



Hardware-Aware Static Optimization of Hyperdimensional Computations

PU (LUKE) YI, Stanford University, USA

SARA ACHOUR, Stanford University, USA

Binary spatter code (BSC)-based hyperdimensional computing (HDC) is a highly error-resilient approximate computational paradigm suited for error-prone, emerging hardware platforms. In BSC HDC, the basic datatype is a *hypervector*, a typically large binary vector, where the size of the hypervector has a significant impact on the fidelity and resource usage of the computation. Typically, the hypervector size is dynamically tuned to deliver the desired accuracy; this process is time-consuming and often produces hypervector sizes that lack accuracy guarantees and produce poor results when reused for very similar workloads. We present HEIM, a hardware-aware static analysis and optimization framework for BSC HD computations. HEIM analytically derives the minimum hypervector size that minimizes resource usage and meets the target accuracy requirement. HEIM *guarantees* the optimized computation converges to the user-provided accuracy target on expectation, even in the presence of hardware error. HEIM deploys a novel static analysis procedure that unifies theoretical results from the neuroscience community to systematically optimize HD computations.

We evaluate HEIM against dynamic tuning-based optimization on 25 benchmark data structures. Given a 99% accuracy requirement, HEIM-optimized computations achieve a 99.2%-100.0% median accuracy, up to 49.5% higher than dynamic tuning-based optimization, while achieving 1.15x-7.14x reductions in hypervector size compared to HD computations that achieve comparable query accuracy and finding parametrizations 30.0x-100167.4x faster than dynamic tuning-based approaches. We also use HEIM to systematically evaluate the performance benefits of using analog CAMs and multiple-bit-per-cell ReRAM over conventional hardware, while maintaining iso-accuracy – for both emerging technologies, we find usages where the emerging hardware imparts significant benefits.

CCS Concepts: • **Hardware** → **Emerging languages and compilers**; Memory and dense storage; • **Software and its engineering** → *Software notations and tools*.

Additional Key Words and Phrases: unconventional computing, emerging hardware technologies, program optimization

ACM Reference Format:

Pu (Luke) Yi and Sara Achour. 2023. Hardware-Aware Static Optimization of Hyperdimensional Computations. *Proc. ACM Program. Lang.* 7, OOPSLA2, Article 222 (October 2023), 30 pages. <https://doi.org/10.1145/3622797>

1 INTRODUCTION

Over the years, researchers have developed many emerging memory technologies (e.g., FeRAM, ReRAM, STT-MRAM) that offer non-volatility, better write endurance, and faster write speeds, and support integration into monolithic 3D integrated circuits due to low annealing temperatures. [Halawani et al. 2021; Imani et al. 2017b, 2019c; Karunaratne et al. 2020; Poduval et al. 2021; Rahimi et al. 2017; Wu et al. 2018] Moreover, resistive memories, such as ReRAM, have also been used to build

Authors' addresses: Pu (Luke) Yi, Computer Science Department, Stanford University, 450 Jane Stanford Way, Stanford, California, 94305, USA, lukeyi@stanford.edu; Sara Achour, Computer Science and Electrical Engineering Departments, Stanford University, 450 Jane Stanford Way, Stanford, California, 94305, USA, sachour@stanford.edu.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2023 Copyright held by the owner/author(s).

2475-1421/2023/10-ART222

<https://doi.org/10.1145/3622797>

analog in-memory computing fabrics that eliminate data movement by performing computation directly within memory cells, and as memory units in monolithic systems that employ other emerging device technologies (CNFETs) to realize extremely communication dense, next-generation hardware platforms. These new technologies offer unprecedented benefits but have not seen broad adoption, as they have much higher bit corruption rates than conventional hardware. These hardware errors often arise due to intrinsic properties in the involved materials, and therefore remain a significant problem despite investments from the devices community. [Imani et al. 2017b; Shulaker et al. 2014]

Challenges with Approximate Classical Computation. Practitioners from the hardware and software communities have developed a range of techniques for statically and dynamically optimizing classical computations to execute reliably on error-prone hardware. [Achour and Rinard 2015; Misailovic et al. 2014; Sharif et al. 2021] All of these methods grapple with two truths of classical approximate computing: (1) certain bits are essential to the computation and must be retained accurately (e.g., exponent vs. mantissa bits), (2) some compute operations (e.g., branching) need to execute accurately to obtain a usable result. As a result, these techniques typically require computations and data to be partitioned into precise/approximation-amenable regions, where precise data and compute are either run separately on reliable hardware or run with a number of protection mechanisms (e.g., redundancy, ECC) that guard against bit corruptions. These requirements complicate the architectural designs of these platforms and introduce overheads into the computation. These inefficiencies and the relative difficulty associated with statically propagating error through programs without over-approximation make it exceedingly difficult to soundly and efficiently perform computation on emerging hardware.

1.1 Hyperdimensional Computing / Vector Symbolic Architectures

Hyperdimensional Computing (HDC), or alternatively Vector Symbolic Architectures (VSA), is an approximate computing paradigm well-matched to error-prone, emerging computing platforms. The basic unit of data is a hypervector – a large numeric or binary vector – that distributes program information evenly across bits/values. There are many variants or dialects of HDC; this work focuses on the binary spatter code (BSC) variant of HDC that works with dense binary hypervectors. [Kanerva et al. 1997] BSC HDC offers three key computational characteristics which together enable sound and robust approximate computing on emerging technologies:

- ▶ **Distributed Data Representation.** All hypervector bits are *equally important*, and all bit errors have the same effect on the hypervector result. Therefore, a single-bit flip has the same effect on the computed result, regardless of where it occurs in memory, or within the computational pipeline.
- ▶ **Distance-Based Computation.** Information is encoded over the relative similarity/dissimilarity of hypervectors, where the similarity of two hypervectors is computed with the Hamming distance metric. The Hamming distance is highly resilient to bit errors, as many bit corruptions are required to substantially influence the calculated distance.
- ▶ **Simple Operators.** The basic HD operators are implemented with circular shift and bit-wise XOR and majority operations. These operators are both hardware-efficient and amenable to analysis.

In contrast, in classical computation, a single bit error can have an outsized effect on a computational result, the sensitivity of the final result to error is workload-dependent, and it is typically difficult to statically analyze the propagation of bit errors through the program without over-approximation.

Applications. HD computing has enjoyed increased attention in the hardware and software research communities. [Imani et al. 2019a, 2018, 2019b; Kim et al. 2020] Practitioners have devised HD computations to build a range of data structures, including database records, graphs, trees, and finite-state automata, [Osipov et al. 2017; Yerxa et al. 2018] and to perform a number of processing tasks, including signal and language classification, information retrieval, workload balancing, and analogical reasoning. [Eggimann et al. 2021; Gayler and Levy 2009; Heddes et al. 2022; Jones and

Mewhort 2007; Kanerva 2010; Karunaratne et al. 2020; Kleyko et al. 2022, 2020; Pashchenko et al. 2020; Plate 2000; Rachkovskij and Slipchenko 2012; Simpkin et al. 2019] HD computation has also been successfully used in recent years to improve the accuracy and efficiency of edge ML models, and to embed intuition about problem structure into ML training tasks. [Imani et al. 2017a; Rahimi et al. 2016, 2018; Schlegel et al. 2021, 2022; Theiss et al. 2022]

Challenges with VSA/HDC. One drawback to HD computing is that large hypervectors are usually required to encode information reliably. The hypervector *size* strongly affects the accuracy of the implemented HD computation and determines the amount of information that can be reliably encoded with the hypervector. Lower dimensional bit-vectors consume less space but potentially reduce one’s ability to retrieve information reliably. Typically, practitioners either leave the hypervector size unoptimized or dynamically tune the hypervector size by running Monte Carlo simulations for each parametrization of the target computation. [Kanerva 2009, 2014, 2018; Montagna et al. 2018; Rahimi et al. 2017] Dynamic tuning is time-consuming and has a tendency to overfit – the chosen hypervector sizes do not generalize well when minor adjustments are made to the computation.

1.2 The HEIM Optimizer

We present HEIM, the first (to our knowledge) static analysis and optimization tool for BSC HD computations. To summarize, the HDC paradigm enables robust computation on error-prone hardware, and HEIM delivers accuracy guarantees even in the presence of hardware error. Given a hardware error model and an accuracy specification, HEIM derives the smallest hypervector size that meets the specified accuracy requirements on the target hardware platform:

- ▶ **Analysis.** HEIM deploys a precise and sound static analysis that guarantees the convergence of the accuracy of the HD computation to the desired accuracy on expectation. The analysis uses several novel theoretical results to soundly derive the expected accuracy for HD computations (see Table 1).
- ▶ **Hardware-Aware.** HEIM optimizes HD computations to execute accurately on hardware platforms that use error-prone and emerging device technologies. HEIM’s analysis procedure works with a hardware error model and delivers accuracy guarantees in the presence of hardware error.
- ▶ **Robust Optimization.** The HEIM-derived parametrization is guaranteed to deliver the desired accuracy for all HD computations captured by the accuracy specification. HEIM analytically derives important HDC program parameters, including distance thresholds and HD operation-specific hypervector sizes, which together are used to optimize the computation.

1.3 Contributions

- ▶ **HEIM Accuracy Analysis.** We present a novel accuracy analysis that employs novel theoretical results to derive the expected accuracy for a set of BSC computations on an emerging hardware technology. The accuracy analysis works with an accuracy specification that supports the description of HD computations and their associated accuracy constraints.
- ▶ **HEIM Optimizer.** We present an algorithm that uses the above accuracy analysis to statically derive thresholds and hypervector sizes that minimize resource usage for a given HD computation while satisfying the accuracy constraints provided in the HEIM accuracy specification.
- ▶ **Evaluation.** We evaluate HEIM against dynamic tuning-based optimization on 25 data structures. Given a 99% accuracy requirement, HEIM-optimized computations achieve a 99.2%-100.0% median accuracy, up to 49.5% higher than dynamic tuning, and achieve 1.15x-7.14x reductions in hypervector size compared to iso-accuracy dynamically tuned executions. HEIM also finds parametrizations 30.0x-100167.4x faster than dynamic tuning-based approaches. We also use HEIM to find optimized computations at iso-accuracy for two emerging hardware technologies and find usages where the emerging hardware imparts significant benefits over the classical HDC implementation.

2 HYPERDIMENSIONAL COMPUTING

Hyperdimensional computing (HDC) is a highly error-resilient brain-inspired computational paradigm. In HDC, information is encoded by computing over randomly generated vectors corresponding to symbols (e.g., letters, colors, objects) in the application domain using binding, bundling, and permutation operations. A hypervector is a numerical vector which may contain binary, modular integer, complex, or real values depending on the HDC. Information is retrieved from HD-encoded data by computing the *distance* ($dis(hv, hv')$) between hypervectors, where hypervectors with small distances are similar. The *distance threshold* (thr) determines the cutoff point that separates a "small" and a "large" distance. Because information is evenly distributed across hypervector bits, and distance calculations are resilient to error, HD queries can complete successfully even when bit corruptions occur.

BSC HDC. This work targets the binary spatter code (BSC) HDC variant, which works with dense binary hypervectors. Random hypervectors are generated by sampling bits from a $p=0.5$ Bernoulli distribution, and permutation, binding, and bundling operations are implemented with circular shifts, bit-wise exclusive-OR (XOR), and bit-wise majority operations, respectively. The Hamming distance $dis(hv, hv') = \frac{1}{n} \sum_i hv_i \wedge hv'_i$ is the BSC distance measure, where n is the hypervector size. The bit-wise majority operation computes whether there are a majority of "1" or "0" bits at each bit position, and can be alternatively interpreted as an element-wise sum, followed by a thresholding operation.


Operators. The binding and bundling operations ($hv \odot hv'$ and $hv + hv'$) respectively produce hypervectors that are *dissimilar* and *similar* to the input hv , hv' hypervectors. The permutation operation $hv' = \rho_i(hv)$ produces a hypervector hv' that is dissimilar to the input hypervector hv , where the original hypervector can be recovered by inverting the permutation ($hv = \rho_{-i}(hv')$), where i is an integer value. Generally, binding and permutation operations distribute over bundling, and for BSC HDC, bundling/bundling are commutative and associative, and binding is invertible.

Codebooks. These operations are performed over a "codebook" of *basis hypervectors*, which represent atomics in the problem domain. Examples of codebooks include letters of the alphabet ($K=26$), numerical digits ($K=10$), graph nodes ($K=\#$ nodes), and primary colors ($K=3$). Each basis hypervector, or code, in the codebook is typically randomly generated; the associated atomics (e.g., letters) are, therefore, dissimilar from one another. Conceptually, this dissimilarity captures the idea that atomics are distinct – the letter **A** is distinct from the letter **C**, for example. The HD permutation, binding, and bundling operations are then applied to encode information using these atomics. For example, bundling the "A" and "B" basis hypervectors ($hv = hv_A + hv_B$) produces a basis hypervector similar to both **A** and **B**.

2.1 Data Structure/Query Interpretation of HD Computing

An HD computation can be thought of as a *data structure encoding* operation that produces a *data structure* hypervector \mathbf{ds} that can be queried by computing its distance from a *query* hypervector \mathbf{q} . Because hypervectors are lossy information encodings, a query against an HD data structure may occasionally return an incorrect result – the *accuracy* of a query is the probability that a query returns the correct result. The hypervector size and the distance thresholds together control the computation's accuracy.

Data Structures. The bundling operation conceptually creates a set of elements where the membership of an element or subset can be tested with a distance calculation. For example, for the $\mathbf{ds} = hv_A + hv_B + hv_C \approx \{A, B, C\}$, The membership of a subset $\mathbf{q} = hv_A + hv_B \approx \{A, B\}$ is tested by computing the distance $dis(hv_A + hv_B, hv_A + hv_B + hv_C)$. If the distance falls below a distance threshold thr , the set contains the subset; this is referred to as a *match*. A false positive occurs when a data structure hypervector falsely *matches* the query, and a false negative occurs when a data structure hypervector falsely fails to match the query. The binding operation conceptually creates a record of elements ($hv_A \odot hv_B \sim \langle A, B \rangle$), where each record is only similar to other matching records, and the permutation operation is used to encode positional information into the HD data structure.



(a) Student knowledge graph.

```

1 //data structure: one hypervector for each node in item memory
2 ds[jack] = act⊙(likes⊙(banana+apple+mary)+plays⊙tennis)//jack
3 ds[apple] = target⊙likes⊙(jack+mary)//apple
4 ds[mary] = act⊙(likes⊙apple+plays⊙tennis)+target⊙likes⊙jack//mary
5 ds[tennis] = target⊙plays⊙(jack+mary)//tennis
6 ds[banana] = target⊙likes⊙jack//banana
7 query = act⊙likes⊙apples//how many people like apples query
8 return count(dis(query,ds[node]) <= thr[node])//execute query on graph
  
```

(b) HDC encoding of student knowledge graph and *apples* query.

