Novelty Assessment Report

Paper: Answering Counterfactual Queries on Graph Databases

PDF URL: https://openreview.net/pdf?id=8QdA8mUfYF

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2026-01-08

Abstract

Counterfactual analysis on graph data is central to causal reasoning and interpretability, yet existing graph-based methods rely on ad hoc perturbations and remain tied to model behavior rather than underlying data. To address this challenge, we introduce Counterfactual Graph Database (CF-GDB) queries, the first query-based framework for counterfactual reasoning on graphs that grounds counterfactuals in verifiable database instances. Our approach abstracts graphs into semantically meaningful concepts and compares them using a hypergraph-based distance that integrates local structure with global semantics. To ensure efficiency and scalability, we propose two complementary indices: the Concept Distribution Index (CDI), a histogram that provides certified lower bounds, and the Concept Semantic Index (CSI), a continuous embedding that provides upper bounds. These indices yield provably tight sandwich guarantees and enable efficient candidate pruning while preserving the fidelity of counterfactual retrieval. Using 8 read data sets across 4 domains, CF-GDB improves accuracy by over 20% and achieves up to 20× faster performance, demonstrating both fidelity and scalability.

Core Task Landscape

This paper addresses: **Counterfactual Reasoning on Graph Databases**

A total of **50 papers** were analyzed and organized into a taxonomy with **13 categories**.

Taxonomy Overview

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

- **Counterfactual Explanation Methods for Graph Neural Networks**
- **Counterfactual Fairness and Bias Mitigation on Graphs**
- **Counterfactual Learning and Data Augmentation on Graphs**
- **Causal Inference and Reasoning Frameworks for Graphs**
- **Counterfactual Query Systems and Database Integration**
- **Application-Specific Counterfactual and Causal Methods**

Complete Taxonomy Tree

- Counterfactual Reasoning on Graph Databases Survey Taxonomy
- Counterfactual Explanation Methods for Graph Neural Networks
  - Instance-Level Counterfactual Explainers (5 papers)
  - [5] Robust Counterfactual Explanations on Graph Neural Networks (Mohit Bajaj, 2021) View paper
  - [7] Counterfactual Learning on Heterogeneous Graphs with Greedy Perturbation (Qiang Yang, 2023) View paper
  - [8] CF-GNNExplainer: Counterfactual Explanations for Graph Neural Networks (Lucic, 2022) View paper
  - [31] Counterfactual Graphs for Explainable Classification of Brain Networks (Carlo Abrate, 2021) View paper
  - [45] MEG: Generating Molecular Counterfactual Explanations for Deep Graph Networks (Danilo Numeroso, 2021) View paper
  - Global and Inductive Counterfactual Explainers (4 papers)
  - [4] Gcfexplainer: Global counterfactual explainer for graph neural networks (Mert Kosan, 2025) View paper
  - [6] Global Counterfactual Explainer for Graph Neural Networks (Zexi Huang, 2023) View paper
  - [12] InduCE: Inductive counterfactual explanations for graph neural networks (S Verma, 2024) View paper
  - [33] CLEAR: Generative Counterfactual Explanations on Graphs (Ma Jing, 2022) View paper
  - Robustness and Evaluation of Counterfactual Explanations (3 papers)
  - [15] A survey on graph counterfactual explanations: definitions, methods, evaluation, and research challenges (Mario Alfonso Prado-Romero, 2024) View paper
  - [30] Adapting to Change: Robust Counterfactual Explanations in Dynamic Data Landscapes (Bardh Prenkaj, 2023) View paper
  - [35] GRETEL: Graph Counterfactual Explanation Evaluation Framework (Mario Alfonso Prado-Romero, 2022) View paper
- Counterfactual Fairness and Bias Mitigation on Graphs
  - Counterfactual Fairness in Graph Neural Networks (4 papers)
  - [10] Graph fairness via authentic counterfactuals: Tackling structural and causal challenges (Zichong Wang, 2025) View paper
  - [18] Advancing graph counterfactual fairness through fair representation learning (Zichong Wang, 2024) View paper
  - [39] Toward fair graph neural networks via real counterfactual samples (Zichong Wang, 2024) View paper
  - [46] Contrastive learning for fair graph representations via counterfactual graph augmentation (Chengyu Li, 2024) View paper
  - Counterfactual Fairness in Recommendation Systems (1 papers)
  - [16] Counterfactual graph augmentation for consumer unfairness mitigation in recommender systems (Ludovico Boratto, 2023) View paper
- Counterfactual Learning and Data Augmentation on Graphs (6 papers)
  - [1] Counterfactual Learning on Graphs: A Survey (Zhimeng Guo, 2023) View paper
  - [14] Counterfactual Graph Learning for Anomaly Detection on Attributed Networks (Chunjing Xiao, 2023) View paper

