

RESEARCH

Open Access

A multi-modal network architecture for knowledge discovery

Craig M Vineyard^{1*}, Stephen J Verzi¹, Michael L Bernard¹, Shawn E Taylor¹, Irene Dubicka¹ and Thomas P Caudell²

Abstract

The collection and assessment of national security related information often involves an arduous process of detecting relevant associations between people, events, and locations—typically within very large data sets. The ability to more effectively perceive these connections could greatly aid in the process of knowledge discovery. This same process—pre-consciously collecting and associating multimodal information—naturally occurs in mammalian brains. With this in mind, this effort sought to draw upon the neuroscience community's understanding of the relevant areas of the brain that associate multi-modal information for long-term storage for the purpose of creating a more effective, and more automated, association mechanism for the analyst community. Using the biology and functionality of the hippocampus as an analogy for inspiration, we have developed an artificial neural network architecture to associate k-tuples (paired associates) of multimodal input records. The architecture is composed of coupled unimodal self-organizing neural modules that learn generalizations of unimodal components of the input record. Cross modal associations, stored as a higher-order tensor, are learned incrementally as these generalizations are formed. Graph algorithms are then applied to the tensor to extract multi-modal association networks formed during learning. Doing so yields a potential novel approach to data mining for intelligence-related knowledge discovery. This paper describes the neurobiology, architecture, and operational characteristics, as well as provides a simple intelligence-based example to illustrate the model's functionality.

Introduction

Currently, intelligence analysts are hampered by the need to sift through very large amounts of constantly changing data in order to forage for “nuggets” of information that may support or discredit an existing hypothesis. The collection and assessment of national security related information often involves an arduous process of “connecting the dots” within very large data sets. This process has proven to be extremely difficult, especially when analysts need to piece together information cues associated with various individuals, groups, events, and places, along with such items as communication and transportation logs [1]. The ability to more effectively perceive connections among events, locations, and people could greatly aid in the process of knowledge discovery.

Recent data mining and fusion tools have become much more effective in uncovering evidence of potential

threats by sifting through Internet traffic, financial and communications records, as well as transcripts of audio streams for patterns of interest. While this type of capability is useful in understanding general patterns of behaviour, it is typically limited to one type of information domain (e.g. textual) and must rely on a large number of statistically related links to uncover relevant patterns. In addition, systems that utilize video sources to analyse video surveillance information to classify video footage are typically achieved without the ability to infer common relationships among related video events or actors.

Regardless of the information source, a significant problem faced by existing approaches is the immense difficulty in finding an information signal that is indicative of specific adversary behaviours and associating it with other meaningful signals in a vast expanse of noise. That is, current statistical database approaches, by themselves, are generally ill equipped to detect meaningful associations across a spectrum of information sources. Consequently, existing systems are generally considered poorly equipped to actively assist in the marshalling and assessment of multi-source information. Developing a

* Correspondence: cmviney@sandia.gov

¹Sandia National Laboratories, Albuquerque, NM, USA

Full list of author information is available at the end of the article

system that assists analysts with knowledge discovery by helping to uncover associations, as well as help marshal evidence by assembling individual pieces of evidence into a single context, would be a great advancement to the analyst community. This is particularly true with the increasing need to more rapidly detect associations across various information modes for threat identification and determination in real-time, security-related contexts—for example, in situations involving time critical targets of national importance where rapid assessments must be made as to the type and degree of threat that may or may not exist.

In response to this need, an internally-funded effort sponsored by Sandia National Laboratories is seeking to advance the field of knowledge discovery by exploring both traditional statistics-based approaches as well as a neurologically-based, or “neuromorphic”, approach to auto-associate information similar to the way a mammalian brain processes and associates multi-sourced information. This process of collecting and storing information naturally occurs in an awake mammalian brain. While a system that can fully auto-associate relevant, multi-modal information as described above is still in the future, we assert that an effort to replicate the associative processes of the brain to an appropriate degree, has the promise to greatly advance the process of knowledge discovery. This process would more effectively generate threat determinations in support of rapid decision-making in security-related contexts by filtering through a large corpus of multi-source/multi-mode information to uncover relevant associations.

The focus of this effort, termed the Augmented COgnition for Rapid Decision making (ACORD) capability, is to explore how to model relevant neurological processes in the brain that naturally associate information from different modalities for long-term storage as a memory episode. Recent advances in knowledge pertaining to the processes underlying associative memory have made it possible to model these processes at a level of fidelity that is applicable to knowledge discovery. This discussion will emphasize our neuromorphic approach underlying the ACORD capability.

Neurological underpinnings of the ACORD effort

The brain receives a variety of sensory input signals such as visual, auditory, and olfactory. Although each input stream does receive its share of focused individual processing, additional insight comes from the converged processing of all input modalities. Such an occurrence takes place within the Medial Temporal Lobe (MTL) region of the brain, and more specifically within the hippocampus. Beyond receiving a convergence of sensory inputs, the hippocampus is essentially involved in episodic memory formation. Rather than simply being a

mechanism for storing information, episodic memory associates information such as the spatial and temporal contexts of an event.

