
Enabling multi-agent collaboration in knowledge graph environments

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Knowledge graphs (KGs) are critical for grounding large language models and
2 providing them with persistent memory. However, the development of agents
3 capable of collaboratively building and maintaining these KGs is hindered by a
4 lack of suitable environments. Current tools often obscure the construction process,
5 lack standardized editing APIs, and offer poor support for concurrent, multi-agent
6 collaboration. To address this, we introduce GRAPHWORLD, an environment
7 for agents to build and edit graphs. GRAPHWORLD provides agents with tools
8 to create, update, or delete nodes, edges, or properties of KGs, with support for
9 Python, TypeScript, and Rust. GRAPHWORLD relies on DIAMOND, a compact,
10 property-preserving storage format that permits scaling to multi-million-edge KGs.
11 All changes are version-controlled by integrating graph-aware diffing and merging
12 directly into Git, enabling multi-agent and human-agent collaboration through
13 standard branching workflows.

14 1 Introduction

15 Knowledge graphs (KGs) are relational databases that use a graph-based data model to encode
16 knowledge-informed interactions between different objects [1, 2]. Formally, a KG is defined by
17 a set of nodes as well as a set of edges that describe relationships between the nodes. In modern
18 heterogeneous KGs, nodes and edges have different types, and may also contain extra properties
19 or information [3]. KGs are rapidly becoming a core component of modern large language model
20 (LLM)-based systems, enabling retrieval [4], grounded reasoning, and persistent memory. Integrating
21 KGs with LLMs can equip LLMs with rich, structured factual knowledge and traceable information
22 provenance, addressing common challenges faced by LLMs, including hallucination, indecision,
23 poor interpretability, and lack of domain-specific knowledge [5]. Further, KGs offer a mechanism
24 for agents to build and maintain a persistent memory of their interactions and learned knowledge
25 [6]. LLMs and agents excel at generating datasets from structured or unstructured sources, including
26 textual datasets (*e.g.*, LangChain’s Tuna [7] and Alibaba’s DataJuicer [8, 9]) and graph datasets (*e.g.*,
27 Graphiti’s Zep [10]); however, to date, most KGs are created and maintained through manual or
28 semi-automated, non-agentic data pipelines. Tools to build and maintain KG environments both
29 for and with agents remain underdeveloped. Current methods for KG editing often do not expose
30 intermediate edits, lack a standard, multi-language edit API, and provide weak versioning and merge
31 semantics for concurrent work. These obstacles make agent-driven KG editing challenging.

32 To address these challenges, we introduce GRAPHWORLD, an environment for agents to build and
33 edit knowledge graphs. GRAPHWORLD provides a set of atomic edit tools that allow agents to create,
34 update, or delete nodes, edges, and knowledge properties. These tools are accessible from Python,
35 TypeScript, and Rust, ensuring broad interoperability across development ecosystems. To support
36 distributed and concurrent multi-agent or human-agent collaboration, GRAPHWORLD integrates

graph-aware diffs and three-way merges directly with Git. For scaling to multi-million-edge graphs, we further introduce DIAMOND, a compact, property-preserving storage format optimized for labeled property graphs. In experiments, DIAMOND achieves up to $34.1\times$ compression over widely used LPG formats, reducing storage demands while preserving all graph structure and metadata. Benchmarking on synthetic and real-world KGs demonstrates that DIAMOND consistently outperforms JSON, JSON Lines, and PG-JSON, especially as graphs grow in property density and size. The Git integration in GRAPHWORLD produces semantic diffs that reveal node- and edge-level changes rather than opaque binary differences, enabling reproducible and auditable histories of KGs. Applications of GRAPHWORLD include automated validation of human edits, natural language editing tools for domain experts, and ontology alignment workflows, all of which highlight GRAPHWORLD as a human-AI environment for KG development. In biomedical settings, for example, GRAPHWORLD compresses PrimeKG [11] to less than 9% of its original size while supporting transparent version control, facilitating large-scale analysis and sharing. More broadly, GRAPHWORLD supports multi-agent pipelines that continuously propose, review, and merge edits in knowledge graphs.

2 Related work

2.1 Tools or environments for agent-based graph interaction

Prior systems address specific features of the KG construction and agent integration workflow, converting text into graph structures using LLMs [12], managing agent memory in real time [10, 13], enabling retrieval-augmented generation (RAG) over graphs [14], or supporting versioning and provenance [15–17]. By contrast, GRAPHWORLD treats KG construction itself – observable, version-controlled, multi-language editing of property-rich labeled property graphs (LPGs) – as the first-class object.

Single-agent KG construction. Several methods exist to convert unstructured text into graph representations. LangChain’s LLMGraphTransformer, KGGen [12], and related approaches build graphs via prompting heuristics or extraction pipelines. Methods also exist to allow single agents to build a KG; for example, Graphiti [10] provides an interface for an agent to build and query graphs. While effective for initially populating a KG, these approaches do not expose per-edit observability, structured diffs, or merge semantics, and are often bound to a single runtime or framework. By contrast, GRAPHWORLD provides a standard set of atomic graph editing operations across languages and under version control, allowing multiple agents, rather than a single LLM, to collaboratively edit a graph. Moreover, single-agent KG construction systems can adopt the observable edit tools, branching or merging semantics, and graph storage of GRAPHWORLD as a backend.

Graph retrieval for LLM reasoning. Graph-based RAG systems such as Neo4j [18], GraphRAG [19], GNN-RAG [20], KG²RAG [21], KET-RAG [22], KG-RAG [23], and LlamaIndex KG indices are consumers of graphs, improving grounding and multi-hop reasoning. GRAPHWORLD sits upstream, allowing agents to build and maintain the graphs that retrieval systems depend on.

2.2 Graph data models

Various data models exist to represent graphs. GRAPHWORLD includes a version control system that operates on top of a storage format, DIAMOND, and represents graphs using a specific data model. Therefore, to explain the design choices of GRAPHWORLD, it is necessary to discuss the significant body of existing work on graph data models, storage structures, and versioning approaches. We highlight components that were adopted in GRAPHWORLD as well as features of GRAPHWORLD that differ from previous work.

Machine-readable representations of KGs often use the Resource Description Framework (RDF) data model [24], introduced at the World Wide Web Consortium (W3C) in 1999. Under RDF, a graph is represented by a collection of triples. These triples follow the structure of subject-predicate-object (*e.g.*, (Horacio, likes, cars)). A resource is any subject, predicate, or object, and is identified through a Unique Resource Identifier, or URI. Subjects in a triple can be a URI or a blank node, while predicates must be a URI, and objects can be a URI, a blank node, or a literal (*e.g.*, an integer, a float, a string). A graph \mathcal{R} is then represented as a set of RDF triples (also called semantic triples):

$$\mathcal{R} \subseteq (\text{URI} \cup \text{blank}) \times \text{URI} \times (\text{URI} \cup \text{blank} \cup \text{literal})$$

87 where `URI` is the set of all URIs, and `blank` and `literal` are the sets of blank nodes and literals,
88 respectively.

89 Often KGs require the attachment of non-structural information to a node or edge in the graph. RDF
90 does not natively support this, and instead relies on reification. Reification achieves this through the
91 use of metadata predicate types that enable writing triples about other triples, but comes at the cost of
92 larger graph sizes, and has been a source of criticism for the format [25].