Fig. 1. Illustrative Example: Knowledge Graph

Several complex data structures that compose sets, sequences, and records can be built from these basic operations. For example, $hv_A + \rho_1(hv_A)$ builds the sequence $[A, A]$ that can be indexed at index 1 by computing $\rho_{-1}(hv_A + \rho_1(hv_A))$, the $hv_A \odot hv_B + hv_C \odot hv_D$ encoding builds the set of records $\{\langle A, B \rangle, \langle C, D \rangle\}$ that can be queried with record subsets, and the $hv_A \odot \rho_1(hv_B)$ encoding builds a tuple $\langle A, B \rangle$ of ordered elements, such that $\langle A, B \rangle \neq \langle B, A \rangle$.

Item Memories. Complex data structures, such as graphs and databases, can be encoded through the use of an *item memory*, a key-hypervector data store that maps identifiers to hypervector item memory rows hv_i , which implements HD data structures. For example, an HD graph item memory maps nodes to hypervectors, and each hypervector "value" encodes the set of edges connected to the associated "key" node. Item memory-based data structures support two queries: (1) threshold-based queries and (2) w -winner winner-take-all queries. For both queries, the distance between the query hypervector and each item memory row $dis(q, hv_i)$ is computed. In threshold-based querying, all item memory rows that match a query hypervector q are returned. In winner-take-all querying, the w item memory rows closest to the query hypervector are returned. Threshold queries require a distance thr to operate, which may be individually set for each item memory row, while winner-take-all queries take no additional parameters. See Section 10 for more discussion on these two query types.

3 ILLUSTRATIVE EXAMPLE: KNOWLEDGE GRAPH

Knowledge graphs capture networks of real-world entities (objects, people, situations), and model relationships between them. An outgoing edge indicates the originating node is *acting* on a target node, and an incoming edge indicates the receiving node is the *target* of another node. Nodes map to concepts (e.g., **apple**, **mary**, **tennis**), and edges are labeled with relations (e.g., **plays**, **likes**, **hates**).

HD Knowledge Graph. Figure 1b presents the HD encoding of the student knowledge graph from Figure 1a (Lines 2-6). This encoding works with codebooks that specify the relations **{likes, plays}**, the interactions **{act, target}**, and concepts **{jack, mary, banana, apple, tennis}** that may appear in the student knowledge graph. Each node's edge information is then encoded as a hypervector $ds[node]$ in item memory. A hypervector that encodes each incoming and outgoing edge is constructed by binding together interaction, relation, and concept tuples (e.g., $act \odot likes \odot apples$). The hypervector that encodes the set of edges connected to a given node is constructed by bundling (+) all of the edge hypervectors containing the target node together. In BSC, binding distributes over bundling.

Each edge set hypervector is then stored in an *item memory*. The keys in the knowledge graph's item memory map to concepts, and the hypervectors are the constructed edge list. In this example, the item memory hypervectors are stored in 2-bit-per-cell resistive RAM (ReRAM). This storage medium is 2x denser than conventional binary storage but has a 2.15% chance of corrupting a bit in memory – this error rate is collected from a real ReRAM array (Section 8). Therefore, the emerging memory may sporadically corrupt bits at random positions in the hypervector data structure.

```

1: procedure TESTACC(n,accReq,tests,nCbs,nTraces)
2:   data = []
3:   for q,ds,label in tests do
4:     for i in range(nCbs×nTraces) do
5:       itemMemDist = execMonteCarlo(q,ds,n)
6:       data.append((itemMemDist,label))
7:   dists = sort(getDists(data))
8:   thrs = set((dist[i]+dist[i+1])/2 for 0...|dists|-1)
9:   thr,acc = bruteSearch(thrs,data)
10:  return (thr, acc >= accReq )
11: procedure DYN_TUNE(nMax,accReq,tests,nCbs,nTraces)
12:  fn =  $\lambda$ .n: testAcc(n,accReq,tests,nCbs,nTraces)
13:  return binSearch(0,nMax,fn)

```

(a) Dynamic tuning algorithm

```

1 spec {
2   abs-data query = prod(itypes,rels,concepts);
3   abs-data ds = sum(4,prod(itypes,rels,concepts));
4   thr-query(query, ds, 1, 0.99, 0.01, 0.01); }

```

(b) Knowledge graph specification

```

1 hardware-model {
2   mem codebook = 0.00; mem item-mem = 0.0215;
3   op bind = 0.00; op bundle = 0.00; }

```

(c) Hardware error model

Fig. 2. Dynamic tuning algorithm and HEIM specifications.

Queries. We now want to query the knowledge graph to answer the following question: "How many students like apples?". To answer this question, we would count how many nodes have outgoing edges with the **likes** relation label pointing to apple – this will be referred to as the *apples* query. We want the *apples* query to complete with 99% accuracy, *even in the presence of hardware error*.

This query can be dispatched on the hypervisor representation of the student knowledge graph. We construct the *apples* query hypervisor by binding together the relation, concept, and interaction hypervectors together (Line 8) – this is the same encoding used for the edge in the knowledge graph. We then determine if a node hypervisor contains the query tuple by computing the Hamming distance **dis(query, ds[node])** between the query hypervisor and the node hypervectors in item memory (line 10). Node hypervectors with a distance below some node-specific distance threshold *thr* contain **(act, likes, apples)** query tuple in its edge set, and the corresponding key is returned as a match. This query cannot be expressed as a winner-take-all query as the number of matches is unknown.

3.1 Naive Query Optimization

We are interested in *minimizing* the hypervisor size to reduce the memory usage while still attaining this 99% accuracy target. Typically, this is done by dynamically tuning the size and distance thresholds to execute the provided query with acceptable accuracy. Figure 2a presents a dynamic tuning algorithm that finds a minimum hypervisor size between zero and *nMax* and the associated distance threshold that achieves a query accuracy of *accReq* over a test set of labelled matching/non-matching HD queries and data structures (*tests*). The algorithm performs a binary search over hypervisor sizes. For each size, the algorithm executes each test query and data structure for *nCbs*×*nTraces* Monte Carlo trials to build up the dataset (*data*) of query-item memory distances. The algorithm then uses a brute-force search to find the distance threshold that maximizes the accuracy, or the fraction of correctly classified samples over the constructed dataset. In total, the dynamic tuning algorithm executes $\log(\mathbf{nMax}) \times |\mathbf{tests}| \times \mathbf{nCbs} \times \mathbf{nTraces}$ Monte Carlo trials of the computation.

Accuracy. We parameterize the dynamic tuner with a test set containing the "apples" query and knowledge graph data structures, over 30 random codebooks, and 30 error traces – in total, 900 trials. The dynamic tuner completes in 35.4 seconds (averaged over ten runs) and finds a hypervisor size of 117 bits and threshold of **0.402** that empirically delivers a query accuracy of **99.044%** on the test set. With these dynamically tuned parameters, we attain an 85.47x reduction in the hypervisor size, compared to the unoptimized 10,000-bit hypervisor size used in previous literature – dynamic tuning therefore significantly reduces the memory footprint of the knowledge graph item memory. [Kanerva 2009, 2014, 2018; Montagna et al. 2018; Rahimi et al. 2017]

While this parameterization delivers the desired accuracy for the *apples* query, it does not generalize well to other knowledge graphs or queries of comparable complexity. To demonstrate this, we randomly construct 1000 knowledge graphs containing five nodes with a maximum degree of 4, generate five random edge queries for each graph, and then evaluate the accuracy of the edge queries over 900 trials. Over these 5000 randomly generated data structure-query combinations, we find that the accuracy for 0-degree nodes, 1-degree nodes, 2-degree nodes, 3-degree nodes, and 4-degree nodes are 100.0%, 98.8%, 98.8%, 98.9%, 97.6% respectively. Notably, queries over 1-4 degree nodes fail to meet the 99% accuracy target when the dynamically tuned parametrization is used. This issue can be addressed by dynamically tuning over randomly generated queries and data structures; however, doing so will drastically increase parameter tuning time.

3.2 Optimizing the Apples Query with HEIM

With HEIM, we can *statically* optimize the size and distance thresholds for a given HD computation without performing any simulation. HEIM delivers precise accuracy guarantees for data structure queries that generally hold, on expectation, over different query and data structure instantiations and can deliver these guarantees even in the presence of hardware error.

HD Specifications. HEIM works with an *specification* of the HD computation that describes the accuracy requirements for a set of data structure queries. Figure 2b presents a HEIM specification that verifies that all edge queries over 5-node knowledge graphs with a maximum degree of 4 achieve a query accuracy of at least 99%. The *apples* query on the student knowledge graph presented in Figure 1a is an example of a concrete data structure query that adheres to this specification. Line 2 defines a query as a product of interactions (**itypes**), relations (**rels**), and concepts (**concepts**), and Line 3 defines a node hypervector in item memory as a bundle (**sum**) of up to 4 edge hypervectors, where each edge tuple is binding of an interaction, relation, and concept. The accuracy assertion on Line 4 requires all tuple queries made to a node hypervector to return the correct result at least 99% of the time, with maximum false positive (incorrect match) and false negative (incorrect not match) rates of 1%. See supplementary materials for more information on the analysis-amenable knowledge graph data structure.

HEIM considers the effect of hardware error during optimization and works with a hardware error model that captures the error rates in the target device. Figure 2c presents the hardware error model for the 2 BPC ReRAM-based accelerator we are targeting in this example; this model defines the per-bit corruption probability for data in item memory as 0.0215. All other operations are error-free.

HEIM-Optimized Parameters. We use HEIM to identify an optimal threshold and hypervector size for the specification in Figure 2b and the hardware error model in Figure 2c. HEIM completes its analysis in 13.58 milliseconds (2606.8x faster than dynamic tuning) and returns a hypervector size of 173 and distance thresholds of 0.4116, 0.3795, 0.3795, and 0.1744 for nodes with degrees 4, 3, 2, and 1. The HEIM-optimized *apples* query attains an accuracy of 99.944%. HEIM, therefore, meets the 99% accuracy target and delivers a 57.80x reduction in hypervector size over unoptimized 10,000 element vectors, reducing the number of 2 BPC ReRAM cells required to store each node hypervector from 5,000 cells to just 87 cells. The HEIM hypervectors are 1.48x larger than the dynamically tuned hypervector size but much more reliably deliver the desired accuracy across data structures. In fact, HEIM guarantees that the derived threshold and hypervector size will classify edge queries on node hypervectors (with node degree ≤ 4) with at least 99% accuracy on expectation. We evaluate the HEIM-optimized HD computation over the random knowledge graphs and queries described in Section 3.1 and find HEIM empirically attains an accuracy of 100.0%, 100.0%, 99.9%, 99.9%, 98.9% for 0-degree, 1-degree, 2-degree, 4-degree and 5-degree nodes respectively, all close to or higher than the target accuracy of 99%. Therefore, while HEIM's hypervector size is larger than the dynamically tuned hypervector size, the HEIM-tuned computation more reliably meets the 99% accuracy constraint.

$$\begin{array}{l}
 x, x', x'' \in \text{Reals}, v \in \text{Literals}, i, k, w, m, t \in \text{Integers} \\
 \text{SExp} ::= v \mid \text{perm}(i, v) \mid \text{prod}(\text{SExp}^*) \\
 \text{CExp} ::= \text{sum}(i, \text{SExp}^*) \mid \text{prod}(\text{sum}(i, \text{SExp}^*)^*) \\
 \text{Expr} ::= \text{SExp} \mid \text{CExp} \\
 \text{Stmt} ::= \text{thr-query}(\text{Expr}, \text{Expr}', k, x, x', x'') \\
 \quad \mid \text{wta-query}(\text{Expr}, \text{Expr}', k, w, m, x, t, x') \\
 \quad \mid \text{abs-data } v = \text{Expr} \\
 \text{Spec} ::= \text{spec} \{ \text{Stmt}^* \}
 \end{array}
 \qquad
 \begin{array}{l}
 x \in \text{Reals} \\
 \text{HOp} ::= \text{bundle} \mid \text{bind} \mid \text{perm} \\
 \text{MLoc} ::= \text{codebook} \mid \text{item-mem} \mid \text{query} \\
 \text{Stmt} ::= \text{HOp} = x \mid \text{mem MemLoc} = x \\
 \text{Mdl} ::= \text{hardware-model} \{ \text{Stmt}^* \}
 \end{array}$$

Fig. 3. Program grammars - HEIM accuracy specification language (*Spec*) and hardware model (*Mdl*).

4 HEIM SPECIFICATION LANGUAGES

The HEIM accuracy specification language *Spec* enables practitioners to specify the structure of HD programs to optimize. The specification language supports the specification of *abstract programs* that capture a set of HD computations. The language supports describing HD computations with the following statements:

HD Expression. Each **abs-data** $v = \text{Expr}$ statement maps an HD expression *Expr* to a variable *v*. The HD expression *Expr* statically describes the structure of the HD computation to analyze. We break up HD expressions into simple (*SExp*) and complex (*CExp*) HD expressions. A simple HD expression can be, a code, a permutation of a code, where the basic permutation operator **perm** specifies the number of times to apply or unapply the permutation (*i*), or a tuple of (permutation of) codes. A complex HD expression is either bundle of simple expressions **sum**(*i,SExp**), or binding of several bundles of simple expressions **prod**(**sum**(*i,SExp**)*). All **sum** expressions specify the maximum number of hypervectors that will be summed together (*i*) – this is necessary for HEIM to complete its analysis.

Thresholding Query Accuracy Constraint. The **thr-query**(*Expr, Expr', k, x, x', x''*) statement imposes the requirement that the thresholding query $|\text{Expr} \cap \text{Expr}'| \geq k$ produces an accurate result with a probability of at least *x*. Intuitively, this formulation checks that thresholding on $\text{dis}(\text{Expr}, \text{Expr}')$ can determine whether at least **k** elements in the query *Expr* are contained within the data structure *Expr'* correctly with a probability of at least *x*. The statement also defines the maximum probability of a false positive (*x'*) and false negative (*x''*) occurrence.

Winner-take-all Query Accuracy Constraint. The **wta-query**(*Expr, Expr', k, w, m, x, t, x'*) statement imposes the requirement that a WTA query produces an accurate result with a probability of at least *x*. Specifically, the WTA query has a query in the form of *Expr* and *m* data structures in the form of *Expr'* in item memory, out of which only *w* match the query, satisfying that $|\text{Expr} \cap \text{Expr}'| \geq k$. The WTA query returns the *w* in *m* data structures with the smallest distances to the query. We define the result as *accurate* when the returned *w* ones are exactly the *w* positives. The statement also specifies a softer constraint that the *w* true matches are contained within the top *t* lowest distances to the query ($t \geq w$), with probability *x'*.

4.1 Hardware Error Model

HEIM works with a hardware error model (*Mdl*) that describes the error rates for the basic HD computational operators. The **HOp** = *x* statements define the per-bit error rates for the bundling (**bundle**), binding (**bind**), and permutation (**perm**) operators. The **MLoc** = *x* defines the per-bit error rate associated with storing data in item memory, in codebook memory, or the query buffer. The item memory data storage location supports in-memory distance calculations, the query buffer stores the query to apply to item memory, and the codebook memory stores the basis vectors for the codebooks.

