

# Molecular Classification Using Hyperdimensional Graph Classification

Pere Vergés, Igor Nunes, Mike Heddes, Tony Givargis and Alexandru Nicolau  
University of California Irvine

pvergesb@uci.edu, igord@uci.edu, mheddes@uci.edu, givargis@uci.edu, nicolau@ics.uci.edu

**Abstract**—Our work introduces an innovative approach to graph learning by leveraging Hyperdimensional Computing. Graphs serve as a widely embraced method for conveying information, and their utilization in learning has gained significant attention. This is notable in the field of chemoinformatics, where learning from graph representations plays a pivotal role. An important application within this domain involves the identification of cancerous cells across diverse molecular structures. We propose an HDC-based model that demonstrates comparable Area Under the Curve results when compared to state-of-the-art models like Graph Neural Networks (GNNs) or the Weisfieler-Lehman graph kernel (WL). Moreover, it outperforms previously proposed hyperdimensional computing graph learning methods. Furthermore, it achieves noteworthy speed enhancements, boasting a 40x acceleration in the training phase and a 15x improvement in inference time compared to GNN and WL models. This not only underscores the efficacy of the HDC-based method, but also highlights its potential for expedited and resource-efficient graph learning.

**Index Terms**—Hyperdimensional Computing, Graph Kernels, Classification, Machine Learning

## I. INTRODUCTION

Molecular classification is a crucial task in pharmaceutical research [1], [2], gaining increased prominence in recent years. Various techniques, including statistical models, decision trees [3], [4], Bayesian models [5], [6], kernel-based methods [7]–[9], neural networks [3], [10], and sub-graph mining techniques [11], [12], are employed in this field. The quantitative structure-activity relationship (QSAR) [6] serves as the foundation for these techniques, leveraging the assumption that molecular structure reflects multiple chemical properties. In Figure 1, we show an example of the conversion of benzene to a graph. Analyzing the graph structure of molecules provides valuable insights into their relationships and similarities.




Figure 1. Benzene to graph representation [12].

Among the molecular classification approaches, methods based on feature vector representation have emerged, par-

ticularly utilizing chemical descriptors and fingerprints. Fingerprints [13]–[16], for instance, generate fixed-width bit-vectors by enumerating cycles and linear paths, offering a comprehensive representation of molecular features.

Another set of approaches involves mining frequent substructures based on biological and chemical properties of molecules [17], [18]. Such methods generate feature vectors indicating the presence of distinct properties, requiring a frequency threshold to differentiate between frequent and infrequent substructures. Notably, studies suggest that infrequent substructures are more effective in discriminating between different molecule types [12].

This work integrates the insights of infrequent substructures learned from graph mining works [12] with Hyperdimensional Computing (HDC) [19]. As Machine Learning gains popularity for solving complex problems [20], HDC has emerged as a favored solution for resource-constrained scenarios [21], [22], spanning voice and gesture recognition [23], [24], natural language processing [25], graph learning [26], [27], and bioinformatics [28]. HDC appeal lies in its ability to abstract information efficiently.

Beyond HDC, there is a growing interest in graph learning techniques, specifically those focused on graph similarity [29]. Graph kernels, commonly coupled with support vector machines (SVM) [30], have gained attention. Nevertheless, their computational cost scales quadratically, rendering them unsuitable for online learning and real-time applications [31]. Another popular approach is the use of Graph Neural Networks (GNN). However, GNNs present challenges due to their computational time complexity and energy costs, thereby limiting their applicability in real-world scenarios, such as embedded or IoT applications [32]–[34] that are subject to strict design constraints.

Our research introduces a novel method for detecting cancerous cells by comparing the cell structures of cancerous and non-cancerous cells and employing graph classification based on HDC. The approach is based on sub-graph kernel matching [35], taking into account infrequent cell patterns to refine the learned associative memory using HDC adaptive learning techniques. The primary objective of our method is to improve efficiency while ensuring a level of accuracy comparable to state-of-the-art graph learning methods in the context of cancerous cell detection. The applicability of our work extends to various real-time scenarios requiring online learning for timely execution [12], [36], [37], such as the

evaluation of evolving cells in human liver diseases [38] or gliomas [39]; and in anomaly detection for molecular data [40], [41], where approaches like GNNs or Weisfeiler-Lehman (WL) kernels may prove too expensive for real-time conditions.

## II. RELATED WORK

In this section, we present related works on graph learning methods, covering techniques for molecular classification.

### A. Graph Learning Methods

*a) Kernel methods:* are popular machine learning techniques for comparing data points within a graph using similarity metrics. The literature encompasses various kernel methods, each focusing on different properties of a graph. Notable examples include the Random Walk Kernel [42], Tree Pattern Kernel [43], Shortest-Path Kernel [44], Optimal Assignment Kernel [45], Graphlet Kernel [46], Weisfeiler-Lehman Kernel [36], Subgraph Matching Kernel [12], among others. From this extensive set, we will focus on comparing our results with the state-of-the-art Weisfeiler-Lehman (WL) kernel.

The WL kernel iteratively refines label assignments for graph nodes based on their neighbors' structure. The iterative algorithm updates node labels until no further changes occur. The execution cost depends on the neighborhood's number of hops, increasing with larger values. An enhancement to the WL algorithm involves combining it with Optimal Assignment [47], determining node pairs with minimal cost in graph comparisons according to a similarity metric.

