

A semi-holographic hyperdimensional representation system for hardware-friendly cognitive computing

A. Serb*, I. Kobzyev†, J. Wang*, T. Prodromakis*

*Zepler Institute, University of Southampton, SO17 1BJ, UK

† David R. Cheriton School of Computer Science, University of Waterloo, N2L 3G1, Canada

Corresponding author email: A.Serb@soton.ac.uk

Abstract—One of the main, long-term objectives of artificial intelligence is the creation of thinking machines. To that end, substantial effort has been placed into designing cognitive systems; i.e. systems that can manipulate semantic-level information. A substantial part of that effort is oriented towards designing the mathematical machinery underlying cognition in a way that is very efficiently implementable in hardware. In this work we propose a ‘semi-holographic’ representation system that can be implemented in hardware using only multiplexing and addition operations, thus avoiding the need for expensive multiplication. The resulting architecture can be readily constructed recycling standard microprocessor elements and ;Something about hardware performance;. Our proposed ‘cognitive processing unit’ (CoPU) is intended as just one (albeit crucial) part of much larger cognitive systems where artificial neural networks of all kinds and associative memories work in concord to give rise to intelligence.

I. INTRODUCTION

The explosive scale of research output and investment in the field of artificial intelligence (AI) and machine learning (ML) testify to the tremendous impact of the field to the world. Thus far this has manifested itself as a mass-scale proliferation of artificial neural network-based (ANN) algorithms for data classification. This covers multiple data modalities such as most prominently images [1] and speech/sound [2], and relies on a number of standard, popular ANN architectures, most notably multi-layer perceptrons [3], recurrent NNs (in particular, LSTM [4] and GRU [5]) and convolutional NNs [6] amongst many others [7], [8].

Thus far the vast majority of market-relevant ANN-based systems belong to the domain of statistical learning, i.e. perform tasks which can be generally reduced to some sort of pattern recognition and interpolation (in time, space, etc.). This, though demonstrably useful, is akin to memorising every answer to every question plus some ability to cope with uncertainty. In contrast, higher level intelligence must be able to support fluid reasoning and syntactic generalisation, i.e. applying previous knowledge/experience to solve novel problems. This requires the packaging of classified information generated by traditional ANNs into higher level variables (which we may call ‘semantic objects’), which can then be fluently manipulated at that higher level of abstraction. A number of cognitive architectures have been proposed to perform such post processing, most notably the ACT-R architecture [9] and the semantic pointer architecture (SPA) [10], which is an effort to manipulate symbols using neuron-based implementations.

Handling the complex interactions/operations between semantic objects requires both orderly semantic object representations and machinery to carry out useful object manipulation operations. Hyperdimensional vector-based representation systems [11] have emerged as the de facto standard approach and are employed in both the SPA and ACT-R. Their mathematical machinery typically includes generalised vector addition (combine two vectors in such way that the result is as similar to both operands as possible), vector binding (combine two vectors in such way that the result is as dissimilar to both operands as possible) and normalisation (scale vector elements so that overall vector magnitude remains constant). These operations may be instantiated in holographic (all operands and results have fixed, common length) or non-holographic manners. Non-holographic systems have employed convolution [12] or tensor products [13] as binding. Holographic approaches have used circular convolution [11] and element-wise XOR [14]. Meanwhile, element-wise addition tends to remain the vector addition operation of choice across the board.

Finally, whichever computational methodology is adopted for cognitive computing must be implementable in hardware with extremely high power efficiency in order to realise its full potential for practical impact. This is the objective pursued by a number of accelerator architectures spanning from limited precision analogue neuron-based circuits [15], through analogue/digital mixtures [16] to fully analogue chips seeking to emulate the diffusive kinetics of real synapses [17]. More recently memristor-based architectures has also emerged [18].

In this work, we summarise an existing, abstract mathematical structure for carrying out semantic object manipulation computations and propose an alternative, hardware-friendly instantiation. Our approach uses vector concatenation and modular addition as its fundamental operations (in contrast to the more typical element-wise vector addition and matrix-vector multiplication respectively). Crucially, the chosen set of operations no longer forms a holographic representation system. This trades away some ‘expressivity’ (ability to form semantic object expressions within limited resources) in exchange for compression: Unlike holographic representations semantic object vector length depends on its information content. Furthermore, the proposed system avoids use of multiplication completely, thus allowing for both fast and efficient processing in hardware (avoiding both expensive multipliers

and relatively slow spiking systems). Finally, we illustrate how the proposed system can be easily mapped onto a simple vector processing unit and provide some preliminary, expected performance metrics based on simulations in Cadence on a 65nm technology.

II. MATHEMATICAL FOUNDATIONS AND MOTIVATION

Generalising the series of work on models of associative memory, many of them inspired from the world of optics [11], [13], [14], [19]–[24], one may inspect the most abstract algebraic formulation of it. All we need is a commutative ring R with a distance metric dist .

In order to give this mathematical machinery sufficient power to describe cognitive tasks, one must initially specify the ring operations and impose some restrictions on them. The primary operation (addition, denoted by $+$), enables superposition¹; that is the combination of two elements in such way that the result is equidistant from its operands under the metric dist (i.e. for $a, b \in R$, one has $\text{dist}(a + b, a) \approx \text{dist}(a + b, b)$). The secondary operation (multiplication, denoted by $*$) enables binding; that is the combination of two elements in such manner that the result is ideally completely different from both operands. Next, one needs to store a (finite) set of elements of R including both invertible elements, which we call ‘pointers’ (or ‘roles’), and not necessarily invertible ones, which we call ‘fillers’.