Formulation	reference	description
WTA- <i>acc</i> , $w=1$ (6)	[Frady et al. 2018]	WTA accuracy for exactly one winner $w=1$
WTA- <i>acc</i> (10)	Section 5.4	WTA accuracy for more than one winner $w>1$
WTA- <i>prob</i> (12)	Section 5.4	probability of the w positives being in top t .
QDS I (14)	[Kanerva et al. 1997]	single-element sum-of-product set membership
QDS II (15)	[Kleyko et al. 2016]	subset sum-of-product set membership
QDS III (17)	Section 6.6	single-element product-of-sum set membership
Hardware Error (20)	Section 6.8	incorporation of hardware error

Table 1. Summary of Theoretical Formulations




Fig. 4. Visualization of WTA/threshold query over match/no-match distance distributions. Points map to sampled match and no-match query-item memory row distances (\bullet and \blacktriangleright) for a 10-element item memory, where match/no-match distances are sampled from match/no-match distance distributions (\blacksquare and \blacktriangleright). Circled points (\circ) map to correct row matches for query.

5 HEIM ACCURACY ANALYSIS

At the heart of HEIM is a novel static accuracy analysis that derives the query accuracy for threshold-based and winner-take-all data structure queries. The analytically derived accuracy is both precise and sound – HEIM’s analysis guarantees that the query under study will converge to the computed accuracy on expectation. This accuracy analysis works with analytical models of the query-data structure distance distributions that are parametrized over hypervector size, the hardware error model, and the query and data structure expressions – these models are used to analytically derive the accuracy of each type of query. For the HEIM analysis to be sound, HEIM requires certain *mutual independence* constraints to hold over the query and dataset. Section 5.1-5.2 overviews the relationship between distance distributions and query accuracy, Sections 5.3-5.4 present the accuracy analysis, and Section 6 presents the analytical distance distribution models. Table 1 summarizes the novel and previously published theoretical results used in this analysis.

HEIM Optimizer. The HEIM accuracy analysis is used by the HEIM optimizer to find the minimum hypervector size for a given HEIM specification. The HEIM optimizer returns a set of query-specific thresholds that can be used to more accurately query from the data structure, and a query-specific hypervector size that can be used to soundly do partial computation. HEIM also returns a set of query-specific mutual independence constraints that must hold for the above analysis to be fully sound; these constraints can be optionally checked at runtime with the HEIM mutual independence checking algorithm. Section 7 presents the HEIM optimization algorithm, and Section 7.1 presents the algorithm for checking that the mutual independence constraints hold over concrete data structures and queries.

5.1 Intuition: Accuracy of Threshold-Based Queries

Consider the following HEIM accuracy constraint over a threshold query $\text{thr-query}(q, ds, k, x, x', x'')$, where q is the query HD expression, ds is the HD expression of the rows in a data structure’s item memory. Each item memory row is a *match* if it contains at least k elements in the query expression, and a

not match otherwise. Figure 4 presents the match (blue) and not-match (red) query-data structure distance distributions. Each distance distribution is normally distributed, with a mean $\mu(q, ds, k, hw)$ that depends on the query and data structure expressions, the number of matching elements, the hardware error model, and a standard deviation $\sigma(\mu(q, ds, k), n)$ is a function of the mean and hypervector size.

Example. In Figure 4, the points on the match and not-match distributions correspond to the query-data structure distances for a ten-element item memory with five match rows (blue, circled) and five not-match rows (red), given a sampled set of codebooks and error trace. Each point on the distance distribution corresponds to a distance between the query hypervector and a row hypervector in item memory. In threshold-based querying, points to the left of the distance threshold (grey line) are returned as a match, and points to the right are returned as not a match. The above example correctly returns four of five matches and incorrectly returns one not-matching distance as a match – corresponding to 1 false positive and one false negative. The *accuracy* of the threshold query corresponds to the probability of correctly classifying each row in the data structure’s item memory. In the above example, 8/10 item memory rows are correctly classified, yielding an 80% accuracy.

Expected Accuracy. The associated match and not-match distributions can be analyzed to compute the expected accuracy. The overlapping area between the match and not-match distributions is the probability that a query is erroneously misclassified. Because the distance distributions overlap, there is ambiguity on whether the query matches a given data structure row. The accuracy is, therefore, the probability that a distance is sampled from the non-overlapping regions of the match and not-match distributions. The false positive rate is the portion of the overlapping area left of the chosen threshold, and the false negative rate is the portion of the overlapping area to the right of the selected threshold. The degree of separation between the two distributions depends on how the distributions’ mean and standard deviations are parameterized.

5.2 Intuition: Accuracy of Winner-Take-All Queries

Consider a winner-take-all (WTA) query accuracy constraint $\text{wta-query}(q, ds, k, w, m, x, t, x')$ with w winners for an m -row item memory. The q , ds , and k arguments correspond to the query and data structure expressions and the number of query elements that must be contained in an item memory row to be considered a match. In w -winner WTA queries, there are precisely w true matches sampled from the match distribution. We note (1) that within matches sampled from the match distribution, there is no order to sampled matches, and (2) WTA queries are not used to query not-matching elements.

Example. Figure 4b shows a 5-winner WTA query over a ten-element item memory. The above accuracy constraint requires all true winners ω (circled) to be contained within the top five lowest distances (left of **top $w=5$** line) with probability x , and all true winners ω to be contained within the **$t=8$** lowest distances (left of **top $t=8$** line) with probability x' . Intuitively, the former requirement requires only rows that are true matches to be returned for a WTA query with probability x . The latter requirement is a soft requirement that ensures all true matches are contained within the top t lowest distances with probability x' , where $t \geq k$. The above figure violates the hard constraint since the top 5 lowest distances contain one not-matching row, and satisfies the soft constraint since all true matches are in the top 8 lowest distances.

Expected Accuracy. Next, we provide an intuitive explanation of the expected accuracy. First, we draw w distances from the match distribution and $m - w$ distances from the not-match distribution. Intuitively, the expected accuracy of a w -winner WTA query corresponds to the probability that all w matching distances are to the left of all $m - w$ not-matching distances. The probability the soft accuracy requirement is satisfied corresponds to the probability that the top t distances contain all w matching distances. The WTA distance distributions have the same μ and σ parameters as the threshold-based distance distributions.

5.3 Threshold-Based Query Accuracy Analysis (**thrAccAnalysis**)

Given a hypervector size n , a hardware specification hw , and a threshold query constraint **thr-query**(q, ds, k, n, hw), the analysis queries the analytical distance model $M(q, ds, k, n, hw)$ to retrieve the corresponding match and not-match distance distributions Φ_M, Φ_{NM} and the associated independence constraint *indepCstr*. HEIM finds the optimal distance threshold thr_{opt} maximizing *acc* for the given match (Φ_M) and not-match distributions (Φ_{NM}) and upper bounds on false negative *reqFn* and false positive *reqFp* rates. The analysis returns the computed threshold, the independence constraint, and whether the derived threshold satisfies the provided accuracy requirements.

Optimal Threshold Derivation. For a thr , denoting fp and fn as the false positive and false negative rates, we have $(CDF(f)(x))$ is the cumulative distribution function of distribution f evaluated on x)

$$fn = 1 - CDF(\Phi_M)(thr), fp = CDF(\Phi_{NM})(thr) \quad (1)$$

and with the increase of thr , fn decreases and fp increases. Therefore, requirements $fp \leq reqFp$ and $fn \leq reqFn$ can be translated into bounds on thr .

$$CDF(\Phi_M)^{-1}(1 - reqFn) = thr_l \leq thr \leq thr_h = CDF(\Phi_{NM})^{-1}(reqFp) \quad (2)$$

If $thr_l > thr_h$, no thr satisfies the requirements and we set *success* to **False**. Otherwise, we aim to maximize the accuracy $acc = 1 - \frac{1}{2}(fp + fn)$ (assuming balanced positive and negative queries).

$$\max_{thr_l \leq thr \leq thr_h} acc = \frac{1}{2}(1 + CDF(\Phi_M)(thr) - CDF(\Phi_{NM})(thr)) \quad (3)$$

We take the derivative of *acc* as follows ($PDF(f)(x)$ is the probability density function of distribution f evaluated on x)

$$\frac{\partial acc}{\partial thr} = \frac{1}{2}(PDF(\Phi_M)(thr) - PDF(\Phi_{NM})(thr)) \quad (4)$$

We have $\frac{\partial acc}{\partial thr} > 0$ when $\mu_M < thr < x$, and $\frac{\partial acc}{\partial thr} < 0$ when $x < thr < \mu_{NM}$, where x is the intersection of two PDF curves in range $\mu_M < x < \mu_{NM}$ (Figure 4). Solving x is trivial as $PDF(\Phi_M)(x) = PDF(\Phi_{NM})(x)$ is a quadratic equation of x . Therefore, the optimal thr is the point closest to x in range $[thr_l, thr_h]$, i.e.,

$$thr_{opt} = \max(thr_l, \min(thr_h, x)) \quad (5)$$

The analysis reports a satisfying threshold was found iff thr_{opt} achieves the accuracy requirement and returns the optimal threshold and the independence constraint *indepCstr* on success.

5.4 Winner-Take-All Query Accuracy Analysis (**wtaAccAnalysis**)

Given a hypervector size n , a hardware specification hw , and a WTA query constraint, **wta-query**(q, ds, k, w, m, x, t, x'), the analysis queries the analytical distance model $M(q, ds, k, n, hw)$ to retrieve the corresponding match and not-match distance distributions Φ_M, Φ_{NM} and the associated independence constraint *indepCstr*. The algorithm computes the expected hard and soft accuracies (*acc* and *prob*) from the Φ_M, Φ_{NM} distance distributions and the w, m , and t WTA query parameters. The expected accuracy is then compared with the provided accuracy requirement ($acc \geq x, prob \geq x'$) to determine whether the hypervector size is sufficiently large. On success, the algorithm returns the independence constraint *indepCstr*.

Winner-Take-All Query Accuracy. We present how the hard accuracy (*acc*) and soft accuracy (*prob*) are computed in this section. We start with *acc*. We denote the WTA query accuracy as $acc(w, m, \Phi_M, \Phi_{NM})$. Frady et al. [Frady et al. 2018] developed a perception theory that gives the expected WTA accuracy when $w = 1$ as follows

$$acc(1, m, \Phi_M, \Phi_{NM}) = \int_{-\infty}^{\infty} PDF(\Phi_M)(x) (1 - CDF(\Phi_{NM})(x))^{m-1} dx \quad (6)$$

An intuitive explanation of the above equation is that when the one positive vector has distance x to ds (with probability density $PDF(\Phi_M)(x)$), the result is accurate iff the other $m - 1$ distractor vectors

all have a distance greater than x (each independently with probability $1 - CDF(\Phi_{NM})(x)$), and the result is the integral of it for all possible x over distribution Φ_M .

Accuracy with Multiple Winners. We extend the theory to handle the general cases when $w > 1$. Following the intuition of (6), if we have $PDF(MAX_{w,\Phi_M})$, the probability density function of the maximum of distances of w positive vectors, then

$$acc(w, m, \Phi_M, \Phi_{NM}) = \int_{-\infty}^{\infty} PDF(MAX_{w,\Phi_M})(x) (1 - CDF(\Phi_{NM})(x))^{m-w} dx \quad (7)$$

because the results are accurate (i.e., the returned w vectors are all positives) iff the maximum distance of the w positives is no greater than the minimum of the $m - w$ negatives. We then derive $PDF(MAX_{w,\Phi_M})$. First, we have

$$CDF(MAX_{w,\Phi_M})(x) = CDF(\Phi_M)^w(x) \quad (8)$$

because $MAX_{w,\Phi_M} \leq x$ iff all w positive distances are no greater than x , each independently with probability $CDF(\Phi_M)(x)$. Then, with the relation of CDF and PDF and the chain rule, we have

$$PDF(MAX_{w,\Phi_M})(x) = CDF(MAX_{w,\Phi_M})'(x) = (CDF(\Phi_M)^w(x))' = w CDF(\Phi_M)^{w-1}(x) \cdot PDF(\Phi_M)(x) \quad (9)$$

With (9), we now conclude that

$$acc(w, m, \Phi_M, \Phi_{NM}) = \int_{-\infty}^{\infty} w CDF(\Phi_M)^{w-1}(x) PDF(\Phi_M)(x) (1 - CDF(\Phi_{NM})(x))^{m-w} dx \quad (10)$$

Analysis of Soft Accuracy Constraint. Besides specifying the desired accuracy, HEIM also enables users to specify *soft accuracy constraint* for WTA queries, with the following form: with probability x' the distances of the w positives are all in the top- t smallest. In other words, there can be at most $t - w$ in the other $m - w$ vectors in the codebook that have a distance smaller than any positive, with probability x' . To formulate this new constraint, we denote $prob(w, m, t, \Phi_M, \Phi_{NM})$ as the probability that w positives all have at least top- t smallest distances. To calculate $prob(w, m, t, \Phi_M, \Phi_{NM})$, we enumerate the number of negatives that have a smaller distance than any positives (i), which should be no more than $t - w$. For each i , and $MAX_{w,\Phi_M} = x$, the probability is that exactly i negatives have smaller distances than x , each independently with probability $CDF(\Phi_{NM})(x)$, and the other $m - w - i$ negatives have greater distances than x , each independently with probability $1 - CDF(\Phi_{NM})(x)$. Therefore, we have

$$prob(w, m, t, \Phi_M, \Phi_{NM}) = \sum_{i=0}^{t-w} \binom{m-w}{i} \cdot \int_{-\infty}^{\infty} PDF(MAX_{w,\Phi_M}) CDF(\Phi_{NM})^i(x) (1 - CDF(\Phi_{NM})(x))^{m-w-i} dx \quad (11)$$

Substituting $PDF(MAX_{w,\Phi_M})$ with (9), we have

$$prob(w, m, t, \Phi_M, \Phi_{NM}) = \sum_{i=0}^{t-w} \binom{m-w}{i} \cdot \int_{-\infty}^{\infty} w CDF(\Phi_M)^{w-1}(x) PDF(\Phi_M)(x) CDF(\Phi_{NM})^i(x) (1 - CDF(\Phi_{NM})(x))^{m-w-i} dx \quad (12)$$

6 ANALYTICAL MODEL (M)

The HEIM analytical model derives the match distribution parameters $\langle \mu_M, \sigma_M \rangle$, the not-match distribution parameters $\langle \mu_{NM}, \sigma_{NM} \rangle$ for the provided query, and the mutual independence constraint *indepCstr* that must hold for the analysis to be valid:

$$\begin{aligned} \langle MeanDist_M, MeanDist_{NM}, indepCstr \rangle &= QDS(q, ds, k) \\ \langle \mu_M, \sigma_M \rangle &= ToNormal(HwErr(hw, MeanDist_M), n) \\ \langle \mu_{NM}, \sigma_{NM} \rangle &= ToNormal(HwErr(hw, MeanDist_{NM}), n) \end{aligned}$$

Sections 6.2-6.6 describes how the match and not-match mean distances are derived from the query and data structure (*QDS*), Section 6.7 describes how hardware error is incorporated into the error-free mean distance (*HwErr*), and Section 6.8 derives the standard deviation from the mean of the same distance distribution (*ToNormal*).