The episodic memory capability of the brain enables us to encode personal experiences as converged neural activations across cortical areas using diverse sensory modalities. In doing so, we are able to remember a large number of events including detailed sequences of events comprising experience as well as the temporal and spatial context of each event in the sequence [2]. One brain area, the hippocampus, is critically involved in remembering the spatial and temporal context of an event. The location of the hippocampus within the human brain may be seen in Figure 1. The MTL, where the hippocampus is located, is the recipient of inputs from widespread areas of the cortex and supports the ability to bind together cortical representations [2]. A key component of episodic memory is association of diverse types of information [3]. This capability allows humans to relate knowledge pertaining to elements of an event such as who, what, when, and where. In this paper we present an artificial neural network architecture that learns these types of association inspired by the hippocampus. We start with an overview of hippocampus biology; then explain the design of our neural network architecture; follow this with sample experimental results to motivate its use; then present a real world terrorist network scenario, and conclude with a discussion of the significance of our architecture and directions for future research.

Cortical inputs to MTL arrive from various sensory modalities, with different emphases depending upon the mammalian species. For instance, rats receive a significant olfactory influence whereas bats receive a strong auditory influence [4]. Nevertheless, across species, most of the neocortical inputs to the perirhinal cortex come from cortical areas which process unimodal sensory information about qualities of objects, called the “what” stream, and most of the neocortical inputs to the parahippocampal cortex come from cortical areas which process polymodal spatial information, called the “where” stream [4,2]. There are some connections between the two streams, however overall processing in each stream remains largely segregated until they converge within the hippocampus [5,6].

Extensive neuroscience research typically describes the anatomy of the hippocampus as consisting of a loop (see Figure 1) beginning with the dentate gyrus (DG), proceeding to CA3, followed by CA1 and propagating through the subiculum out to the input streams [7]. The hippocampus receives its inputs from the entorhinal cortex (EC) and passes its outputs back to the EC. The EC receives inputs from both perirhinal (dorsal) and parahippocampal (ventral) cortices. The perirhinal cortex is

involved in object recognition, and the parahippocampal cortex is involved in recognizing scenes. The sub regions of the hippocampus and surrounding cortex, specifically related to the motivation for our neural model, will be addressed individually as follows (also see columns 1 and 2 in Table 1).

The DG receives the conjoined multimodal sensory signals from EC. Anatomically, DG consists of a large number of neurons with a relatively sparse neural activation code at any given instant. Effectively, this behaviour suggests that the DG creates non-overlapping sparse codes for unique events [8]. In this case an event consists of simultaneous neural activation leading into (afferent to) the hippocampus (specifically the DG in this case) within a short span of time. The sparse DG outputs serve as the input for CA3.

The CA3 region of the hippocampus consists of extensive recurrent connections. The CA3 region also receives direct input from the EC. The sparse encoding of the DG allows the CA3 to uniquely encode EC activation patterns as specific events within an episode as well as facilitating later semantic encoding. These neural processes enable CA3 to perform auto-association. Anatomically, the output of CA3 proceeds to CA1 and subiculum as the major output regions of the hippocampus [9]. While the exact functionality of the subiculum is largely unknown, CA1 functionality is typically identified as learning relational information for temporal sequences and connecting episodic encodings from CA3 with the original EC sensory activations. We have used some of these functional properties of the hippocampus as the basis for an artificial neural network architecture

for learning and forming associations. Table 1 depicts the relationship between these anatomical regions and the corresponding computational implementation, which we will describe next.

Computational architecture

In general, an association is a relationship between entities where they share some degree of commonality. For example, an individual is associated with his/her name, or two individuals may be associated with a common workplace. All entities are trivially related to themselves. The simplest non-trivial association is between two entities, but in general, k individual entities may be associated with one another. The question arises as to how relationships are learned and encoded as memory?

Numerous domain specific rules or heuristics may be utilized to discover commonality among entities based upon criteria such as distance metrics or shared feature counts. In contrast, our architecture inspired by the hippocampus builds relational codes by associating multiple modal specific entities with their mutual context, analogous to the dorsal and ventral partitioning in EC sensory input signals. In its simplest form, our approach associates what and where information based upon their shared frame of reference. For example, multiple people may be associated with the house in which they live.

In order to create associations, the network must first create representations (or neural codes) of the individual unimodal sensory perceptions of the entities. Prior to entering hippocampus, sensory signals pass through numerous layers of cortex. Throughout these layers a distributed representation for entities is gradually




Figure 1 Hippocampus and associated structures.