Let us now give an example of how such mathematical machinery may give rise to simple cognition. Assume that we have a ring R with the distance dist and the operations, satisfying the desired properties. Also assume that we fixed five elements of R : “*obj*” and “*col*” are invertible, “*red*”, “*green*” and “*car*” are any elements. We can now construct a new element: $s = \text{obj} * \text{car} + \text{col} * \text{red}$, which can be interpreted as a semantic object “red car”. Now one can ask: what colour is this car? The answer can be accessed by performing an algebraic operation: $\text{col}^{-1} * s = \text{red} + \text{col}^{-1} * \text{obj} * \text{car}$. Then if the term $\text{col}^{-1} * \text{obj} * \text{car}$ is either close to zero or in some other way does not interfere with the computation of dist , the stored memory element closest to the result of the query is *red*. Mathematically, the query is $\text{argmin}_{r \in R} (\text{dist}(\text{col}^{-1} * s, r)) = \text{red}$. Thus, we observe that the mathematical foundation of AI is underpinned by a solid computational/information processing foundation whose functionality must be preserved in any proposed alternative representation system, even if not necessarily via a distance-equipped commutative ring.

The classical realisation of the commutative ring-based cognition principle is the holographic-like memory [11]. In this case R is defined as follows: the set is a collection of n -dimensional real vectors (n -vectors) \mathbb{R}^n . The ring operations are the element-wise addition and circular convolution. The distance metric is the simple Euclidean. To define a pointer or a filler one just needs to independently sample each entry of the vector from the normal distribution $\mathcal{N}(0, 1/n)$.

¹In the literature this is typically called ‘chunking’, but this term by itself does not allude strongly enough to the desired simultaneous similarity between operands and result.




Fig. 1. Summary of key terms used throughout this text.

Finally, the operations of the system must be ideally implementable in hardware in a way that minimises power and area requirements. In practice this means that the fundamental superposition and binding operations must rely on energetically cheap building block operations such as thresholding (an inverter), shifts (flip-flop chain), addition (sum of currents on a wire or digital adder) or possibly analogue multiplication (memristor + switch) [25]. Implementation details will ultimately determine the actual cost of each operation. The main approaches so far either use too many multiply-accumulate (MAC) operations (circular convolution-based binding from [11] requires $\approx n^2$ MACs/binding), or are applicable only to binary vectors (radix $k = 2$) [14].

III. PROPOSED SEMI-HOLOGRAPHIC REPRESENTATION SYSTEM

In this section we provide an intuitive overview followed by a rigorous mathematical explanation of the proposed architecture interwoven with pointers on how our design decisions aim towards hardware efficiency. Overall, in order to achieve a more hardware-friendly cognitive algebra realisation we trade away some of the mathematical simplicity from the previous section for implementability. The algebraic structure we are using for cognition is no longer a ring, but a rather exotic construction. It consists of an underlying set and two binary operations (superposition, binding).

A. Building a set of semantic objects

In our proposed system, the set of semantic objects is perhaps best understood in terms of two subsets: i) Fixed-length ‘base items’, each consisting of y integer elements in the range $[0, p - 1]$. The choices of p and y link to desired memory capacity, i.e. the number of semantic objects the system is capable of representing reliably - see section IV). ii) Variable-length ‘item chains’ consisting of multiple concatenated base elements. The maximum length for chains is d base items for a total of n numerical elements, where d is determined by the hardware design² and affects the capacity of the system to hold/express multiple basic items at the same time. The number of base items in a chain is defined as the rank of the chain. The terminology is summarised in figure 1

Some observations about our implementation: i) Base items are generally intended for encoding the fundamental vocabu-

²However, note that much akin to standard computers being able to process numbers more than 32 or 64 bits, there is no reason why chains longer than d base items cannot be processed using similar techniques.

TABLE I

SUMMARY OF RELATIONS BETWEEN BASIC MATHEMATICAL OBJECTS USED IN THIS WORK. AS AN EXAMPLE OF HOW TO READ THE TABLE, THE TOP LEFT ENTRY STATES THAT: "IN EVERY ELEMENT THERE ARE $p = 2^l$ STATES". THE TERM 'CHAIN' REFERS TO A MAXIMUM SIZE CHAIN. THE TERM 'ITEM' REFERS TO BASE ITEMS. ALL PARAMETERS ARE INTEGERS.