Formalization of HD Computation. A code c is a randomly generated hypervector that maps to a distinct atomic symbol (e.g., a letter). A code set C is a vector that is the superposition (bundle, $+$) of a set of codes. We denote $c \in C$ if c is a code in the set, and $C = \{c_1, c_2, \dots, c_m\}$ if C is a superposition of codes $c_1 + c_2 + \dots + c_m$. We denote a code tuple t as a product (\odot) of two or more codes and create code tuples by binding codes, e.g., $c \odot c' = t$. A code tuple set T is a set of code tuples. The analysis works with codes, code tuples, code sets, and code tuple sets. All permutation operators over codes $\rho_k(c)$ are represented as distinct codes c' in our formalization. This transformation can be applied because the dependency of c and $\rho_k(c)$ does not affect distance computation.

6.1 Mutual Independence

The HEIM analytical model assumes that both the data structure and query contain *mutually independent* codes or tuples. In the analytical model, HEIM identifies the mutual independence constraints that must hold. These mutual independence constraints are dynamically checked when constructing the data structure and query. The specific mutual independence constraint depends on the type of analytical procedure used to perform the analysis. We next present two types of mutual independence constraints. We discuss the implications of the mutual independence constraint in Section 6.10, present an efficient dynamic independence checker and prove the equivalence of mutual independence constraints to statistical independence in Section 7.1.

Independent Set We define independent sets as a set of codes or tuples that are *mutually independent*. Formally, given any HD expression, it can be flattened into the superposition (set) of code tuples $T = \{t_1, t_2, \dots, t_l\}$ ($expr = \sum_{i=1}^l t_i$). T is an independent set if and only if for any $t_i, 1 \leq i \leq l$, there exists no subset $T' \subset T$ other than the subset $\{t_i\}$ that contains only t_i , where t_i is the binding of tuples in T' (i.e., $t_i = \odot_{t \in T'} t$). For example, $expr = (a+b) \odot c + a \odot b = a \odot c + b \odot c + a \odot b$ is not an independent set, because $a \odot b = (a \odot c) \odot (b \odot c)$. Recalling a tuple in an independent tuple set is similar to recalling a code from a simple code set (they fall into the same QDS type, see Section 6.4).

Independent Product. A product of sets $expr = \odot_{i=1}^l s_i$ is called an independent product if and only if the multiplicand sets are disjoint and the sum of all the multiplicand sets $\sum_{i=1}^l s_i$ is an independent set. For example, $(a+b) \odot (c+d)$ is an independent product because these two sets are disjoint and a, b, c, d are mutually independent. $(a \odot b + a \odot c) \odot (b \odot c + b \odot d)$ is not an independent product because $a \odot b, a \odot c, b \odot c, b \odot d$ are not mutually independent, even though the two multiplicand sets are both independent sets. Note that a product of sets can also be flattened into a set of tuples, and the flattened set is usually not independent, e.g., $(a+b) \odot (c+d) = a \odot c + a \odot d + b \odot c + b \odot d$ is not an independent set because $(b \odot c) \odot (b \odot d) \odot (a \odot d) = a \odot c$.

6.2 Query-Data Structure (QDS) Predicates

We introduce the concept of a query-data structure (QDS) predicate, a unifying formalization that enables HEIM to implement an analysis that both leverages theoretical results from previous literature [Kleyko et al. 2022, 2021], and the novel derivations (Section 6.6). A query-data structure predicate is a set membership expression with the formulation $|q \cap ds| \geq k$ for which the symbolic match and not-match mean distances have been analytically derived. The HEIM analysis supports three forms of QDS predicates (tuple t and tuple set T can also be simply code c and code set C):

- $|\{t\} \cap T| \geq 1$ **Type I, Single Element, Independent Tuple-Set** [Section 6.4]
- $|T \cap T'| \geq k$ **Type II, Subset, Independent Tuple-Set** [Section 6.5]
- $|\{t\} \cap \odot_i T_i| \geq 1$ **Type III, Single Element, Independent Product** [Section 6.6]

Type I QDS predicates test if a code/tuple is in an independent code/tuple set, type II QDS predicates test if k elements of a code/tuple set is a subset of an independent code/tuple set, and type

III predicates test if a tuple is in an independent product. Each QDS type is associated with an independence constraint on the data structures. Type I QDS can be seen as a special case of type II QDS, but it is listed as a separate QDS as it is most frequently used. Note that these 3 QDS predicates systematically cover all cases when the data structure is independent, except for the subset query in the independent product. Subsets of independent products may have dependent tuples, which are theoretically challenging [Clarkson et al. 2023]. This kind of queries are rarely used in data structures [Kleyko et al. 2022, 2021]. We leave incorporating this kind of queries as future work.

QDS Classification. Given a query and data structure, HEIM classifies the query into 3 supported QDS predicates by inspecting the data structure and the query expression form. If the data structure is in the form of product of sums (bound tuple sets), then it must be a type III query and the query must be a tuple. Otherwise, the data structure must be an independent code/tuple set, in the form of sum or products (bundle of code tuples). Then if the query is a code set or code tuple set, it must be a Type II query. If the query is a code or code tuple, then it must be a Type I query.

Mutual Independence Constraints. Depending on the QDS Type, HEIM returns a mutual independence constraint that must hold over the data structure for the corresponding analysis to be valid. For sum-of-product formed data structure $expr$, the independence constraints $indepCstr = iset(expr)$ requires that $expr$ is an independent set. For product-of-sum formed data structure $expr$, the independence constraint $indepCstr = iproduct(expr)$ requires that $expr$ is an independent product. The definitions of independent set and product are in Section 6.1.

6.3 Most Common Not-Match Distribution - Independent Vectors

We start with the simplest and the most commonly used not-match distribution - the distance distribution between two independent/unrelated vectors $expr_1$ and $expr_2$. We denote the mean distance between two vectors $expr_1$ and $expr_2$ as $M(expr_1, expr_2) = E[dis(expr_1, expr_2)]$. In the following analysis, we assume all the code sets and code tuple sets are of odd size. The reason is that when bundling even number of vectors, the common practice is to add one more randomly generated vector as an operator to prevent the potential ties for majority [Kleyko et al. 2021].

Because Hamming distance is the average distance across all vector dimensions, the mean distance between two vectors is then the probability of them to differ in any one dimension. Since no correlation exists between two independent vectors, each dimension of one is equally likely (with probability 0.5) to be 0/1 (same/different) from the perspective of the other. The mean distance is therefore:

$$M(expr_1, expr_2) = \frac{1}{2}, \forall \text{ independent } expr_1, expr_2 \quad (13)$$

6.4 Type I (Single Element - Independent Set) Analytical Model

In this QDS type, the query is a code c /tuple t and the data structure is a code set C /tuple set T . Since the distance distributions in this QDS are the same for $|\{c\} \cap C| \geq 1$ or $|\{t\} \cap T| \geq 1$, we describe the $|\{c\} \cap C| \geq 1$ case. The query asks whether $c \in C$. Suppose $|C| = m$. In the not-match case $|\{c\} \cap C| < 1$, i.e., $c \notin C$, the mean distance is (13) as the two vectors have no correlation. The mean distance of the match case $c \in C$ has been derived by Kanerva [Kanerva et al. 1997]:

$$M(\{c_1\}, \{c_1, c_2, \dots, c_m\}) = \frac{1}{2} - \frac{\binom{m-1}{\frac{m-1}{2}}}{2^m} \quad (14)$$

For Type I QDS queries, equations (14) and (13) are the match mean distance $MeanDist_M$, and not-match mean distance ($MeanDist_{NM}$) respectively. The independence constraint ($indepCstr$) for Type I queries requires all codes be mutually independent. ($iset\{c_1, c_2, \dots, c_m\}$)

6.5 Type II (Subset - Independent Set) Analytical Model

QDS Type II queries are a generalization of Type I QDS. In this QDS type, the query is a code set C /tuple set T , and the data structure is a code set C' /tuple set T' . Again, in this QDS type, the distance distributions are the same for $|C \cap C'| \geq k$ or $|T \cap T'| \geq k$, we describe the $|C \cap C'| \geq k$ case. Assume $C' = \{c_1, c_2, \dots, c_m\}$, $C = \{c_1, c_2, \dots, c_n, c'_1, c'_2, \dots, c'_p\}$ has l ($l \leq m$) codes c_1, c_2, \dots, c_l also in C' and the p other codes c'_1, c'_2, \dots, c'_p not in C' . The mean distance of C and C' in this case has been derived by Kleyko et al. [Kleyko et al. 2016]:

$$M(\{c_1, c_2, \dots, c_l, c'_1, c'_2, \dots, c'_p\}, \{c_1, c_2, \dots, c_m\}) = 1 - \frac{1}{2^{p+m-1}} \sum_{i=0}^{\min(\frac{l+p-1}{2}, \frac{m-1}{2})} \binom{l}{i} \sum_{j=0}^{l+p-1-i} \binom{p}{j} \sum_{k=0}^{m-1-i} \binom{m-1}{k} \quad (15)$$

For Type II QDS queries, equation (15) implements the match mean distance ($MeanDist_M$) when $p = k$, and the not-match mean distance ($MeanDist_{NM}$) when $p = k - 1$. The Type II independence constraint (*indepCstr*) for Type II queries requires both subset and set codes be mutually independent ($iset\{c_1, c_2, \dots, c_m\} \wedge iset\{c_1, c_2, \dots, c_l, c'_1, c'_2, \dots, c'_p\}$).

6.6 Type III (Tuple-Set) Analytical Model

In this QDS type, the query vector is a code tuple t and the data-structure vector is a product-of-sum formed code tuple set $T = \odot_{i=1}^w T_i$, i.e., T is a binding of several code tuple sets. We derive for this QDS type because binding code sets is a common operation used in constructing data structures, e.g., analogical database [Kanerva 2010], finite-state-automata [Pashchenko et al. 2020]. Although T is also a tuple set (it can be flattened to sum-of-product form), this QDS type differs from type I in that the tuples in T have dependencies, while QDS type I assumes independence of tuples in the set. For example, consider $T = (c_1 + c_2) \odot (c_3 + c_4) = c_1 \odot c_3 + c_1 \odot c_4 + c_2 \odot c_3 + c_2 \odot c_4$. The first tuple $c_1 \odot c_3$ is the binding of the other three. The dependencies make the distance distributions different.

Since this analysis requires that $T = \odot_{i=1}^w T_i$ is an independent product (Section 6.1), enabling us to view tuples in $\cup_{1 \leq i \leq w} T_i$ as independent codes, in the following analysis we assume $T = \odot_{i=1}^w C_i$, i.e. T is a product of code sets, and the results generalize to the $T = \odot_{i=1}^w T_i$ case. We first consider the simple case where $T = C \odot C'$ is the product of two code sets. Assume $C = \{c_1, c_2, \dots, c_l\}$ is of size l and $C' = \{c'_1, c'_2, \dots, c'_m\}$ is of size m . The not-match case is $|\{t\} \cap T| < 1$, i.e., $t \notin T$, then the two independent vectors have mean distance (13). Otherwise, suppose $t = c_1 \odot c'_1$, we derive that:

$$M(c_1 \odot c'_1, \{c_1, c_2, \dots, c_l\} \odot \{c'_1, c'_2, \dots, c'_m\}) = \frac{1}{2^{l+m-2}} \left[\sum_{i=0}^{\frac{l-1}{2}} \binom{l-1}{i} \binom{\frac{m-3}{2}}{i} + \sum_{i=0}^{\frac{l-3}{2}} \binom{l-1}{i} \binom{\frac{m-1}{2}}{i} \right] \quad (16)$$

The derivation is as follows. Since binding is commutative, we have:

$$\begin{aligned} dis(c_1 \odot c'_1, \{c_1, c_2, \dots, c_l\} \odot \{c'_1, c'_2, \dots, c'_m\}) &= \frac{1}{n} |(c_1 \odot c'_1) \odot (\{c_1, c_2, \dots, c_l\} \odot \{c'_1, c'_2, \dots, c'_m\})| \\ &= \frac{1}{n} |(c_1 \odot \{c_1, c_2, \dots, c_l\}) \odot (c'_1 \odot \{c'_1, c'_2, \dots, c'_m\})| \\ &= dis(c_1 \odot C_1, c'_1 \odot C'_1) \end{aligned}$$

Therefore, for one dimension of t and T to differ, either in the dimension c_1 and C are the same while c'_1 and C' differ, or the c'_1 and C' are the same while c_1 and C differ. The probability of c_1 and C to be the same in a dimension is $\frac{1}{2^{l-1}} \sum_{i=0}^{\frac{l-1}{2}} \binom{l-1}{i}$ because it requires less than half (at most $\frac{l-1}{2}$) of c_2, c_3, \dots, c_l to differ from c_1 in the dimension, and the number of possible choices satisfying this is $\sum_{i=0}^{\frac{l-1}{2}} \binom{l-1}{i}$ and there are 2^{l-1} choices for c_2, c_3, \dots, c_l in total. For c_1 and C to be different in a dimension, there has to be at most $\frac{l-3}{2}$ of c_2, c_3, \dots, c_n to be the same as c_1 in the dimension, with probability

$\frac{1}{2^{l-1}} \sum_{i=0}^{\frac{l-3}{2}} \binom{l-1}{i}$. By symmetry, we also get the probability for c'_1 and C' to be the same or differ in one dimension. Combining them together gives (16). Note that the computation of (16) and (15) can be sped up by pre-computing the binomial coefficients and prefix sums of them with Pascal's triangle. More generally, T can be the binding of $w, \forall w \geq 2$ sets. Assume $T = \odot_{i=1}^w C_i$, and $C_i = \{c_{i1}, c_{i2}, \dots, c_{il_i}\}$ is a code set of size l_i . In this general case, we have:

$$M(\odot_{i=1}^w c_{i1}, \odot_{i=1}^w \{c_{i1}, c_{i2}, \dots, c_{il_i}\}) = \frac{1}{2^{\sum_{i=1}^w l_i - w}} \sum_{\sum_{i=1}^w e_i \text{ is odd}} \left[\prod_{j=1}^w \left(\sum_{k=0}^{\frac{l_j-1}{2} - e_j} \binom{l_j-1}{k} \right) \right] \quad (17)$$

The derivation is similar to (16). By commutativity of binding, we have:

$$\frac{1}{n} |(\odot_{i=1}^w c_{i1}) \odot (\odot_{i=1}^w \{c_{i1}, c_{i2}, \dots, c_{il_i}\})| = \frac{1}{n} |(\odot_{i=1}^w (c_{i1} \odot \{c_{i1}, c_{i2}, \dots, c_{il_i}\}))|$$

Therefore, denoting e_i as the value of one dimension of $C_{i1} \odot \{c_{i1}, c_{i2}, \dots, c_{il_i}\}$ (0 or 1), for $t \in T$ to be 1 in the dimension, odd number of e_i s should be 1, i.e., $\sum_{i=1}^w e_i$ is odd. The probability for each e_i to be 0 or 1 has been derived for (16). Adding the probability of all the independent cases gives (17). Note that for large w , (17) is non-trivial to compute, as there are exponential number of cases where $\sum_{i=1}^w e_i$ is odd. However, common HDC computations do not involve binding of more than 2 sets. We leave the problem of computing (17) more efficiently to future work. For Type III QDS queries, equation (17), or (16) when $w=2$ computes the match mean distance ($MeanDist_M$). Equation (13) computes the not-match mean distance ($MeanDist_{NM}$). The mutual independence constraint (*indepCstr*) for this QDS query is *iproduct*($\odot_{i=1}^w \{c_{i1}, c_{i2}, \dots, c_{il_i}\}$).