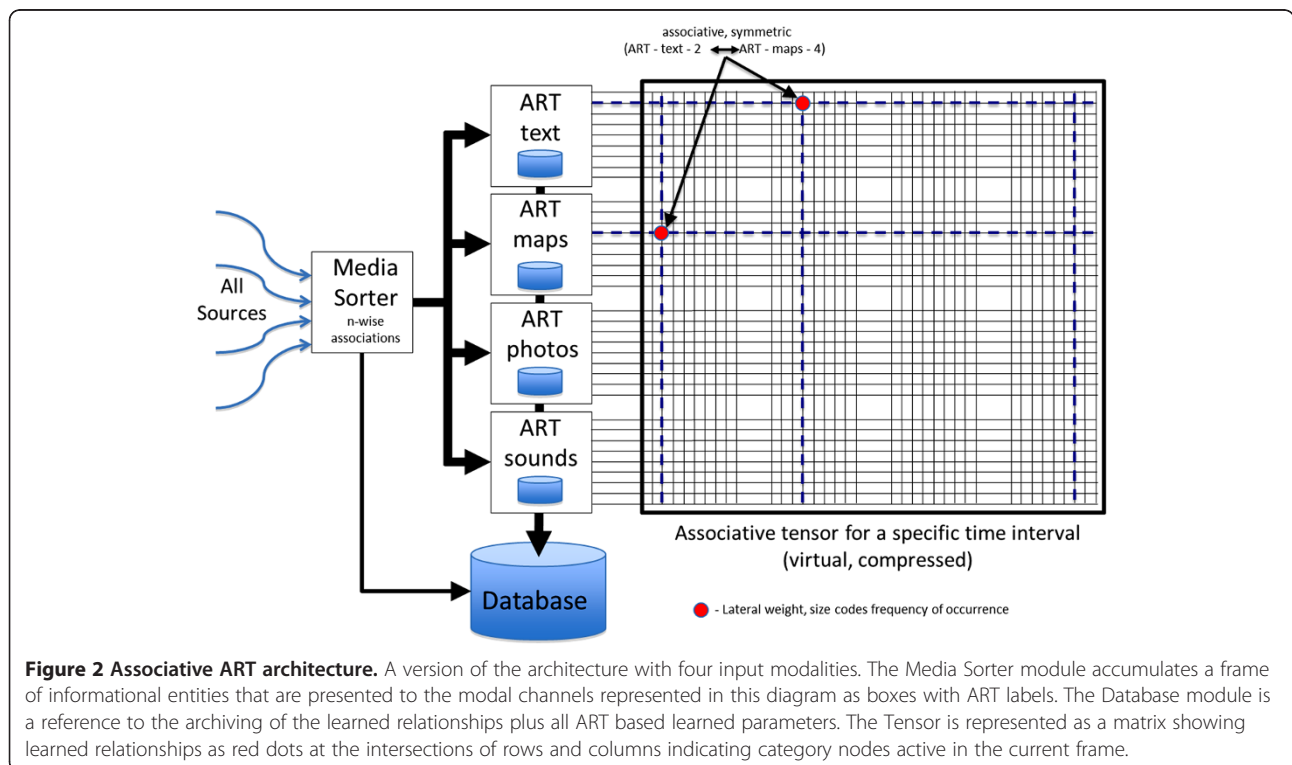
Table 1 Relation between anatomical functionality and model representation

Associative brain anatomy	Function	Model representation
Hippocampal Formation	Associates and consolidates multi-modal, event related information into long term memory	Associative-ART
Entorhinal Cortex (EC)	Aggregates sensory inputs	Fuzzy-ART modules inputs
Dentate Gyrus (DG)	Provides sparse coding of dense neuron population to enable pattern separation	Fuzzy-ART modules
Cornu Ammonis Subarea (CA3)	Acts as the auto-association network within the hippocampus	Tensor mapfield

constructed. Eventually, within the hippocampus, the DG is believed to create unique sparse encodings for unique multimodal sensory perceptions allowing it to either learn new associations or recall existing ones. Through the use of self-organizing neural networks, our architecture performs similar operations. It can detect entities that it has previously experienced and therefore reinforce existing associations, or detect novel entities necessitating a new association encoding.

Our architecture, shown in Figure 2, addresses this capability by using fuzzy-Adaptive Resonance Theory (Fuzzy ART) artificial neural network modules. Developed by Carpenter and Grossberg, the ART family of neural networks are online, unsupervised (self organizing) neural networks, which are excellent at rapid category formation [10]. The Fuzzy-ART version used in our architecture operates upon real valued inputs. Given a vector of real valued numbers as input, Fuzzy ART

performs pattern categorization through winner-take-all competition, yielding a unique category code as output. A single parameter (vigilance) regulates mathematically how similar inputs must be in order to be grouped within the same category. A vigilance parameter value of unity requires the inputs in the same group to be identical. Lowering the vigilance parameter towards zero allows for generalization such that similar, but not exactly identical, inputs may be grouped together. If no existing category is sufficiently close to capture an input, then Fuzzy ART automatically creates a new category code. In the neurophysiology, DG creates nearly unique encoding for novel inputs. Fuzzy-ART creates and maintains representative categories for inputs. Repeated presentation of previous inputs activates the same categorical representation whereas novel inputs are represented by new category encodings. Categorical activations are the basis for learning associations in our architecture.