*b) Graph Neural Network (GNN):* Recent advances in Graph Neural Networks (GNNs) involve iteratively updating representations by exchanging information with their neighbors, and once trained, they are used to perform classification. Despite the popularity of GNNs, graph kernels remain competitive. In fact, it has been shown that GNNs are, at most, as powerful as the Weisfeiler-Lehman (WL) test in graph discernment [48].

*c) Fingerprints:* These are well-known techniques in cheminformatics for representing molecules using feature vectors [15], [16]. These fingerprints are generated by enumerating all molecular graph substructures and extracting predefined relevant substructures. Feature vectors can then be compressed using hashing techniques. In order to perform classification using fingerprints one can use similarity metrics to define graph similarity. This is typically done using the Tanimoto coefficient [49].

*d) Subgraph mining:* This is another prevalent approach applicable to large graphs. Various techniques focus on mining frequent patterns, with GraphSig [12] yielding superior results. GraphSig mines significant subgraphs, translating them into the feature vector space by evaluating statistical significance patterns, which helps to use the most relevant features of the graph when performing graph classification, which is performed using similarity measures. This technique highlights that infrequent patterns are more significant than frequent ones.

## III. HYPERDIMENSIONAL COMPUTING (HDC)

In this section, we provide a brief introduction to HDC, explain how classification is performed, and introduce two graph encoding approaches proposed in the literature.

### A. Information Representation

In HDC, information is encoded as points in a high-dimensional space (*hyperspace*)  $\mathcal{H}$ , known as *hypervectors*. Given input data  $x \in \mathcal{X}$ , a mapping  $\phi: \mathcal{X} \rightarrow \mathcal{H}$  produces the *encoding*  $\phi(x)$ . The core concept is that similar inputs in the original space map to similar vectors in the hyperspace.

### B. Operations

Once information is in the hyperspace, operations like *superposition*, *binding*, and *permutation* are applied to create composite representations. These operations preserve hypervector dimensionality, generating a hypervector in the same hyperspace as the operands. Their composability allows for versatile encoding tailored to diverse applications, capturing data compositionality effectively. The *similarity* metric evaluates hypervector relations and resemblance.

*a) Similarity metric:* The similarity between hypervectors is given by  $\delta: \mathcal{H} \times \mathcal{H} \rightarrow \mathbb{R}$ , and is used to determine relations, crucial for classification problems [50]. Usually, the similarity metric is implemented using cosine similarity, dot product, or hamming distance.

*b) Superposition:* The superposition operation aggregates information into a single hypervector representation using element-wise vector addition. This operation produces a vector maximizing similarity to the original operands. In certain hyperspaces, an inverse operation to bundling exists, removing an item from the bundled representation by applying superposition with the additive inverse [51], [52].

*c) Binding:* The binding operation associates two hypervectors, yielding a hypervector dissimilar to both inputs. Widely employed for information association, similar to assigning values to variables, binding has an inverse operation (unbinding) involving the multiplicative inverse. Depending on the hyperspace, the binding operation can be self-inverse or may require a different operation [51], [52].

*d) Permutation:* The permutation operator is used for imposing order on hypervectors. This operation yields a dissimilar hypervector, enabling the assignment of specific orders within the hyperspace. Importantly, the inverse operation allows exact retrieval of the original input hypervector, ensuring reversibility in the permutation process [53].

### C. Hyperdimensional Classification

Hyperdimensional classification [54] comprises two stages. Firstly, the training stage involves learning from the provided input data. Secondly, the inference stage involves the model predicting the class of unseen samples.

a) *Training*: The initial training step involves encoding input data samples into hypervectors by employing the HDC operations. Subsequently, the encoded sample is added to the associative memory, which is comprised of hypervectors that represent the prototypes of the different classes. This addition entails element-wise addition of the encoded sample with the corresponding class hypervector, identified by its label.

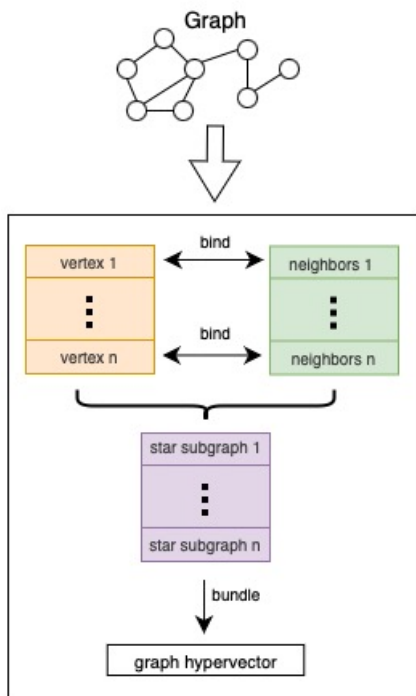


Figure 2. Proposed graph encoding.

b) *Inference*: The inference process consists of three steps. Similar to training, the initial inference step involves encoding input data into the hyperdimensional space. Ensuring consistency and accuracy, it is crucial to use the same encoding for both phases. Following encoding, the next step computes the similarity between the encoded sample and associative memory. This yields a list of distance measures between the sample and different classes. Finally, the class most similar to the encoded sample is selected from the list of similarities for predicting the input data class.

#### D. Hyperdimensional graph encodings

The first work that proposed encoding graphs into hyperspace is due to Gayler and Levy [55]. Poduval et al. [27] built up on it and demonstrated different properties of the encoding. Another work, named GraphHD [26], proposed creating an ordering of the nodes using a node centrality metric when assigning each one its hypervector for classification. This work shows that using HDC on certain datasets gives comparable results to state-of-the-art models such as GNNs or WL kernels, while having faster training and inference times.

