Novelty Assessment Report

Paper: Welfarist Formulations for Diverse Similarity Search
PDF URL: https://openreview.net/pdf?id=UWhOUrsgkA
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2026-01-01

Abstract

Nearest Neighbor Search (NNS) is a fundamental problem in data structures with wide-ranging applications, such as web search, recommendation systems, and, more recently, retrieval-augmented generations (RAG). In such recent applications, in addition to the relevance (similarity) of the returned neighbors, diversity among the neighbors is a central requirement. In this paper, we develop principled welfare-based formulations in NNS for realizing diversity across attributes. Our formulations are based on welfare functions---from mathematical economics---that satisfy central diversity (fairness) and relevance (economic efficiency) axioms. With a particular focus on Nash social welfare, we note that our welfare-based formulations provide objective functions that adaptively balance relevance and diversity in a query-dependent manner. Notably, such a balance was not present in the prior constraint-based approach, which forced a fixed level of diversity and optimized for relevance. In addition, our formulation provides a parametric way to control the trade-off between relevance and diversity, providing practitioners with flexibility to tailor search results to task-specific requirements. We develop efficient nearest neighbor algorithms with provable guarantees for the welfare-based objectives. Notably, our algorithm can be applied on top of any standard ANN method (i.e., use standard ANN method as a subroutine) to efficiently find neighbors that approximately maximize our welfare-based objectives. Experimental results demonstrate that our approach is practical and substantially improves diversity while maintaining high relevance of the retrieved neighbors.

Core Task Landscape

This paper addresses: **Balancing Diversity and Relevance in Nearest Neighbor Search**

A total of **50 papers** were analyzed and organized into a taxonomy with **32 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Welfare-Based and Principled Formulations for Diverse Search**
- **Heuristic and Constraint-Based Diversity Methods**
- **Graph-Based Indexing and Diverse Neighbor Graphs**
- **Hashing and Inverted File Indexing for Diverse Search**
- **Exploration-Exploitation Trade-Off in Interactive and Adaptive Search**
- **Recommendation Systems with Diversity-Accuracy Trade-Offs**
- **Approximate Diverse k-Nearest Neighbor Algorithms**
- **Feature Weighting and Distance Metric Learning for k-NN**
- **Ensemble and Diversity-Based k-NN Methods**
- **Exploration-Exploitation in Non-Search Domains**
- ... and 5 more categories

Complete Taxonomy Tree

- Balancing Diversity and Relevance in Nearest Neighbor Search Survey Taxonomy
- Welfare-Based and Principled Formulations for Diverse Search
  - Welfare Function Approaches ★ (1 papers)
  - [0] Welfarist Formulations for Diverse Similarity Search (Anon et al., 2026) View paper
  - Graph-Based Diverse Similarity Search (1 papers)
  - [3] Graph-based algorithms for diverse similarity search (Indyk, 2025) View paper
- Heuristic and Constraint-Based Diversity Methods
  - Spatial and Geometric Diversity Metrics
  - Angle-Based and Spatial Diversity (1 papers)
    - [30] Finding diverse neighbors in high dimensional space (Qi Guo, 2018) View paper
  - Category-Constrained Spatial Diversity (2 papers)
    - [32] Towards Spatially- and Category-Wise k-Diverse Nearest Neighbors Queries (Camila F. Costa, 2017) View paper
    - [50] Spatial search for K diverse-near neighbors (Gregory Ference, 2013) View paper
  - Attribute and Semantic Diversity Approaches
  - Attribute-Weighted and Two-Stage Neighbor Selection (1 papers)
    - [1] Improving the accuracy and diversity of personalized recommendation through a two-stage neighborhood selection (Junpeng Guo, 2025) View paper
  - Dissimilarity and Furthest Neighbor Models (2 papers)
    - [19] Improving novelty and diversity of nearest-neighbors recommendation by exploiting dissimilarities (Pablo SÃ¡nchez, 2025) View paper
    - [29] Increasing diversity through furthest neighbor-based recommendation (Alan Said, 2012) View paper