		In every...		
		Element	Item	Chain
There are this many...	States	$p = 2^l$	2^{l+z}	2^{l+z+m}
	Elements	1	$y = 2^z$	$n = 2^{z+m}$
	Items	N/A	1	$d = 2^m$

lary items of the system (e.g. 'red', 'apple', 'colour') and possible bindings, including the classical 'pointer-filler' pairings (e.g. 'colour'*'red': the value of the colour 'attribute' is 'red'). In contrast, chains are intended for simultaneously holding (superpositions of) multiple base items in memory (e.g. composite descriptions of objects such as: *colour * red + object * apple* ('a red apple'), or collections of unrelated items such as: *shape * circle + shape * square* ('a circle and a square'). The order in which the superposed items are kept in memory does not bear any functional significance; for the purposes of our system items are either present or absent from a chain. Cognitive systems that are order- or even position-dependent can be, of course, conceived; all that is necessary is for each item to have some mechanism (e.g. a position indicator) for marking its location within a chain. ii) Setting p, y, d, n as powers of 2 offers the attribute of naturally advantageous implementation in digital hardware. This is the approach we choose in this work, as shown in table I. The choice of p is not necessarily obvious as what constitutes a 'good' choice of p will depend on the specific implementations of superposition and binding. iii) Any chain can be zero-padded until it forms a maximum-length chain.

Mathematically the above can be described as follows: Fix natural numbers p and y as above. Then the set of base items is a group $B = (\mathbb{Z}/p)^y$ (under element-wise mod p summation). The way to form item chains is by executing a direct product of copies of B . Then we say that any element of $B^r = \prod_{i=1}^r B$ has rank r . The chain of maximal length will be an element of B^n , and $n = d \cdot y$.

B. Superposition and binding

Next, we define our set of basic operations. The superposition operation '+' is defined as follows: If a and b are semantic objects, then:

$$a + b = (a, b) \quad (1)$$

which is a standard direct sum. The result contains both a and b operands preserved completely intact. This can be contrasted with superposition implemented as regular element-wise summation, where each operand is 'blurred' and merged into the result. Superpositions of semantic objects whose

combined ranks exceed d are not allowed³.

Formally speaking, given $a \in B^{d_1}$, $b \in B^{d_2}$, the superposition $a + b$ is just an element (a, b) in a direct products of the groups $B^{d_1+d_2}$. If $d_1 + d_2 > n$, the operation is not defined.

Next, the binding operation '*' is defined as a variant of a tensor product between semantic objects where the individual pairings are subsequently subjected to element-wise addition modulo p . Mathematically, for given natural numbers d_1 and d_2 , such that $d_1 \cdot d_2 < n$, one can define the binding operation $*$: $B^{d_1} \times B^{d_2} \rightarrow B^{d_1 \cdot d_2}$ by the formula:

$$(a_1, a_2, \dots, a_{d_1}) * (b_1, b_2, \dots, b_{d_2}) = (a_1 + b_1, a_2 + b_1, \dots, a_{d_1} + b_1, a_1 + b_2, \dots, a_{d_1} + b_{d_2}) \quad (2)$$

where $+$ is the group operation in $B = (\mathbb{Z}/p)^y$. One can see that any element from B (base item) is invertible under the binding⁴.

One should notice that modular addition is losslessly reversible: we may indefinitely add and subtract n-vectors, and therefore can perfectly extract any individual term from any multi-term binding combination if we bind with the modulo p summation inverses of all other terms. We also remark that within the context of the order-independence property any binding of chains with length greater than 1 item is effectively a convenient shorthand for describing multiple base item bindings and adds no further computational (or indeed semantic) value.

We conclude this section by highlighting that our superposition operation is not length-preserving but our binding is when one of the operands consists of 1 basic item. Thus we describe our system as semi-holographic. Interestingly, this is the opposite of the classical convolution-based system from [12], where the binding operation is not length-preserving but superposition (element-wise average) is.

C. Similarity metric

Let us define a distance. First, we use a "circular distance" on \mathbb{Z}/p : for $a \in \mathbb{Z}/p$, one has $\text{dist}_o(a, 0) = \min(|a|, |p - a|)$, here we also denoted by $a = a + 0 \cdot p$ the corresponding representative in \mathbb{Z} . For example, for $4 \in \mathbb{Z}/5$, $\text{dist}(4, 0) = 1$. Analogously one defines a distance $\text{dist}_o(a, b)$ for any $a, b \in \mathbb{Z}/p$ as $\min(|b - a|, p - |b - a|)$. For two vectors $a, b \in B$, one defines the distance as:

$$\text{dist}_B(a, b) = \sum_{i=1}^y \text{dist}_o(a_i, b_i) \quad (3)$$

³In a practical hardware implementation we would either: i) raise an exception and forbid the operation, ii) truncate the result to size d and raise a warning flag or iii) raise a flag and trigger a software sequence (program) designed to handle overlength chains - equivalent to branching to a different subroutine in Assembly language.

⁴This is where the consequences of our choice of p become apparent: Consider the item consisting of all elements equal to $p/2$. Binding this item to itself twice results in the original item. This becomes problematic if we wish to define a sequence of items as a succession of bindings, e.g. if we define the semantic object '2' as $1 * 1$, '3' as $2 * 1$ etc. If, on the other hand p is prime, then for any integer $x \neq 0$ there is a guarantee that if $(k \cdot x) \bmod p = x$, the next greatest solution after $k = 1$ is $k = p + 1$; this may allow the construction of longer, non-tautological self-bindings vs. non-prime p systems. Morale: the choice of p is not always obvious.

For $a \in B$ and $b = (b_1, \dots, b_r) \in B^r$ we define:

$$\text{dist}(a, b) = \min_i \text{dist}_B(a, b_i) \quad (4)$$

One can note that $\text{dist}(a + b, a) = 0$ for any $a \in B$.