6.7 Hardware Error-Aware Mean Distance Model ($HwErr(hw, MeanDist)$)

HDC is a suitable computing paradigm for emerging hardware platforms because it is highly resilient against noises in them [Halawani et al. 2021; Imani et al. 2017b, 2019c; Karunaratne et al. 2020; Poduval et al. 2021]. HEIM incorporates the noise present in hardware, which works simultaneously for all the distance distributions above. We consider the bit-flip error model, where the probability of each bit in hyper-vectors to flip is p . Bit flips change the mean distance between two vectors – we denote $M'(expr_1, expr_2)$ as the mean distance of two vectors considering possible bit flips:

$$M'(expr_1, expr_2) = (p^2 + (1-p)^2)M(expr_1, expr_2) + 2p(1-p)(1 - M(expr_1, expr_2)) \quad (18)$$

The derivation is as follows. $M(expr_1, expr_2)$ is the probability of two vectors to be the same in one dimension, and $M'(expr_1, expr_2)$ is the probability considering bit flips. There are two cases where they are the same in one dimension with possible bit flips. First, they can be the same before possible bit flips with probability $M(expr_1, expr_2)$, and then two vectors either both have a bit flip, or both have no bit flip in the dimension, with probability $p^2 + (1-p)^2$. Second, they differ before possible bit flips with probability $1 - M(expr_1, expr_2)$, and a bit flip occurs only to one of the two vectors in this dimension, with probability $2p(1-p)$.

Hardware Errors Increase the Expected Distance Between Vectors. In all cases we consider, $M(expr_1, expr_2) \leq \frac{1}{2}$. The maximum $M(expr_1, expr_2)$ is $\frac{1}{2}$ when $expr_1$ and $expr_2$ are unrelated, as shown in (13), and relatedness of vectors makes their mean distance smaller. We show that possible bit flips increase the mean distances between vectors.

$$M'(expr_1, expr_2) - M(expr_1, expr_2) = 2p(1-p)(1 - 2M(expr_1, expr_2)) \geq 0 \quad (19)$$

Larger mean distance means closer to distribution of unrelated vector, implying loss of relation information encoded in the hypervectors. The implication is that hardware noise decreases the information resolution, which is intuitive. We note that information loss increases with p in a reasonable noise range, as in (19) $2p(1-p)$ increases monotonically for $0 < p < 0.5$. This enables us to use an upper bound of p in our analysis and deliver a sound accuracy guarantee.

Bit-Flip Probability. We derive the bit flip error probability from the hardware specification hw . As a standard practice, raw bit flip error rate is commonly used to characterize hardware noises. [Grossi et al. 2019; Le et al. 2021; Li et al. 2016] Given the hypervector operators and memory locations $op \in Op = \{\text{bind, bundle, codebook, item_mem, query}\}$ from the hardware specification, we denote the error of an operator as $err(op)$. We can compute the p as follows:

$$p = 1 - \prod_{op \in Op} (1 - err(op))$$

p captures the probability of at least one bit flip happens in certain operator or memory location. Note that p is a probability upper bound of a bit flip occurs in query or data structure during distance calculations, and using p delivers sound accuracy analysis as the information loss increases with p in a reasonable range $0 < p < 0.5$ (shown in (19)).

6.8 Correspondance between Mean Distance and Distance Distributions (ToNormal)

Given a mean distance $M'(expr_1, expr_2)$ considering bit flips and hypervector size n , we can derive the standard deviation to get the corresponding normal distribution $N(\mu, \sigma)$. Denote x as the value of one dimension in $expr_1 \odot expr_2$. Note that $dis(expr_1, expr_2) = \frac{1}{n} |expr_1 \odot expr_2|$. Since x is either 0 or 1, the following holds:

$$E[x^2] = E[x] = M'(expr_1, expr_2), Var[x] = E[x^2] - E^2[x] = M'(expr_1, expr_2)(1 - M'(expr_1, expr_2))$$

Since all n dimensions in $expr_1 \odot expr_2$ are independent and symmetric, we have

$$Var[dis(expr_1, expr_2)] = \frac{1}{n} Var[x] = \frac{1}{n} M'(expr_1, expr_2)(1 - M'(expr_1, expr_2))$$

To sum up, the distance distribution is determined by the mean distance and hypervector size n .

$$dis(expr_1, expr_2) \sim N\left(M'(expr_1, expr_2), \sqrt{\frac{1}{n} M'(expr_1, expr_2)(1 - M'(expr_1, expr_2))}\right) \quad (20)$$

6.9 Discussion on Model Simplifications

Use of Normal Distributions. We next justify the use of a normal distributions to model match and not-match distances. Recall, the distance metric is the Hamming distance, essentially the average distance in all dimensions. Since all the dimensions are symmetric, the distance of each dimension follows the same distribution. Therefore, the distance is the average of many i.i.d. variables. Furthermore, the hypervectors are long in HD computation, meaning that the number of i.i.d. variables averaged is large. By the central limit theorem, the distance distributions can be well approximated by normal distributions. In fact, the Kolmogorov–Smirnov difference (supremum of absolute distance) of the cumulative distribution functions (CDF) of the binomial distribution and its corresponding normal distribution is bounded by $\Omega(n^{-\frac{1}{2}})$ [Nagaev and Chebotarev 2011]. Approximation with normal distributions is also a standard practice by theoreticians in this field [Fraday et al. 2018], and we note that using binomial distributions to model hypervector distances is computationally expensive (costly to compute PDF and CDF, compared with normal distributions). Therefore, we view the distance distributions as normally distributed with a standard deviation and mean that both depend on the hypervector dimension, query and data structure sizes, and bit error probability.

Elimination of Permutation Operations. Because each bit of a code c is independently randomly generated, each bit of c and the corresponding bit of $\rho_k(c)$ are independent, and thus the distance distribution between c and permutations of it $\rho_k(c)$ are exactly the same as that of two independently generated codes, as in (13), unless k is equal to n (number of dimensions) or multiple of n times. Therefore, we may treat permuted codes as an independent code of the original codes.

```

1: procedure OPTIMIZE(hwModel,heimSpec,maxN)
2:   optDim = 0, queryParams = {}
3:   for queryCstr in getConstraints(heimSpec) do
4:     match queryCstr do
5:       case queryCstr is threshold-query
6:         success,indepCstr,minDim,thrs = thrAccAnalysis.binsearch(hw,0,maxN,queryCstr)
7:         queryParams = queryParams  $\cup$  {queryCstr,indepCstr,minDim,thrs}
8:       case queryCstr is wta-query
9:         success,indepCstr,minDim = wtaAccAnalysis.binsearch(hw,0,maxN,queryCstr)
10:        queryParams = queryParams  $\cup$  {queryCstr,indepCstr,minDim,[]}
11:   optDim = max(optDim,minDim)
12:   assert(success, "failed to find size that satisfies query accuracy requirements")
13:   return (optDim,queryParams)

```

Fig. 5. HEIM accuracy analysis

6.10 Discussion on Mutual Independence

Heim requires the HDC data structure and query to satisfy mutual independence constraints for the analysis to hold. We show that hypervectors that satisfy mutual independence can represent set, knowledge graph, analogical database, and NFA (see supplementary materials). Besides, a number of data structures, including stacks, sequences, and 2D images can be also be represented. These data structures are useful for signal and language classification, information retrieval, workload balancing, and analogical reasoning workloads. [Kleyko et al. 2021] In data structures with correlated information, independence can be induced by partitioning the data structure across multiple hypervectors, where each hypervector encodes mutually independent elements. Because hypervector size linearly increases with the number of stored elements (Figure 9a), the independent sub-hypervectors take up almost the same amount of space. We note this technique does not work well for applications where information loss induced by the bundling operation is a feature, such as feature encoding for machine learning applications, and cannot be used on queries with correlated elements.

Correlated Data Structures. The analysis of models with correlation is known to be hard and is an open problem in the HDC community. [Clarkson et al. 2023] This work establishes a core HDC analysis that is precise and sound. In the future, the analysis can be extended to directly support data structures that have correlations – these extensions would likely need to use overapproximations or use empirically derived information, and will not deliver the same guarantees as Heim’s core analysis.

Example. It is possible to implement data structures with correlations and satisfy HEIM’s mutual independence constraints. Consider the edge set $\{a \odot b, a \odot c, a \odot d, b \odot d, c \odot d\}$ for a 4-node graph. If we encode this set as one vector $G = a \odot b + a \odot c + a \odot d + b \odot d + c \odot d$, the expected distance of vectors of G and $a \odot b$ is $\frac{1}{2}$ (obtained by enumerating all 2^4 value combinations), totally indistinguishable from the distance of two independent vectors, although $a \odot b$ is a member of the set. In this case, no threshold can have a $> 50\%$ accuracy (random guessing). However, one can decompose a dependent set into a number of independent sets, each stored in one vector. Therefore, instead of storing the graph as a set of all edges G , we can store the set of incident edges of each node as one vector, similar to adjacency lists. This way, each vector is a mutually independent set, and a query falls into the QDS I predicate (Section 6.4).

7 HEIM OPTIMIZATION FRAMEWORK

Figure 5 presents the HEIM optimization algorithm. The optimizer takes a hardware error specification (**hwModel**), a HEIM specification (**heimSpec**), and a maximum hypervector size (**maxN**) as input, and returns both the smallest hypervector size that satisfies all accuracy constraints (**optDim**), query-optimized collection of hypervector sizes, independence constraints, and distance thresholds (**queryParams**) (line 13). HEIM iterates over each query accuracy constraint in the HEIM specification,

and derives the minimum hypervector size `minDim` required for the query, the mutual independence constraints `indepCstr` that must hold for the analysis to be sound, and a set of distance thresholds to use for threshold-based queries `thr` (lines 6-7, 9-10). The `wtaAccAnalysis.binsearch` and `thrAccAnalysis.binsearch` routines derive the minimum hypervector size `minDim` and associated thresholds and independence constraints for the given query by performing a binary search over hypervector sizes `0..maxN` and invoking the appropriate accuracy analysis for each candidate size. The final hypervector size returned by is the maximum required dimension across all queries. The `HEIM` returns early with an error if the accuracy analysis fails to find an appropriate size for any one query.

7.1 Dynamic Independence Checker

`HEIM` offers an optional dynamic checking algorithm that validates the concrete data structure and query hypervectors that meet all independence constraints associated with the selected query. The independence checker helps users build data structures that satisfy `HEIM`'s independence requirements and can be disabled if the user is sure these requirements are met. The algorithm tests if a concrete HD expression contains *mutually independent* tuples or codes. Intuitively, a set of elements is mutually independent if the existence of each set element does not depend on the existence of other set elements.

Algorithm. We present an efficient dynamic checking algorithm for validating the mutual independence of products and sets of tuples. Since the definition of product independence is derived from set independence, an efficient set independence checker can also check product independence. To avoid checking independence with a time-intensive exhaustive search, we derive an equivalent condition that can be efficiently checked to ensure independence. Given a tuple set $expr = \sum_{i=1}^n t_i$, denote c_1, c_2, \dots, c_m as the codes that are a factor of some t_i . For each tuple $t_i, 1 \leq i \leq n$ in $expr$, we create a binary vector V_i of length m , where the j -th element is 1 if c_j is a factor of t_i , and 0 otherwise. The equivalent condition of independence of set $expr = \sum_{i=1}^n t_i$ is that V_1, V_2, \dots, V_n are linearly independent in $GF(2)$. In other words, if we make a $n \times m$ matrix M , where the i -th row is V_i , the equivalent condition is that the M has rank n in $GF(2)$. For example, for $expr = a \odot c + b \odot c + a \odot b$, if c_1, c_2, c_3 are a, b, c respectively, the vectors for $a \odot c, b \odot c, a \odot b$ are $[1,0,1], [0,1,1], [1,1,0]$ respectively. These 3 vectors are not linearly independent because $[1,0,1] + [0,1,1] = [1,1,0]$ in $GF(2)$, so $expr$ is not an independent set. Intuitively, this is because the binary vector addition in $GF(2)$ represents the binding of tuples, e.g., $1+1=0$ in the previous example's last vector element is because binding $a \odot c$ and $b \odot c$ cancels out c . Calculating the rank of a $n \times m$ matrix can be done with a $O(n^2m)$ or $O(m^2n)$ Gaussian elimination algorithm.

Correctness. A binding of tuples corresponds to the sum of their V s in $GF(2)$. And for an index set $s \subset \{1, 2, \dots, n\}$, $t_i \odot_{j \in s} t_j$ is equivalent to a linear equation $V_i + \sum_{j \in s} V_j = \mathbf{0}$ in $GF(2)$. Therefore, when C is an independent set, there exist no such linear equations, which is equivalent to linear independence of V s.

Connection to Statistical Independence. From the formalism of the independence checker, we can derive that T is an independent set that is equivalent to that the bit values of $t_i \in T$ are statistically mutually independent. We consider the value of one dimension for all the vectors, and all the other dimensions follow the same arguments. Denote the values of the code vectors as \mathbf{x} , i.e., \mathbf{x}_i is the value of c_i 's vector. \mathbf{x} can be any of the 2^m values, each with probability 2^{-m} , as codes are independently randomly generated. For an assignment of the tuple vectors \mathbf{y} , we have $\mathbf{y} = M\mathbf{x}$. There are 2^n possible \mathbf{y} s. For each given assignment \mathbf{y} , the probability mass of it is 2^{-m} times the number of solutions \mathbf{x} of equation $M\mathbf{x} = \mathbf{y}$. Since M is rank n , it is guaranteed that there is at least one solution, and after Gaussian elimination, the Echelon form has $m-n$ free variables in \mathbf{x} , each of which can be either 0 or 1. This means that there are exactly 2^{m-n} solutions of \mathbf{x} , and thus the probability mass of this assignment \mathbf{y} is $2^{m-n} \cdot 2^{-m} = 2^{-n}$. Thus, the joint distribution of the vector bit values of t_i is a uniform distribution over all possible 2^n values (each bit is 0 or 1, with equal 50% probability). This is exactly the same as the joint distribution of n randomly generated codes, so we can analyze them as if they were independent codes.

Table 2. Summary of randomly generated data structure and query characteristics as a function of data structure size (k or (k, m)). Non-standard queries are described in grey. WTA queries only support matches.

benchmark	query type	data structure and query sizes				
set	threshold	50k-100k element sets, 1 element/query				
db-match	threshold	5k-10k fields/record, 50-100 records, m fields/query				
kgraph	threshold	1-100k edges/concept, 100k+10 concepts, 800k-1000k edges, 1 edge/query, 2 relations				
nfa	threshold	recognizes str with length k , 1- k character strings/query, 26 letters/alphabet				
<i>query: matches are substrings of str, non-matches are partial substrings of str</i>						
db-analogy	winner-take-all (WTA)	$k/2$ - k fields/record, 50m-100m records, 1 analogy query				
<i>query: select rows a, b where $\langle k, v \rangle \in a, \langle k, v' \rangle \in b$, infer v from item memory and v'</i>						
benchmark	size parameters	benchmark sizes				
set	k	1	2	3	4	5
db-match	(k, m)	(1,2)	(2,4)	(3,6)	(4,8)	(5,10)
kgraph	k	1	2	3	4	5
nfa	k	6	8	10	12	14
db-analogy	(k, m)	(4,1)	(8,2)	(12,3)	(16,4)	(20,5)

7.2 Discussion

Use of Binary Search. HEIM’s optimization algorithm exploits the fact that the accuracy of the HD computation increases monotonically with hypervector size to parametrize the HD computation efficiently. This has been shown theoretically [Fraday et al. 2018; Gallant and Okaywe 2013; Kleyko et al. 2022], and we also verify it in our evaluation (Figures 9a-9d). Because the accuracy is monotonic with respect to hypervector size, we can perform a binary search over hypervector sizes to identify the smallest hypervector size that satisfies a minimum accuracy requirement.