    - Semantic Compression and Graph-Augmented Retrieval (1 papers)
        - [38] Beyond Nearest Neighbors: Semantic Compression and Graph-Augmented Retrieval for Enhanced Vector Search (Rahul Raja, 2025) View paper
    - Constraint-Based and Greedy Diversity Enforcement
    - Minimum Distance Diversity Constraints (2 papers)
        - [23] Diverse near neighbor problem (Sofiane Abbar, 2013) View paper
        - [34] Providing Diversity in K-Nearest Neighbor Query Results (Jain, 2022) View paper
    - Post-Processing and Learned Cutoff Filtering (2 papers)
        - [12] Efficient diversity-aware search (Albert Angel, 2011) View paper
        - [42] LotusFilter: Fast Diverse Nearest Neighbor Search via a Learned Cutoff Table (Yusuke Matsui, 2025) View paper
- Graph-Based Indexing and Diverse Neighbor Graphs
    - k-Diverse Nearest Neighbor Graphs (1 papers)
    - [28] Fast approximate nearest neighbor search via k-diverse nearest neighbor graph (Yan Xiao, 2018) View paper
    - Cluster-Guided and Supervised Indexing (2 papers)
    - [22] Cluster Guided Truncated Hashing for Enhanced Approximate Nearest Neighbor Search (Mingyang Liu, 2024) View paper
    - [25] Semi-supervised inverted file index approach for approximate nearest neighbor search (Anton Bazdyrev, 2023) View paper
- Hashing and Inverted File Indexing for Diverse Search (1 papers)
    - [4] EI-LSH: An early-termination driven I/O efficient incremental c-approximate nearest neighbor search (Wanqi Liu, 2021) View paper
- Exploration-Exploitation Trade-Off in Interactive and Adaptive Search
    - Reinforcement Learning for Adaptive Ranking (2 papers)
    - [13] Balancing exploration and exploitation: Empirical parameterization of exploratory search systems (Kumaripaba Ahukorala, 2015) View paper
    - [18] Beyond relevance: Adapting exploration/exploitation in information retrieval (Kumaripaba Athukorala, 2016) View paper
    - Relevance Feedback and Session Utility Optimization (2 papers)
    - [41] A learning approach to optimizing explorationâ□□exploitation tradeoff in relevance feedback (Maryam Karimzadehgan, 2013) View paper
    - [43] Exploration-exploitation tradeoff in interactive relevance feedback (Maryam Karimzadehgan, 2010) View paper
    - Query Decomposition and Document Retrieval (1 papers)
    - [49] Query Decomposition for RAG: Balancing Exploration-Exploitation (Murray, 2025) View paper
- Recommendation Systems with Diversity-Accuracy Trade-Offs
    - Matching and Ranking with Diversified Preferences (1 papers)
    - [14] Improving accuracy and diversity in matching of recommendation with diversified preference network (Ruobing Xie, 2021) View paper
    - Graph Neural Network Recommendation (1 papers)
    - [16] BDARec: Balancing Diversity and Accuracy of Recommendation Model with Graph Neural Networks (Li Mengmeng, 2025) View paper
    - Anytime Nearest Neighbor Classification (1 papers)
    - [15] Adding diversity to rank examples in anytime nearest neighbor classification (C. Lemes, 2014) View paper
- Approximate Diverse k-Nearest Neighbor Algorithms (2 papers)
    - [36] Approximate Diverse k-nearest Neighbor Search in Vector Database (Zhao Jia-chen, 2025) View paper
- Feature Weighting and Distance Metric Learning for k-NN
    - Feature Importance and Minkowski Distance (2 papers)
    - [5] A new approach to K-nearest neighbors distance metrics on sovereign country credit rating (Ali Ä°hsan Ã□etin, 2025) View paper
    - [6] Local means-based fuzzy k-nearest neighbor classifier with Minkowski distance and relevance-complementarity feature weighting (Mahinda Mailagaha Kumbure, 2024) View paper
    - Similarity Metric Optimization (1 papers)
    - [21] A recommender system based on a new similarity metric and upgraded crow search algorithm (Bam Bahadur Sinha, 2020) View paper
- Ensemble and Diversity-Based k-NN Methods (2 papers)
    - [27] Measuring diversity and accuracy in ANN ensembles (M. Sesmero, 2018) View paper
    - [35] Nearest neighbor ensemble (C. Domeniconi, 2004) View paper
- Exploration-Exploitation in Non-Search Domains
    - Web Content Discovery and Crawling (1 papers)
    - [2] Improving the exploration/exploitation trade-off in web content discovery (PETER SCHULAM, 2023) View paper
    - Molecule Generation and Creative Foraging (2 papers)
    - [9] Molecule generation with fragment retrieval augmentation (Karsten Kreis, 2024) View paper
    - [10] Creative Foraging and the Explore-Exploit Trade-off in Knowledge Networks (Kara Kedrick, 2025) View paper
    - IoT Interaction Analysis (1 papers)
    - [7] Exploration and Exploitation in Consumer Automation: Visualizing IoT Interactions with Topological Data Analysis (Thomas Novak, 2024) View paper
- Optimization and Efficiency in Approximate Nearest Neighbor Search (2 papers)
    - [11] Optimizing Memory Bandwidth for Efficient Approximate Nearest Neighbor Search (Zhang, 2025) View paper
    - [37] ANN-Benchmarks: A Benchmarking Tool for Approximate Nearest Neighbor Algorithms (AumÃ¼ller, 2022) View paper
- Domain-Specific Applications of Nearest Neighbor Methods
    - Credit Rating and Employment Recommendation (1 papers)
    - [39] Hybrid Employment and Entrepreneurship Recommendation System based on K-Nearest Neighbors (Qin Wang, 2025) View paper
    - Psychosis Research and Image Retrieval (2 papers)
    - [45] Semantic Search in Psychosis: Modeling Local Exploitation and Global Exploration (Nancy B. Lundin, 2020) View paper
    - [46] Exploration and search-by-similarity in cbir (M. Cord, 2003) View paper