D. Basic properties

In terms of fundamental mathematical properties: The superposition operation is not closed in general, but it acts as closed when our restriction on the sum of the ranks of the operands is met. It is associative but not commutative. It has an identity element (the empty string), but no inverse operation as such.

The binding operation is not closed, but acts as closed when the restriction on the product of the ranks of the operands is met. This is always the case when one of the operands is a basic item, i.e. $a \in B$. If a is a basic item, then for any $b \in B^d$, we have the commutativity: $a * b = b * a$. If $a \in B^{d_1}, b \in B^{d_2}, c \in B^{d_3}$, and at least one of $d_i = 1$, then we have associativity: $(a * b) * c = a * (b * c)$. In general it is neither associative nor commutative, however, modulo permutation group on basic item components, it has those properties.

Finally, one has a distributivity in case of a basic item: $a \in B, b \in B^{d_1}, c \in B^{d_2}$, then $a * (b + c) = a * b + a * c$. In general, as above this property no longer holds (unless we don't care about the order of terms and factorise by the action of permutation group).

The identity element is the zero element of B . All basic elements are invertible under binding.

These properties form a good start for building a cognitive system.

IV. CAPACITY

In terms of higher level properties, a key metric is memory capacity: the maximum number of basic elements storable given some minimum upper bound for memory recall reliability. Each rank 1 semantic object (base item), the smallest type of independent semantic objects, must be uniquely identifiable. As a result, there can be no more than $Q = p^y$ basic memories in total without guaranteeing at least one ambiguous recall, i.e. Q is the maximum memory capacity⁵. However, an additional sparsity requirement is necessary in order to guarantee that the system is capable of unambiguously answering queries. Returning to the example from section II, in order for the term $col^{-1} * obj * car$ to be culled from any semantic pointer or filler from our vocabulary it should not coincide with a valid object from the fixed fundamental vocabulary. In order to achieve that, we may impose that our memory safely stores only up to Q_s vocabulary objects, where $s \in \mathbb{R}$ is the desired sparsity factor, and the following formula holds:

$$Q_s = \frac{Q}{s}. \quad (5)$$

⁵It is very expedient if any semantic object that needs to be stored for quick recall is constructed as a basic object, not in the least because binding any operand with a basic object does not lengthen the operand. For that reason we only consider basic elements when computing memory capacity.

A lower bound for s is given by calculating the number of basic items J that the system can generate given a set of Q_s vocabulary items and allowed complexity. These will all need to be accommodated unambiguously for guaranteeing reliable recall. In our proposed system the only operation that can generate basic items from combinations of vocabulary items is the binding operation. Therefore for Q_s vocabulary items we obtain $\frac{Q_s^2}{2}$ derived items arising from all the possible unordered (to account for the commutativity) pairwise bindings. This rises to $\frac{Q_s^\gamma}{\gamma!}$ for exactly γ allowed bindings, and in general the system can generate:

$$J = \sum_{i=0}^{\Gamma} \frac{Q_s^i}{i!} \approx \frac{Q_s^\Gamma}{\Gamma!}, \text{ for } \frac{Q_s}{\Gamma} \gg 1 \quad (6)$$

basic items, if we allow anything between 0 and Γ bindings in total. Ideally we want to account for all possible basic items from the fundamental vocabulary via bindings, so $J = Q (= p^y)$, and therefore we can transform equation 6 into:

$$Q_s \approx \sqrt[\Gamma]{\Gamma!} \cdot p^{\frac{y}{\Gamma}} \quad (7)$$

revealing how expressivity is traded against capacity, at least in the absence of any further allowances to combat possible uncertainty in the encoding, decoding or recall of semantic objects. Whether this boundary can be reached in practice requires further study as the particular encodings of each basic item will determine whether specific bindings coincide with pre-learned vocabulary or other bindings. Let us observe that the more binding is allowed in the system, the less fundamental vocabulary it can memorize (hint: $\lim_{x \rightarrow \infty} \sqrt[x]{x!} = \frac{x}{e}$). This is an example of a trade-off between capacity and complexity.

Example: if we choose $p = 16, y = 128$ and we allow the system to have at most $\Gamma = 20$ bindings, then the upper bound on the length of the core dictionary we can encode is 422 million items.

V. ADDITIONAL SEMANTIC OBJECT MANIPULATIONS

In order to complete the description of the proposed system we need to cover three further issues: i) How does the system cope with uncertainty? ii) Since the system is semi-holographic how does the system map multi-item chains to single base items when necessary? In this work we provide some cursory answers as these questions merit substantially deeper study in the own right.

Dealing with uncertainty: The implementation of de-noising will strongly depend on the form of the uncertainty present in the system. We may define uncertainty as a probability distribution that encodes how likely it is to obtain semantic object x' when in fact the ground truth is x . For example, if the probability density only depends on the 'circular distance' (eq. 3) between the x and x' objects⁶ We may use an adaptation of element-wise average for de-noising. The average is computed as the mid-point along the geodesic. In particular, for $a \in \mathbb{Z}/p$

⁶This alludes to the radial basis functions (RBFs) used in radial basis neurons [26].

let also denote by the same symbols $a = a + 0 \cdot p$ its representative in \mathbb{Z} . Also denote by $\Delta = \text{ceil}(\frac{\text{dist}_0(a,b)}{2})$. For $a, b \in \mathbb{Z}/p$, if $|b - a| \leq p - |b - a|$, pick a smaller representative with respect to the standard ordering on \mathbb{Z} (say, it is a), then $\text{avg}(a, b) = a + \Delta$. If the alternative inequality happens, pick the greater representative (say, it is b), then $\text{avg}(a, b) = (b + \Delta) \bmod p$. In general, for items $a, b \in B$, we define the average as the element-wise average.