Our method uses hyperdimensional classification method for encoding the graph into hypervectors. The approach takes

Table I  
CENTRALITY METRICS AUC PERFORMANCE USING GRAPHHD.

Centrality metric	AUC
None (Gayler & Levy)	<b>61.12</b>
PageRank (GraphHD)	59.57
Degree centrality	59.70
Closeness centrality	58.99
Betweenness centrality	59.57
Load centrality	59.84
Subgraph centrality	59.38
Harmonic centrality	58.83

on the insights that are learnt about infrequent patterns from subgraph mining techniques, by creating the encoding representation using star subgraphs and giving more importance to the graph infrequent patterns by using RefineHD [56]. This allows our method to have similar execution times to HDC approaches while achieving comparable accuracy results to the state-of-the-art graph learning methods.

Our approach employs a hyperdimensional classification method, which involves encoding the graph into hypervectors. By leveraging the knowledge gained from subgraph mining techniques, particularly focusing on infrequent patterns, and subgraph matching kernels, our method employs star subgraphs in the encoding representation. To emphasize the importance of graph infrequent patterns, we utilize RefineHD [56]. This strategy ensures that our approach maintains execution times similar to HDC approaches, while also delivering accuracy results comparable to state-of-the-art graph learning methods.

#### IV. OUR METHOD

The encoding process in HDC poses a significant challenge and holds paramount importance as it dictates how data is captured. Our primary focus in this study is on defining a method to encode the graph of cells effectively, with the goal of capturing maximal information. This capability is crucial for distinguishing whether a molecule exhibits an active or inactive cell line. Consequently, our input comprises structural graphs, and we must subsequently map this information into the hyperdimensional space.

Before introducing our method, we conducted an evaluation of centrality metrics to be applied to the GraphHD encoding, since their work showed great results when using PageRank centrality metric [57]. Our findings indicate that centrality metrics are not effective in capturing the graph’s information when applied to HDC and molecule graphs. Table I illustrates the average AUC (Area Under the Curve), on the PTC\_FM, MUTAG, NCI1, ENZYMES, and PROTEINS datasets, of GraphHD using different centrality metrics. Notably, the first scenario, employing random hypervectors without considering centrality metrics, outperforms all cases where nodes are ordered using such metrics. This observation underscores the limited efficacy of centrality metrics in enhancing the AUC of the encoding process.

Our approach is inspired by subgraph-kernel matching. One approach that uses subgraph matching kernel [35] counts the

---

**Algorithm 1** Graph classification algorithm: Training

---

**Input:** Training set of graphs and their class labels  $\mathcal{G}$ **Output:** Memory  $\mathcal{M}$  containing the set of class hypervectors $C_1, \dots, C_k$ **for**  $(g, i) \in \mathcal{G}$  **do** $H_g \leftarrow \emptyset$ **for** node  $v \in V(g)$  **do** $H_v \leftarrow \phi(v)$ **for**  $u \in \text{neighbours}(v)$  **do** $H_v \leftarrow \text{bind}(H_v, \phi(u))$ **end for** $H_g \leftarrow \text{bundle}(H_g, H_v)$ **end for** $\mathcal{M} \leftarrow \text{RefineHD}(H_g, i, \mathcal{M})$ **end for**

---

number of matching subgraphs between two graphs. Taking this idea into account, we generated a graph encoding that creates a histogram of star subgraphs. Figure 2 depicts the proposed graph encoding.

By performing a histogram of star subgraphs, we make our representation more complex compared to the previously proposed methods. This helps to make a more drastic distinction between different graph patterns. Moreover, our method builds on the idea shown in [12], where they found that when classifying molecules using their structure, in this study [12] was shown that subgraphs that appear in the graph with low frequency carry more information than frequent graphs. For this reason, we decided to apply adaptive learning to our classification algorithm. This will help in giving less weight to frequent graphs and subgraphs and giving more value to those that are more infrequent.

Table II  
ADAPTIVE LEARNING AUC.

Model	Add	AdaptHD	OnlineHD	RefineHD
AUC	76.32	89.36	89.41	<b>90.51</b>

In HDC, adaptive learning is a common technique to improve AUC while maintaining good efficiency, although this technique can also be applied to perform iterative training, which augments its training time considerably. There are three proposed models: AdaptHD [58], OnlineHD [59], and RefineHD [56]. We have evaluated all these models, and we have seen that RefineHD is the one that gives the best AUC for this application. In Table II, we show the comparison between the different methods of adding information to the associative memory. The first method (Add) adds all graph encodings to the memory. The second method (AdaptHD) only adds misclassified patterns to the memory. The third method (OnlineHD) only adds misclassified patterns but in a weighted manner, taking into account the cosine similarity of the sample and the class. The last method (RefineHD) adds both classified and misclassified samples in a weighted manner but only adds a sample when the cosine similarity

Table III  
REFINEHD THRESHOLD VALUE.

Threshold	1	1.5	1.8	2
AUC	90.51	91.00	<b>91.09</b>	90.98

to the correct class is below a threshold, which is partially computed using the similarities of the incorrectly classified samples. From these methods, we see that RefineHD is the one that behaves the best, having a +1% improvement compared to the other adaptive learning methods.

Since RefineHD was the method that behaved the best, we decided to conduct a study and see what threshold performed better. In Table III, we see in the results that having a threshold value between 1.5 and 2 gives the best results.

We show the training pseudocode of our graph classification in Algorithm 1, for each class we will encode all graphs using the proposed star subgraph encoding that uses bind as the HDC to build the subgraphs, in order to add a graph to its class we use the RefineHD method.