The DG encodings in the hippocampus propagate to the CA3 region that is believed to be heavily composed of recurrent connections supporting the formation of associations. In a sense, the CA3 acts like an association “mapfield” where the simultaneous arrival of signals at CA3 neurons from DG neurons representing mixed modal entities strengthens their ability to fire in the future. Similarly, in our architecture, the activations of category codes from k unimodal Fuzzy ART modules are connected in a fully connected mapfield containing synaptic weights encoding associations among k -tuples of inputs. This association map has the structure of a k^{th} rank tensor with variable dimensions and will be henceforth referred to as the “tensor mapfield”.

Existing ART based associative neural network architectures, such as ARTMAP [11] and LAPART [12], link two ART modules using a mapfield that effectively associate category outputs from the two ARTs together. This class of architecture connects an ART module to each axis of a matrix of synaptic weights-- the intersecting grid lines of which encode a connection between the two ART modules [11]. These models are usually used for supervised learning or function approximation applications that require unidirectional many-to-one associations from the first ART to the second (see Figure 1 in [11]).

Our architecture consists of an arbitrary number of unimodal Fuzzy ART modules symmetrically connected through an association tensor mapfield to encode arbitrary associations between unimodal entities (Figure 2). As with ARTMAP and LAPART, the output category layer of each unimodal module is connected to an axis of the tensor. Unlike these architectures, the category layers of each module are buffered and connected to a mirror axis (see the top axis in Figure 2) of the tensor, thus allowing associations between entities of the same modality. All tensor elements, which can be thought of as synapses between modalities, are initialized to zero prior to learning.

During training, the system receives a sequence of data records that can contain any mixture of modal data components. Upon entry into the system, the components of the record are placed in a queue from which each unimodal component is directed to the corresponding unimodal ART module sorted by its modal type. This module performs its categorization, activating the corresponding output node and its gridline in the tensor mapfield. For each grid intersection in the tensor where there exists at least two current activations, the tensor element (modelled synapse) is strengthened. In the version of this architecture described in this paper, the synapse strength as represented by the tensor element is immediately set to unity. After learning has occurred, the active node is mirrored and buffered for the

remainder of the processing of the record, and the next modal data component is drawn from the queue and directed to the appropriate unimodal module. Through the processing of the sequence of data records, the tensor mapfield learns symmetric binary associations between pairs of unimodal module output representations.

Fuzzy ART has excellent learning properties for this type of application. Configured with its choice parameter set near zero and with the use of complement encoding on the input vector [10], this module exhibits single pass learning. That is, given a finite set of training patterns, the number of learned categories and all internal synaptic weights converges to their final values in one training epoch. A training epoch is the process of presenting each and every member of the training set to the module once and only once. During the second presentation of the training set, it is possible that individual patterns will change category membership, but this will cease in subsequent presentations. Scaling studies have shown that for higher dimensional input patterns, membership change is unlikely during the second presentation epoch. As mentioned above, the vigilance parameter for Fuzzy ART determines the ultimate number of categories learned during the first presentation. When this parameter is unity, the number of categories equals the number of unique training patterns, thus memorizing the training set. When this parameter is near zero, the number of categories will approach one, thus over generalizing over the training set. The choice of this parameter will strongly affect the dimensionality of the tensor mapfield.

Extension of trained ACORD system to scenario analysis

Once we have trained an ACORD system using k -tuples of domain-specific records, it can then be used with analyst direction to traverse the associations learned. Given a particular unimodal record of interest, such as the image of a person, vehicle or residence, or a name (text) of an individual or place, the ACORD system can help guide the analyst through the levels of transitive associations it represents. In this mode of operation there is no learning, thus when a record of a particular modality is presented to the system, it will activate the closest matching category (in its unimodal ART), which through the tensor (mapfield) will activate any (and all) matched categories of any modality which can then be traversed further in either depth or breadth first fashion. In this mode of ACORD operation, which we call re-resonance, the analyst provides feedback and direction to the reinforcement learning. This mode of operation is called re-resonance because the initial (unimodal) input record resonates with the best matching category in its modal ART, which is then allowed to generate a cascade of re-resonances determined by the associations learned in

the tensor specific to the input record presented at the start of this operation. Note that this portion of the model has not been fully implemented, but we will explain its potential in the simple example described next and return to it in the future work section below.

Experiments

Initial association study

As an introduction to the associative capability of our Associative-ART architecture we have constructed an initial image association experiment with 20 unique inputs and 22 associations among the inputs. The parametric configuration we used for each Fuzzy-ART module is β set to 1 (fast learning), a choice parameter α of 0.01, and a vigilance of 0.99. The values selected for α and β are standard choices [10]. The vigilance value

specified can be lowered as desired to allow for greater generalization of information. However, we have selected a large value to ensure accurate entity identification in a domain as sensitive as intelligence analysis. As the base case, we have set k equal to two so that the associations are pairs. While ART is capable of processing any vectorized inputs, for this experiment we have presented our architecture with images of uniformly sized circled letters and numbers as shown in Figure 3. Each row in the figure portrays an associative pairing and the column depicts the individual input which was presented to the architecture.