93 Labeled Property Graphs (LPGs) have recently emerged as a more flexible and efficient alternative to
94 RDF [3]. They combine edge-labeled graphs, which consist of graphs where nodes are connected
95 by directed edges that contain a label, or type, with property graphs, which allow arbitrary property
96 information to be added to nodes and edges. LPGs have been adopted as the official storage format
97 of the leading graph databases, including Neo4j and ArangoDB, which has driven to further adoption
98 in production settings [26]. It is worth noting that not all databases adopted the same exact definition
99 of LPGs, and each provider has slightly different variants. For example, Neo4j allows any number of
100 labels on nodes but only one label per edge, while ArangoDB supports only one label per node and
101 one label per edge [26].

102 GRAPHWORLD adopts a multi-label definition of LPGs with flat properties that allow agents trained
103 on it to interact with modern KG tooling and databases.

104 2.3 Storage formats and compression

105 While these databases offer avenues to store graphs, they do not do so in a stateless manner. To load,
106 process, and serve the data, the database server must be on. A file format is needed to serialize
107 KGs and enable the exchange of information in a standardized format which can be loaded by the
108 database system or program of the user’s choice. This is akin to relational database management
109 systems (RDBMs) and the “comma-separated values” (CSV) file format that stores columnar data in
110 a text-based encoding. To address this need, Chiba *et al.* [27] have proposed the “Property Graph
111 Exchange Format” (PG) as a standard for representing LPGs. PG is a text-based format that encodes
112 the graph structure as well as the node and edge properties in a human-readable format. The format
113 specification also contains the definition of PG-JSON, a JSON-based serialization of PG that is easy
114 to parse, and PG-JSONL, a newline-delimited JSON format. The format is designed to be easy to
115 read and write, making it suitable for both humans and machines.

116 While versatile, the format can make graph files grow quickly in size, making it impractical to store
117 on git-backed server, which oftentimes have file size restrictions. Thus, GRAPHWORLD supports the
118 PG-JSON family of formats and also introduces DIAMOND, an LPG format that heavily compresses
119 graphs into a fraction of the size of their PG-JSON equivalents.

120 There is a significant body of literature on algorithms for lossless RDF graph compression, which has
121 been reviewed by Besta & Hoefler [28]. For example, some methods apply text compression methods
122 to text-based graph representations [29, 30]. The seminal WebGraph framework uses lexicographic
123 locality and reference encoding to greatly reduce storage per edge [31]. Brisaboa *et al.* [32] proposed
124 k^2 trees, a succinct data structure for graph adjacency that models the graph as a tree, then recursively
125 partitions and stores the adjacency matrix to capitalize on large empty regions in sparse graphs [33,
126 34]. However, these methods only target topology, or the adjacency structure or derived indices. Edge
127 or node attributes are usually not considered or stored in separate, uncompressed arrays. We could
128 identify no compression tools compatible with storing node and edge properties, which is necessary
129 for our use case. That is, no binary encoding format or other efficient representation exists for LPGs,
130 motivating the need for DIAMOND. DIAMOND is described in more detail in Appendix A.2.

131 2.4 Version control systems

132 Unlike text or code, where line-based version control (à la Git) works well, KGs contain complex
133 structured data that requires specialized versioning strategies. In particular, KG versioning approaches
134 encounter the following challenges:

- 135 • **Conflict management and resolution.** While many existing solutions can record linear KG edit
136 histories, they lack support for branching and merging graph versions, and thus cannot support
137 parallel development or KG contextualization [17]. Reconciling KG merge conflicts entails more

than a simple union of all changes: for example, one branch may delete an edge that another branch modifies.

- **Static and queryable storage sizes.** Naïvely storing multiple versions of large KGs can impose significant storage demands. This is especially true of modern KGs that consist of millions, or billions, of nodes and edges. For example, GenomicKG contains 347 million nodes, 1.36 billion edges, and 3.9 billion node and edge properties [35]. GenomicKB was constructed from over 30 genomic datasets and annotations which are regularly updated, necessitating versioning of this large KG. Instead of representing each version as a standalone file, which does not scale with the number of edits, binary encoding and compression techniques can be used to both encode graph structure and eliminate redundancy by only storing version changes. However, graph look-up, membership, or other queries often cannot be executed directly over a compressed KG, leading to undesirable query performance degradation.
- **Lack of standards and interoperability.** Currently, there does not exist a unified “Git for KGs” standard; existing systems often sacrifice one aspect (*e.g.*, merge support or query performance) to gain another (*e.g.*, storage efficiency or simplicity).

Several approaches have been proposed to address these challenges. Existing methods can be categorized as follows [15–17]:

- **Independent copies.** Each graph version is managed and stored as a separate dataset. This method is straightforward to implement but highly redundant and inefficient, both in terms of query time and storage.
- **Change-based versioning.** Only differences, or deltas, between versions are stored. Broadly, change-based methods reduce storage requirements but incur query performance overhead when reconstructing past versions. Many existing change-based approaches like Frommhold *et al.* [36], TailR [37], and RDF-adapted Darcs [38] lack support for branching [17]. Those that do support branching, like R43ples [39], lack streamlined merge support. Thus, current methods are not suited for parallel development.
- **Timestamp-based versioning.** Each edge is annotated with temporal validity information in a single, unified KG. Query performance of timestamp-based methods like R&Wbase [40], ConVer-G [41], and Aion [42] suffers as queries need to filter by version and time, and storage requirements also increase due to the need to preserve temporal metadata. For example, ConVer-G, which uses PostgreSQL as the storage, introduces in-graph provenance information in the graph to facilitate querying, increasing the storage demands.

Finally, inspired by Git, some systems have implemented distributed version control mechanisms for KGs. Most notably, QuitStore [36] and Git4Voc [43] adapt Git-like operations such as branching, merging, and commit tracking to RDF data. However, QuitStore [17] relies on NQuads [44, 45], an inefficient text-based storage format, to store the change information in a manner that Git tools can easily interpret and display. In fact, in the Discussion section of QuitStore, Arndt *et al.* [17] write, “The evaluation of more advanced compression systems for the stored RDF dataset or the employment of a binary RDF serialization format, such as HDT [46] is still subject to future work.” Thus, to date, no single system for versioning large-scale KGs has successfully combined efficient branching and merging with scalable storage and high-performance queries. We sought to achieve this goal in GRAPHWORLD by developing a novel versioning system for graphs, which we describe in Section 3.2.2. This version control system underlies the ability of GRAPHWORLD to support multi-agent concurrent KG editing.

3 The GRAPHWORLD environment

3.1 Theoretical framework

We model GRAPHWORLD as an environment whose state is a fully observable LPG together with its version history. At each time step, the agent proposes an atomic edit to the graph; the environment applies this edit deterministically, producing a new LPG and a new commit in the version history. Formally, GRAPHWORLD is defined by the tuple $\mathcal{M} = (S, A, \delta, R, O)$ where S is the space of states (G, H) with G a labeled property graph and H its commit history, A is the set of atomic LPG edits, δ is the deterministic transition function, R is an externally specified reward function, and $O = S$ since the agent can access the entire LPG structure together with the version history.

191 **State and observation.** At time t , the environment is in state $s_t = (G_t, H_t)$, where $G_t =$
 192 $(V, E, L, l_V, l_E, K, W, p_V, p_E)_t$ is an LPG as defined in Appendix A.1, and H_t is the sequence
 193 of commits leading to G_t . The observation o_t is identical to the state s_t for every t .