- ◦ [17] Learning from counterfactual links for link prediction (Zhao Tong, 2022) View paper
  - ◦ [22] Generating Counterfactual Hard Negative Samples for Graph Contrastive Learning (Haoran Yang, 2023) View paper
  - ◦ [27] GraphCA: Learning from graph counterfactual augmentation for knowledge tracing (Xinhua Wang, 2023) View paper
  - ◦ [44] Social Recommendation via Graph-Level Counterfactual Augmentation (Yinxuan Huang, 2025) View paper
- Causal Inference and Reasoning Frameworks for Graphs
  - ◦ Causal Inference Foundations and Theory (3 papers)
  - ◦ [3] Causal inference (Peter, 2022) View paper
  - ◦ [26] Causal Inference with Deep Causal Graphs (Parafita, 2022) View paper
  - ◦ [41] Individual Causal Inference with Structural Causal Model (Chang, 2025) View paper
  - ◦ Causal Discovery and Graph Structure Learning (3 papers)
  - ◦ [38] Causal Discovery in Physical Systems from Videos (Li, 2022) View paper
  - ◦ [42] Causalaffect: Causal discovery for facial affective understanding (Guanyu Hu, 2025) View paper
  - ◦ [43] Causal discovery with endogenous context variables (Andreas Gerhardus, 2024) View paper
  - ◦ Causal Reasoning with Large Language Models (3 papers)
  - ◦ [9] Causal Reasoning and Large Language Models: Opening a New Frontier for Causality (KÄ±cÄ±man, 2023) View paper
  - ◦ [21] Counterfactual causal inference in natural language with large language models (Gendron GaÃ«l, 2024) View paper
  - ◦ [25] Causal Cartographer: From Mapping to Reasoning Over Counterfactual Worlds (Gendron GaÃ«l, 2025) View paper
- Counterfactual Query Systems and Database Integration ★ (5 papers)
  - ◦ [0] Answering Counterfactual Queries on Graph Databases (Anon et al., 2026) View paper
  - ◦ [11] What If: Causal Analysis with Graph Databases (Amedeo Pachera, 2024) View paper
  - ◦ [13] A framework to improve causal inferences from visualizations using counterfactual operators (Arran Zeyu Wang, 2025) View paper
  - ◦ [20] Causal what-if and how-to analysis using hyper (Fangzhu Shen, 2023) View paper
  - ◦ [23] SIERRA: A Counterfactual Thinking-based Visual Interface for Property Graph Query Construction (Jiebing Ma, 2024) View paper
- Application-Specific Counterfactual and Causal Methods
  - ◦ Counterfactual Reasoning on Knowledge Graphs (4 papers)
  - ◦ [2] Counterfactual Reasoning with Knowledge Graph Embeddings (Roth, 2024) View paper
  - ◦ [28] Difficulty-controllable question generation over knowledge graphs: A counterfactual reasoning approach (Sheng Bi, 2024) View paper
  - ◦ [34] Causal Inference for Knowledge Graph Based Recommendation (Yinwei Wei, 2022) View paper
  - ◦ [36] Causal Reinforcement Learning for Knowledge Graph Reasoning (Li Dezhi, 2024) View paper
  - ◦ Causal and Counterfactual Methods in Recommendation (2 papers)
  - ◦ [19] Attention Is Not the Only Choice: Counterfactual Reasoning for Path-Based Explainable Recommendation (Yicong Li, 2024) View paper
  - ◦ [47] A survey on causal inference for recommendation (Huishi Luo, 2024) View paper
  - ◦ Specialized Domain Applications (8 papers)
  - ◦ [24] Causal Inference for banking finance and insurance a survey (Kumar Satyam, 2023) View paper
  - ◦ [29] Introducing causal inference in the energy-efficient building design process (Chen Xia, 2022) View paper
  - ◦ [32] Counterfactual Data Generation Method for Fault Diagnosis of Complex Electromechanical Systems (Chong Wang, 2025) View paper
  - ◦ [37] Toward interpretable and actionable data analysis with explanations and causality (Sudeepa Roy, 2022) View paper
  - ◦ [40] Structure your data: Towards semantic graph counterfactuals (Lymperaiou, 2024) View paper
  - ◦ [48] Causal graph ode: Continuous treatment effect modeling in multi-agent dynamical systems (Zi-Jie Huang, 2024) View paper
  - ◦ [49] Graph Neural Networks for Vulnerability Detection: A Counterfactual Explanation (Chu Zhaoyang, 2024) View paper
  - ◦ [50] Learning and Evaluating Graph Neural Network Explanations based on Counterfactual and Factual Reasoning (Tan, 2022) View paper