---

**Algorithm 2** Graph classification algorithm: Inference

---

**Input:** Graph  $g$  and trained memory  $\mathcal{M}$  containing the set of class hypervectors  $C_1, \dots, C_k$ **Output:** Predicted label  $\hat{y}$  $H_g \leftarrow \emptyset$ **for** node  $v \in V(g)$  **do** $H_v \leftarrow \phi(v)$ **for**  $u \in \text{neighbours}(v)$  **do** $H_v \leftarrow \text{bind}(H_v, \phi(u))$ **end for** $H_g \leftarrow \text{bundle}(H_g, H_v)$ **end for** $\hat{y} \leftarrow \arg \max_i (\delta(H_g, C_i))$ 

---

For inference, it is necessary to encode the graph using the hyperdimensional vectors that represent each node. After encoding, perform the cosine similarity operation against the previously trained memory. This will return a value between 0 and 1 for each class. The resulting prediction will be the one that has the value closest to one, indicating higher similarity to the corresponding class.

## V. EXPERIMENTS

To evaluate our method, we utilized eleven anticancer screen datasets from the TUDataset [29], which were extracted from PubChem<sup>1</sup>. The method was implemented using TorchHD [60], a Python library dedicated to HDC. We compared the performance of our model against both previous HDC graph classification methods and state-of-the-art graph classification methods. We chose to compare it against the 1-WL sparse graph kernel, as the dense 1-WL and WL-OA are too computationally expensive for cancer cells detection. Additionally, we conducted a comparison with a GNN model

<sup>1</sup><https://pubchem.ncbi.nlm.nih.gov>

Table IV  
ANTICANCER SCREEN DATASETS.

Dataset	Size	Description
MCF-7	28972	Breast
MOLT-4	41810	Leukemia
NCI-H23	42164	Non-small Cell lung
OVCAR-8	42386	Ovarian
PC-3	28679	Prostate
P388	46440	Leukemia
SF-295	40350	Central nervous system
SN12C	41855	Renal
SW-620	42405	Colon
UACC-257	41864	Melanoma
Yeast	83933	Yeast anticancer

based on the Graph Isomorphism Network (GIN) [61]. The configuration of these two methods adhered to the default settings from the benchmark evaluation of the TUDataset. All experiments were executed on a machine with 2 x 6-core Intel Xeon X5680 @ 3.33GHz, 96GB RAM, 64-bit CentOS Linux 7. For each experiment, we utilized 10,000 dimensions and conducted 10 repetitions.

#### A. Dataset

The datasets were extracted from PubChem, which contains a compilation of biological activities of different molecules. We utilized the records of anticancer screen data tested against various cancer cell lines. Each of these datasets contains molecules that have been tested against cancer cells, along with the results indicating whether they are active or inactive. A summary of the datasets we have used is presented in Table IV.

Table V  
ADAPTIVE LEARNING AUC RESULTS.

Dataset	Baseline	AdaptHD	OnlineHD	RefineHD
MCF-7	54.64	86.81	86.63	<b>88.43</b>
MOLT-4	54.15	87.25	86.57	<b>88.24</b>
NCI-H23	57.74	91.20	91.73	<b>92.64</b>
OVCAR-8	56.88	91.40	91.44	<b>92.44</b>
PC-3	56.49	90.33	90.44	<b>91.41</b>
P388	56.52	91.00	91.14	<b>92.14</b>
SF-295	57.97	91.70	91.81	<b>92.27</b>
SN12C	58.01	91.77	91.82	<b>92.70</b>
SW-620	55.87	90.10	89.67	<b>91.27</b>
UACC257	57.00	92.60	92.70	<b>93.78</b>
Yeast	51.14	78.84	79.53	<b>80.33</b>

#### B. Adaptive learning

The first experiment we conducted aimed to determine which adaptive learning algorithm would perform best for detecting cancerous cells. As explained in Section III-C, we compared the hyperdimensional classification baseline to AdaptHD, OnlineHD, and RefineHD. As shown in Table V, RefineHD performed the best in terms of AUC results for all datasets. In Table VI, we present the average time per sample for both training and inference. Despite RefineHD being slower compared to the other methods, we chose it for

our method since it yielded the best performance across all other methods.

Table VI  
ADAPTIVE LEARNING TECHNIQUES TRAIN AND INFERENCE TIMES PER SAMPLE.

VSA	Baseline	AdaptHD	OnlineHD	RefineHD
Train time	0.42	0.57	0.58	0.62
Inference time	0.53	0.65	0.64	0.69

#### C. Vector Symbolic Architectures

The second experiment we conducted involved evaluating our proposed graph encoding using RefineHD and trying different Vector Symbolic Architectures (VSA) [62]. We experimented with Multiply Add Permute (MAP), Fourier Holographic Reduced Representation (FHRR), and Vector-Derived Transformation Binding (VTB).

Table VII shows the performance results with the three different VSA, and we observe that in terms of AUC, the resulting values are very similar. However, in terms of execution time, the MAP architecture is the most efficient one.

Table VII  
VECTOR SYMBOLIC ARCHITECTURES PERFORMANCE.

VSA	MAP	FHRR	VTB
AUC	90.51	90.78	90.33
Train time	0.82	0.92	5.36
Inference time	0.87	0.97	5.38

#### D. Methods comparison

The third experiment involved comparing our method to both hyperdimensional graph classification methods (GraphHD and GraphHD) and state-of-the-art graph models (1-WL sparse and GIN). We attempted to execute against dense 1-WL and WLOA; however, the execution time was too large to be applied to cancer cell detection, since the execution of just one dataset was taking several days.