To this we add the following observations: i) The purpose of the de-noising average is to reconcile multiple, corrupted versions of a single semantic object vector, not combine different vectors into new semantic objects (i.e. a_i is expected to be reasonably close to b_i most of the time). Nevertheless, when used with radically different semantic objects as inputs, it is inescapable to observe that the operation acts very similarly to binding. The effects of using a binding-like operation for denoising (a task usually handled by superposition) are an interesting subject for further study. ii) Different uncertainty descriptors (probability distribution functions) may lend themselves to different de-noising strategies. So will different metrics. iii) Even with fixed underlying probability distribution assumptions, de-noising may be carried out using multiple alternative strategies. Examples applicable to our assumptions would be majority voting (select element-wise mode instead of mean - works best for large number of input sample terms) or median selection.

Compressing long chains into basic items: Ideally any cognitive system should be able to take any expression and collapse it into a new memory that can be stored, recalled and used with the facileness that basic items enjoy. In our case this requires compressing chains into the size of a basic item. In principle, any compression algorithm will suffice. Examples could be applying genetic algorithm-like methods [27] on the items of a chain or combining said items using any multiplication (e.g., circular convolution etc).

We conclude by remarking that the operation of creating a new semantic object can be reasonably expected to be executed orders of magnitude less frequently than any of the other operations. As such, it is possible to dedicate hardware that is both more complex (luxury of using relatively heavy computation) and more remotely located from the core of the semantic object processor (luxury of preventing the layout footprint of the semantic object generator from impacting the layout efficiency of the processor core).

VI. HARDWARE IMPLEMENTATION

In this section we examine how the mathematical machinery can be mapped onto a hardware module which we call the ‘Cognitive Processing Unit’ (CoPU). The system receives chains as input operands and generates new chains at its output after executing the requested superposition and/or binding operations. The CoPU is based on a common block-level design blueprint which can then be instantiated as specific CoPU designs. It is at the point of instantiating a particular CoPU design that the values of key parameters p, y, d are decided upon.

A. Hardware system design strategy

The proposed holographic representation machinery can be implemented as a fully digital system in a very straightforward manner. The underlying set will be implicitly determined by the bit-width used. The inverses of each n -vector element under element-wise modular addition are simply their 2’s complements. Full representation of any semantic object can therefore consist of d , $\log_2 p$ -bit words, plus x flag bits for tracking the number of items in any given chain.

The superposition operation can be handled by the hardware as ‘APPEND’ operations (akin to linked lists); the system need only know the operands and the state of their flag bits. In practice this would be implemented as d ‘SELECT’ operations, which directly map onto a simple $l \cdot n$ -width⁷ multiplexer/demultiplexer (MUX/DEMUX) pair. A small digital circuit determines the appropriate, successive configurations of the MUX/DEMUX structure depending on the flag bits of the operands. Finally, the circuit also sets the flag bits of the resulting chain. The block diagram of this system is shown in Figure [FIG]. The hardware-level complexity of our proposed system can be contrasted with the standard element-wise addition approach, which requires n times $z = \text{ceil}(l)$ -level ‘ADD’ operations (cost: n , z -bit adders, or one time-shared z -size adder or valid trade-off solutions in between).

The binding operation can be carried out by n element-wise addition/subtractions (ADD/SUB), implementable as n , z -bit ADD/SUB modules. Because of the modular arithmetic rules overflow bits are simply ignored. The ADDSUB terminal of each module can directly convert one of the operands into its 2’s complement inverse as is standard. This is illustrated in Figure [FIG]. The complexity of (a maximum of) n , z -bit additions can be contrasted to the computational cost of circular convolution, which would involve n^2 multiplication and $n \cdot (n - 1)$ additions ($= n \cdot (n - 1)$ MACs + n multiplications). On top of this, the additional hardware cost of shifting a chosen operand of the circular convolution n times in its entirety must also be considered.

Naturally, alternative hardware implementations are also possible. This might include fully analogue ones, e.g. using analogue multiplexers for superposition and current-steering-based binding [28]. Alternatively it might include ‘packet’-based ones where chains are packaged into e.g. TCP-like (Transmission Control Protocol) packets and communicated across an internet-like router structure. Each packet could contain a header detailing the number of items within the packet and a payload, a technique similar to the protocol used in neuromorphic systems communications over the internet [29]. The proposed implementation is chosen because it naturally maps onto easily synthesisable digital hardware. The most efficient implementation technique in any given system, however, will naturally depend on the rest of the system, e.g. on whether the broader environment operates in mainly analogue or digital.

A block diagram of a full CoPU is shown in Figure [FIG].

⁷ n ‘bundles’ of l binary lines.

B. Example instantiation of CoPU and performance evaluation

-Concept of introducing hardware-relevant performance metrics such as number and type of ops. needed to carry out bindings.