Narrative

Core task: counterfactual reasoning on graph databases. The field encompasses a diverse set of approaches that apply counterfactual and causal thinking to graph-structured data. At the highest level, the taxonomy organizes work into several major branches: methods for explaining graph neural network predictions via counterfactual examples (e.g., Gcfexplainer[4], CF-GNNExplainer[8]), techniques addressing fairness and bias mitigation through counterfactual notions (e.g., Authentic Counterfactuals Fairness[10], Counterfactual Fairness Representation[18]), frameworks for counterfactual learning and data augmentation on graphs (e.g., Counterfactual Learning Graphs Survey[1], GraphCA[27]), general causal inference and reasoning systems (e.g., Causal Inference[3], Deep Causal Graphs[26]), integration of counterfactual queries with database systems (e.g., What If Databases[11], Counterfactual Graph Queries[0]), and application-specific methods spanning domains such as recommendation, anomaly detection, and knowledge graphs (e.g., Causal Knowledge Graph Recommendation[34], Counterfactual Anomaly Detection[14]).

A particularly active line of work focuses on explainability for GNNs, where researchers seek minimal graph edits that flip model predictions, balancing interpretability with robustness (Robust Counterfactual GNN[5], Global Counterfactual GNN[6]). In contrast, the database integration branch explores how counterfactual queries can be formulated and executed over structured graph repositories, enabling users to ask "what if" questions directly within query languages. Counterfactual Graph Queries[0] sits squarely in this latter branch, alongside What If Databases[11] and Counterfactual Visualization Framework[13], emphasizing the operational and system-level challenges of embedding counterfactual reasoning into database workflows. Compared to neighbors like SIERRA[23] or Causal Hyper[20], which lean toward causal discovery or hypergraph reasoning, the original paper prioritizes query expressiveness and integration with existing database infrastructure, reflecting a systems-oriented perspective on counterfactual reasoning rather than purely algorithmic or fairness-driven concerns.