- Metaheuristic and Optimization Algorithms with Similarity-Based Components (2 papers)
  - [24] The Discrete Carnivorous Plant Algorithm with Similarity Elimination Applied to the Traveling Salesman problem (Pan-Li Zhang, 2022) View paper
  - [26] A similarity-based neighbourhood search for enhancing the balance exploration-exploitation of differential evolution (Eduardo Segredo, 2020) View paper
- Specialized Learning and Retrieval Frameworks
  - Multi-Fidelity and Active Learning (2 papers)
  - [17] Interactive multi-fidelity learning for cost-effective adaptation of language model with sparse human supervision (Zhang JiaXin, 2023) View paper
  - [47] Deep similarity-based batch mode active learning with exploration-exploitation (Changchang Yin, 2017) View paper
  - Model Serving and Meta-Learning (2 papers)
  - [20] Mlta: a meta-learning toolbox for automl (Gellecom, 2024) View paper
  - [48] Resource-Efficient Model Serving in Multi-Tenant GPU Clusters (Lingyun Yang, 2025) View paper
- Unrelated or Peripheral Applications (4 papers)
  - [8] Security Enhancement in 5G Networks by Identifying Attacks Using Optimized Cosine Convolutional Neural Network (Premalatha Santhanamari, 2025) View paper
  - [31] Balancing exploration and exploitation: An image-based approach to item retrieval with enhanced diversity (A. Pereira, 2020) View paper
  - [33] StereoGest-SNN: Robust Gesture Detection With Stereo Acoustic Setup Using Spiking Neural Networks (Andrew Gigie, 2023) View paper
  - [44] Hierarchical affinity landscape navigation through learning a shared pocket-ligand space (Bin Feng, 2025) View paper

