, the, sugar, has, dissolved]
[take, the, pan, off, the, heat, and, add, half, the, raspberries]
[using, a, hand, blender, blend, the, sauce, until, smooth]
[pass, the, sauce, through, a, sieve, into, a, bowl, and, stir, in, the, remaining, fruit]
[to, serve, turn, each, panna, cotta, out, onto, a, serving, plate]
[spoon, over, the, sauce, and, garnish, with, a, sprig, of, mint]
[dust, with, icing, sugar]

Notes

1. In statistics, a pivotal quantity is a function of observations and unobservable parameters. It can be used in normalisation, to allow data from different data sets to be compared.
2. http://www.scholarpedia.org/article/Santiago_Ramón_y_Cajal.

References

1. Anderson, J.A., Silverstein, J.W., Ritz, S.A. and Jones, R.A. (1977). Distinctive Features, Categorical Perception, and Probability Learning: Some Applications of a Neural Model, Psychological Review, Vol. 84, No. 5.
2. ARC Prize. (2025). https://arcprize.org/. (last accessed 30/4/25).
3. Barabasi A.L. and Albert R. (1999). Emergence of scaling in random networks, Science, 286:509-12.
4. Brown, P.F., Della Pietra, V.J., Desouza, P.V., Lai, J.C. and Mercer, R.L.. (1992). Class-based n-gram models of natural language. Computational linguistics, 18(4), pp.467-480.
5. Buffart, H. (2017). A formal approach to Gestalt theory, Blurb, ISBN: 9781389505577.
6. Cavanagh, J. P. (1972). Relation between the immediate memory span and the memory search rate. Psychological Review, 79, pp. 525 - 530.
7. Cover, T.M. and Joy, A.T. (1991). Elements of Information Theory, John Wiley & Sons, Inc. Print ISBN 0-471-06259-6 Online ISBN 0-471-20061-1.
8. Dobrynin, V., Sherman, M., Abramovich, R., and Platonov, A. (2024). A Sparsifier Model for Efficient Information Retrieval, AICT'24.
9. Dobson, S. and Fields, C. (2023). Constructing condensed memories in functorial time. Journal of Experimental & Theoretical Artificial Intelligence, pp.1-25.
10. Dong, M., Yao, L., Wang, X., Benatallah, B. and Zhang, S. (2018). GrCAN: Gradient Boost Convolutional Autoencoder with Neural Decision Forest. arXiv preprint arXiv:1806.08079.
11. Dorigo, M., Bonabeau, E. And Theraulaz, G. (2000). Ant algorithms and stigmergy, Future Generation Computer Systems, Vol. 16, pp. 851 - 871.
12. Eliasmith, C., Stewart, T.C., Choo, X., Bekolay, T. DeWolf, T., Tang, Y. and Rasmussen, D. (2012). A Large-Scale Model of the Functioning Brain, Science, 338(6111), pp. 1202 - 1205.
13. Feng, S., Sun, H., Yan, X., Zhu, H., Zou, Z., Shen, S. and Liu, H.X. (2023). Dense reinforcement learning for safety validation of autonomous vehicles. Nature, 615(7953), pp. 620 - 627.
14. Fink, G.A. (2014). Markov models for pattern recognition: from theory to applications. Springer Science & Business Media.
15. Friedman, R. (2021). Cognition as a Mechanical Process. NeuroSci, 2, 141–150. https://doi.org/10.3390/neurosci2020010.
16. Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. and Bengio, Y. (2014). Generative adversarial nets. Advances in neural information processing systems, 27.
17. Grassé P.P. (1959). La reconstruction dun id et les coordinations internidividuelles chez Bellicositermes natalensis et Cubitermes sp., La théorie de la stigmergie: essais d'interprétation du comportement des termites constructeurs, Insectes Sociaux, Vol. 6, pp. 41-84.
18. Greer, K. (2022). Neural Assemblies as Precursors for Brain Function, NeuroSci, 3(4), pp. 645 - 655. https://doi.org/10.3390/neurosci3040046. Also published in Eds. Parnetti, L., Paoletti, F.P. and Gallart-Palau, X., Feature Papers in NeuroSci : From Consciousness to Clinical Neurology, July 2023, pages 256. ISBN 978-3-0365-7846-0 (hardback); ISBN 978-3-0365-7847-7 (PDF) https://doi.org/10.3390/books978-3-0365-7847-7.
19. Greer, K. (2021). New Ideas for Brain Modelling 7, International Journal of Computational and Applied Mathematics & Computer Science, Vol. 1, pp. 34-45.
20. Greer, K. (2021). Is Intelligence Artificial? Euroasia Summit, Congress on Scientific Researches and Recent Trends-8, August 2-4, The Philippine Merchant Marine Academy, Philippines, pp. 307 - 324. Also available on arXiv at https://arxiv.org/abs/1403.1076.