194 **Actions.** Each action $a_t \in A$ corresponds to a single mutation of the 9-tuple structure, for example
 195 adding nodes `add_node($v, l_V(v), p_V(v)$)`, edges `add_edge($e, l_E(e), p_E(e)$)`, updating properties
 196 `update_property(x, k, w)` and deleting nodes `delete_node(v)` and edge `delete_edge(e)`. All
 197 actions are recorded in the commit history H_t , preserving an auditable trail of modifications.

198 **Transition dynamics.** The transition function is deterministic $s_{t+1} = \delta(s_t, a_t)$, where δ is the
 199 graph edit operator. A valid edit updates the LPG tuple G_t and appends a commit to H_t . Invalid edits
 200 (e.g., deleting a non-existent node) produce a no-op or transition to an error state, depending on the
 201 configuration of the task.

202 **Rewards.** GRAPHWORLD is task-agnostic. The reward function $R(s_t, a_t, s_{t+1})$ is external to the
 203 environment and can be specified according to downstream objectives. For example, rewards may
 204 encourage edits that improve graph completeness, penalize violations of ontology constraints, or favor
 205 the addition of provenance information. This separation allows evaluation along two complementary
 206 axes: construction quality metrics (task-dependent) and process metrics (environment-dependent,
 207 e.g., edit efficiency or branching performance).

208 **Episodes.** An episode begins with an initial state (G_0, H_0) and proceeds until a stopping condition
 209 is met, such as reaching a target benchmark graph, exhausting an edit budget, or receiving an explicit
 210 termination signal. Because the environment incorporates Git-style versioning, episodes may include
 211 branching and merging trajectories, enabling the study of collaborative and multi-agent editing
 212 strategies.

213 3.2 System overview

214 GRAPHWORLD is designed to provide agents with a practical and extensible environment for building
 215 and maintaining knowledge graphs. Its design follows three guiding principles: standardization,
 216 production readiness, and observability.

217 **Standardization.** GRAPHWORLD offers a uniform interface for graph editing, allowing agents to
 218 interact with knowledge graphs through a consistent set of tools.

219 **Production readiness.** By relying on proven libraries and widely used infrastructure, any agent
 220 developed for GRAPHWORLD can be moved into production with minimal overhead.

221 **Observability.** The framework records each agent action as an auditable step, enabling fine-grained
 222 evaluation of graph construction processes (e.g., scoring agents based on the efficiency of their editing
 223 strategies).

224 At the system level, GRAPHWORLD consists of modular components that can be used independently
 225 or in combination. The two core components are a language-agnostic software development kit
 226 (SDK) for graph compression and manipulation, and a command-line interface (CLI) for versioning
 227 and collaboration.

228 3.2.1 SDK for KG compression with DIAMOND

229 A core component of GRAPHWORLD is DIAMOND, a compact, property-preserving graph storage
 230 format. DIAMOND provides high-performance graph compression and decompression routines,
 231 implemented in Rust for efficiency, with bindings for Python and Node.js via `PyO3` and `napi-rs`.
 232 This multi-language design enables developers to embed the same compression methods into agents
 233 written in different ecosystems, while maintaining consistency across them. Example programs
 234 demonstrating usage in each language are provided in Appendix A.5.

235 The SDK supports two primary capabilities. First, it enables compression and decompression of
 236 property graphs from common formats such as JSON Lines and serialized into the `.diamond` binary
 237 encoding described in Appendix A.2. Second, it provides in-memory graph representations via the

PropertyGraph class, which allows developers to manipulate graphs directly and convert them across formats without loss of information.

3.2.2 CLI for KG versioning

While the SDK addresses scalability, effective multi-agent editing requires version control. To this end, GRAPHWORLD extends Git to support property graphs by defining custom drivers for filtering, diffing, and merging. This integration leverages Git’s mature branching, archiving, and collaboration workflows while adapting them to the semantics of graph data.

The interaction between GRAPHWORLD and Git is coordinated through the `.gitattributes` file, which specifies how different file types should be handled during serialization, comparison, and merging. When a repository is configured for use with GRAPHWORLD, graph files matching user-defined patterns (*e.g.*, `.jsonl` or `.diamond`) are automatically associated with these drivers. This allows contributors to work with a familiar Git interface, while GRAPHWORLD manages the underlying representation of graphs.

Filter drivers handle the translation between formats stored in the repository and those exposed in the working tree. For example, a graph may be archived in the compressed DIAMOND format to save space and ensure consistency, while being presented in a human-readable format when checked out. This dual representation enables contributors to inspect or edit graphs locally without sacrificing the efficiency of binary storage.

Diff drivers define how Git compares graph files across commits. Since graphs are often stored in compressed binary form, traditional line-based diffs are meaningless. The custom diff driver decompresses and interprets the files, producing a semantic comparison in terms of nodes, edges, and properties. This allows reviewers to see meaningful changes, such as a node label update or an edge addition, rather than an opaque binary. Figure 1 provides an example diff between two graphs.

Merge drivers extend this functionality to conflict resolution. Git’s default merge algorithm assumes text-based files, where order matters. For graphs, order is irrelevant, and conflicts should only arise when the same structural element is edited differently across branches. The custom merge driver reconciles graphs by structure, reducing spurious conflicts and enabling order-independent merging (Figure 2). This makes it possible for multiple agents or human collaborators to work concurrently without frequent manual intervention.

```
~ Node (node_101)
  labels: ["person"]
~ country: ["United States", "Japan"]
~ name: ["Alice", "Carol", "Juan"]
+ nicknames: ["Jan"]

~ Edge (edge_101_103_directed)
  from: 101
  to: 103
  directed: true
  labels: ["likes"]
  since: ["2015"]
  engaged: [false, true]
```

Figure 1: A human-readable Git diff of an LPG used in GRAPHWORLD. Red elements indicate deletions, yellow elements indicate modifications, and green elements indicate additions.

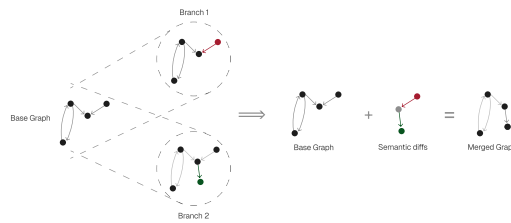


Figure 2: Three-way merge of an LPG in GRAPHWORLD. Starting from a base graph, two branches introduce independent edits. GRAPHWORLD computes semantic diffs and reconciles them through a three-way merge to produce the merged graph.

These drivers integrate with Git’s existing workflows. Contributors can branch, commit, and merge as usual, while GRAPHWORLD ensures that the version control semantics respect the underlying graph structure.

270 4 Applications of GRAPHWORLD environment

271 4.1 Multi-agent collaboration in GRAPHWORLD environment

272 The GRAPHWORLD environment is designed to model agent-agent and agent-human collaboration.
 273 It abstracts away the other participants in the collaboration, such that each contributor must make
 274 no assumptions about the features or functionalities of other collaborators. The agents that operate
 275 within GRAPHWORLD can vary widely in their goals, design, and implementation. We provide two
 276 examples of agents that could exist within the GRAPHWORLD environment.

277 **Change-proposing agent.** An agent might specialize in proposing improvements to a graph, both
 278 in the form of new nodes or edges, or in an edit proposal for an existing entity within the graph. There
 279 are many ways in which these suggestions can be made. The agent could train a graph neural network
 280 (GNN) on the graph with a link prediction objective, predict the edges that are more likely to exist,
 281 and pick the top $k \in \mathbb{N}$ as new edges to add. Alternatively, it could read unstructured documents to
 282 identify known graph entities and look for relationships referenced in the text that are not present in
 283 the current version of the graph. It could also take as input a particular node of interest in a sparse
 284 region of the graph and look for information online that suggests the existence of currently missing
 285 edges. Regardless of how these proposals arise, the agent can create a new branch, perform the
 286 changes, and open a pull request (PR) to merge the new branch into the main one. Figure 3a depicts
 287 the structure of such an agent.