Narrative

Core task: Balancing diversity and relevance in nearest neighbor search. The field addresses the challenge of retrieving results that are both similar to a query and sufficiently varied to avoid redundancy. The taxonomy reveals a rich landscape spanning principled formulations, algorithmic techniques, and application domains. Welfare-Based and Principled Formulations (including works like Welfarist Diverse Search[0]) ground diversity objectives in formal frameworks, while Heuristic and Constraint-Based Diversity Methods offer practical trade-offs. Graph-Based Indexing (e.g., Graph Diverse Search[3], K-Diverse Neighbor Graph[28]) and Hashing approaches (Early Termination LSH[4]) tackle scalability through specialized data structures. Meanwhile, Exploration-Exploitation Trade-Off in Interactive and Adaptive Search (Exploratory Search Parameterization[13], Interactive Relevance Tradeofff[43]) and Recommendation Systems with Diversity-Accuracy Trade-Offs (Diversified Preference Network[14], BDARec[16]) emphasize user-centric adaptation. Approximate Diverse k-Nearest Neighbor Algorithms (Approximate Diverse KNN[36][40], Anytime Neighbor Diversity[15]) and Ensemble methods (ANN Ensemble Diversity[27]) focus on computational efficiency, while Feature Weighting and Distance Metric Learning refine similarity measures themselves.

Several active lines of work highlight contrasting priorities. Graph-based indexing methods pursue efficient diverse neighbor retrieval at scale, often trading exact optimality for speed, whereas welfare-based formulations emphasize principled objective functions that capture user utility more rigorously. Interactive and recommendation-oriented branches explore dynamic adaptation to user feedback, balancing exploration of novel items against exploitation of known preferences. The original paper, Welfarist Diverse Search[0], sits squarely within the Welfare Function Approaches, offering a theoretically grounded perspective that contrasts with more heuristic strategies like Efficient Diversity Search[12] or constraint-driven methods. Compared to graph-centric works such as Graph Diverse Search[3] or Two-Stage Neighborhood Selection[1], Welfarist Diverse Search[0] prioritizes formal welfare criteria over indexing efficiency, positioning it as a principled foundation for understanding diversity-relevance trade-offs rather than an engineering-focused solution.

Related Works in Same Category

No sibling papers were found in the same taxonomy leaf. A taxonomy-subtopic-level comparison will be produced instead.

Taxonomy-Level Summary

Both subtopics address the challenge of balancing diversity and relevance in nearest neighbor search, but through fundamentally different mathematical frameworks. Welfare Function Approaches draw on mathematical economics to create adaptive, query-dependent trade-offs, while Graph-Based Diverse Similarity Search uses structural graph properties and two-stage retrieval pipelines. These represent complementary principled methods for the same core problem.

Similarities: - Both aim to balance diversity and relevance in search results rather than optimizing for relevance alone - Both provide principled, theoretically-grounded approaches to the diversity-relevance trade-off - Both are query-dependent methods that adapt to specific search contexts

Differences: - Welfare Function Approaches use economic utility theory and welfare functions, while Graph-Based methods leverage topological graph structures - Graph-Based methods employ two-stage retrieval-then-filtering pipelines, while Welfare Function Approaches integrate diversity considerations directly into the ranking objective - The mathematical foundations differ fundamentally: economic optimization theory versus graph-theoretic algorithms

Suggested Search Directions: - Hybrid approaches combining welfare functions with graph-based filtering stages - Comparative empirical studies evaluating welfare-based versus graph-based diversity methods - Computational complexity comparisons between economic optimization and graph traversal approaches

Sibling Subtopics

- **Graph-Based Diverse Similarity Search** (leaves: 1, papers: 1)
- Scope: Algorithms leveraging graph structures to retrieve diverse and relevant neighbors through two-stage retrieval and filtering.
- Exclude: Excludes welfare-based formulations; see Welfare Function Approaches.

Contributions Analysis

Overall novelty summary. The paper introduces welfare-based formulations for diverse nearest neighbor search, grounding diversity-relevance trade-offs in mathematical economics. It resides in the 'Welfare Function Approaches' leaf, which contains only this paper within the broader 'Welfare-Based and Principled Formulations for Diverse Search' branch. This positioning reflects a sparse research direction: while the taxonomy includes 50 papers across diverse methods, the welfare-theoretic approach appears underexplored. The paper's focus on Nash social welfare and axiomatic foundations distinguishes it from the more populated heuristic and constraint-based branches, suggesting it occupies a relatively novel conceptual space within the field.