Metadata for Independence Checker. Many usage patterns involve building a data structure once and then querying the data structure hypervector. Once the data structure is checked for independence and built as hypervectors, no independence metadata about the data structure needs to be stored. Only the subset queries must be checked for independence when the data structure is queried. This can be done with the algorithm we described. We note it may be possible to directly embed this check in the encoding computation, which could be more lightweight.

8 EVALUATION ON ERROR-FREE HARDWARE

We evaluate HEIM on five analysis-amenable HDC-based data structures over five different data structure complexities (k or (k, m)) – Table 2 summarizes the complexity of the randomly generated data structures and queries at each size.¹ For example, the **set-100** benchmark has complexity $k = 100$, and is evaluated over random sets containing 50-100 elements and single element queries over the set. Each HEIM data structure parametrization is evaluated over 100 randomly generated data structures and 20 randomly generated match/not-match queries, where half the queries evaluate to "match". The accuracy of each data structure-query computation is evaluated over ten randomly sampled codebooks. All baseline and HEIM executions are evaluated over the same randomly sampled data structures, queries, and codebooks to reduce the effect of variance on the evaluation.

Query Accuracy Metric. Given \mathbf{P} matching query executions and \mathbf{N} not-matching query executions that produce \mathbf{TP} true positive and \mathbf{TN} true negative results, the accuracy of each benchmark is defined as one minus the average of the true positive and true negative rates ($\frac{1}{2}(\frac{\mathbf{TP}}{\mathbf{P}} + \frac{\mathbf{TN}}{\mathbf{N}})$). We employ a balanced strategy where false positive and negative rates are equally important since the real distributions of positive and negative queries depend on the target applications. HEIM supports unbalanced false positive and false negative rates, so unbalanced query distributions can also be

¹See supplementary materials for implementation details

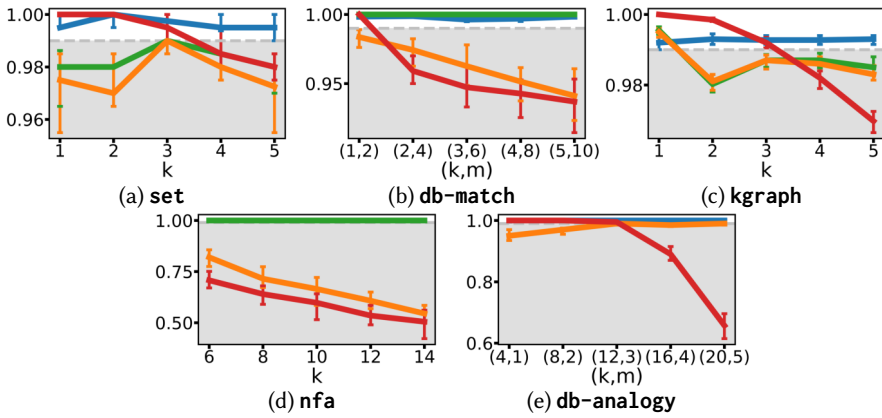


Fig. 6. Y-axis is median of the reported accuracies (0.0-1.0), error bars are 25% and 75% percentiles. Shaded area is below the target 99% accuracy. ■ for **Heim**, ■ for **dt-all**, ■ for **dt-par**, ■ for **dt-hybrid**.

handled. We also report the accuracy ratio (**rat**) for benchmark applications, corresponding to the percentage of random data structure instantiations satisfying the target accuracy.

Baselines. We evaluate HEIM-optimized hypervector size and threshold parametrizations against dynamic tuning-based baselines. These comparisons isolate the accuracy and performance benefits delivered by HEIM over traditional parameter tuning approaches. Each parameterization is optimized to deliver a target query accuracy of 99%, all dynamic tuning baselines use the dynamic tuning algorithm from Figure 2a, and all described baselines have error injection disabled:

- ▶ **Heim**: HEIM is used to statically optimize the distance thresholds and hypervector size for each benchmark with HEIM to get 99% accuracy. For queries composed of sets of elements (**db-match**, **nfa**), HEIM computes distance thresholds for different query set sizes.
- ▶ **dt-par**: The hypervector size is fixed at 10,000 bits and a single distance threshold is dynamically tuned to attain 99% query accuracy. WTA queries accept no settable query parameters and, therefore, do not have **dt-par** executions.
- ▶ **dt-all**: The hypervector size is dynamically tuned to find the smallest size between 1-100,000 bits that attains a query accuracy of 99%. For each candidate size, a single distance threshold is dynamically tuned to maximize accuracy. The search saturates at 100,000 bits.
- ▶ **dt-hybrid**: The hypervector size is dynamically tuned similarly to **dt-all**, but HEIM is used to find the theoretically optimal distance thresholds for each candidate hypervector size. This baseline isolates the effect of using HEIM-derived thresholds for hypervector queries.

8.1 Query Accuracy Comparison

Figure 6 compares the query accuracy of HEIM-optimized programs against the baseline executions. The plot charts the median (timeseries), Q1, and Q3 (vertical bar) for each execution, and the query accuracies that violate the 99% accuracy requirement are shaded grey. HEIM achieves 99.2%-100.0% median accuracy across all benchmarks – the required accuracy target of 99% is therefore met on expectation. Qualitatively, HEIM-optimized executions have low variance in accuracy (vertical bars) across trials and generally deliver consistent accuracy across different benchmark sizes compared to the dynamically tuned and statically sized baselines. Therefore, HEIM-optimized data structures reliably satisfy the desired query accuracy targets and consistently deliver the expected accuracy. We note that 80% of the executed trials exceed the 99% accuracy target (80% **rat**) across all benchmarks – we do not observe adherence to the accuracy constraint 100% of the time because HEIM’s guarantees hold on expectation.

	s-100	s-200	s-300	s-400	s-500	m-10-2	m-20-4	m-30-6	m-40-8	m-50-10	k-100	k-200	k-300
HEIM	3407	6807	10208	13608	17009	1376	4224	8633	14601	22129	3407	6807	10208
dt-par	0.34	0.68	1.02	1.36	1.70	0.14	0.42	0.86	1.46	2.21	0.34	0.68	1.02
dt-all	1.69	1.76	1.27	1.63	1.81	0.86	0.04	0.09	0.15	0.22	0.87	1.39	1.20
	2016	3865	8049	8324	9394	1604	100000	100000	100000	100000	3909	4908	8496
	k-400	k-500	n-6	n-8	n-10	n-12	n-14	a-4-100	a-8-200	a-12-300	a-16-400	a-20-500	
HEIM	13608	17009	3594	6517	10387	15235	21088	1509	6119	14157	25804	41186	
dt-par	1.36	1.70	0.36	0.65	1.04	1.52	2.11	0.15	0.61	1.42	2.58	4.12	
dt-all	1.21	1.29	0.04	0.07	0.10	0.15	0.21	2.22	1.96	1.52	1.62	1.51	
	11254	13195	100000	100000	100000	100000	100000	681	3125	9318	15897	27241	

Fig. 7. Baselines' hypervector size in each benchmark. **s**, **m**, **k**, **n**, **a** abbreviate **set**, **db-match**, **kgraph**, **nfa**, **db-analogy** respectively. Row 1, 4 present **dt-par**, **dt-all** hypervector sizes (in bits), and rows 2,3 present ratio of HEIM size to **dt-all** and **dt-par** size. **dt-par** uses hypervector size 10000 in all benchmarks. ■ and ■ cells have a **rat** of less than 50% and 80% respectively.

In contrast, the dynamically tuned (**dt-all**) baseline delivers 54.5%-99.5% median query accuracy, where only four benchmark evaluations (4 points) have median accuracies that meet the accuracy target of 99%. The **dt-all** benchmark evaluations also experience more significant fluctuations in query accuracy (vertical bars) than HEIM-optimized HD computations and experience degradations in accuracy as the benchmark size increases, likely because the empirically derived parametrizations do not generalize well, especially as the queries and data structures grow in complexity. Statically fixing the hypervector size to 10,000 bits and dynamically tuning only the hypervector threshold (**dt-par**) attains higher median accuracies than full dynamic tuning and HEIM when the data structure is small. For 3 of 5 benchmarks, **dt-par** achieves at least 99% median accuracy for the smaller 1-2 benchmark executions. The median accuracy of **dt-par** evaluations substantially degrades as the size of the data structures increases – over all executions, **dt-par** optimized programs attain a 50.5%-100.0% median accuracy. This phenomenon occurs because threshold-only tuning cannot expand the hypervector size to accommodate hypervectors that implement larger data structures and encode more information.

Hybrid Optimization. We also evaluate the query accuracy of a hybrid optimization approach (**dt-hybrid**) that uses HEIM to find thresholds, given dynamically tuned query size. For 24 of the 25 benchmark executions, **dt-hybrid** executions attain median accuracies that are 0.1%-45.5% higher than **dt-all**.² Notably, the **dt-hybrid** executions attain substantially better accuracy than **dt-all** on the **nfa** and **db-match** benchmarks, which supports the claim that specializing the distance threshold to the query size is important for queries containing multiple elements. We observe the **db-match** and **nfa** benchmarks use queries containing multiple elements and, therefore, likely work best when the distance thresholds are selected based on the query size. Because the dynamic tuning baselines only tune one threshold, the threshold may not work well across different query sizes. While it is possible to tune multiple thresholds dynamically, this would be prohibitively expensive to do with dynamic tuning.

8.2 Hypervector Size Comparison

Figure 7 compares the hypervector sizes of HEIM-optimized executions against dynamically tuned (**dt-all**) and statically sized (**dt-par**) executions. The red-shaded cells fail to meet the 99% median accuracy requirement, and the grey shaded cells meet the 99% accuracy requirement less than 80% of the time. For the 4 of 25 **dt-all** benchmark executions which achieve 99% median accuracy, **dt-all** finds 1.27-1.52x smaller hypervector sizes than HEIM for three executions, and a 1.15x larger hypervector size than HEIM for one execution. Of the benchmarks that do not meet the accuracy requirement,

²The **dt-hybrid** execution reports a 0.1% lower median accuracy than **dt-all** on the remaining execution due to sample randomness.




Fig. 8. Y-axis is measured runtime (s) for tuning the hypervector size and threshold, averaged over 10 runs on a single-core machine. The ■ is **Heim**, ■ is **dt-all**, ■ is **dt-par**. The **dt-par** trendline is omitted in the **db-analogy** figure, because the benchmark uses a WTA query and does not require a threshold.

9 of the 25 executions max out at the largest hypervector size. Therefore, dynamic tuning (**dt-all**) typically produces smaller hypervectors than **HEIM** when it can find a parametrization that meets the desired accuracy target, as **HEIM** employs a conservative strategy, but **dt-all** is rarely able to find a parametrization that reliably delivers a 99% query accuracy on average. In contrast, **HEIM** chooses larger hypervectors than dynamic tuning but reliably meets the accuracy target on expectation with low variance in all cases. In cases where dynamic tuning cannot meet the target query accuracy (size=100k), the dynamic tuning algorithm selects hypervector sizes 4.55-25x larger than **HEIM**. **HEIM** can analytically derive smaller hypervectors by searching over a larger parameter space that includes thresholds and sizes tailored to specific queries and data structures.

For the 10 of 25 **dt-par** benchmark executions which achieve 99% median accuracy, **HEIM** finds hypervectors that are 1.47x-7.14x smaller than 10k bits for 7 of 10 executions, and finds hypervectors that are 1.02x-1.42x larger than 10k bits for 3 of 10 executions, where 10k is the statically configured hypervector size. We also find **dt-par**'s accuracy degrades for larger benchmarks and selects unnecessarily large hypervectors for small benchmark executions. These issues arise because **cdt-par** cannot flexibly adjust the size to accommodate larger or smaller benchmark executions. Therefore, **HEIM** more consistently meets the target accuracy requirement than static allocation strategies while also delivering space savings for executions that can execute with smaller numbers of bits.

8.3 Optimization Time Comparison

Figure 8 compares the optimization runtimes for **HEIM** against the dynamic tuning baselines. **HEIM** completes its analysis in 85-3210 milliseconds, while **dt-all** takes 40 seconds to 22.72 hours to optimize the hypervector size and threshold. Fixing the hypervector size and dynamically tuning only the threshold (**cdt-par**) is substantially faster than full dynamic tuning, taking between 3.9 seconds and 168.8 seconds to compile. The **HEIM** optimizer is 303.0x-100167.4x faster than full dynamic tuning (**dt-all**) and 30.0x-874.4x faster than threshold-only dynamic tuning (**dt-par**), and generally scales better as the benchmark size increases. Because both dynamic tuning approaches are simulation-based, execution time scales poorly as the number of parameters to tune and the complexity of the data structure increases. In contrast, **HEIM**'s static analysis procedure is model-based and computes the optimal threshold and hypervector size in constant time. The performance of the parameter derivation algorithm is insensitive to the size of the optimized data structure, enabling scalable analysis.

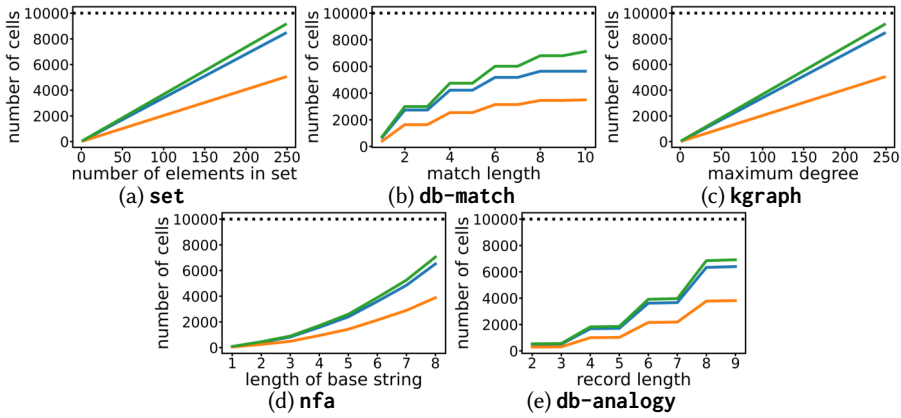


Fig. 9. Memory cells required to store each hypervector with **99%** accuracy. The \blacksquare is unoptimized hypervector size (10,000 bits, 1-bit-per-cell), \blacksquare is Heim, \blacksquare is Heim-2bpc, \blacksquare is Heim-3bpc. db-match records use 20 fields, db-analogy databases have 300 records.