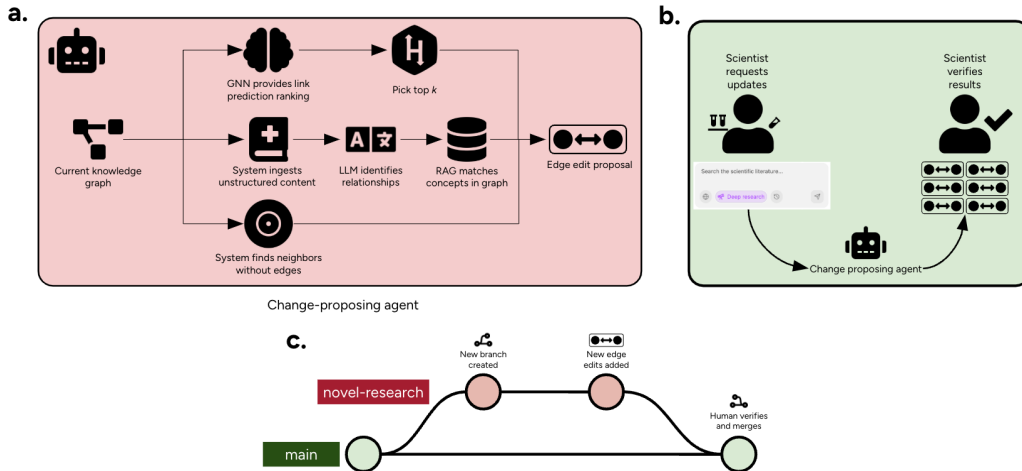


Figure 3: Overview of a change-proposing agent. (a) The agent has access to a toolbox of methods, including GNN-based link prediction on the current KG, relation extraction from unstructured text, and RAG over the literature to score candidates and then synthesize the top- k into edge-edit proposals. (b) An expert scientist states an intent in natural language. The agent materializes it as concrete graph edits on an isolated branch, opens a PR, and returns a reviewable checklist the scientist can accept, modify, or reject. (c) The new agent-created branch does not interfere with the main KG and provides an auditable, reproducible history of KG evolution.

288 **Evidence-gathering agent.** An evidence-gathering agent might review open PRs in the repository
 289 that contain edge proposals and look for information on the Internet, in the scientific literature, in
 290 proprietary databases, or perform independent tests to support or oppose the creation of the new edges
 291 (Figure 4a). This agent plays the same role as peer review in the academic world, but it is automated
 292 and scalable. Scientific research agents show strong performance in synthesizing existing scientific
 293 knowledge from multiple sources [47, 48] and, therefore, may be adept at this role; however, they do
 294 not produce consistent, structured outputs unless appropriately prompted; graphs provide a universal
 295 schema to record their outputs in the context of existing scientific knowledge.

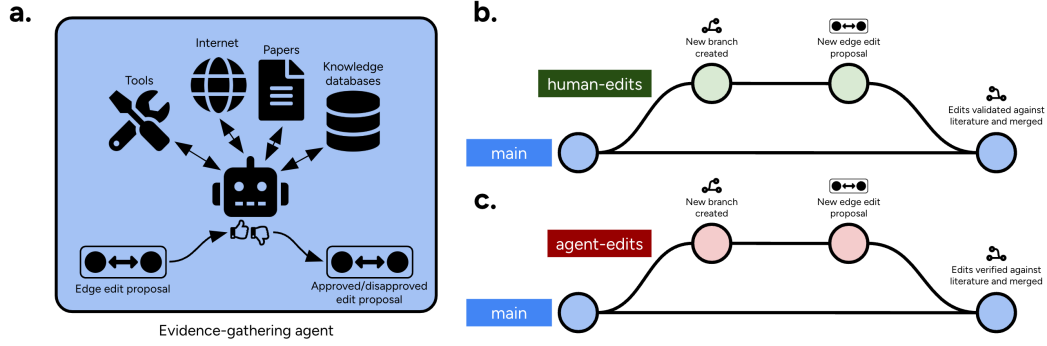


Figure 4: Overview of an evidence-gathering agent. (a) The agent consults multiple external sources, including online tools, scientific papers, and structured knowledge databases, to validate edit proposals. Each candidate edit is either approved or disapproved based on supporting evidence. (b) Human-generated edits proceed through the same process; they are isolated in a new branch, validated against the literature, and then merged into the mainline. (c) Similarly, agent-generated edits, such as those from the change-proposing agent in Figure 3, are proposed on a separate branch and merged once validated.

4.2 Human-AI co-creation of KGs in GRAPHWORLD

Beyond the design of individual agents, GRAPHWORLD supports complex workflows that integrate human input with automated reasoning. These workflows combine the strengths of human domain expertise with the scalability of agents. We describe three representative examples.

Automated validation of human edits. Human maintainers frequently seek to extend or refine a KG. However, detecting semantic inconsistencies or unsupported claims is challenging without computational assistance. In GRAPHWORLD, a change-proposing agent can continuously monitor open pull requests (PRs) submitted by human editors. The agent verifies proposed edits against the available literature or other structured resources and can take several actions: independently approving and merging the PR, providing detailed feedback, or suggesting alternative edits (Figure 4b). This workflow establishes a cycle in which human edits are systematically verified by machine reasoning. The result is improved accuracy of human-generated graphs and reduced reliance on costly manual validation. For graphs deployed in safety-critical applications, such as healthcare or logistics [49–51], the verification workload can be dynamically scaled according to the estimated importance of the affected nodes and edges.

Natural language editing for human domain experts. Domain experts often wish to contribute to KG construction but may lack the technical expertise to interface directly with graph databases or APIs. In GRAPHWORLD, agents serve as interpreters and actuators that translate high-level human intentions into executable graph edits. A user specifies a desired change in natural language. The agent supplements this instruction with information retrieved from structured or unstructured data sources and proposes a set of candidate edits. These edits are applied in an isolated branch, committed, and returned to the user for inspection (Figure 3b, Figure 3c). This workflow allows domain experts to contribute without using graph query languages. As agent capabilities in GRAPHWORLD expand, these systems can move beyond execution of text instructions toward co-pilots that engage in multi-turn discussions with experts about the validity and scope of proposed knowledge changes [52].

Ontology alignment and entity deduplication. Large-scale KG construction often introduces redundant nodes when integrating diverse sources or ontologies [11, 53]. These duplications can fragment the graph: metadata and updates may apply to one duplicate node but not to others, and the topology can become misleading if neighbors are split across copies [54]. In GRAPHWORLD, an agent can be configured to inspect graph branches (such as `main`) and detect such redundancies. Detection leverages semantic signals, including lexical similarity of node labels, synonym resolution, and consistency of attached properties [55]. Upon identifying a likely duplication, the agent can either (i) perform a merge of the redundant nodes, preserving their combined metadata and edges, or (ii)

raise an alert to human maintainers for confirmation. This workflow ensures that entity resolution is continuous and systematic, improving graph integrity as the KG evolves. Multiple agents can monitor graph development in parallel [56, 57], providing early detection of conflicts and contributing to a high-quality and reliable final knowledge base.