Related Works in Same Category

The following **4 sibling papers** share the same taxonomy leaf node with the original paper:

1. What If: Causal Analysis with Graph Databases

Authors: Amedeo Pachera, Mattia Palmiotto, Angela Bonifati, Andrea Mauri | **Year/Venue**: 2024 | **URL**: View paper

Abstract

Graphs are powerful abstractions for modeling relationships and enabling data science tasks. In causal inference, Directed Acyclic Graphs (DAGs) serve as a key formalism, but they are typically handcrafted by experts and rarely treated as first-class data artifacts in graph data management systems. This paper presents a novel vision to align causal analysis with property graphsâthe foundation of modern graph databasesâby rethinking graph models to incorporate hypernodes, structural equations...

Relationship Analysis

Both papers belong to the Counterfactual Query Systems and Database Integration category, focusing on enabling counterfactual reasoning through database query mechanisms. They overlap in their goal of grounding counterfactual analysis in database instances rather than model-specific perturbations, and both leverage graph structures for causal reasoning. However, the original paper focuses on retrieving counterfactual graphs from databases using concept-based distances and optimal transport, while the candidate paper proposes integrating causal DAGs directly into property graph databases with extensions to query languages (GQL/Cypher) for causal analysis, mediation queries, and structural equation modeling.

---

2. A framework to improve causal inferences from visualizations using counterfactual operators

Authors: Arran Zeyu Wang, D. Borland, David Gotz, David Borland | **Year/Venue**: 2025 | **URL**: View paper

Abstract

Exploratory data analysis of high-dimensional datasets is a crucial task for which visual analytics can be especially useful. However, the ad hoc nature of exploratory analysis can also lead users to draw incorrect causal inferences. Previous studies have demonstrated this risk and shown that integrating counterfactual concepts within visual analytics systems can improve usersâ understanding of visualized data. However, effectively leveraging counterfactual concepts can be challenging, with on...

Relationship Analysis

Both papers belong to the Counterfactual Query Systems and Database Integration category, focusing on integrating counterfactual reasoning with data systems. The original paper develops a query-based framework for retrieving counterfactual graphs from graph databases using concept-based distance measures and indices, while the candidate paper proposes a framework and library (Co-op) for integrating counterfactual operators into visual analytics systems to improve causal inferences from exploratory data analysis. The key difference is that the original paper focuses on graph database queries with formal distance metrics and retrieval guarantees, whereas the candidate paper addresses visualization systems and user interaction workflows for exploratory analysis.

---

3. Causal what-if and how-to analysis using hyper

Authors: Fangzhu Shen, Kayvon Heravi, Oscar Gomez, Ãscar GÃ³mez, Sainyam Galhotra, et al. (8 authors total) | **Year/Venue**: 2023 | **URL**: View paper

Abstract

What-if and How-to queries are fundamental data analysis questions that provide insights about the effects of a hypothetical update without actually making changes to the database. Traditional systems assume independence across differÂ¬ent tuples and non-updated attributes of the database. However, different attributes and tuples are generally dependent in real-world scenarios. We propose to demonstrate HypeR, a novel system to efficiently answer what-if and how-to queries while capturing causal...

Relationship Analysis

Both papers belong to the Counterfactual Query Systems and Database Integration category, focusing on enabling counterfactual reasoning within database systems. While the original paper introduces CF-GDB for answering counterfactual queries on graph databases using concept-based distance measures and hypergraph representations, the candidate paper (HypeR) addresses what-if and how-to queries on relational databases by capturing causal dependencies across tuples and attributes. The key difference lies in their data models: the original paper operates on graph-structured data with node-edge relationships, whereas HypeR targets traditional relational databases with tuple-attribute structures.

---

4. SIERRA: A Counterfactual Thinking-based Visual Interface for Property Graph Query Construction

Authors: Jiebing Ma, Sourav S. Bhowmick, Lester Tay, S. Bhowmick, Byron Choi | **Year/Venue**: 2024 | **URL**: View paper

Abstract

Attractive visual query interfaces (VQIs) have great potential to democratize the usage of property graph databases as they facilitate user-friendly query formulation without demanding the need to learn a property graph query language e.g.Cypher. Existing VQIs, however, do not embrace HCI principles and psychology theories to inform their design and as a result may limit their potential due to usability andaesthetics challenges. In this demonstration, we present a novelcounterfactual thinking -b...