In Figure 3, we observe the AUC of all the methods. As we can see, the two hyperdimensional classification models have considerably lower AUC values, lower by 20% and 25% compared to our proposed method. Compared to the state-of-the-art models, we have comparable results, lower by 3% and 4%. In terms of efficiency, our method is 1.25 times faster than (Gayler & Levy encoding) and 1.33 times slower than GraphHD during training, with similar times for inference. Compared to the state-of-the-art models, we achieve large speedups in terms of training. Against WL, we get a 31x speedup, and a 41x speedup compared to GIN. In terms of inference, we have a speedup of 15x compared to WL and a slight speedup of 1.07x against GIN. All these results are shown in Tables VIII and IX.

#### E. Scalability

Finally, the last experiment consisted of evaluating the scalability performance of this method. It aimed to assess both




Figure 3. Method comparison techniques AUC.

Table VIII  
METHOD COMPARISON TRAIN TIME.

	Gayler & Levy	GraphHD	Ours	1-WL	GIN
MCF-7	0.18	0.10	0.13	3.73	5.67
MOLT-4	1.30	0.72	0.95	32.28	40.74
NCI-H23	1.75	0.96	1.27	42.44	55.18
OVCAR-8	1.15	0.63	0.84	26.80	36.43
PC-3	1.21	0.66	0.89	23.85	37.15
P388	0.61	0.39	0.46	15.25	12.80
SF-295	1.36	0.75	0.99	31.94	45.72
SN12C	1.18	0.65	0.87	27.79	37.76
SW-620	0.12	0.06	0.09	2.95	3.93
UACC257	0.01	0.01	0.01	0.44	0.60
Yeast	0.03	0.02	0.02	1.25	0.66

Table IX  
METHOD COMPARISON INFERENCE TIME.

	Gayler & Levy	GraphHD	Ours	1-WL	GIN
MCF-7	0.20	0.11	0.14	1.65	0.18
MOLT-4	1.45	0.83	1.01	15.95	1.12
NCI-H23	1.94	1.12	1.36	21.60	1.39
OVCAR-8	1.28	0.73	0.89	14.19	0.91
PC-3	1.35	0.77	0.95	10.49	1.31
P388	0.70	0.45	0.50	7.9	0.51
SF-295	1.51	0.87	1.06	16.76	1.09
SN12C	1.33	0.76	0.93	14.54	0.95
SW-620	0.13	0.07	0.09	1.54	0.09
UACC257	0.02	0.01	0.01	0.23	0.02
Yeast	0.03	0.02	0.02	0.74	0.02

Table X  
AVERAGE AUC AND TRAIN/INFERENCE TIMES COMPARISON.

Methods	AUC	Train time	Inference time
Gayler & Levy	68.73 (x0.75)	0.81 (x1.25)	0.91 (x1.42)
GraphHD	66.76 (x0.73)	0.45 (x0.75)	0.53 (x0.82)
Ours	90.93 -	0.60 -	0.64 -
WL	94.97 (x1.04)	18.98 (x31.6)	9.60 (x15)
GIN	93.68 (x1.03)	25.15 (x41.91)	0.69 (x1.07)

the AUC improvement achievable by adding more dimensions and the effect on training time and inference time. In Figure 4, we can observe that adding dimensions up to 50,000 results in an improvement in AUC. However, going beyond that range does not provide further improvement, as the AUC plateaus. Additionally, the training and inference times increase linearly with the number of dimensions.

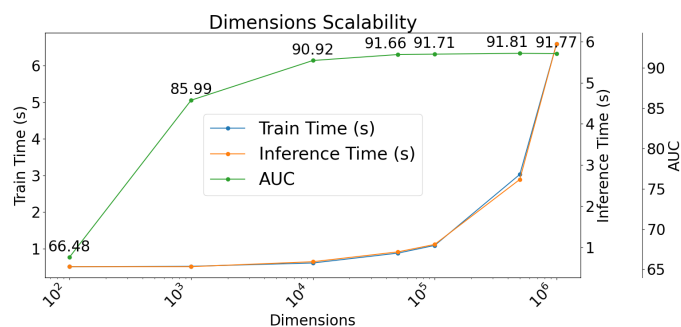


Figure 4. Method scalability.

## VI. CONCLUSIONS

This work introduces a novel graph encoding using HDC, employing star-subgraphs to represent graph information in hyperspace, and utilizing adaptive learning through the RefineHD approach to prioritize infrequent subgraph patterns. The focus is on representing graphs of molecules where a cancer cell is either active or inactive. The comparison with two graph encodings proposed in the literature using HDC demonstrates an improvement in accuracy of up to 25% while maintaining a similar execution time. Additionally, our method is compared to state-of-the-art graph learning approaches such as Graph Neural Networks (GNNs) and the Weisfeiler-Lehman kernel, achieving comparable accuracy results and significant speedups of 40x during training and 15x during inference.