4.3 Limitations and future work

Inefficient memory use of library bindings. Internally, the Rust core stores the decompressed graphs in memory as Polars tables. This is extremely memory efficient, but it is difficult to transfer across the bridge to the host binding language. Currently, the graph is converted into a simpler PG-JSON-like representation before it is sent, but switching to this representation forces a copy of memory, uses a space-inefficient layout, and prevents vectorized operations. To address this, we will expose table-level APIs in all bindings and develop zero-copy exchange PyPolars and NodePolars.

Unnecessary deserialization cost. Every file read is currently canonicalized into an in-memory PG-JSON representation before compression, duplicating the memory requirements to read the graph and introducing an unnecessary extra pass. We will add streaming front-ends for every supported format that will perform early type inference to amortize the ingestion overhead and will reduce the need to maintain the entire graph at once in a memory-inefficient representation.

Property type constraints. DIAMOND, following the PG-JSON standard, currently assumes flat optional properties whose values are homogeneous lists of primitive types (Boolean, String, Number). Mixed-type lists, nested objects, and unions are not supported. We will extend schema inference to nested types and perform an intermediate flattening step that enables support for nested types while preserving compact encodings and round-trip fidelity.

These limitations largely reflect the unoptimized nature of the first version of our compression library. We aim to reduce memory consumption through a mix of streaming, pipelining, and the elimination of unnecessary steps, all while improving developer experience.

5 Conclusion

We presented GRAPHWORLD, an environment that makes the construction of KGs observable, collaborative, and reproducible. GRAPHWORLD combines a language-agnostic manipulation library for labeled property graphs, semantic graph diffs and merges built on top of Git, and DIAMOND, a compact, property-preserving columnar graph encoding that scales to multi-million-edge KGs. Together, these components let agents and humans propose, review, and merge edits with auditable histories. We hope GRAPHWORLD will serve as the interoperability layer upon which multi-agent KG systems are built, compared, and improved.

References

1. Ehrlinger, L. & Wöß, W. Towards a definition of knowledge graphs. *SEMANTiCS (Posters, Demos, SuCCESS)* **48**, 2 (2016).
2. Hogan, A. *et al.* Knowledge Graphs. *ACM Comput. Surv.* **54**, 71:1–71:37. doi:10.1145/3447772 (2021).
3. Angles, R., Arenas, M., Barceló, P., Hogan, A., Reutter, J. & Vrgoč, D. Foundations of Modern Query Languages for Graph Databases. *ACM Comput. Surv.* **50**, 68:1–68:40. doi:10.1145/3104031 (2017).
4. Wu, S. *et al.* STaRK: Benchmarking LLM Retrieval on Textual and Relational Knowledge Bases 2024. doi:10.48550/arXiv.2404.13207.
5. Pan, S., Luo, L., Wang, Y., Chen, C., Wang, J. & Wu, X. Unifying Large Language Models and Knowledge Graphs: A Roadmap. *IEEE Transactions on Knowledge and Data Engineering* **36**, 3580–3599. doi:10.1109/TKDE.2024.3352100 (2024).
6. Park, J. S., O’Brien, J., Cai, C. J., Morris, M. R., Liang, P. & Bernstein, M. S. *Generative Agents: Interactive Simulacra of Human Behavior in Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (Association for Computing Machinery, New York, NY, USA, 2023), 1–22. doi:10.1145/3586183.3606763.

- 378 7. Gao, A. K. *Introducing Tuna - A Tool for Rapidly Generating Synthetic Fine-Tuning Datasets*
379 2023.
- 380 8. Chen, D. *et al.* *Data-Juicer: A One-Stop Data Processing System for Large Language Models*
381 2023. doi:10.48550/arXiv.2309.02033.
- 382 9. Chen, D. *et al.* *Data-Juicer 2.0: Cloud-Scale Adaptive Data Processing for and with Foundation*
383 *Models* 2025. doi:10.48550/arXiv.2501.14755.
- 384 10. Rasmussen, P., Paliychuk, P., Beauvais, T., Ryan, J. & Chalef, D. *Zep: A Temporal Knowledge*
385 *Graph Architecture for Agent Memory* 2025. doi:10.48550/arXiv.2501.13956.
- 386 11. Chandak, P., Huang, K. & Zitnik, M. Building a knowledge graph to enable precision medicine.
387 *Scientific Data* **10**, 67. doi:10.1038/s41597-023-01960-3 (2023).
- 388 12. Mo, B. *et al.* *KGGen: Extracting Knowledge Graphs from Plain Text with Language Models*
389 2025. doi:10.48550/arXiv.2502.09956.
- 390 13. Zhao, X. *et al.* *AGENTiGraph: An Interactive Knowledge Graph Platform for LLM-based*
391 *Chatbots Utilizing Private Data* 2024. doi:10.48550/arXiv.2410.11531.
- 392 14. Larson, J. & Truitt, S. *GraphRAG: A new approach for discovery using complex information*
393 2024.
- 394 15. Tzitzikas, Y., Theoharis, Y. & Andreou, D. *On Storage Policies for Semantic Web Repositories*
395 *That Support Versioning in The Semantic Web: Research and Applications* (eds Bechhofer, S.,
396 Hauswirth, M., Hoffmann, J. & Koubarakis, M.) (Springer, Berlin, Heidelberg, 2008), 705–719.
397 doi:10.1007/978-3-540-68234-9_51.
- 398 16. Fernández, J. D., Polleres, A. & Umbrich, J. Towards Efficient Archiving of Dynamic Linked
399 Open Data. *DIACRON@ ESWC* **1377**, 34–49 (2015).
- 400 17. Arndt, N., Naumann, P., Radtke, N., Martin, M. & Marx, E. Decentralized Collaborative
401 Knowledge Management using Git. *Journal of Web Semantics* **54**, 29–47. doi:10.1016/j.
402 websem.2018.08.002 (2019).
- 403 18. Miller, J. J. *Graph database applications and concepts with Neo4j in Proceedings of the*
404 *southern association for information systems conference, Atlanta, GA, USA* **2324** (2013), 141–
405 147.
- 406 19. Edge, D. *et al.* From local to global: A graph rag approach to query-focused summarization.
407 *arXiv:2404.16130* (2024).
- 408 20. Mavromatis, C. & Karypis, G. GNN-RAG: Graph neural retrieval for large language model
409 reasoning. *arXiv:2405.20139* (2024).
- 410 21. Zhu, X., Xie, Y., Liu, Y., Li, Y. & Hu, W. Knowledge graph-guided retrieval augmented
411 generation. *NAACL* (2025).
- 412 22. Huang, Y., Zhang, S. & Xiao, X. *KET-RAG: A cost-efficient multi-granular indexing framework*
413 *for graph-rag in KDD* (2025), 1003–1012.
- 414 23. Sanmartin, D. KG-RAG: Bridging the gap between knowledge and creativity. *arXiv:2405.12035*
415 (2024).
- 416 24. Wood, D., Lanthaler, M. & Cyganiak, R. RDF 1.1 concepts and abstract syntax. *W3C Recom-*
417 *mendation, W3C* (2014).
- 418 25. Zaho, Z., Han, S. K. & Kim, J. R. LPG Representation of the Reification of RDF. *International*
419 *Journal of Engineering and Technology* **7**, 562–566. doi:10.14419/ijet.v7i3.34.19382
420 (2018).
- 421 26. Besta, M. *et al.* Demystifying Graph Databases: Analysis and Taxonomy of Data Organization,
422 System Designs, and Graph Queries. *ACM Comput. Surv.* **56**, 31:1–31:40. doi:10.1145/
423 3604932 (2023).
- 424 27. Chiba, H., Yamanaka, R. & Matsumoto, S. *Property Graph Exchange Format* 2019. doi:10.
425 48550/arXiv.1907.03936. arXiv: 1907.03936 [cs].
- 426 28. Besta, M. & Hoefler, T. *Survey and Taxonomy of Lossless Graph Compression and Space-*
427 *Efficient Graph Representations* 2019. doi:10.48550/arXiv.1806.01799.
- 428 29. Navarro, G. Compressing web graphs like texts. *Dept. Comput. Sci., Univ. Chile, Santiago,*
429 *Chile, Tech. Rep. TR/DCC-2007-2* (2007).
- 430 30. Claude, F. & Navarro, G. Fast and Compact Web Graph Representations. *ACM Trans. Web* **4**,
431 16:1–16:31. doi:10.1145/1841909.1841913 (2010).