Relationship Analysis

Both papers belong to the Counterfactual Query Systems and Database Integration category, focusing on enabling counterfactual or what-if queries on graph databases. While the original paper (Answering Counterfactual Queries on Graph Databases) develops a retrieval-based framework for finding counterfactual graph instances from databases using concept-based distance measures and optimal transport, SIERRA focuses on the user interface aspect by providing a visual query construction system informed by counterfactual thinking theory and HCI principles. The key difference is that the original paper addresses the computational and algorithmic challenges of counterfactual retrieval, whereas SIERRA addresses the usability and interaction design challenges of query formulation for property graph databases.

Contributions Analysis

Overall novelty summary. The paper introduces CF-GDB, a query-based framework for counterfactual reasoning on graph databases that grounds counterfactuals in verifiable database instances rather than model perturbations. It resides in the 'Counterfactual Query Systems and Database Integration' leaf, which contains only five papers total (including this one). This is a relatively sparse research direction within the broader taxonomy of 50 papers, suggesting that database-centric counterfactual query systems remain an emerging area compared to the more crowded GNN explainability branches.

The taxonomy reveals that most counterfactual graph research concentrates on GNN explanation methods (12 papers across instance-level and global explainers) and fairness applications (5 papers). The original paper's leaf sits alongside work on what-if databases and counterfactual visualization frameworks, but diverges from neighboring branches focused on causal discovery, LLM-based causal reasoning, and application-specific methods in recommendation or knowledge graphs. The scope_note emphasizes integration with query languages and database systems, distinguishing this work from theoretical causal frameworks and GNN-centric explanation techniques that dominate sibling categories.

Among 17 candidates examined across three contributions, no refutable prior work was identified. The CF-GDB framework examined 10 candidates with zero refutations, the C2GQ method examined 2 candidates with zero refutations, and the dual indexing scheme examined 5 candidates with zero refutations. This suggests that within the limited search scope—focused on top-K semantic matches and citation expansion—the specific combination of concept-based abstraction, hypergraph distance, and dual indexing for counterfactual graph queries appears distinct from examined prior work.

Based on the limited literature search of 17 candidates, the work appears to occupy a relatively novel position at the intersection of counterfactual reasoning and database query systems. However, the sparse population of the target leaf (5 papers) and the modest search scope mean this assessment reflects only the examined neighborhood, not an exhaustive survey of all potentially relevant database, graph query, or counterfactual reasoning literature.

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

Contribution 1: Counterfactual Graph Database (CF-GDB) framework

Description: The authors propose CF-GDB, a novel framework that reframes counterfactual reasoning as a query problem over graph databases. Unlike prior approaches that generate perturbed graphs to flip model predictions, CF-GDB retrieves dataset-grounded, domain-valid counterfactuals anchored in verifiable instances.

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. GRETEL: Graph Counterfactual Explanation Evaluation Framework

URL: View paper

Brief Assessment

GRETEL[35] is an evaluation framework for testing graph counterfactual explanation methods, not a query-based retrieval system. The original paper proposes CF-GDB as a database query framework that retrieves dataset-grounded counterfactuals, whereas GRETEL[35] focuses on developing and testing explanation techniques across different settings.

---

2. Counterfactual fairness with partially known causal graph

URL: View paper

Brief Assessment

Partial Causal Fairness[54] addresses counterfactual fairness in machine learning with partially known causal graphs (mpdags), focusing on feature selection for fair predictions. The original paper proposes CF-GDB for counterfactual reasoning as database queries over graph instances, a fundamentally different problem domain and methodology.

---

3. Counterfactual-Based Root Cause Analysis for Misconfigurations in Autonomous Driving Systems

URL: View paper

Brief Assessment

Counterfactual Root Cause[58] focuses on root cause analysis for misconfigurations in autonomous driving systems through counterfactual analysis of driving scenarios, not on counterfactual reasoning as a query problem over graph databases with verifiable instances.