## REFERENCES

- [1] A. Bender and R. C. Glen, "Molecular similarity: a key technique in molecular informatics," *Org. Biomol. Chem.*, vol. 2, pp. 3204–3218, 2004. [Online]. Available: <http://dx.doi.org/10.1039/B409813G>
- [2] P. Willett, "Searching techniques for databases of two- and three-dimensional chemical structures," *Journal of Medicinal Chemistry*, vol. 48, no. 13, pp. 4183–4199, 2005, pMID: 15974568. [Online]. Available: <https://doi.org/10.1021/jm0582165>
- [3] A. Rusinko, M. W. Farnen, C. G. Lambert, P. L. Brown, and S. S. Young, "Analysis of a large structure/biological activity data set using recursive partitioning," *Journal of Chemical Information and Computer Sciences*, vol. 39, no. 6, pp. 1017–1026, 1999, pMID: 10614024. [Online]. Available: <https://doi.org/10.1021/ci9903049>
- [4] D. Zmuidinavicius, R. Didziapetris, P. Japertas, A. Avdeef, and A. Petrauskas, "Classification structure-activity relations (c-sar) in prediction of human intestinal absorption," *Journal of Pharmaceutical Sciences*, vol. 92, no. 3, pp. 621–633, 2003. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0022354916311996>
- [5] P. Watson, "Naïve bayes classification using 2d pharmacophore feature triplet vectors," *Journal of Chemical Information and Modeling*, vol. 48, no. 1, pp. 166–178, 2008, pMID: 18183968. [Online]. Available: <https://doi.org/10.1021/ci7003253>