31. Boldi, P. & Vigna, S. *The webgraph framework I: compression techniques* in *Proceedings of the 13th international conference on World Wide Web* (Association for Computing Machinery, New York, NY, USA, 2004), 595–602. doi:10.1145/988672.988752.
32. Brisaboa, N. R., Ladra, S. & Navarro, G. *k2-Trees for Compact Web Graph Representation in String Processing and Information Retrieval* (eds Karlgren, J., Tarhio, J. & Hyvärinen, H.) (Springer, Berlin, Heidelberg, 2009), 18–30. doi:10.1007/978-3-642-03784-9_3.
33. Claude, F. & Ladra, S. *Practical representations for web and social graphs* in *Proceedings of the 20th ACM international conference on Information and knowledge management* (Association for Computing Machinery, New York, NY, USA, 2011), 1185–1190. doi:10.1145/2063576.2063747.
34. Brisaboa, N. R., Ladra, S. & Navarro, G. Compact representation of Web graphs with extended functionality. *Information Systems* **39**, 152–174. doi:10.1016/j.is.2013.08.003 (2014).
35. Feng, F. *et al.* GenomicKB: a knowledge graph for the human genome. *Nucleic Acids Research* **51**, D950–D956. doi:10.1093/nar/gkac957 (2023).
36. Frommhold, M., Piris, R. N., Arndt, N., Tramp, S., Petersen, N. & Martin, M. *Towards Versioning of Arbitrary RDF Data* in *Proceedings of the 12th International Conference on Semantic Systems* (Association for Computing Machinery, New York, NY, USA, 2016), 33–40. doi:10.1145/2993318.2993327.
37. Meinhardt, P., Knuth, M. & Sack, H. *TailR: a platform for preserving history on the web of data* in *Proceedings of the 11th International Conference on Semantic Systems* (Association for Computing Machinery, New York, NY, USA, 2015), 57–64. doi:10.1145/2814864.2814875.
38. Cassidy, S. & Ballantine, J. Version Control for RDF Triple Stores. *ICSOFT (ISDM/EHST/DC)* **7**, 5–12 (2007).
39. Graube, M., Hensel, S. & Urbas, L. *R43ples: Revisions for triples* in *Proceedings of the 1st Workshop on Linked Data Quality co-located with 10th International Conference on Semantic Systems (SEMANTiCS 2014)* (2014).
40. Vander Sande, M., Colpaert, P., Verborgh, R., Coppens, S., Mannens, E. & Van de Walle, R. R&Wbase: Git for triples. *LDOW* **996** (2013).
41. Gil, J. P., Coquery, E., Samuel, J. & Gesquiere, G. *ConVer-G: Concurrent versioning of knowledge graphs* 2024. doi:10.48550/arXiv.2409.04499.
42. Theodorakis, G., Clarkson, J. & Webber, J. *Aion: Efficient Temporal Graph Data Management* in (Paestum, Italy, 2024). doi:10.48786/EDBT.2024.43.
43. Halilaj, L., Grangel-González, I., Coskun, G. & Auer, S. *Git4Voc: Git-based Versioning for Collaborative Vocabulary Development* 2016. doi:10.48550/arXiv.1601.02433.
44. *RDF 1.1 N-Quads* 2014.
45. *RDF Binary using Apache Thrift* 2025.
46. Fernandez, J. D., Martínez-Prieto, M. A., Gutiérrez, C., Polleres, A. & Arias, M. *Binary RDF Representation for Publication and Exchange (HDT)* SSRN Scholarly Paper. Rochester, NY, 2013. doi:10.2139/ssrn.3198999.
47. Skarlinski, M. D. *et al.* *Language agents achieve superhuman synthesis of scientific knowledge* 2024. doi:10.48550/arXiv.2409.13740.
48. *Ai2. Introducing Ai2 Paper Finder* 2025.
49. Alber, D. A. *et al.* Medical large language models are vulnerable to data-poisoning attacks. *Nature Medicine* **31**, 618–626 (2025).
50. Alsentzer, E. *et al.* Few shot learning for phenotype-driven diagnosis of patients with rare genetic diseases. *npj Digital Medicine* **8**, 380 (2025).
51. Yang, J. *et al.* Poisoning medical knowledge using large language models. *Nature Machine Intelligence* **6**, 1156–1168 (2024).
52. Gao, S. *et al.* *Empowering Biomedical Discovery with AI Agents* 2024. doi:10.48550/arXiv.2404.02831.
53. Lobentanzer, S. *et al.* Democratizing Knowledge Representation with BioCypher. *Nature Biotechnology* **41**, 1056–1059. doi:10.1038/s41587-023-01848-y (2023).
54. Callahan, T. J. *et al.* An open source knowledge graph ecosystem for the life sciences. *Scientific Data* **11**, 363 (2024).

- 486 55. Johnson, R. *et al.* ClinVec: Unified Embeddings of Clinical Codes Enable Knowledge-Grounded
487 AI in Medicine. *medRxiv*, 2024–12 (2024).
- 488 56. Lu, Y. & Wang, J. KARMA: Leveraging Multi-Agent LLMs for Automated Knowledge Graph
489 Enrichment. *arXiv:2502.06472* (2025).
- 490 57. Liu, B., Zhang, J., Lin, F., Yang, C., Peng, M. & Yin, W. *SymAgent: A neural-symbolic self-
491 learning agent framework for complex reasoning over knowledge graphs* in *Proceedings of the
492 ACM on Web Conference* (2025), 98–108.
- 493 58. Chiba, H. & Voß, J. Property Graph Exchange Format (PG). doi:10.5281/zenodo.13859531
494 (2024).
- 495 59. *pola-rs/polars* 2025.
- 496 60. *apache/arrow-rs* 2025.
- 497 61. Vohra, D. in *Practical Hadoop Ecosystem: A Definitive Guide to Hadoop-Related Frameworks
498 and Tools* (ed Vohra, D.) 325–335 (Apress, Berkeley, CA, 2016). doi:10.1007/978-1-4842-
499 2199-0_8.
- 500 62. *apache/parquet-java* 2025.
- 501 63. Collet, Y. & Kucherawy, M. *Zstandard Compression and the application/zstd Media Type*
502 Request for Comments RFC 8478 (Internet Engineering Task Force, 2018). doi:10.17487/
503 RFC8478.
- 504 64. Chandak, P., Huang, K. & Zitnik, M. Building a Knowledge Graph to Enable Precision Medicine.
505 *Scientific Data* **10**, 67. doi:10.1038/s41597-023-01960-3 (2023).