---

4. User-friendly, interactive, and configurable explanations for graph neural networks with graph views

URL: View paper

Brief Assessment

Interactive GNN Explanations[55] focuses on explaining GNN predictions through interactive visualization and graph views for individual instances, not on database-scale counterfactual retrieval. The candidate does not address query-based counterfactual reasoning over graph databases or propose frameworks for retrieving dataset-grounded counterfactuals from collections of graphs.

---

5. Leveraging structured biological knowledge for counterfactual inference: a case study of viral pathogenesis

URL: View paper

Brief Assessment

Biological Counterfactual Inference[56] focuses on converting qualitative biological knowledge graphs into quantitative structural causal models for medical interventions, not on query-based counterfactual retrieval from graph databases with verifiable instances as proposed in the original paper.

---

6. PerfCE: Performance Debugging on Databases with Chaos Engineering-Enhanced Causality Analysis

URL: View paper

Brief Assessment

PerfCE[57] applies counterfactual analysis to database performance debugging using chaos engineering, not to graph-based machine learning explanations. The candidate focuses on diagnosing performance anomalies in databases by identifying root cause KPIs through causal graphs and structural equation models, whereas the original work retrieves semantically similar counterfactual graphs from databases for ML model explanation.

---

7. Design of an Automated Construction Platform for Advanced Mathematics Content Integrating Knowledge Graph and Generative Artificial Intelligence

URL: View paper

Brief Assessment

Mathematics Knowledge Graph[60] focuses on automated construction of mathematics educational content using knowledge graphs and generative AI. It does not address counterfactual reasoning on graphs or query-based frameworks for retrieving counterfactuals from graph databases.

---

8. What If: Causal Analysis with Graph Databases

URL: View paper

Brief Assessment

What If Databases[11] focuses on integrating causal analysis with graph databases through structural causal models and do-calculus operators, not on counterfactual retrieval from databases to flip model predictions as in CF-GDB.

---

9. Actual causality canvas: a general framework for explanation-based socio-technical constructs

URL: View paper

Brief Assessment

Actual Causality Canvas[59] focuses on socio-technical constructs and general explanation frameworks using actual causality theory, not specifically on counterfactual reasoning as a query problem over graph databases with verifiable instances.

---

10. AGENTICTS:Robust Text-to-SPARQL via Agentic Collaborative Reasoning over Heterogeneous Knowledge Graphs for the Circular Economy

URL: View paper

Brief Assessment

AGENTICTS[53] focuses on text-to-SPARQL generation for question answering over heterogeneous knowledge graphs in the circular economy domain, not on counterfactual reasoning or graph database queries for causal analysis.

Contribution 2: Concept-Based Counterfactual Graph Query (C2GQ) method

Description: The authors introduce C2GQ, which abstracts graphs into semantically meaningful concepts serving as prototypes that cluster structurally similar subgraphs. Differences are measured using a hypergraph-based concept distance grounded in unbalanced optimal transport, jointly capturing fine-grained local changes and global distributional shifts.

This contribution was assessed against **2 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. One Concept at a Time: Subspace-Constrained Causal Inference for High-Dimensional Treatments

URL: View paper

Brief Assessment

Subspace Causal Inference[51] focuses on high-dimensional treatment variables in causal inference using subspace constraints and activation steering in pretrained models, not on graph databases or counterfactual queries over graph structures. The candidate addresses causal inference with text/image treatments, while the original addresses counterfactual retrieval from graph databases.

2. Conceptual Graph Counterfactuals

URL: View paper

Brief Assessment

Conceptual Graph Counterfactuals[52] focuses on counterfactual explanations for visual classifiers using scene graphs and GNN-based graph similarity, not on counterfactual database queries with optimal transport-based concept distances for general graph data.