The taxonomy reveals neighboring work in 'Graph-Based Diverse Similarity Search' (one sibling paper using two-stage retrieval) and extensive activity in 'Heuristic and Constraint-Based Diversity Methods' (spanning spatial metrics, attribute-based approaches, and greedy constraints). The paper diverges from these by replacing fixed diversity constraints with adaptive welfare functions that balance relevance and diversity in a query-dependent manner. While graph-based methods prioritize indexing efficiency and constraint-based approaches enforce hard diversity thresholds, this work offers a parametric framework rooted in economic principles. The taxonomy's

scope notes explicitly exclude welfare-theoretic foundations from heuristic branches, underscoring the conceptual gap this paper addresses.

Among 11 candidates examined across three contributions, no clearly refuting prior work was identified. The core welfare-based formulation examined zero candidates, suggesting limited direct precedent in the search scope. The algorithmic contribution examined one candidate without refutation, while the parametric trade-off mechanism examined ten candidates, none providing overlapping prior work. This limited search scale (11 papers, not hundreds) means the analysis captures top semantic matches and citations but cannot claim exhaustive coverage. The absence of refutations within this scope suggests the welfare-theoretic framing and parametric control represent relatively fresh angles, though broader literature may contain related ideas not surfaced here.

Given the constrained search scope and sparse taxonomy positioning, the work appears to introduce a principled perspective underrepresented in the examined literature. The welfare-based formulation and adaptive balancing mechanism contrast with prevalent heuristic and constraint-driven methods. However, the analysis reflects top-30-scale semantic search, not comprehensive field coverage, leaving open the possibility of related welfare-theoretic or economic approaches in adjacent communities (e.g., fair ranking, social choice) not captured by this taxonomy.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: Welfare-based formulations for diverse similarity search

Description: The authors introduce novel formulations for nearest neighbor search that incorporate diversity using welfare functions from mathematical economics, specifically Nash social welfare and p-mean welfare. These formulations balance relevance and diversity in a query-dependent manner without requiring fixed diversity quotas.

This contribution was assessed against **0 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

Contribution 2: Efficient algorithms with provable guarantees for welfare-based objectives

Description: The authors develop polynomial-time algorithms for optimizing their welfare-based objectives. For single-attribute settings, they provide an exact greedy algorithm that can leverage any ANN method as a subroutine. For multi-attribute settings, they provide a 0.63-approximation algorithm despite NP-hardness.

This contribution was assessed against **1 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Optimized Artificial Neural Network Techniques to Improve Cybersecurity of Higher Education Institution.

URL: View paper

Brief Assessment

Cybersecurity Neural Network[51] focuses on intrusion detection in higher education institutions using deep neural networks and sparrow search algorithms, not on welfare-based objectives or approximate nearest neighbor search algorithms.

Contribution 3: Parametric control of relevance-diversity trade-off via p-mean welfare

Description: The authors extend their framework to p-mean welfare functions that interpolate between pure relevance (p equals 1) and pure diversity (p approaches negative infinity), giving practitioners a tunable parameter to adjust the balance between these objectives based on application needs.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. A Spatially-Aware Search Engine for Textual Content in Images

URL: View paper

Brief Assessment

Spatially-Aware Image Search[59] focuses on spatial text localization in images using geometric bounding boxes and n-gram matching, not on parametric control of relevance-diversity trade-offs in nearest neighbor search or welfare-based formulations.

---

2. Diversity based relevance feedback for time series search

URL: View paper

Brief Assessment

Time Series Diversity[58] focuses on diversity-based relevance feedback for time series search using MMR and cluster-based methods. It does not address parametric control of relevance-diversity trade-offs through p-mean welfare functions or provide a tunable parameter that interpolates between pure relevance and pure diversity objectives.

---

3. Image search reranking with relevance, diversity and topic coverage

URL: View paper

Brief Assessment

Image Reranking Coverage[57] focuses on image search reranking with visual similarity and topic coverage, not on parametric control of relevance-diversity trade-offs through p-mean welfare functions in nearest neighbor search.