Although this is a generic, fictitious example, it can be conceptually compared with the type of task presented to data analysts. These analysts are tasked with processing large quantities of information and forming

Input α	Input β	Input α	Input β
1	a	8	r
a	m	s	t
b	2	s	r
3	7	t	r
3	v	5	g
7	c	g	9
c	v	c	3
3	c	k	2
3	m	2	x
2	8	v	7
6	2	z	x

Figure 3 Initial association study inputs.

associations for a variety of reasons including but not limited to knowledge discovery, discovering groups and individuals of interest, and analyzing criminal or terrorist networks. Thus, for this initial example it may be understood that numbers are representative of people and letters representative of locations respectively. As such, people may be associated with other people or locations such as businesses and addresses. Likewise, locations may be associated with people for the same reasons just stated, or they may also be associated with other locations. Locations may be associated with other locations for a variety of reasons such as representing concepts like a business partnerships or geographic hierarchies such as cities within a subsuming state or country.

Association results

Upon processing the inputs shown in Figure 3, the tensor mapfield of our architecture has encoded all of the symmetric binary associations between inputs. For this initial study, the mapfield encodes the associations between pairs of images. The tensor mapfield may be queried in order to generate an equivalent, but often more visually intuitive, association graph. The graph of this example is illustrated in Figure 4. As portrayed by this figure, although simple pairings were presented to the architecture, the net result is a more sophisticated network, which may aid an analyst in comprehending the structure of criminal networks or identify connections previously unperceived. For instance, as may be seen in Figure 4 the resultant association network is disjoint and from the 22 input pairings 3 sub networks emerge. The topological layout shown is arbitrary, but the connectivity allows for potential knowledge discovery such as by means of transitive association path analysis. As an example, individuals '1' and '7' were never presented

together in a single intelligence snippet^a, but by traversing the association graph they may be connected through a third party (individual '3') and two locations ('a' and 'm'). This example, although hypothetical, illustrates the extension of a trained ACORD system to analysis with analyst feedback to help the reinforcement learning of the most relevant transitive associations to the specific task at hand.

Degree centrality is a quantitative network analysis technique to assess the relative significance of a node within the network. Degree centrality is computed as the summation of edges incident with a node normalized by the number of nodes within the network [13]. The degree centralities for all of the nodes in this experiment are listed in Table 2 in decreasing order. This analytical technique quantitatively captures than individual '2' is the most connected entity within the association network, and conversely several other nodes such as individual '1' only have a single association. Information such as degree centrality can be used to help the analyst navigate very large trained ACORD association networks by focusing on nodes (or links) with a specific value (such as high degree centrality). Actual intelligence scenarios are typically much more complex yielding much more ambiguous association networks and requiring more sophisticated analysis techniques beyond the scope of this paper. However, as follows is a second scenario illustrating the associative capabilities of the ACORD architecture.

Benchmark comparison with K-means clustering

The K-means clustering algorithm is a widely used, popular unsupervised category formation methodology. It is an iterative method which strives to partition the given data into K clusters such that each data point is




Figure 4 Initial association study association graph.

affiliated with the cluster with the closest mean [14]. We have analysed our initial association study using K-means clustering as a benchmark comparison.

The K-means algorithm requires that the desired number of clusters, K, must be specified *a priori* whereas our architecture does not. Consequently, we have performed a parameter sweep experimenting with K values from two to seven. This range was selected based upon knowing the structure of the associations in the example. The partition results that K-means comes up with can depend upon the randomized selection of initial cluster means. For this particular dataset, as K increased, the variability in the partitioning increased as well. With three clusters (K = 3), there are only two unique partitions. These two partitions are shown in Figure 5 with the ovals representing clusters. If clusters are perceived as connected sub-graphs, then our architecture identified three clusters as shown in Figure 4. One method of formally assessing the similarity between graphs is to compute the edit distance, which is the minimum cost to transform one graph to another. The edit distances between the result given by our architecture and the two partitions yielded by K-means (A and B of Figure 5 respectively) are 9 and 10. Considering the small size of the example, these are relatively high edit distances (approximately half of the nodes in the graphs had to be

modified) expressing the poor performance by the K-means algorithm.

The clusters K-means identifies are typically not connected sub-components of the overall graph, but rather consist of several disjoint groups. This disconnectedness brings into question the validity of the clusters K-means identified since inter-cluster components may have no apparent reason to be grouped together. Furthermore, by examining the individual nodes to try and identify any sort of semantic coherence within a cluster one may identify the same nodes are frequently located within different clusters. This is a consequence of K-means errantly misinterpreting repeated presentation of the same data element.

Real world network example

To analyse our architecture with a more realistic data set, we used an actual terrorist scenario. However, the events pertaining to the scenario were anonymized, stripping out the actual names, places, and events. These were replaced with generic descriptions (for example Villain1, Villain2). This scenario focuses upon the central character Villain1 and its known associations. The scenario was designed to investigate Villain1's known criminal network and reveal potential interconnections between Villain1 and criminals directly linked to a terrorist act. The intent was to embody traits associated with real terrorist networks as opposed to artificial network types that may or may not be realistic. While the scenario portrayed in this example is relatively simple compared to a complete scenario, it is able to demonstrate this architecture's ability to operate upon a much larger scenario.