Contribution 3: Dual indexing scheme with certified bounds

Description: The authors propose two complementary indices for scalable counterfactual queries: CDI provides certified lower bounds via histogram-based concept counts, while CSI provides upper bounds through continuous embeddings. These indices yield provably tight sandwich guarantees and enable efficient candidate pruning while preserving retrieval fidelity.

This contribution was assessed against **5 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Efficient Exact Subgraph Matching via GNN-based Path Dominance Embedding (Technical Report)

URL: View paper

Brief Assessment

Path Dominance Technical[62] focuses on exact subgraph matching using GNN-based path embeddings with dominance relationships, not counterfactual graph database queries with histogram-based concept counts and continuous embeddings for certified bounds.

2. HFGNN: Efficient Graph Neural Networks Using Hub-Fringe Structures

URL: View paper

Brief Assessment

HFGNN[65] focuses on hub-fringe structures for scalable graph neural network message passing, not on dual indexing schemes for counterfactual queries with certified bounds using histogram-based concept counts and continuous embeddings.

3. Graph homomorphism revisited for graph matching

URL: View paper

Brief Assessment

Graph Homomorphism Matching[63] focuses on graph matching via homomorphism extensions (p-homomorphism) for structural similarity in web site matching, not on counterfactual query indexing with certified bounds for database retrieval.

4. Efficient exact subgraph matching via gnn-based path dominance embedding

URL: View paper

Brief Assessment

Path Dominance Embedding[61] focuses on exact subgraph matching using GNN-based path embeddings with dominance relationships, not counterfactual queries on graph databases with certified histogram-based bounds and continuous embeddings for retrieval.

5. Efficient frequent subtree mining beyond forests

URL: View paper

Brief Assessment

Frequent Subtree Mining[64] focuses on mining frequent subtree patterns from graph databases using spanning tree sampling and forest representations, not on indexing schemes with certified bounds for counterfactual queries. The candidate addresses a fundamentally different problem domain (pattern mining) compared to the original paper's focus on counterfactual database queries with provable retrieval guarantees.

Appendix: Text Similarity Detection

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

References

- [0] Answering Counterfactual Queries on Graph Databases View paper
- [1] Counterfactual Learning on Graphs: A Survey View paper
- [2] Counterfactual Reasoning with Knowledge Graph Embeddings View paper
- [3] Causal inference View paper
- [4] Gcfexplainer: Global counterfactual explainer for graph neural networks View paper
- [5] Robust Counterfactual Explanations on Graph Neural Networks View paper