VII. DISCUSSION

The starting point of this work is the observation that any system consisting of a length n vector with p states per element (corresponding to some fixed number of digital signal lines) can only represent p^n uniquely identifiable vectors. This is effectively a hardware resource constraint and imposes a number of trade-offs warranting design decisions.

Trade-off 1 - expressivity vs. capacity: In the classical holographic representation systems all semantic object vectors are of equal length no matter how many times semantic objects are combined together through superposition or binding. By contrast, in our proposed system some objects will be base items and others will be chains of various lengths. This introduces some constraints into which combinations of semantic objects are allowable, yet the system retains the capability of representing p^n states overall. This seems to be a manifestation of a fundamental trade-off. Cognitive systems may either:

- 1) Operate on relatively few basic semantic objects (objects stored in memory as meaningful/significant) but allow many possible combinations between them, i.e. be expressive but low capacity.
- 2) Operate on relatively many basic semantic objects but only accommodate certain possible combinations between them. This is the regime in which our proposed system operates.

We note that the question of the optimum balance between expressivity and capacity is highly complex and requires further study in its own right. In our proposed system capacity and expressivity are to some extent decoupled: p, y affect capacity and expressivity in a trade-off manner whilst d affects only capacity.

Trade-off 2 - ‘holographic’ vs. compression: Cognitive systems can be conceived at different levels of ‘holographic’ as determined by the percentage of operations that are operand length-preserving. For fixed maximum semantic object length the choice lies between the extreme of always utilising the full length of n elements in order to represent every possible semantic object (full-holographic), or allowing some semantic objects to be shorter (non-holographic). This significantly impacts the amount of information each numerical element carries. In a fully holographic representation transmitting or processing even a single-item-equivalent semantic object requires handling of n elements; the same as transmitting/processing the equivalent of a long chain. The semantic information per element may dramatically differ in each situation. In our proposed system, however, superpositions of fewer items are represented by shorter chains. This illustrates how less holographic systems generally offer the option of operating on more compressed information, i.e. closer to the signal-to-noise ratio (SNR) limit.

Naturally there is a price to pay for compression: when creating new semantic objects for storage it is extremely useful if these new objects can be mapped onto minimum-length units (the semantic object basis of any cognitive system). Mechanisms for mapping any arbitrary chain onto such units need to be supported, adding to system complexity. Furthermore, in a non-holographic system any circuitry designed to support the last items of a chain may be utilised only infrequently. This is expected to strongly affect hardware design decisions.

Trade-off 3 - long vectors with few states per element vs short vectors with many states per element: If we have a fixed number of binary lines (i.e. $l \cdot y = C$), we have a choice of treating C as either: i) one single, large identifier number, ii) a collection of binary bits independent of one another or iii) certain possibilities in between. For example, for $C = 16$ we can have $\{l, y\} \in \{(1, 16), (2, 8), (4, 4), (8, 2), (16, 1)\}$. The number of states we can represent remains fixed at 2^{ly} , but:

- The distance relationships between semantic objects will be different in each case. In the case (1,16) our item consists of a vector of 16x 1-bit elements, and therefore there are 16 nearest neighbours each item (all items that differ from the base object at exactly one position). In the case (16,1) our item is a single 16-bit number which has exactly two nearest neighbours (the elements/items different from the base object by one unit of distance). Note that the case (1,16) corresponds tightly to the spatter code system proposed by Kanerva [14] since modular addition now reduces to a simple XOR.
- The degree of modularity achievable in hardware may be impacted in each case. The (1,16) case requires 16x XOR gates in order to perform one item-item binding whilst in the (16,1) case requires a single 16-bit adder. In the case of large values of C there may be an additional impact on speed (how viable is to make a 512-bit adder that computes an answer in one clock cycle/step? - 512x XOR gates on the other hand will compute 512 outputs in one step). This subject requires further, dedicated study.

Trade-off 4 - operation complexity vs. property attractiveness: As a rule of thumb operations with more attractive mathematical properties tend to introduce computational and implementational difficulties. This is perhaps well exemplified by examining different binding operations:

- Convolution commutes, ‘scrambles’ the information well⁸ and preserves information. However, it lengthens the vectors that it processes and it is computationally heavy (many MACs).
- Circular convolution commutes and scrambles. Lengthening no longer occurs, but information is lost and the operation is still heavy on MACs.
- Modular arithmetic commutes. Lengthening does not occur and the operation is MAC-lightweight, but information is lost and the scrambling properties are similar to those of superposition by element-wise addition, so the similarity requirements for defining two semantic

⁸The result bears in general very little resemblance to either of the operands.

objects as corrupted versions of each other have to be substantially tightened.

Ultimately, a complex mix of factors/specs in all trade-off directions will determine the best cognitive system implementation. This may depend on the overall cognitive capabilities required of the system. In this work we have focussed on a partially holographic system based on effectively multiplexing and addition as the system operations. The advantage of this implementation vs. the holographic approach that we have used as standard and inspiration is that both operations have been simplified in hardware: superposition became a multiplexing operation instead of addition whilst binding became element-wise addition instead of circular convolution. The balance of these advantages vs. the attributes that had to be traded-away (mathematical elegance, full holographic, etc.) needs to be considered very carefully. In general, however the system is designed for occasions where we have partially restricted expressivity (notable cap on chain length - effective number of successive superpositions allowed) but enables extreme implementational simplicity and high energy efficiency.