- [6] P. Labute, "Binary qsar: a new method for the determination of quantitative structure activity relationships," *Pacific Symposium on Biocomputing*, no. 1, pp. 444–55, 1999, pMID: 18183968. [Online]. Available: [https://doi.org/10.1142/9789814447300\\_0044](https://doi.org/10.1142/9789814447300_0044)
- [7] H. Fröhlich, J. K. Wegner, F. Sieker, and A. Zell, "Optimal assignment kernels for attributed molecular graphs," *Proceedings of the 22nd international conference on Machine learning*, 2005. [Online]. Available: <https://api.semanticscholar.org/CorpusID:1459319>
- [8] G. Harper, J. Bradshaw, J. Gittins, D. Green, and A. Leach, "Prediction of biological activity for high-throughput screening using binary kernel discrimination," *Journal of chemical information and computer sciences*, vol. 41, no. 5, p. 1295–1300, 2001. [Online]. Available: <https://doi.org/10.1021/ci000397q>
- [9] K.-R. Müller, G. Rätsch, S. Sonnenburg, S. Mika, M. Grimm, and N. Heinrich, "Classifying 'drug-likeness' with kernel-based learning methods," *Journal of Chemical Information and Modeling*, vol. 45, no. 2, pp. 249–253, 2005, pMID: 15807485. [Online]. Available: <https://doi.org/10.1021/ci049737o>
- [10] G. Keserü, L. Molnár, and I. Greiner, "A neural network based virtual high throughput screening test for the prediction of cns activity," *Combinatorial chemistry high throughput screening*, vol. 3, no. 6, p. 535–540, December 2000. [Online]. Available: <https://doi.org/10.2174/1386207003331346>
- [11] X. Yan, H. Cheng, J. Han, and P. S. Yu, "Mining significant graph patterns by leap search," in *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, ser. SIGMOD '08. New York, NY, USA: Association for Computing Machinery, 2008, p. 433–444. [Online]. Available: <https://doi.org/10.1145/1376616.1376662>
- [12] S. Ranu and A. K. Singh, "Graphsig: A scalable approach to mining significant subgraphs in large graph databases," in *2009 IEEE 25th International Conference on Data Engineering*, 2009, pp. 844–855.
- [13] N. Brown, "Chemoinformatics—an introduction for computer scientists," *ACM Comput. Surv.*, vol. 41, no. 2, feb 2009. [Online]. Available: <https://doi.org/10.1145/1459352.1459353>
- [14] G. W. Adamson and J. A. Bush, "A method for the automatic classification of chemical structures," *Information Storage and Retrieval*, vol. 9, no. 10, pp. 561–568, 1973. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0020027173900594>
- [15] A. Capecchi, D. Probst, and J.-L. Reymond, "One molecular fingerprint to rule them all: Drugs, biomolecules, and the metabolome," *Journal of Cheminformatics*, vol. 12, 06 2020.
- [16] L. David, A. Thakkar, R. Mercado, and O. Engkvist, "Molecular representations in ai-driven drug discovery: a review and practical guide," *Journal of Cheminformatics*, vol. 12, 09 2020.
- [17] A. Mrzic, P. Meysman, W. Bittremieux, P. Moris, B. Cule, B. Goethals, and K. Laukens, "Grasping frequent subgraph mining for bioinformatics applications," *BioData Mining*, vol. 11, 2018. [Online]. Available: <https://api.semanticscholar.org/CorpusID:52184639>
- [18] S. Naulaerts, P. Meysman, W. Bittremieux, T. N. Vu, W. Berghe, B. Goethals, and K. Laukens, "A primer to frequent itemset mining for bioinformatics," *Briefings in bioinformatics*, vol. 16, 10 2013.
- [19] P. Kanerva, "Hyperdimensional computing: An introduction to computing in distributed representation with high-dimensional random vectors," *Cognitive Computation*, vol. 1, pp. 139–159, 2009. [Online]. Available: <https://api.semanticscholar.org/CorpusID:733980>
- [20] A. Fawzi, M. Balog, A. Huang, T. Hubert, B. Romera-Paredes, M. Barekatin, A. Novikov, F. J. R. Ruiz, J. Schrittwieser, G. Swirszcz, D. Silver, D. Hassabis, and P. Kohli, "Discovering faster matrix multiplication algorithms with reinforcement learning," *Nature*, vol. 610, no. 7930, pp. 47–53, 2022.
- [21] P. Vergés, M. Heddes, I. Nunes, T. Givargis, and A. Nicolau, "Hdcc: A hyperdimensional computing compiler for classification on embedded systems and high-performance computing," 2023.
- [22] A. Abusnaina, A. Khormali, H. Alasmay, J. Park, A. Anwar, and A. Mohaisen, "Adversarial learning attacks on graph-based iot malware detection systems," in *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*, 2019, pp. 1296–1305.
- [23] A. Rahimi, S. Benatti, P. Kanerva, L. Benini, and J. M. Rabaey, "Hyperdimensional biosignal processing: A case study for emg-based hand gesture recognition," in *International Conference on Rebooting Computing (ICRC)*. IEEE, 2016, pp. 1–8.
- [24] M. Imani, D. Kong, A. Rahimi, and T. Rosing, "Voicehd: Hyperdimensional computing for efficient speech recognition," in *International Conference on Rebooting Computing (ICRC)*. IEEE, 2017, pp. 1–8.
- [25] A. Rahimi, P. Kanerva, and J. M. Rabaey, "A robust and energy-efficient classifier using brain-inspired hyperdimensional computing," in *International Symposium on Low Power Electronics and Design (ISLPED)*, 2016, pp. 64–69.
- [26] I. Nunes, M. Heddes, T. Givargis, A. Nicolau, and A. Veidenbaum, "Graphhd: Efficient graph classification using hyperdimensional computing," in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2022.
- [27] P. Poduval, H. Alimohamadi, A. Zakeri, F. Imani, M. H. Najafi, T. Givargis, and M. Imani, "Graphd: Graph-based hyperdimensional memorization for brain-like cognitive learning," *Frontiers in Neuroscience*, vol. 16, 02 2022.
- [28] Y. Kim, M. Imani, N. Moshiri, and T. Rosing, "Geniehd: Efficient dna pattern matching accelerator using hyperdimensional computing," in *2020 Design, Automation Test in Europe Conference Exhibition (DATE)*, 2020, pp. 115–120.
- [29] C. Morris, N. M. Kriege, F. Bause, K. Kersting, P. Mutzel, and M. Neumann, "Tudataset: A collection of benchmark datasets for learning with graphs," in *ICML 2020 Workshop on Graph Representation Learning and Beyond (GRL+ 2020)*, 2020. [Online]. Available: [www.graphlearning.io](http://www.graphlearning.io)
- [30] S. S. Keerthi, O. Chapelle, D. DeCoste, K. P. Bennett, and E. Parrado-Hernández, "Building support vector machines with reduced classifier complexity," *J. Mach. Learn. Res.*, vol. 7, pp. 1493–1515, 2006. [Online]. Available: <https://api.semanticscholar.org/CorpusID:13375961>
- [31] Y. Zhan and D. Shen, "Increasing the efficiency of support vector machine by simplifying the shape of separation hypersurface," in *Computational and Information Science*, J. Zhang, J.-H. He, and Y. Fu, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 732–738.
- [32] R. Khan, S. U. Khan, R. Zaheer, and S. Khan, "Future internet: The internet of things architecture, possible applications and key challenges," in *2012 10th International Conference on Frontiers of Information Technology*, 2012, pp. 257–260.
- [33] X. Qi and C. Liu, "Enabling deep learning on iot edge: Approaches and evaluation," in *2018 IEEE/ACM Symposium on Edge Computing (SEC)*, 2018, pp. 367–372.
- [34] P. Verges, I. Nunes, M. Heddes, T. Givargis, and A. Nicolau, "Accelerating permute and n-gram operations for hyperdimensional learning in embedded systems," in *2023 IEEE 29th International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA)*. Los Alamitos, CA, USA: IEEE Computer Society, sep 2023, pp. 253–260. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/RTCSA58653.2023.00037>
- [35] N. Kriege and P. Mutzel, "Subgraph matching kernels for attributed graphs," in *Proceedings of the 29th International Conference on International Conference on Machine Learning*, ser. ICML'12. Madison, WI, USA: Omnipress, 2012, p. 291–298.
- [36] N. Shervashidze, P. Schweitzer, E. J. van Leeuwen, K. Mehlhorn, and K. M. Borgwardt, "Weisfeiler-lehman graph kernels," *Journal of Machine Learning Research*, vol. 12, no. 77, pp. 2539–2561, 2011. [Online]. Available: <http://jmlr.org/papers/v12/shervashidze11a.html>
- [37] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, "A comprehensive survey on graph neural networks," *CoRR*, vol. abs/1901.00596, 2019. [Online]. Available: <http://arxiv.org/abs/1901.00596>
- [38] V. Cernigliaro, R. Peluso, B. Zedda, L. Silengo, E. Tolosano, R. Pellicano, F. Altruda, and S. Fagoonee, "Evolving cell-based and cell-free clinical strategies for treating severe human liver diseases," *Cells*, vol. 9, no. 2, 2020. [Online]. Available: <https://www.mdpi.com/2073-4409/9/2/386>
- [39] O. Keunen, T. Taxt, R. Grüner, M. Lund-Johansen, J.-C. Tonn, T. Pavlin, R. Bjerkvig, S. P. Niclou, and F. Thorsen, "Multimodal imaging of gliomas in the context of evolving cellular and molecular therapies," *Advanced Drug Delivery Reviews*, vol. 76, pp. 98–115, 2014, targeted imaging. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0169409X14001513>
- [40] X. Luo, J. Wu, J. Yang, S. Xue, H. Peng, C. Zhou, H. Chen, Z. Li, and Q. Z. Sheng, "Deep graph level anomaly detection with contrastive learning," *Scientific Reports*, vol. 12, no. 1, p. 19867, 2022.
- [41] L. Buck, T. Schmidt, M. Feist, P. Schwarzfischer, D. Kube, P. J. Oefner, H. U. Zacharias, M. Altenbuchinger, K. Dettmer, W. Gronwald, and R. Spang, "Anomaly detection in mixed high-dimensional molecular data," *Bioinformatics*, vol. 39, no. 8, p. btad501, 08 2023. [Online]. Available: <https://doi.org/10.1093/bioinformatics/btad501>