- [6] Global Counterfactual Explainer for Graph Neural Networks View paper
- [7] Counterfactual Learning on Heterogeneous Graphs with Greedy Perturbation View paper
- [8] CF-GNNExplainer: Counterfactual Explanations for Graph Neural Networks View paper
- [9] Causal Reasoning and Large Language Models: Opening a New Frontier for Causality View paper
- [10] Graph fairness via authentic counterfactuals: Tackling structural and causal challenges View paper
- [11] What If: Causal Analysis with Graph Databases View paper
- [12] InduCE: Inductive counterfactual explanations for graph neural networks View paper
- [13] A framework to improve causal inferences from visualizations using counterfactual operators View paper
- [14] Counterfactual Graph Learning for Anomaly Detection on Attributed Networks View paper
- [15] A survey on graph counterfactual explanations: definitions, methods, evaluation, and research challenges View paper
- [16] Counterfactual graph augmentation for consumer unfairness mitigation in recommender systems View paper
- [17] Learning from counterfactual links for link prediction View paper
- [18] Advancing graph counterfactual fairness through fair representation learning View paper
- [19] Attention Is Not the Only Choice: Counterfactual Reasoning for Path-Based Explainable Recommendation View paper
- [20] Causal what-if and how-to analysis using hyper View paper
- [21] Counterfactual causal inference in natural language with large language models View paper
- [22] Generating Counterfactual Hard Negative Samples for Graph Contrastive Learning View paper
- [23] SIERRA: A Counterfactual Thinking-based Visual Interface for Property Graph Query Construction View paper
- [24] Causal Inference for banking finance and insurance a survey View paper
- [25] Causal Cartographer: From Mapping to Reasoning Over Counterfactual Worlds View paper
- [26] Causal Inference with Deep Causal Graphs View paper
- [27] GraphCA: Learning from graph counterfactual augmentation for knowledge tracing View paper
- [28] Difficulty-controllable question generation over knowledge graphs: A counterfactual reasoning approach View paper
- [29] Introducing causal inference in the energy-efficient building design process View paper
- [30] Adapting to Change: Robust Counterfactual Explanations in Dynamic Data Landscapes View paper
- [31] Counterfactual Graphs for Explainable Classification of Brain Networks View paper
- [32] Counterfactual Data Generation Method for Fault Diagnosis of Complex Electromechanical Systems View paper
- [33] CLEAR: Generative Counterfactual Explanations on Graphs View paper
- [34] Causal Inference for Knowledge Graph Based Recommendation View paper
- [35] GRETEL: Graph Counterfactual Explanation Evaluation Framework View paper
- [36] Causal Reinforcement Learning for Knowledge Graph Reasoning View paper
- [37] Toward interpretable and actionable data analysis with explanations and causality View paper
- [38] Causal Discovery in Physical Systems from Videos View paper
- [39] Toward fair graph neural networks via real counterfactual samples View paper
- [40] Structure your data: Towards semantic graph counterfactuals View paper
- [41] Individual Causal Inference with Structural Causal Model View paper
- [42] Causalaffect: Causal discovery for facial affective understanding View paper
- [43] Causal discovery with endogenous context variables View paper
- [44] Social Recommendation via Graph-Level Counterfactual Augmentation View paper
- [45] MEG: Generating Molecular Counterfactual Explanations for Deep Graph Networks View paper
- [46] Contrastive learning for fair graph representations via counterfactual graph augmentation View paper
- [47] A survey on causal inference for recommendation View paper
- [48] Causal graph ode: Continuous treatment effect modeling in multi-agent dynamical systems View paper
- [49] Graph Neural Networks for Vulnerability Detection: A Counterfactual Explanation View paper
- [50] Learning and Evaluating Graph Neural Network Explanations based on Counterfactual and Factual Reasoning View paper
- [51] One Concept at a Time: Subspace-Constrained Causal Inference for High-Dimensional Treatments View paper
- [52] Conceptual Graph Counterfactuals View paper
- [53] AGENTICTS:Robust Text-to-SPARQL via Agentic Collaborative Reasoning over Heterogeneous Knowledge Graphs for the Circular Economy View paper
- [54] Counterfactual fairness with partially known causal graph View paper
- [55] User-friendly, interactive, and configurable explanations for graph neural networks with graph views View paper
- [56] Leveraging structured biological knowledge for counterfactual inference: a case study of viral pathogenesis View paper
- [57] PerfCE: Performance Debugging on Databases with Chaos Engineering-Enhanced Causality Analysis View paper
- [58] Counterfactual-Based Root Cause Analysis for Misconfigurations in Autonomous Driving Systems View paper
- [59] Actual causality canvas: a general framework for explanation-based socio-technical constructs View paper
- [60] Design of an Automated Construction Platform for Advanced Mathematics Content Integrating Knowledge Graph and Generative Artificial Intelligence View paper
- [61] Efficient exact subgraph matching via gnn-based path dominance embedding View paper
- [62] Efficient Exact Subgraph Matching via GNN-based Path Dominance Embedding (Technical Report) View paper
- [63] Graph homomorphism revisited for graph matching View paper
- [64] Efficient frequent subtree mining beyond forests View paper
- [65] HFGNN: Efficient Graph Neural Networks Using Hub-Fringe Structures View paper