Finally, we envision that our proposed CoPU will form a core component of larger systems with cognitive capability. Much like in a traditional computer, our CPU-equivalent will need a memory to which it can communicate as well as peripheral structures. Work in that general direction has very recently begun to gain traction [18], [30]. Relating this back to biological brains we see the closest analogue of our CoPU in the putative attentional systems of the brain; the contents of the input buffers at any given time could be interpreted as the semantic objects in the machine's 'conscious attention'. In conclusion, we envisage that future thinking machines will be complex systems consisting of multiple, heterogeneous modules including ANNs, memories (bio-inspired or standard digital look-up tables), sensors, possibly even classical microprocessors and more; all working together to give rise to cognitive intelligence. We hope that our CoPU will play a central role in this 'hyperarchitecture' structure by acting as the equivalent of the CPU in a classical computer, and that it will do so with the energy efficiency required for enabling widespread adaptation of cognitive computers.

ACKNOWLEDGEMENTS

The authors would like to thank Prof. Chris Eliasmith whose work provided much of the inspiration for this work. We also thank Prof. Jesse Hoey for his support and fruitful discussions.

REFERENCES

[1] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," *Advances In Neural Information Processing Systems*, pp. 1–9, 2012.

[2] A. Hannun, C. Case, J. Casper, B. Catanzaro, G. Diamos, E. Elsen, R. Prenger, S. Satheesh, S. Sengupta, A. Coates, and A. Y. Ng, "Deep Speech: Scaling up end-to-end speech recognition," dec 2014. [Online]. Available: <http://arxiv.org/abs/1412.5567>

[3] Y. Lecun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015. [Online]. Available: <https://www.nature.com/nature/journal/v521/n7553/pdf/nature14539.pdf><http://arxiv.org/abs/1606.01781><http://arxiv.org/abs/1603.05691>

[4] K. Greff, R. K. Srivastava, J. Koutnik, B. R. Steunebrink, and J. Schmidhuber, "LSTM: A Search Space Odyssey," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 28, no. 10, pp. 2222–2232, oct 2017. [Online]. Available: <http://ieeexplore.ieee.org/document/7508408/>

[5] C. W. Wang, P. C. Guo, X. Wang, Paerhati, L. B. Li, and J. P. Bai, "Autologous peroneus brevis and allogeneic tendon to reconstruct lateral collateral ligament of the ankle joint," *Chinese Journal of Tissue Engineering Research*, vol. 19, no. 30, pp. 4908–4914, sep 2015. [Online]. Available: <http://arxiv.org/abs/1409.1259>

[6] C. J. Spoeer, P. McClure, and N. Kriegeskorte, "Recurrent Convolutional Neural Networks: A Better Model of Biological Object Recognition," *Frontiers in Psychology*, vol. 8, p. 1551, sep 2017. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/28955272><http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC5600938><http://journal.frontiersin.org/article/10.3389/fpsyg.2017.01551/full>

[7] G. B. Kaplan and C. Güzelis, "Hopfield networks for solving Tower of Hanoi problems," *Ari*, vol. 52, no. 1, pp. 23–29, 2001. [Online]. Available: <https://www.researchgate.net/profile/Cueneyt>{_\}Guezelis/publication/225181806{_\}Hopfield{_\}networks{_\}for{_\}solving{_\}Tower{_\}of{_\}Hanoi{_\}problems/links/54520f750cf24884d8873da2.pdfhttps://ejournal.csiro.au/cgi-bin/sciserv.pl?collection=journals{\&}journal=14345641{\&}issue=v5

[8] F. Schurmann, K. Meier, and J. Schemmel, "Edge of Chaos Computation in Mixed-Mode VLSI - A Hard Liquid," *Proc. of NIPS*, 2005.

[9] J. R. Anderson, M. Matessa, and C. Lebiere, "ACT-R: A Theory of Higher Level Cognition and Its Relation to Visual Attention," *Human-Computer Interaction*, vol. 12, no. 4, pp. 439–462, dec 1997. [Online]. Available: <http://www.tandfonline.com/doi/abs/10.1207/s15327051hci1204>{_\}5

[10] C. Eliasmith, *How to build a brain : a neural architecture for biological cognition*.

[11] T. A. Plate, "Holographic Reduced Representations," *IEEE Transactions on Neural Networks*, vol. 6, no. 3, pp. 623–641, may 1995. [Online]. Available: <http://ieeexplore.ieee.org/document/377968/>

[12] P. H. Schönemann, "Some algebraic relations between involutions, convolutions, and correlations, with applications to holographic memories," *Biological Cybernetics*, vol. 56, no. 5–6, pp. 367–374, jul 1987. [Online]. Available: <http://link.springer.com/10.1007/BF00319516>

[13] P. Smolensky, "Tensor product variable binding and the representation of symbolic structures in connectionist systems," *Artificial Intelligence*, vol. 46, no. 1, pp. 159–216, 1990.

[14] P. Kanerva, "Fully Distributed Representation," *Proceedings of 1997 Real World Computing Symposium*, no. c, pp. 358–365, 1997.