The parametric configuration used for each Fuzzy-ART module is β set to 1 (fast learning), a choice parameter α of 0.01, and a vigilance of 0.99. The values selected for α and β are standard choices [10]. The vigilance value specified can be lowered as desired to allow for greater generalization of information. However, we have selected a large value to ensure accurate entity identification in a domain as sensitive as intelligence analysis.

As a more sophisticated example, rather than constraining the inputs to be simple pairs the *k*-tuple inputs varied in size from 2 to 4 entities being presented to the architecture simultaneously. Overall, this example was comprised of 179 tuples constructed from 189 unique inputs.

Real world network results

As noted with the previous example, the tensor mapfield yielded from processing the input tuples encodes the binary associations between inputs. Also, the tensor mapfield may be queried to generate an association

Table 2 Degree centrality measures for initial association experiment

Input	Degree centrality
2	0.263158
3	0.210526
r	0.157895
c	0.157895
v	0.157895
7	0.157895
x	0.105263
g	0.105263
t	0.105263
s	0.105263
8	0.105263
m	0.105263
a	0.105263
z	0.052632
k	0.052632
9	0.052632
5	0.052632
6	0.052632
b	0.052632
1	0.052632

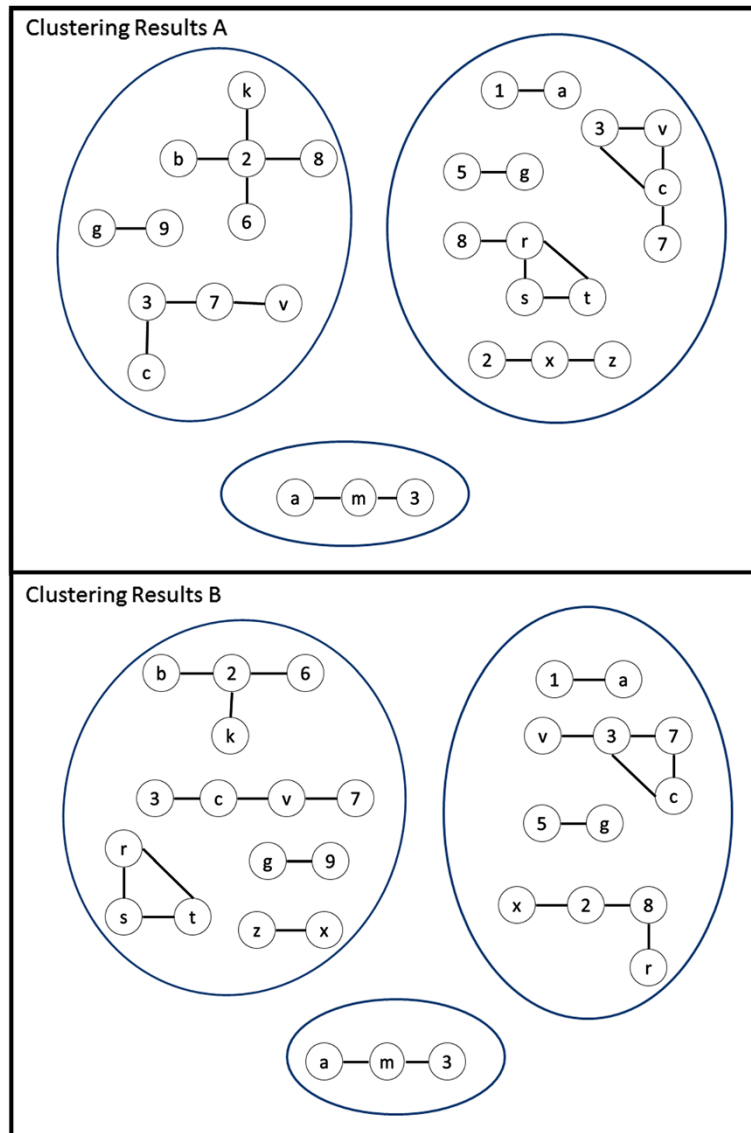


Figure 5 K-means benchmark comparison association graphs.

graph. The association graph of this real world example is captured in Figure 6. Just as the input tuples are more complicated both in terms of quantity and length in this real world example, the resulting association graph is much richer. In this case, the graph shown is nearly fully connected, but that is a consequence of the highly focused nature of the input data, whereas in a general intelligence scenario it is also possible that an analyst may receive a barrage of unrelated information. Despite the indecipherable appearance of this association graph, it provides a means for an analyst to investigate individual associations and facilitates a larger overall understanding of the mass quantity of information. For example, Figure 7 illustrates a subset of the overall scenario, focusing upon Villain1 and highlighting all of the

interconnections between people, places, and events associated with the intelligence data. Not only does it allow for connections between crimes and specific individuals to be identified, but all the additional associations allow an analyst to attain a greater understanding of the overall scenario, such as perceiving the structure of the larger criminal network carrying out a specific event, as well as how the criminal network operates.

