Hyperdimensional Computing for Filtered Vector Search

Sebastián Bugedo, Stijn Vansummeren

DSI UHasselt, Belgium
{sebastian.bugedo,stijn.vansummeren}@uhasselt.be

Approximate nearest-neighbour (ANN) search is a key primitive in modern data systems, powering semantic search; image and video retrieval; and entity resolution. Given a query vector, the goal is to find the k most similar data vectors. Because retrieving this set of k vectors exactly requires a linear data scan, ANN methods choose to trade a small loss in accuracy for significantly faster processing by returning approximately the k nearest neighbours. Given the importance of ANN, a wide range of index structures have been proposed to support ANN efficiently.

In many applications, however, in particular enterprise RAG, the returned vectors must additionally satisfy predicates on structured attributes, a setting known as *filtered ANN* (FANN). Current systems tackle FANN in two ways. (I) By performing attribute filtering and ANNS as separate steps, in which case existing ANN indexes can be reused at the expense of also needing to build a traditional attribute-based index, and needing to merge query results originating from both indexes. (II) By building a new index structure that also incorporates the attribute information, in which case querying can be done in a single step.

In this ongoing work, we explore the use of Hyperdimensional Computing (HDC) for FANN search. HDC [1] is an approach to neural computing that exploits the properties of randomization in tandem with simple algebraic vector operations to allow symbolic reasoning on structured data (e.g. attribute filters) to be expressed by means of nearest-neighbour search in vector space. A particular challenge that we solve to support FANN by means of HDC, is to modify the classical HDC encoding of structured data: traditional HDC assumes that all input vectors are randomly generated, whereas FANN needs to starting from a given set of data vectors. The result is a generic method for reducing the Filtered ANN problem to the normal ANN problem. Consequently, our approach is index-agnostic, allowing to reuse existing ANN methods for FANN search without requiring additional attribute indexes or expensive merge steps (in contrast to I above), and without requiring to engineer completely new index structures (in contrast to II above).

We present a preliminary empirical evaluation of our approach on a real-world dataset of paper abstracts [2], finding that HDC improves on type-I methods in terms of QPS vs Recall trade-off while being competitive with type-II methods.

References

[1] Sebastián Bugedo and Stijn Vansummeren. Hyperdimensional computing for structured querying on tabular data embeddings. In *Proceedings of the VLDB 2026 Workshop on Tabular Data Analysis (TaDA)*. VLDB Endowment, 2026.

[2] Patrick Iff, Paul Brügger, Marcin Chrapek, David Kochergin, Maciej Besta, and Torsten Hoefler. Benchmarking filtered approximate nearest neighbor search algorithms on transformer-based embedding vectors. In *Proceedings of the 49th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '26)*, Melbourne, Australia, 2026. ACM.