[15] F. Akopyan, J. Sawada, A. Cassidy, R. Alvarez-Icaza, J. Arthur, P. Merolla, N. Imam, Y. Nakamura, P. Datta, G. J. Nam, B. Taba, M. Beakes, B. Brezzo, J. B. Kuang, R. Manohar, W. P. Risk, B. Jackson, and D. S. Modha, "TrueNorth: Design and Tool Flow of a 65 mW 1 Million Neuron Programmable Neurosynaptic Chip," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 34, no. 10, pp. 1537–1557, oct 2015. [Online]. Available: <http://ieeexplore.ieee.org/document/7229264/>

[16] A. Necker, S. Fok, B. V. Benjamin, T. C. Stewart, N. N. Oza, A. R. Voelker, C. Eliasmith, R. Manohar, and K. Boahen, "Braindrop: A Mixed-Signal Neuromorphic Architecture With a Dynamical Systems-Based Programming Model," *Proceedings of the IEEE*, vol. 107, no. 1, pp. 144–164, jan 2019. [Online]. Available: <https://ieeexplore.ieee.org/document/8591981/>

[17] N. Qiao, H. Mostafa, F. Corradi, M. Osswald, F. Stefanini, D. Sumislawska, and G. Indiveri, "A reconfigurable on-line learning spiking neuromorphic processor comprising 256 neurons and 128K synapses," *Frontiers in Neuroscience*, vol. 9, no. APR, p. 141, 2015. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/25972778><http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC4413675>

[18] A. Rahimi, T. F. Wu, H. Li, J. M. Rabaey, H. S. P. Wong, M. M. Shalaker, and S. Mitra, "Hyperdimensional Computing Nanosystem," nov 2018. [Online]. Available: <http://arxiv.org/abs/1811.09557>

[19] D. Casasent and B. Telfer, "Key and recollection vector effects on heteroassociative memory performance," *Applied optics*, vol. 28, no. 2, pp. 272–283, jan 1989. [Online]. Available: <https://www.osapublishing.org/abstract.cfm?URI=ao-28-2-272><http://www.ncbi.nlm.nih.gov/pubmed/20548469>

- [20] A. D. Fisher, W. L. Lippincott, and J. N. Lee, "Optical implementations of associative networks with versatile adaptive learning capabilities," *Applied Optics*, vol. 26, no. 23, p. 5039, dec 1987. [Online]. Available: <https://www.osapublishing.org/abstract.cfm?URI=ao-26-23-5039>
- [21] E. G. Paek and D. Psaltis, "Optical Associative Memory Using Fourier Transform Holograms," *Optical Engineering*, vol. 26, no. 5, p. 265428, may 1987. [Online]. Available: <http://opticalengineering.spiedigitallibrary.org/article.aspx?doi=10.1117/12.7974093>
- [22] D. Willshaw and P. Dayan, "Optimal Plasticity from Matrix Memories: What Goes Up Must Come Down," *Neural Computation*, vol. 2, no. 1, pp. 85–93, mar 1990. [Online]. Available: <http://www.mitpressjournals.org/doi/10.1162/neco.1990.2.1.85>
- [23] D. W. J. O. B. P. and H. L.-H. C., "Non-holographic associative memory," *Nature*, vol. 222, no. 5197, pp. 960–962, 1969. [Online]. Available: <http://psycnet.apa.org/psycinfo/1970-20044-001>
- [24] D. Aerts, M. Czachor, and B. De Moor, "On Geometric Algebra representation of Binary Spatter Codes," oct 2006. [Online]. Available: <http://arxiv.org/abs/cs/0610075>
- [25] M. Hu, J. P. Strachan, Z. Li, E. M. Grafals, N. Davila, C. Graves, S. Lam, N. Ge, R. S. Williams, J. Yang, and H. P. Labs, "Dot-Product Engine for Neuromorphic Computing: Programming 1T1M Crossbar to Accelerate Matrix-Vector Multiplication," *IEEE Design Automation Conference*, pp. 1—6, 2016. [Online]. Available: <https://www.labs.hpe.com/techreports/2016/HPE-2016-23.pdf>
- [26] J. Park and I. W. Sandberg, "Universal Approximation Using Radial-Basis-Function Networks," *Neural Computation*, vol. 3, no. 2, pp. 246–257, 1991. [Online]. Available: <http://www.mitpressjournals.org/doi/10.1162/neco.1991.3.2.246>
- [27] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," *IEEE Transactions on Evolutionary Computation*, vol. 6, no. 2, pp. 182–197, apr 2002. [Online]. Available: <http://ieeexplore.ieee.org/document/996017/>
- [28] J. Deveugele and M. Steyaert, "A 10-bit 250-MS/s Binary-Weighted Current-Steering DAC," *IEEE Journal of Solid-State Circuits*, vol. 41, no. 2, pp. 320–329, feb 2006. [Online]. Available: <http://ieeexplore.ieee.org/document/1583796/>
- [29] K. Boahen, "Point-to-point connectivity between neuromorphic chips using address events," *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, vol. 47, no. 5, pp. 416–434, may 2000. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=842110>
- [30] A. Graves, G. Wayne, and I. Danihelka, "Neural Turing Machines," oct 2014. [Online]. Available: <http://arxiv.org/abs/1410.5401>