Conclusions and future work

In this paper we first presented an artificial neural network computational architecture with functionality inspired by the neural processes of hippocampus. Specifically, this architecture was based upon the DG and CA3 regions of hippocampus as a means to learn associations

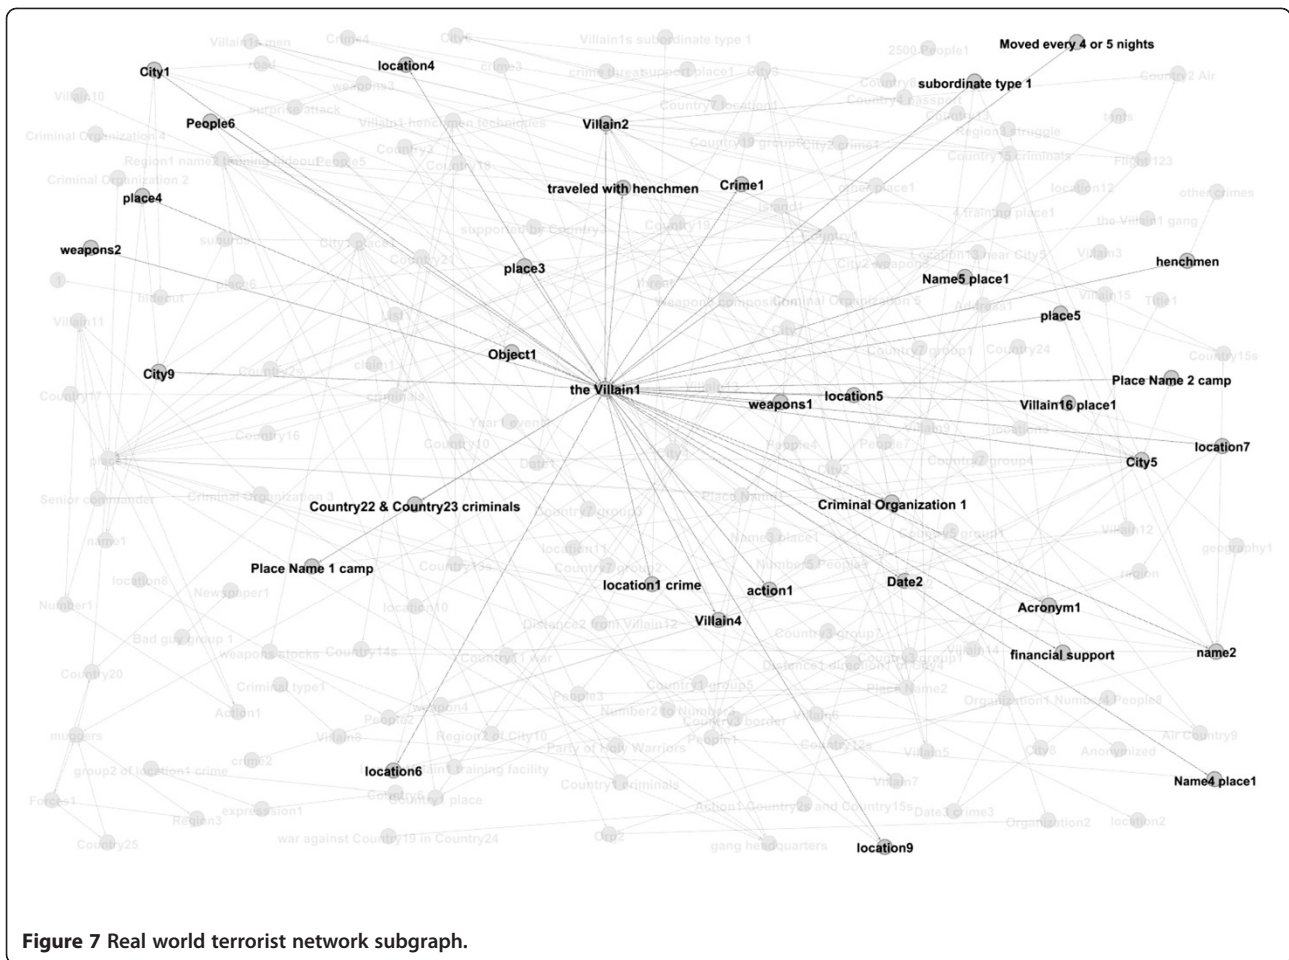


Figure 7 Real world terrorist network subgraph.

trained fully online. If meaningful insights are known about the specific problem domain, performance improvements may be possible by operating in a supervised learning mode. Depending upon the particular application, architecture modifications, such as those described above, could provide great potential for enhanced, further processing, as well as addressing episodic or sequential data. In addition, the resonance mode of operation has the potential to offer semi-automated (possibly even automated) generation of higher order associations, such as transitive chains. The integration of network metrics, such as degree centrality, into the process of generating higher order associations could help focus the analyst's effort on domain areas of metric interest.