A Technical Appendices and Supplementary Material

A.1 Formal definition of labeled property graphs

Formally, LPGs can be represented as a 9-tuple $G = (V, E, L, l_V, l_E, K, W, p_V, p_E)$ [26], where:

- V is the set of nodes,
- E is the set of edges,
- L is the set of labels,
- $l_V : V \rightarrow \mathcal{P}(L)$ is the function that maps nodes to their labels,
- $l_E : E \rightarrow \mathcal{P}(L)$ is the function that maps edges to their labels,
- K is the set of all possible property keys,
- W is the set of all possible property values,
- $p_V : V \rightarrow K \times W$ is the function that maps nodes to their property key-value pairs, and
- $p_E : E \rightarrow K \times W$ is the function that maps edges to their property key-value pairs.

Note that under this formulation l_V and l_E map to the power set of L , and not to L itself. Therefore, every node and edge can have one label, more than one label, or no labels associated.

A.2 Graph compression algorithm

We designed binary encoding and compression techniques to efficiently store and represent labeled property (LP) graphs. This enabled us to encode graph structure and metadata in a single, reduced-size file, and simplified storage in the git server. Initially, we attempted to identify and use a pre-existing binary encoding format for LPG; however, as described in Section 2.3, given the recency of the LPG format, we could not identify any. Therefore, we developed our own binary encoding format for compressing graphs, which we refer to as DIAMOND. The compression steps are as follows (Figure 5).

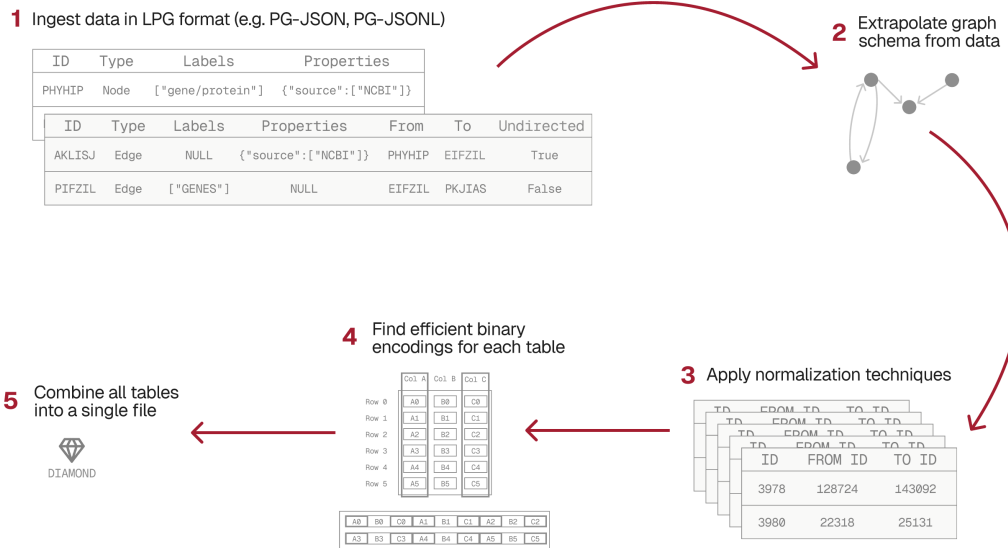


Figure 5: The DIAMOND graph compression algorithm. Each step of the compression algorithm is depicted.

Ingestion. First, we ingest the data in a standard format used to represent LPG graphs, for example, in Graphviz DOT, JSON lines, or PG file formats. If the input format is not PG-JSON, the original file is first converted to PG-JSON. This avoids having to define a separate compression routine for every standard.

532 **Schema extrapolation.** Next, we extrapolate a strongly typed schema for the graph by sifting
533 through all the node and edge records to understand their structure. The first step in the schema
534 extrapolation is to split all the nodes and edges into groups. Groups are understood to potentially
535 have distinct properties. Separation into groups allows us to find an efficient representation for each
536 group that avoids having many nulls in our final binary file.

537 We create these groups by separating graph elements by label combinations, with all nodes with a
538 specific order-agnostic label combination grouped together. The same is done for edges. In practice,
539 we achieve order-agnostic grouping by lexicographically sorting the labels, but the compression
540 mechanism is method-agnostic. Each group within each element type is assigned an identifier in the
541 form of a natural number that will be used to reference the group in other parts of the algorithm. The
542 full pseudocode for GROUP_BY_LABELS can be found in the appendix under Algorithm 1.

543 Once all elements have been partitioned into their respective groups, the types of their properties are
544 inferred (see Algorithm 2). Following the PG-JSON standard [58], we assume that all properties
545 are optional. That is, a node with labels {animal, dog} might have property `nicknames`: ["puppy",
546 "pup"] declared but another node with the same labels might not. According to the PG-JSON
547 standard [58], all properties must have as their value a list of type `Boolean`, `String` or `Number`
548 (floating point). For each element in a group we iterate collecting all the properties present and storing
549 their corresponding types. There are two potential sources of conflict in this scenario.

- 550 1. **Mismatched type within a property value list.** The value for a property list in the PG-
551 JSON format can hold elements of different primitive types (*e.g.*, ["puppy", 1]). This is
552 an unsupported feature in DIAMOND, as much of the space gains come from leveraging the
553 assumption that properties will always have the same type.
- 554 2. **Mismatched type for a property across elements.** Two elements belonging to the same
555 group could have lists composed of different values. Once again, DIAMOND relies on all
556 instances of a property within a group being of the same type, so an error is found if a
557 mismatch occurs.

558 If no conflicts are found the result after this step is finished is a mapping from a group identifier to a
559 map from property name to value list inner type. The full implementation of the type inference for a
560 specific group can be found in the appendix under Algorithm 3.

561 **Table creation.** We create two tables, one for nodes and one edges, that contain metadata about the
562 group. They each have two columns, `id` and `type`. The `type` column contains in each cell a value of
563 type `String[]` that represents a unique label combination and the `id` column contains the numerical
564 identifier assigned to that combination in the group partitioning step. For the detailed algorithm see
565 Algorithm 4.

566 Following the creation of the group metadata tables, we proceed to create one table per element group
567 to hold the information about its elements. For node groups each table consists of an `id` column that
568 holds the node identifier and one column per possible property for that node group (see Algorithm 5).
569 For edge groups each table consists of columns `id`, `from`, `to`, `undirected` plus one column per
570 possible property (see Algorithm 6).

571 **Transform tables to efficient encoding.** We encode all the group tables efficiently in memory by
572 using Polars [59] data frames. Polars data frames are tables consisting of multiple Apache Arrow
573 columns [60]. Arrow is a software framework for dealing with columnar data. It provides efficient
574 in-memory storage for various data types as well as utility functions for operating on that data. For
575 each node and edge group we iterate over their elements and convert each their static (*e.g.*, `id`, `from`)
576 and dynamic (*e.g.*, `nickname`) properties to Arrow series (see Algorithm 7) by finding the appropriate
577 Arrow data type that efficiently represents the underlying values. After all columns for a group are
578 created we group them in a Polars data frame, attaching the appropriate column name to the data
579 frame header.

580 **Serializing the graph.** We obtain a binary file by saving all the data frames created so far to disk
581 using the Apache Parquet format [61, 62] and grouping the files together using `tar`. We save each
582 data frame as a column-oriented Parquet file, compressing every column with Zstandard [63] at level
583 3 – an intermediate setting that offers a favorable trade-off between compression ratio and processing
584 speed – and partition the output into row groups of 1024×1024 rows (roughly one million rows) so

585 that the files remain highly compressible while still allowing efficient parallel reads and selective
 586 access to individual subsets of the data. Finally, all Parquet tables are combined together into a single
 587 on-disk file, which is further compressed to minimize the on-disk file size. The final tar-compressed
 588 file is saved with a `.diamond` extension.