21. Greer, K. (2021). Category Trees - Classifiers that Branch on Category, International Journal of Artificial Intelligence & Applications (IJAIA), Vol. 12, No. 6, pp. 65 - 76.

22. Greer, K. (2020). New Ideas for Brain Modelling 6, AIMS Biophysics, Vol. 7, Issue 4, pp. 308-322. doi: 10.3934/biophy.2020022.

23. Greer, K. (2019). New Ideas for Brain Modelling 3, Cognitive Systems Research, 55, pp. 1-13, Elsevier. DOI: https://doi.org/10.1016/j.cogsys.2018.12.016.

24. Greer, K. (2012). Turing: Then, Now and Still Key, in: X-S. Yang (eds.), Artificial Intelligence, Evolutionary Computation and Metaheuristics (AIECM) - Turing 2012, Studies in Computational Intelligence, 2013, Vol. 427/2013, pp. 43-62, DOI: 10.1007/978-3-642-29694-9_3, Springer-Verlag Berlin Heidelberg.

25. Greer, K. (2011). Symbolic Neural Networks for Clustering Higher-Level Concepts, *NAUN International Journal of Computers*, Issue 3, Vol. 5, pp. 378 – 386, extended version of the WSEAS/EUROPMENT International Conference on Computers and Computing (ICCC'11).

26. Gruber, T. (1993). A translation approach to portable ontology specifications. Knowledge Acquisition, 5, pp. 199 - 220.

27. Gupta, B., Rawat, A., Jain, A., Arora, A. and Dhami, N. (2017). Analysis of various decision tree algorithms for classification in data mining. International Journal of Computer Applications, Vol. 163, No. 8, pp. 15 - 19.

28. Hawkins, J., Lewis, M., Klukas, M., Purdy, S. and Ahmad, S. (2019). A Framework for Intelligence and Cortical Function Based on Grid Cells in the Neocortex, Frontiers in neural circuits, 12, p. 121.

29. Hawkins, J. and Blakeslee, S. On Intelligence. Times Books, 2004.

30. High, R., 2012. The era of cognitive systems: An inside look at IBM Watson and how it works. IBM Corporation, Redbooks, pp.1-16.

31. Hinton, G., 2023. How to represent part-whole hierarchies in a neural network. Neural Computation, 35(3), pp.413-452.

32. Hinton, G.E., Osindero, S. and Teh, Y.-W. (2006). A fast learning algorithm for deep belief nets, Neural computation, Vol. 18, No. 7, pp. 1527 - 1554.

33. Katuwal, R., Suganthan, P.N. (2019). Stacked Autoencoder Based Deep Random Vector Functional Link Neural Network for Classification, accepted: Applied Soft Computing (https: // doi. org/ 10. 1016/ j. asoc. 2019. 105854).

34. Kingma, D.P. and Welling, M., 2019. An introduction to variational autoencoders. Foundations and Trends in Machine Learning, 12(4), pp.307-392.

35. Krizhevsky, A., Sutskever, I. and Hinton, G.E. (2012). Imagenet classification with deep convolutional neural networks. In Advances in neural information processing systems, pp. 1097-1105.

36. Krotov, D. (2023). A new frontier for Hopfield networks. Nature Reviews Physics, 5(7), pp. 366 - 367.

37. Laird, J. (2012). The Soar cognitive architecture, MIT Press.

38. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. Nature 521, 436–444 (2015). https://doi.org/10.1038/nature14539.

39. Lieto, A., Lebiere, C. and Oltramari, A. (2017). The knowledge level in cognitive architectures: Current limitations and possible developments, Cognitive Systems Research.

40. Lynn, C.W., Holmes, C.M. and Palmer, S.E. (2024). Heavy-tailed neuronal connectivity arises from Hebbian self-organization, Nature Physics, 20(3), pp.484-491.