Together, these advancements are intended to provide the national security community with a next-generation knowledge discovery system that associates relevant information across various source modalities. As with many efforts, our goal is to enable analysts to more effectively, and more timely, connect the dots to increase the probability of detecting 9/11-type of events before they are carried out. Not effectively connecting the dots has been seen as a failure of pre-9/11 analysis [16]. Our approach to this effort is to better understand and model the pre-conscious,

associative mechanisms of the mammalian brain, albeit to much simpler degree, in support of rapid decision-making for security-related contexts. While our approach may be considered unconventional compared to most efforts, we believe replicating specific aspects of the brain has the potential to ultimately produce advancements in knowledge discovery that cannot be achieved through current means. We believe the neuromorphic advancements, along with advancements in more conventional, statistically based data filtering have already produced promising results. Ultimately, an appropriate mixture of neuromorphic and statistically based approaches should, in effect, shore up weaknesses in each approach to produce a knowledge discovery system that can more effectively associate relevant information.

Endnotes

^a An intelligence snippet contains specific information concerning an event (or events) related to national or international law enforcement, and this information is presented to the ACORD system as a record (or k-tuple) of events.

Competing interests

The authors declare that they have no competing interests.

Authors' contributions

CMV is the technical lead and participated in the research and development of the model, and led the execution and analysis of experiments on the model. SJV helped in design and implementation of the proposed model. MLB is the project manager and participated in the sequence alignment. SET participated in the design of the study and performed the statistical analysis. ID provided background on hippocampal neuroanatomy and human factors engineering guidance on potential applications to all-source analytical tradecraft. TPC conceived of the study, and participated in its design and coordination. All authors read and approved the final manuscript.

Acknowledgements

We would like to thank Jonathan McClain for his work in developing essential statistical components of the ACORD capability, as well as Wendy Shaneyfelt for her work in structuring multi-model information used for testing the ACORD capability. This research was possible in part by LDRD program support from Sandia National Laboratories. Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Author details

¹Sandia National Laboratories, Albuquerque, NM, USA. ²Department of Electrical and Computer Engineering, University of New Mexico, Albuquerque, NM, USA.

Received: 16 December 2011 Accepted: 7 September 2012

Published: 6 November 2012

References

1. T. Quiggin, *Connecting the Dots: President Obama*, [<http://globalbrief.ca/tomquiggin/2010/01/12/connecting-the-dots-president-obama/>]
2. H. Eichenbaum, The hippocampus and declarative memory: cognitive mechanisms and neural codes. *Behav Brain Res* **127**, 199–207 (2001)
3. N.J. Cohen, J. Ryan, C. Hunt, L. Romine, Hippocampal system and declarative (relational) memory: summarizing the data from functional neuroimaging studies. *Hippocampus* **9**, 83–98 (1999)
4. R.D. Burwell, M.P. Witter, D.G. Amaral, Perirhinal and postrhinal cortices of the rat: a review of the neuroanatomical literature and comparison with findings from the monkey brain. *Hippocampus* **5**, 390–408 (1995)
5. W.A. Suzuki, D.G. Amaral, Perirhinal and parahippocampal cortices of the macaque monkey: cortical afferents. *J Comp Neurol* **350**, 497–533 (1994)
6. W. Suzuki, H. Eichenbaum, The neurophysiology of memory. *Annals of the NY Academy of Sciences* **911**, 175–191 (2000)
7. D. Amaral, P. Lavenex, Ch3. Hippocampal Neuroanatomy, in *The Hippocampus Book*, ed. by P. Anderson, R. Morris, D. Amaral, T. Bliss, J. O'Keefe (Oxford University Press, New York, 2006). ISBN 978-0-19-51002703
8. S. Leutgeb et al., Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science* **309**, 619–623 (2005)
9. H. Eichenbaum, N. Cohen, *From Conditioning to Conscious Recollection: Memory Systems of the Brain* (Oxford University Press, Oxford, 2001)
10. G.A. Carpenter, S. Grossberg, D.B. Rosen, FuzzyART: fast stable learning and categorization of analog patterns by an adaptive resonance system. *Neural Netw* **4**, 759–771 (1991)
11. G.A. Carpenter, S. Grossberg, N. Markuzon, J.H. Reynold, D.B. Rosen, Fuzzy ARTMAP: a neural network architecture for incremental supervised learning of analog multidimensional maps. *IEEE Trans Neural Netw* **3**, 698–713 (1992)
12. M.J. Healy, T.P. Caudell, Acquiring rule sets as a product of learning in the logical neural architecture LAPART. *IEEE Trans Neural Netw* **8**, 461–474 (1997)
13. P. Bonacich, Power and Centrality: A Family of Measures. *Am J Sociol* **5**, 1170–1182 (1987)
14. R.O. Duda, P. Hart, D. Stork, *Pattern Classification* (Wiley-Interscience, New York, 2001)
15. S.J. Verzi, G.H. Heileman, M. Georgiopoulos, Boosted ARTMAP: modifications to Fuzzy ARTMAP motivated by boosting theory. *Neural Netw* **19**(4), 446–468 (2006)
16. 9/11Chair, *Attack Was Preventable*, [http://www.cbsnews.com/2100-18563_162-589137.html]

doi:10.1186/2190-8532-1-20

Cite this article as: Vineyard et al.: A multi-modal network architecture for knowledge discovery. *Security Informatics* 2012 **1**:20.

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- Convenient online submission
- Rigorous peer review
- Immediate publication on acceptance
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► springeropen.com