- [42] M. Sugiyama and K. M. Borgwardt, "Halting in random walk kernels," in *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*, ser. NIPS'15. Cambridge, MA, USA: MIT Press, 2015, p. 1639–1647.
- [43] P. Mahé and J.-P. Vert, "Graph kernels based on tree patterns for molecules," *Machine Learning*, vol. 75, no. 10, 2006.
- [44] G. Nikolentzos, P. Meladianos, F. Rousseau, Y. Stavrakas, and M. Vazirgiannis, "Shortest-path graph kernels for document similarity," in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, M. Palmer, R. Hwa, and S. Riedel, Eds. Copenhagen, Denmark: Association for Computational Linguistics, Sep. 2017, pp. 1890–1900. [Online]. Available: <https://aclanthology.org/D17-1202>
- [45] N. Kriege, P.-L. Giscard, and R. Wilson, *On Valid Optimal Assignment Kernels and Applications to Graph Classification*, 06 2016, pp. 1615–1623.
- [46] N. Shervashidze, S. Vishwanathan, T. Petri, K. Mehlhorn, and K. Borgwardt, "Efficient graphlet kernels for large graph comparison," in *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, D. van Dyk and M. Welling, Eds., vol. 5. Hilton Clearwater Beach Resort, Clearwater Beach, Florida USA: PMLR, 16–18 Apr 2009, pp. 488–495. [Online]. Available: <https://proceedings.mlr.press/v5/shervashidze09a.html>
- [47] H. Fröhlich, J. K. Wegner, F. Sieker, and A. Zell, "Optimal assignment kernels for attributed molecular graphs," in *Proceedings of the 22nd International Conference on Machine Learning*, ser. ICML '05. New York, NY, USA: Association for Computing Machinery, 2005, p. 225–232. [Online]. Available: <https://doi.org/10.1145/1102351.1102380>
- [48] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, "How powerful are graph neural networks?" 2019.
- [49] D. Bajusz, A. Rácz, and K. Héberger, "Why is tanimoto index an appropriate choice for fingerprint-based similarity calculations?" *Journal of Cheminformatics*, vol. 7, no. 05, 2015.
- [50] P. J. Olver, "On multivariate interpolation," *Studies in Applied Mathematics*, vol. 116, 2006. [Online]. Available: <https://api.semanticscholar.org/CorpusID:8314098>
- [51] T. A. Plate, "Holographic reduced representation: Distributed representation for cognitive structures," 2003. [Online]. Available: <https://api.semanticscholar.org/CorpusID:60591699>
- [52] P. Kanerva, *Sparse Distributed Memory*. Cambridge, MA, USA: MIT Press, 1988.
- [53] T. Cohen and D. Widdows, "Bringing order to neural word embeddings with embeddings augmented by random permutations (EARP)," in *Proceedings of the 22nd Conference on Computational Natural Language Learning*. Brussels, Belgium: Association for Computational Linguistics, Oct. 2018, pp. 465–475. [Online]. Available: <https://aclanthology.org/K18-1045>
- [54] P. Vergés, M. Heddes, I. Nunes, T. Givargis, and A. Nicolau, "Classification using hyperdimensional computing: A review with comparative analysis," *Computing in Science & Engineering*.
- [55] R. W. Gayler and S. D. Levy, "A distributed basis for analogical mapping," 2009. [Online]. Available: <https://api.semanticscholar.org/CorpusID:18842042>
- [56] P. Vergés, T. Givargis, and A. Nicolau, "Refinehd: Accurate and efficient single-pass adaptive learning using hyperdimensional computing," EasyChair, 2023.
- [57] L. Page, S. Brin, R. Motwani, and T. Winograd, "The pagerank citation ranking : Bringing order to the web," in *The Web Conference*, 1999. [Online]. Available: <https://api.semanticscholar.org/CorpusID:1508503>
- [58] M. Imani, J. Morris, S. Bosch, H. Shu, G. D. Micheli, and T. Rosing, "Adaphd: Adaptive efficient training for brain-inspired hyperdimensional computing," in *2019 IEEE Biomedical Circuits and Systems Conference (BioCAS)*, 2019, pp. 1–4.
- [59] A. Hernández-Cano, N. Matsumoto, E. Ping, and M. Imani, "Onlinehd: Robust, efficient, and single-pass online learning using hyperdimensional system," in *2021 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2021, pp. 56–61.
- [60] M. Heddes, I. Nunes, P. Vergés, D. Kleyko, D. Abraham, T. Givargis, A. Nicolau, and A. Veidenbaum, "Torchhd: An open source python library to support research on hyperdimensional computing and vector symbolic architectures," *Journal of Machine Learning Research*, vol. 24, no. 255, pp. 1–10, 2023. [Online]. Available: <http://jmlr.org/papers/v24/23-0300.html>
- [61] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, "How powerful are graph neural networks?" 2019.
- [62] R. W. Gayler, "Vector symbolic architectures answer jackendoff's challenges for cognitive neuroscience," in *Joint International Conference on Cognitive Science (ICCS/ASCS)*, 2003, pp. 133–138.