9 EVALUATION OF HEIM ON EMERGING HARDWARE TECHNOLOGIES

HEIM’s accuracy analysis enables sound optimization of HD computations to execute with acceptable accuracy in the presence of hardware error. This capability enables the optimization of HD computations for error-prone emerging hardware technologies. We use HEIM to systematically study the benefits and drawbacks of using different emerging technologies for hyper-dimensional computation. We analyze the performance benefits offered by analog content-addressable memories (CAMs) (Section 9.2), and the storage benefits offered by analog multi-bit storage arrays. (Section 9.1).

9.1 HEIM Storage Density Analysis with MLC ReRAMs

We use HEIM to analyze the storage benefits of using multiple-bit-per-cell (MLC, k BPC) ReRAM-based item memories for HD computation against conventional one-bit-per-cell DRAM-based memory. We use HEIM to minimize the hypervector size while delivering a target query accuracy of **99%** for 2 BPC ReRAM, 3-BPC ReRAM, and conventional DRAM.

MLC ReRAMs. ReRAM is an emerging resistive memory technology that is prone to bit corruption but delivers fast access times, non-volatility, and improved density. We investigate the benefits of 2BPC and 3BPC ReRAM, which have raw bit error rates of 0.0215 and 0.1273 , respectively. All ReRAM error measurements were collected by characterizing a ReRAM storage array fabricated 130nm logic CMOS process in the BEOL with ECC disabled. [Hsieh et al. 2019; Le et al. 2021; Wei et al. 2023]

Analysis. Figure 9 presents the number of memory cells required to store each hypervector as a function of the benchmark size for conventional, 2BPC ReRAM (2bpc), and 3BPC ReRAM (3bpc) hardware platforms. HEIM produces hypervectors that require 76-8400 binary memory cells for conventional memory. HEIM produces 94-10064 bit hypervectors that use 47-5032 memory cells for 2 BPC ReRAM, netting an additional 1.614x-1.677x cell reduction over conventional binary memory. Though the overall hypervector size increases for 2 BPC memory, the 2x improvement in data density for this memory technology subsumes this size increase. HEIM produces 270-27361 bit hypervectors that use 90-9121 memory cells for 3 BPC ReRAM. Though 3 BPC ReRAM is denser than 2 BPC ReRAM, it does not net density improvements because the benchmark HD computations require substantially larger hypervectors to execute accurately in the presence of hardware error.

platform	processor			caches				memory	
	cores	threads/core	frequency	L1D	L1I	L2	L3	main	specialized
micro	2	1	1GHz	32KB/8w	64KB/8w	2MB/16w	16MB/16w	3GB	-
multi	10	2	2.8GHz	32KB/8w	32KB/8w	256KB/4w	20MB/16w	3GB	-
cam	2	1	1GHz	32KB/8w	64KB/8w	2MB/16w	16MB/16w	3GB	Analog CAM


Table 3. Benchmark hardware platforms. All caches are k -way (kw) associative, L1, L2 are private.

Fig. 10. Y-axis is simulated runtime (ms) of a single query with gem5 averaged over 20 runs. The ■ is micro, ■ is multi, ■ is cam.

9.2 HEIM Performance Analysis with Analog CAMs

We evaluate the benchmark applications on embedded (**micro**), multicore (**multi**), and emerging hardware (**cam**) platforms. Each architecture is simulated with the **x86 gem5** simulator using the processor and cache hierarchy presented in Table 3. [Binkert et al. 2011] The embedded and multicore baselines are based on the Intel Atom x7211E and Intel Core i9-10900E x86 architectures, respectively, and the CAM hardware platform uses the embedded architecture coupled with an analog CAM that efficiently performs item memory lookups to implement HD computations, but introduces error into the computation. To ensure an iso-accuracy comparison, we use HEIM to optimize each execution to achieve a 99% accuracy on the respective hardware platform.³ The benchmarks are parallelized and implemented in C, where the query hypervector is built in parallel, and the item memory/query distances are computed in parallel. For the CAM hardware platform, the software instead computes the item memory/query distances by dispatching a single query to the CAM.

The Analog CAM. The **cam** hardware platform has the same baseline characteristics as the embedded hardware platform but also offers a ReRAM-based Analog CAM that uses Ohm’s and Kirchhoff’s laws to perform in-memory, parallel hamming distance calculations against a query bit vector. [Imani et al. 2017b] Because the analog CAM both performs analog computation and uses an emerging device technology, the associated hamming distance calculations are unreliable and complete with a bit error rate of 0.14%. The analog CAM is extremely fast and completes the entire item memory query in 2.74-11.90 nanoseconds, provided 6-100 item memory rows are in use. The associated latency increases with row width, and the bit error rate increases with row width and device density. We parametrize the CAM to use 10k-bit rows, where $> 10k$ -bit hypervectors are split across rows; these row distances are summed in the analog domain with Kirchhoff’s law. The communication costs between the embedded system and the CAM are modeled as a DRAM access, and the threshold query latency

³for **db-analogy** simulation, we set the item memory and codebook sizes to be $\frac{1}{4}$ of the actual size to avoid memory issues with the **gem5-x86** simulator.

is conservatively approximated with the WTA lookup latency. We build performance and hardware error models by fitting regressions to the figures presented in the associated paper. [Imani et al. 2017b]

Analysis Figure 10 presents the simulated runtime of a single query as a function of the benchmark size, averaged over 20 executions. On average, the **multi**, **micro**, and **cam** executions take 0.589-5.418, 0.385-283.226, and 0.000-4.526 milliseconds respectively, depending on the benchmark. We observe that, unsurprisingly, the **multi** executions complete much faster than the **micro** executions for large benchmarks and only execute slightly slower on small benchmarks. The multicore platform has 5x more cores at its disposal and operates at 2.8x the frequency of the embedded system, so it, therefore, can complete HD computations much faster, provided the synchronization overheads amortize.

However, once the embedded platform is paired with a CAM, the embedded platform's performance becomes competitive and, in some cases, better than the multicore system. The **cam** executions are 2.16x-389.51x faster than **micro** executions for all benchmarks and 1.35x-62.55x faster than **multi** executions for 21 out of 25 benchmarks. The **cam** delivers substantial performance improvements because it significantly accelerates the item-memory search, which is usually the bottleneck when the query encoding is simple. We observe that these performance benefits hold, even though HEIM optimizes the CAM executions to use 1.011x-1.012x larger hypervectors to ensure iso-accuracy in the presence of hardware error. The **multi** executions are faster than **cam** in the largest 3 **db-match** benchmarks because the query encoding uses bundling operation and, therefore, becomes the bottleneck.

Discussion While HD computing is, at a glance, more resource-inefficient than classical computation for this class of applications, the HD computing paradigm enables the use of emerging hardware technologies that drastically accelerate computation and enable dense storage, such as CAMs and MLC ReRAM. These emerging technologies only implement highly restrictive subsets of computational operators and are error-prone, which can often lead to unpredictable effects on classical computations. Moreover, because the HDC model uses a distributed information representation, it is also amenable to several program optimizations that do not typically apply to classical programs. For example, the HEIM per-query hypervector sizes and thresholds can be used to compute hypervector distances over smaller sets of bits soundly or to terminate distance calculations early when a match or not-match is guaranteed. Moreover, the binding, bundling, and permutation operation implementations and hypervector sizing can be altered to improve performance, provided the distance relationships between input and output vectors hold. Second, because the encoded information is evenly distributed, algorithms can use highly unusual program transformations. For example, HD data structures can be combined by splicing hypervectors together, distances can be computed over any segment of hypervector bits in any order, queries over data structure hypervectors can be fielded in the middle of a data structure update, and HD queries can be prematurely interrupted to receive a computational result.

10 RELATED WORK

HDC/VSA. Hyperdimensional computing (HDC), or vector symbolic architecture (VSA), is a highly general cognitive computing paradigm that operates on binary, integer, real-valued, or complex-valued hypervectors [Kleyko et al. 2022, 2023b, 2021; Plate 1994, 2003; Yu et al. 2022]. Researchers have implemented VSA computations with emerging hardware platforms [Halawani et al. 2021; Imani et al. 2017b, 2019c; Karunaratne et al. 2020; Langenegger et al. 2023; Poduval et al. 2021]. We focus on HD computing with binary spatter codes [Kanerva et al. 1997] because this computational model is easy to implement in hardware, and amenable to theoretical analysis.

Theoretical Analysis of HD Computation. Kanerva derived the distance distributions for the set-recall [Kanerva et al. 1997], and Kleyko derived distance distributions for subset-recall [Kleyko et al. 2016]. We use Kanerva's and Kleyko's theoretical results to develop the Type I and Type II QDS analysis employed by HEIM. Researchers have also studied theoretical capacity and recall accuracy

of winner-take-all queries over VSA item memories [Frady et al. 2018; Gallant and Okaywe 2013; Kleyko et al. 2023a,c; Plate 1994; Thomas et al. 2021]. We extend the perception theory [Frady et al. 2018] to develop HEIM’s WTA accuracy analysis.

Horizontal Thresholding. Kleyko et al. derived an approach for horizontally thresholding distance distributions for subset-set recall queries. [Kleyko et al. 2016] This thresholding does not apply to Type III QDS queries, discards distance sub-ranges, and does not estimate query accuracy or solve for the optimal distance threshold. In contrast, HEIM analyzes a broader set of data structure queries (Type I-Type 3 QDS), and derives an optimal distance threshold to use. Our approach also uses the entire range of distance values when evaluating a query.

Hypervector Size Minimization for Classification. Prior work has primarily focused on dynamic approaches toward hypervector parameter optimization. Imani et al. dynamically tuned hypervector size to explore the trade-off between computational efficiency and accuracy. [Imani et al. 2018] Morris et al. decomposed computational hypervectors into lower-dimensional vectors of a dynamically selected size. [Morris et al. 2019] Basaklar et al. reduced the hypervector size by tuning the level hypervector construction for classification tasks. [Basaklar et al. 2021] Other works have tuned application-specific hypervector parameters, such as level hypervector chunk sizes, to reduce resource usage while maintaining classification accuracy. [Imani et al. 2019c] All approaches mentioned above are heuristic techniques that leverage dynamic tuning over representative inputs and therefore do not offer static guarantees. In contrast, HEIM employs statically sound analytical methods that deliver accuracy guarantees, even in the presence of hardware error.

11 CONCLUSION

We presented HEIM, a framework for *statically* optimizing HD computation parameters to minimize resource usage in the presence of hardware error. HEIM produced parametrizations that generalize across queries and data structures and satisfy a target accuracy on expectation. HEIM improved on dynamic parameter tuning-based approaches that potentially overfit to test data, provide no guarantees, and take orders of magnitude more time to find parametrizations. We demonstrated that HEIM’s analysis results could be leveraged to perform aggressive space-saving optimizations without compromising result fidelity and to systematically analyze emerging technologies’ benefits and drawbacks while maintaining iso-accuracy. With analysis and programming systems such as HEIM, we can enable the development of principled program optimizations that effectively reduce the resource requirements of HD computations without compromising accuracy.

ACKNOWLEDGMENTS

This research was supported by the Stanford SystemX Alliance and ACCESS – AI Chip Center for Emerging Smart Systems, sponsored by InnoHK funding, Hong Kong SAR.

DATA-AVAILABILITY STATEMENT

The software that supports Sections 8 and 9 is available on Zenodo [Yi and Achour 2023].

REFERENCES

- Sara Achour and Martin C Rinard. 2015. Approximate computation with outlier detection in topaz. *Acm Sigplan Notices* 50, 10 (2015), 711–730. <https://doi.org/10.1145/2858965.2814314>
- Toygun Basaklar, Yigit Tuncel, Shruti Yadav Narayana, Suat Gumussoy, and Umit Y Ogras. 2021. Hypervector design for efficient hyperdimensional computing on edge devices. *arXiv preprint arXiv:2103.06709* (2021). <https://doi.org/10.48550/arXiv.2103.06709>
- Nathan Binkert, Bradford Beckmann, Gabriel Black, Steven K Reinhardt, Ali Saidi, Arkaprava Basu, Joel Hestness, Derek R Hower, Tushar Krishna, Somayeh Sardashti, et al. 2011. The gem5 simulator. *ACM SIGARCH computer architecture news* 39, 2 (2011), 1–7. <https://doi.org/10.1145/2024716.2024718>