---

4. Recommending Paintings in Web Art Gallery with Adjustable Popularity and Diversity

URL: View paper

Brief Assessment

Adjustable Popularity Diversity[53] focuses on user-adjustable controls for popularity-diversity balance in art recommendation systems, not on p-mean welfare functions for nearest neighbor search.

---

5. Jointly optimizing diversity and relevance in neural response generation

URL: View paper

Brief Assessment

Diversity Relevance Response[55] addresses diversity-relevance trade-offs in neural response generation for conversational AI, not nearest neighbor search. The paper does not discuss p-mean welfare functions or parametric control mechanisms for balancing relevance and diversity in retrieval systems.

---

6. A unified framework for recommending diverse and relevant queries

URL: View paper

Brief Assessment

Diverse Relevant Queries[61] focuses on query recommendation in search engines using manifold ranking, not parametric control of relevance-diversity trade-offs in nearest neighbor search via p-mean welfare functions.

---

7. Towards a Similarity-adjusted Surprisal Theory

URL: View paper

Brief Assessment

Similarity-Adjusted Surprisal[60] focuses on linguistic predictability measures in natural language processing, not on nearest neighbor search or relevance-diversity trade-offs in information retrieval systems.

---

8. Diversity vs relevance: A practical multi-objective study in luxury fashion recommendations

URL: View paper

Brief Assessment

Luxury Fashion Multi-Objective[56] focuses on reranking strategies for fashion recommendations using fixed diversification metrics (coverage, serendipity, neighborhood distance), not on parametric p-mean welfare functions that interpolate between relevance and diversity objectives in nearest neighbor search.

---

9. Diversity, serendipity, novelty, and coverage: a survey and empirical analysis of beyond-accuracy objectives in recommender systems

URL: View paper

Brief Assessment

Beyond-Accuracy Survey[54] discusses existing methods with parametric trade-off control (e.g., α parameter in reranking), but does not propose p-mean welfare functions for nearest neighbor search. The survey reviews prior work rather than introducing the specific p-mean interpolation framework.

---

10. Balancing Diversity in Session-Based Recommendation Between Relevance and Unexpectedness

URL: View paper

Brief Assessment

Session Diversity Unexpectedness[52] focuses on session-based recommender systems using serendipity components and determinantal point processes, not on nearest neighbor search with p-mean welfare functions for parametric trade-off control.