41. Meunier, D., Lambiotte, R. and Bullmore, E.T., 2010. Modular and hierarchically modular organization of brain networks. Frontiers in neuroscience, 4, p.200.

42. Miller, G. A. (1956). The magical number seven, plus or minus two: some limits on our capacity for processing information. Psychological Review, 63, pp. 81 - 97.

43. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S. and Dean, J. (2013). Distributed representations of words and phrases and their compositionality. Advances in neural information processing systems, 26.

44. Minaee, S., Mikolov, T., Nikzad, N., Chenaghlu, M., Socher, R., Amatriain, X. and Gao, J., 2024. Large language models: A survey. arXiv preprint arXiv:2402.06196.

45. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D.,

Wierstra, D., Legg, S. and Hassabis, D. (2015). Human-level control through deep reinforcement learning, Nature, Vol. 518, pp. 529-533.

46. Mountcastle, V.B. (1997). The columnar organization of the neocortex, Brain: J Neurol, Vol. 120, pp. 701 - 722.

47. Newell, A. and Simon, H.A. (1976). Computer science as empirical inquiry: Symbols and search, Communications of the ACM, Vol.19, No. 3, pp. 113 - 126.

48. Nguyen, T., Ye, N. and Bartlett, P.L. (2019). Learning Near-optimal Convex Combinations of Basis Models with Generalization Guarantees. arXiv preprint arXiv:1910.03742.

49. OpenAI. 2023. GPT-4 Technical Report. (2023). arXiv:cs.CL/2303.08774.

50. Pulvermüller, F., Tomasello, R., Henningsen-Schomers, M.R. and Wennekers, T. (2021). Biological constraints on neural network models of cognitive function, Nature Reviews Neuroscience, 22(8), pp. 488 - 502.

51. Rock, I. (1977). In defence of unconscious inference. In W. Epstein (Ed.), Stability and constancy in visual perception: mechanisms and processes. New York, N. Y.: John Wiley & Sons.

52. Rubinov, M., Sporns, O., van Leeuwen, C., and Breakspear, M. (2009). Symbiotic relationship between brain structure and dynamics. BMC Neurosci. 10, 55.

53. Sarker, M.K., Zhou, L., Eberhart, A. and Hitzler, P. (2021). Neuro-symbolic artificial intelligence. AI Communications, 34(3), pp. 197 - 209.

54. Shannon, C.E. (1948). A Mathematical Theory of Communication, The Bell System Technical Journal, 27(3), pp. 379 - 423.

55. Spens, E. and Burgess, N. (2024). A generative model of memory construction and consolidation, Nature Human Behaviour, 8(3), pp. 526 - 543, https://doi.org/10.1038/ s41562-023-01799-z.

56. Tkačik, G., Mora, T., Marre, O., Amodei, D., Palmer, S.E., Berry, M.J. and Bialek, W. (2015). Thermodynamics and signatures of criticality in a network of neurons, Proceedings of the National Academy of Sciences, Vol. 112, No. 37, pp. 11508 - 11513.

57. The Gutenberg Project., https://www.gutenberg.org/browse/scores/top. (last accessed 30/4/25).

58. Treves, A. and Rolls, E.T. (1991). What determines the capacity of autoassociative memories in the brain?, Network: Computation in Neural Systems, 2(4), p.371.

59. Tsien, R.Y. Very long-term memories may be stored in the pattern of holes in the perineuronal net. Proc. Natl. Acad. Sci. USA 2013, 110, 12456-12461.

60. Turing, A.M. (1950). Computing machinery and intelligence. Mind, 59, pp. 433 - 460.

61. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L. and Polosukhin, I. (2017). Attention is all you need. Advances in neural information processing systems, 30.

62. Webb, T.W., Frankland, S.M., Altabaa, A., Segert, S., Krishnamurthy, K., Campbell, D., Russin, J., Giallanza, T., O'Reilly, R., Lafferty, J. and Cohen, J.D. (2024). The relational bottleneck as an inductive bias for efficient abstraction. Trends in Cognitive Sciences.

63. Word2Vec – DeepLearning4J. (2025). https://deeplearning4j.konduit.ai/en-1.0.0-beta7/language-processing/word2vec. (last accessed 12/5/25).