- Kenneth L Clarkson, Shashanka Ubaru, and Elizabeth Yang. 2023. Capacity Analysis of Vector Symbolic Architectures. *arXiv preprint arXiv:2301.10352* (2023). <https://doi.org/10.48550/arXiv.2301.10352>
- Manuel Eggimann, Abbas Rahimi, and Luca Benini. 2021. A 5 μw standard cell memory-based configurable hyperdimensional computing accelerator for always-on smart sensing. *IEEE Transactions on Circuits and Systems I: Regular Papers* 68, 10 (2021), 4116–4128. <https://doi.org/10.1109/TCSI.2021.3100266>
- E Paxon Frady, Denis Kleyko, and Friedrich T Sommer. 2018. A theory of sequence indexing and working memory in recurrent neural networks. *Neural Computation* 30, 6 (2018), 1449–1513. https://doi.org/10.1162/neco_a_01084
- Stephen I Gallant and T Wendy Okaywe. 2013. Representing objects, relations, and sequences. *Neural computation* 25, 8 (2013), 2038–2078. https://doi.org/10.1162/NECO_a_00467
- Ross W Gayler and Simon D Levy. 2009. A distributed basis for analogical mapping. In *New Frontiers in Analogy Research; Proc. of 2nd Intern. Analogy Conf*, Vol. 9.
- Alessandro Grossi, Elisa Vianello, Mohamed M Sabry, Marios Barlas, Laurent Grenouillet, Jean Coignus, Edith Beigne, Tony Wu, Binh Q Le, Mary K Wootters, et al. 2019. Resistive RAM endurance: Array-level characterization and correction techniques targeting deep learning applications. *IEEE Transactions on Electron Devices* 66, 3 (2019), 1281–1288. <https://doi.org/10.1109/TED.2019.2894387>
- Yasmin Halawani, Eman Hassan, Baker Mohammad, and Hani Saleh. 2021. Fused RRAM-based shift-add architecture for efficient hyperdimensional computing paradigm. In *2021 IEEE International Midwest Symposium on Circuits and Systems (MWSCAS)*. IEEE, 179–182. <https://doi.org/10.1109/MWSCAS47672.2021.9531748>
- Mike Heddes, Igor Nunes, Tony Givargis, Alexandru Nicolau, and Alex Veidenbaum. 2022. Hyperdimensional hashing: a robust and efficient dynamic hash table. In *Proceedings of the 59th ACM/IEEE Design Automation Conference*. 907–912. <https://doi.org/10.1145/3489517.3530553>
- ER Hsieh, M Giordano, B Hodson, A Levy, SK Osekowsky, RM Radway, YC Shih, W Wan, TF Wu, X Zheng, et al. 2019. High-density multiple bits-per-cell 1T4R RRAM array with gradual SET/RESET and its effectiveness for deep learning. In *2019 IEEE International Electron Devices Meeting (IEDM)*. IEEE, 35–6. <https://doi.org/10.1109/IEDM19573.2019.8993514>
- Mohsen Imani, Samuel Bosch, Sohun Datta, Sharadhi Ramakrishna, Sahand Salamat, Jan M Rabaey, and Tajana Rosing. 2019a. Quanthd: A quantization framework for hyperdimensional computing. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 39, 10 (2019), 2268–2278. <https://doi.org/10.1109/TCAD.2019.2954472>
- Mohsen Imani, Chenyu Huang, Deqian Kong, and Tajana Rosing. 2018. Hierarchical hyperdimensional computing for energy efficient classification. In *Proceedings of the 55th Annual Design Automation Conference*. 1–6. <https://doi.org/10.1145/3195970.3196060>
- Mohsen Imani, Deqian Kong, Abbas Rahimi, and Tajana Rosing. 2017a. Voicehd: Hyperdimensional computing for efficient speech recognition. In *2017 IEEE international conference on rebooting computing (ICRC)*. IEEE, 1–8. <https://doi.org/10.1109/ICRC.2017.8123650>
- Mohsen Imani, Abbas Rahimi, Deqian Kong, Tajana Rosing, and Jan M Rabaey. 2017b. Exploring hyperdimensional associative memory. In *2017 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 445–456. <https://doi.org/10.1109/HPCA.2017.28>
- Mohsen Imani, Sahand Salamat, Saransh Gupta, Jiani Huang, and Tajana Rosing. 2019b. Fach: Fpga-based acceleration of hyperdimensional computing by reducing computational complexity. In *Proceedings of the 24th Asia and South Pacific Design Automation Conference*. 493–498. <https://doi.org/10.1145/3287624.3287667>
- Mohsen Imani, Sahand Salamat, Behnam Khaleghi, Mohammad Samragh, Farinaz Koushanfar, and Tajana Rosing. 2019c. Sparsehd: Algorithm-hardware co-optimization for efficient high-dimensional computing. In *2019 IEEE 27th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*. IEEE, 190–198. <https://doi.org/10.1109/FCCM.2019.00034>
- Michael N Jones and Douglas JK Mewhort. 2007. Representing word meaning and order information in a composite holographic lexicon. *Psychological review* 114, 1 (2007), 1. <https://doi.org/10.1037/0033-295X.114.1.1>
- Pentti Kanerva. 2009. Hyperdimensional computing: An introduction to computing in distributed representation with high-dimensional random vectors. *Cognitive computation* 1, 2 (2009), 139–159.
- Pentti Kanerva. 2010. What we mean when we say "What's the dollar of Mexico?": Prototypes and mapping in concept space. In *2010 AAAI fall symposium series*.
- Pentti Kanerva. 2014. Computing with 10,000-bit words. In *2014 52nd annual Allerton conference on communication, control, and computing (Allerton)*. IEEE, 304–310. <https://doi.org/10.1109/ALLERTON.2014.7028470>
- Pentti Kanerva. 2018. Computing with high-dimensional vectors. *IEEE Design & Test* 36, 3 (2018), 7–14. <https://doi.org/10.1109/MDAT.2018.2890221>
- Pentti Kanerva et al. 1997. Fully distributed representation. *PAT* 1, 5 (1997), 10000.
- Geethan Karunaratne, Manuel Le Gallo, Giovanni Cherubini, Luca Benini, Abbas Rahimi, and Abu Sebastian. 2020. In-memory hyperdimensional computing. *Nature Electronics* 3, 6 (2020), 327–337. <https://doi.org/10.1038/s41565-023-01357-8>

- Yeseong Kim, Mohsen Imani, Niema Moshiri, and Tajana Rosing. 2020. Geniehd: Efficient dna pattern matching accelerator using hyperdimensional computing. In *2020 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 115–120. <https://doi.org/10.23919/DATE48585.2020.9116397>
- Denis Kleyko, Connor Bybee, Ping-Chen Huang, Christopher J Kymn, Bruno A Olshausen, E Paxon Frady, and Friedrich T Sommer. 2023a. Efficient decoding of compositional structure in holistic representations. *Neural Computation* 35, 7 (2023), 1159–1186. https://doi.org/10.1162/neco_a_01590
- Denis Kleyko, Mike Davies, Edward Paxon Frady, Pentti Kanerva, Spencer J Kent, Bruno A Olshausen, Evgeny Osipov, Jan M Rabaey, Dmitri A Rachkovskij, Abbas Rahimi, et al. 2022. Vector Symbolic Architectures as a Computing Framework for Emerging Hardware. *Proc. IEEE* 110, 10 (2022), 1538–1571. <https://doi.org/10.1109/JPROC.2022.3209104>
- Denis Kleyko, Evgeny Osipov, Alexander Senior, Asad I Khan, and Yaşar Ahmet Şekerciogğlu. 2016. Holographic graph neuron: A bioinspired architecture for pattern processing. *IEEE transactions on neural networks and learning systems* 28, 6 (2016), 1250–1262. <https://doi.org/10.1109/TNNLS.2016.2535338>
- Denis Kleyko, Dmitri Rachkovskij, Evgeny Osipov, and Abbas Rahimi. 2023b. A survey on hyperdimensional computing aka vector symbolic architectures, part ii: Applications, cognitive models, and challenges. *Comput. Surveys* 55, 9 (2023), 1–52. <https://doi.org/10.1145/3558000>
- Denis Kleyko, Dmitri A Rachkovskij, Evgeny Osipov, and Abbas Rahimi. 2021. A Survey on Hyperdimensional Computing aka Vector Symbolic Architectures, Part I: Models and Data Transformations. *ACM Computing Surveys (CSUR)* (2021). <https://doi.org/ASurveyonHyperdimensionalComputingakaVectorSymbolicArchitectures>
- Denis Kleyko, Abbas Rahimi, Ross W Gayler, and Evgeny Osipov. 2020. Autoscaling bloom filter: controlling trade-off between true and false positives. *Neural Computing and Applications* 32 (2020), 3675–3684. <https://doi.org/10.1007/s00521-019-04397-1>
- Denis Kleyko, Antonello Rosato, Edward Paxon Frady, Massimo Panella, and Friedrich T. Sommer. 2023c. Perceptron Theory Can Predict the Accuracy of Neural Networks. *IEEE Transactions on Neural Networks and Learning Systems* (2023), 1–15. <https://doi.org/10.1109/TNNLS.2023.3237381>
- Jovin Langenegger, Geethan Karunaratne, Michael Hersche, Luca Benini, Abu Sebastian, and Abbas Rahimi. 2023. In-memory factorization of holographic perceptual representations. *Nature Nanotechnology* (2023), 1–7. <https://doi.org/10.1038/s41565-023-01357-8>
- Binh Q Le, Akash Levy, Tony F Wu, Robert M Radway, E Ray Hsieh, Xin Zheng, Mark Nelson, Priyanka Raina, H-S Philip Wong, Simon Wong, et al. 2021. RADAR: A fast and energy-efficient programming technique for multiple bits-per-cell RRAM arrays. *IEEE Transactions on Electron Devices* 68, 9 (2021), 4397–4403. <https://doi.org/10.1109/TED.2021.3097975>
- Haitong Li, Tony F Wu, Abbas Rahimi, Kai-Shin Li, Miles Rusch, Chang-Hsien Lin, Juo-Luen Hsu, Mohamed M Sabry, S Burc Eryilmaz, Joon Sohn, et al. 2016. Hyperdimensional computing with 3D VRRAM in-memory kernels: Device-architecture co-design for energy-efficient, error-resilient language recognition. In *2016 IEEE International Electron Devices Meeting (IEDM)*. IEEE, 16–1. <https://doi.org/10.1109/IEDM.2016.7838428>
- Sasa Misailovic, Michael Carbin, Sara Achour, Zichao Qi, and Martin C Rinard. 2014. Chisel: Reliability-and accuracy-aware optimization of approximate computational kernels. *ACM Sigplan Notices* 49, 10 (2014), 309–328. <https://doi.org/10.1145/2714064.2660231>
- Fabio Montagna, Abbas Rahimi, Simone Benatti, Davide Rossi, and Luca Benini. 2018. PULP-HD: Accelerating brain-inspired high-dimensional computing on a parallel ultra-low power platform. In *2018 55th ACM/ESDA/IEEE Design Automation Conference (DAC)*. IEEE, 1–6. <https://doi.org/10.1145/3195970.3196096>
- Justin Morris, Mohsen Imani, Samuel Bosch, Anthony Thomas, Helen Shu, and Tajana Rosing. 2019. CompHD: Efficient hyperdimensional computing using model compression. In *2019 IEEE/ACM International Symposium on Low Power Electronics and Design (ISLPED)*. IEEE, 1–6. <https://doi.org/10.1109/ISLPED.2019.8824908>
- SV Nagaev and VI Chebotarev. 2011. On the bound of proximity of the binomial distribution to the normal one. In *Doklady Mathematics*, Vol. 83. Springer, 19–21. <https://doi.org/10.1134/S1064562411010030>
- Evgeny Osipov, Denis Kleyko, and Alexander Legalov. 2017. Associative synthesis of finite state automata model of a controlled object with hyperdimensional computing. In *IECON 2017-43rd Annual Conference of the IEEE Industrial Electronics Society*. IEEE, 3276–3281. <https://doi.org/10.1109/IECON.2017.8216554>
- Dmitry V Pashchenko, Dmitri A Trokoz, Alexey I Martyshev, Mihail P Sinev, and Boris L Svistunov. 2020. Search for a substring of characters using the theory of non-deterministic finite automata and vector-character architecture. *Bulletin of Electrical Engineering and Informatics* 9, 3 (2020), 1238–1250. <https://doi.org/10.11591/eei.v9i3.1720>
- Tony A Plate. 1994. *Distributed representations and nested compositional structure*. Citeseer.
- Tony A Plate. 2000. Analog retrieval and processing with distributed vector representations. *Expert systems* 17, 1 (2000), 29–40. <https://doi.org/10.1111/1468-0394.00125>
- Tony A Plate. 2003. Holographic Reduced Representation: Distributed representation for cognitive structures. (2003).
- Prathyush Poduval, Zhuowen Zou, Hassan Najafi, Houman Homayoun, and Mohsen Imani. 2021. Stochd: Stochastic hyperdimensional system for efficient and robust learning from raw data. In *2021 58th ACM/IEEE Design Automation*

- Conference (DAC)*. IEEE, 1195–1200. <https://doi.org/10.1109/DAC18074.2021.9586166>
- Dmitri A Rachkovskij and Serge V Slipchenko. 2012. Similarity-based retrieval with structure-sensitive sparse binary distributed representations. *Computational Intelligence* 28, 1 (2012), 106–129. <https://doi.org/10.1111/j.1467-8640.2011.00423.x>
- Abbas Rahimi, Simone Benatti, Pentti Kanerva, Luca Benini, and Jan M Rabaey. 2016. Hyperdimensional biosignal processing: A case study for EMG-based hand gesture recognition. In *2016 IEEE International Conference on Rebooting Computing (ICRC)*. IEEE, 1–8. <https://doi.org/10.1109/ICRC.2016.7738683>
- Abbas Rahimi, Sohun Datta, Denis Kleyko, Edward Paxon Frady, Bruno Olshausen, Pentti Kanerva, and Jan M Rabaey. 2017. High-dimensional computing as a nanoscalable paradigm. *IEEE Transactions on Circuits and Systems I: Regular Papers* 64, 9 (2017), 2508–2521. <https://doi.org/10.1109/TCSL.2017.2705051>
- Abbas Rahimi, Pentti Kanerva, Luca Benini, and Jan M Rabaey. 2018. Efficient biosignal processing using hyperdimensional computing: Network templates for combined learning and classification of exg signals. *Proc. IEEE* 107, 1 (2018), 123–143. <https://doi.org/10.1109/JPROC.2018.2871163>
- Kenny Schlegel, Florian Mirus, Peer Neubert, and Peter Protzel. 2021. Multivariate time series analysis for driving style classification using neural networks and hyperdimensional computing. In *2021 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 602–609. <https://doi.org/10.1109/IV48863.2021.9576028>
- Kenny Schlegel, Peer Neubert, and Peter Protzel. 2022. HDC-MiniROCKET: Explicit time encoding in time series classification with hyperdimensional computing. In *2022 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 1–8. <https://doi.org/10.1109/IJCNN55064.2022.9892158>
- Hashim Sharif, Yifan Zhao, Maria Kotsifakou, Akash Kothari, Ben Schreiber, Elizabeth Wang, Yasmin Sarita, Nathan Zhao, Keyur Joshi, Vikram S Adve, et al. 2021. ApproxTuner: a compiler and runtime system for adaptive approximations. In *Proceedings of the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*. 262–277. <https://doi.org/10.1145/3437801.3446108>
- Max M Shulaker, Tony F Wu, Asish Pal, Liang Zhao, Yoshio Nishi, Krishna Saraswat, H-S Philip Wong, and Subhasish Mitra. 2014. Monolithic 3D integration of logic and memory: Carbon nanotube FETs, resistive RAM, and silicon FETs. In *2014 IEEE International Electron Devices Meeting*. IEEE, 27–4. <https://doi.org/10.1109/IEDM.2014.7047120>
- Chris Simpkin, Ian Taylor, Graham A Bent, Geeth de Mel, Swati Rallapalli, Liang Ma, and Mudhakar Srivatsa. 2019. Constructing distributed time-critical applications using cognitive enabled services. *Future Generation Computer Systems* 100 (2019), 70–85. <https://doi.org/10.1016/j.future.2019.04.010>
- Justin Theiss, Jay Leverett, Daeil Kim, and Aayush Prakash. 2022. Unpaired Image Translation via Vector Symbolic Architectures. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXI*. Springer, 17–32. https://doi.org/10.1007/978-3-031-19803-8_2
- Anthony Thomas, Sanjoy Dasgupta, and Tajana Rosing. 2021. Theoretical Foundations of Hyperdimensional Computing. *Journal of Artificial Intelligence Research* 72 (2021), 215–249. <https://doi.org/10.48550/arXiv.2010.07426>
- Anjiang Wei, Akash Levy, Pu (Luke) Yi, Robert Radway, Priyanka Raina, Subhasish Mitra, and Sara Achour. 2023. PBA: Percentile-Based Level Allocation for Multiple-Bits-Per-Cell RRAM. In *ICCAD*.
- Tony F Wu, Haitong Li, Ping-Chen Huang, Abbas Rahimi, Gage Hills, Bryce Hodson, William Hwang, Jan M Rabaey, H-S Philip Wong, Max M Shulaker, et al. 2018. Hyperdimensional computing exploiting carbon nanotube FETs, resistive RAM, and their monolithic 3D integration. *IEEE Journal of Solid-State Circuits* 53, 11 (2018), 3183–3196. <https://doi.org/10.1109/JSSC.2018.2870560>
- Thomas Yerxa, Alexander Anderson, and Eric Weiss. 2018. The hyperdimensional stack machine. *Cognitive Computing* (2018), 1–2.
- Pu (Luke) Yi and Sara Achour. 2023. Artifact for the OOPSLA 2023 Article "Hardware-Aware Static Optimization of Hyperdimensional Computations". <https://doi.org/10.5281/zenodo.8329813>
- Tao Yu, Yichi Zhang, Zhiru Zhang, and Christopher M De Sa. 2022. Understanding hyperdimensional computing for parallel single-pass learning. *Advances in Neural Information Processing Systems* 35 (2022), 1157–1169.

Received 2023-04-14; accepted 2023-08-27