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] Welfarist Formulations for Diverse Similarity Search View paper
- [1] Improving the accuracy and diversity of personalized recommendation through a two-stage neighborhood selection View paper
- [2] Improving the exploration/exploitation trade-off in web content discovery View paper
- [3] Graph-based algorithms for diverse similarity search View paper
- [4] EI-LSH: An early-termination driven I/O efficient incremental c-approximate nearest neighbor search View paper
- [5] A new approach to K-nearest neighbors distance metrics on sovereign country credit rating View paper
- [6] Local means-based fuzzy k-nearest neighbor classifier with Minkowski distance and relevance-complementarity feature weighting View paper
- [7] Exploration and Exploitation in Consumer Automation: Visualizing IoT Interactions with Topological Data Analysis View paper
- [8] Security Enhancement in 5G Networks by Identifying Attacks Using Optimized Cosine Convolutional Neural Network View paper
- [9] Molecule generation with fragment retrieval augmentation View paper
- [10] Creative Foraging and the Explore-Exploit Trade-off in Knowledge Networks View paper
- [11] Optimizing Memory Bandwidth for Efficient Approximate Nearest Neighbor Search View paper
- [12] Efficient diversity-aware search View paper
- [13] Balancing exploration and exploitation: Empirical parameterization of exploratory search systems View paper
- [14] Improving accuracy and diversity in matching of recommendation with diversified preference network View paper
- [15] Adding diversity to rank examples in anytime nearest neighbor classification View paper
- [16] BDARec: Balancing Diversity and Accuracy of Recommendation Model with Graph Neural Networks View paper
- [17] Interactive multi-fidelity learning for cost-effective adaptation of language model with sparse human supervision View paper
- [18] Beyond relevance: Adapting exploration/exploitation in information retrieval View paper
- [19] Improving novelty and diversity of nearest-neighbors recommendation by exploiting dissimilarities View paper
- [20] Mlta: a meta-learning toolbox for automl View paper
- [21] A recommender system based on a new similarity metric and upgraded crow search algorithm View paper
- [22] Cluster Guided Truncated Hashing for Enhanced Approximate Nearest Neighbor Search View paper
- [23] Diverse near neighbor problem View paper
- [24] The Discrete Carnivorous Plant Algorithm with Similarity Elimination Applied to the Traveling Salesman problem View paper
- [25] Semi-supervised inverted file index approach for approximate nearest neighbor search View paper
- [26] A similarity-based neighbourhood search for enhancing the balance exploration-exploitation of differential evolution View paper
- [27] Measuring diversity and accuracy in ANN ensembles View paper
- [28] Fast approximate nearest neighbor search via k-diverse nearest neighbor graph View paper
- [29] Increasing diversity through furthest neighbor-based recommendation View paper
- [30] Finding diverse neighbors in high dimensional space View paper
- [31] Balancing exploration and exploitation: An image-based approach to item retrieval with enhanced diversity View paper
- [32] Towards Spatially- and Category-Wise k-Diverse Nearest Neighbors Queries View paper
- [33] StereoGest-SNN: Robust Gesture Detection With Stereo Acoustic Setup Using Spiking Neural Networks View paper
- [34] Providing Diversity in K-Nearest Neighbor Query Results View paper
- [35] Nearest neighbor ensemble View paper
- [36] Approximate Diverse k-nearest Neighbor Search in Vector Database View paper
- [37] ANN-Benchmarks: A Benchmarking Tool for Approximate Nearest Neighbor Algorithms View paper
- [38] Beyond Nearest Neighbors: Semantic Compression and Graph-Augmented Retrieval for Enhanced Vector Search View paper

- [39] Hybrid Employment and Entrepreneurship Recommendation System based on K-Nearest Neighbors View paper
- [40] Approximate Diverse -nearest Neighbor Search in Vector Database View paper
- [41] A learning approach to optimizing explorationâ□□exploitation tradeoff in relevance feedback View paper
- [42] LotusFilter: Fast Diverse Nearest Neighbor Search via a Learned Cutoff Table View paper
- [43] Exploration-exploitation tradeoff in interactive relevance feedback View paper
- [44] Hierarchical affinity landscape navigation through learning a shared pocket-ligand space View paper
- [45] Semantic Search in Psychosis: Modeling Local Exploitation and Global Exploration View paper
- [46] Exploration and search-by-similarity in cbir View paper
- [47] Deep similarity-based batch mode active learning with exploration-exploitation View paper
- [48] Resource-Efficient Model Serving in Multi-Tenant GPU Clusters View paper
- [49] Query Decomposition for RAG: Balancing Exploration-Exploitation View paper
- [50] Spatial search for K diverse-near neighbors View paper
- [51] Optimized Artificial Neural Network Techniques to Improve Cybersecurity of Higher Education Institution. View paper
- [52] Balancing Diversity in Session-Based Recommendation Between Relevance and Unexpectedness View paper
- [53] Recommending Paintings in Web Art Gallery with Adjustable Popularity and Diversity View paper
- [54] Diversity, serendipity, novelty, and coverage: a survey and empirical analysis of beyond-accuracy objectives in recommender systems View paper
- [55] Jointly optimizing diversity and relevance in neural response generation View paper
- [56] Diversity vs relevance: A practical multi-objective study in luxury fashion recommendations View paper
- [57] Image search reranking with relevance, diversity and topic coverage View paper
- [58] Diversity based relevance feedback for time series search View paper
- [59] A Spatially-Aware Search Engine for Textual Content in Images View paper
- [60] Towards a Similarity-adjusted Surprisal Theory View paper
- [61] A unified framework for recommending diverse and relevant queries View paper