589 By applying these steps – data ingestion, schema inference, normalization, binary encoding, and
 590 bundling – the DIAMOND algorithm losslessly shrinks the size of the original graph data.

591 A.2.1 Algorithms

Algorithm 1 GROUP_BY_LABELS

Input: `elements` ▷ a list of items, each with a `labels()` method
Output: `groups` ▷ a map from a set of labels to (`type_id`, list of items)

```

1: groups ← empty map
2: next_type_id ← 0
3: for all item in elements do
4:   labels ← item.labels()
5:   if labels is not empty then
6:     key ← sort(labels) ▷ sort alphabetically for a stable key
7:     if key not in groups then
8:       groups[key] ← (next_type_id, empty list)
9:       next_type_id ← next_type_id + 1
10:    end if
11:    append item to the list inside groups[key]
12:  end if
13: end for
14: return groups
```

Algorithm 2 INFER_PROPERTY_TYPES_FOR_GROUPS

Input: `groups` ▷ a map from a set of labels to (`type_id`, list of items)
Output: `prop_types_map` ▷ a map from the same label set to (`property name` → `data type`)

```

1: prop_types_map ← empty map
2: for all labels in groups do
3:   items ← groups[labels].items
4:   properties_list ← list of item.properties() for every item in items
5:   status, inferred ← INFERTYPES(properties_list)
6:   if status is Success then
7:     prop_types_map[labels] ← inferred
8:   else if status is MismatchedTypesWithinItem then
9:     raise TypeInferenceFailed(labels, status.property)
10:  else ▷ mismatched types across items
11:    raise MismatchedListInnerType(labels, status.property, status.expected)
12:  end if
13: end for
14: return prop_types_map
```

Algorithm 3 INFER_TYPES_FOR_PROPERTIES_VEC

Input: properties_vec ▷ list of property dictionaries
Output: types_map ▷ map (property name → data type)

```
1: types_map ← empty map
2: for all properties in properties_vec do
3:   for all (key, value_list) in properties do
4:     status, data_type ← INFERTYPE(value_list)
5:     ▷ Iterates over list checking all values are of the same type
6:     if status is ErrorWithinList then
7:       raise MismatchedTypesWithinItem(key, status.data_types)
8:     end if
9:     if key in types_map then
10:      prev_type ← types_map[key]
11:      if prev_type ≠ data_type then
12:        raise MismatchedTypesAcrossItems(key, prev_type, data_type)
13:      end if
14:    end if
15:    types_map[key] ← data_type
16:  end for
17: end for
18: return types_map
```

Algorithm 4 BUILD_TYPES_DATA_FRAME

Input: type_map ▷ map from labels to (type_id, list of items)
Output: df ▷ table with two columns: id (integer) and type (list of labels)

```
1: type_ids ← empty list
2: type_labels ← empty list
3: for all (labels, (tid, _)) in type_map do
4:   append tid to type_ids
5:   append labels to type_labels
6: end for
7: col_id ← CREATECOLUMN("id", type_ids)
8: col_type ← CREATELISTCOLUMN("type", type_labels)
9: df ← CREATEDATAFRAME(col_id, col_type)
10: return df
```

Algorithm 5 BUILD_DATA_FRAME_FOR_NODE_GROUP

Input: nodes ▷ list of node objects
Input: property_types ▷ map (property name → data type)
Output: df ▷ table with an id column plus one column per property

```
1: columns ← empty list
2: append COLUMN("id", list of node.id for each node in nodes) to columns
3: prop_cols ← GETPROPERTYCOLUMNS(nodes, property_types)
4: append every element of prop_cols to columns
5: df ← CREATEDATAFRAME(columns)
6: return df
```

Algorithm 6 BUILD_DATA_FRAME_FOR_EDGE_GROUP

Input: edges ▷ list of edge objects
Input: property_types ▷ map (property name → data type)
Output: df ▷ table with columns id, from, to, undirected, plus one column per property
1: columns \leftarrow empty list
2: append COLUMN("id", list of edge.id for each edge in edges) to columns
3: append COLUMN("from", list of edge.from for each edge in edges) to columns
4: append COLUMN("to", list of edge.to for each edge in edges) to columns
5: append COLUMN("undirected", list of edge.undirected for each edge in edges) to columns
6: prop_cols \leftarrow GETPROPERTYCOLUMNS(edges, property_types)
7: append every element of prop_cols to columns
8: df \leftarrow CREATEDATAFRAME(columns)
9: **return** df

Algorithm 7 GET_DATA_FRAME_PROPERTY_COLUMNS

Input: elements ▷ list of items, each with a properties() map
Input: property_types ▷ map (property name → data type)
Output: columns ▷ list of table columns, one per property
1: columns \leftarrow empty list
2: **for all** (prop_name, dtype) **in** property_types **do**
3: builder \leftarrow CREATELISTBUILDER(dtype)
4: **for all** item **in** elements **do**
5: values \leftarrow item.properties()[prop_name] ▷ may be missing
6: **if** values exists **then**
7: series \leftarrow SERIESFROMVALUES(values, dtype)
8: builder.add(series)
9: **else**
10: builder.add_null()
11: **end if**
12: **end for**
13: col \leftarrow COLUMN(prop_name, builder.finish())
14: append col to columns
15: **end for**
16: **return** columns

592 A.3 Benchmarking DIAMOND

593 A.3.1 Benchmarking on diverse synthetic graphs

594 To evaluate the performance of the DIAMOND library under various graph characteristics and scales,
595 we conducted a benchmarking study using synthetically generated graphs. This approach allowed us
596 to control specific graph properties and observe their impact on compression efficiency and memory
597 consumption.

598 First, we designed a synthetic graph generator parameterized by three variables: $\mu \in \mathbb{R}$, $\sigma \in \mathbb{R}$, and
599 $p \in [0, 1]$. The parameters μ and σ governed the number of properties associated with a group of
600 nodes or edges, which followed a Normal distribution with mean μ and standard deviation σ . The
601 parameter p represented the probability that a given node or edge would have a non-null value for a
602 specific property. By adjusting μ , σ , and p , we could simulate graphs with varying levels of property
603 density.

604 Using this generator, we created LPGBENCH, a dataset of diverse graphs with varying numbers of
605 nodes, edges, properties, and labels per element. The configurations in LPGBENCH included:

- 606 • **Micro:** 10 nodes, 100 edges, maximum 1 label per element, $\mu = 2.0$, $\sigma = 1.0$, $p = 0.3$,
607 maximum 1 value per selected property.

- **Small:** 1,000 nodes, 10,000 edges, maximum 1 label per element, $\mu = 3.0$, $\sigma = 1.0$, $p = 0.5$, maximum 2 values per selected property.
- **Medium:** 100,000 nodes, 1,000,000 edges, maximum 1 label per element, $\mu = 4.0$, $\sigma = 1.0$, $p = 0.7$, maximum 2 values per selected property.
- **Large:** 1,000,000 nodes, 10,000,000 edges, maximum 2 labels per element, $\mu = 5.0$, $\sigma = 1.0$, $p = 0.9$, maximum 3 values per selected property.

We evaluated DIAMOND on LPGBENCH to understand the performance characteristics of the DIAMOND library and compare the .diamond format against other popular graph representations. Specifically, our benchmarking analyses were as follows.

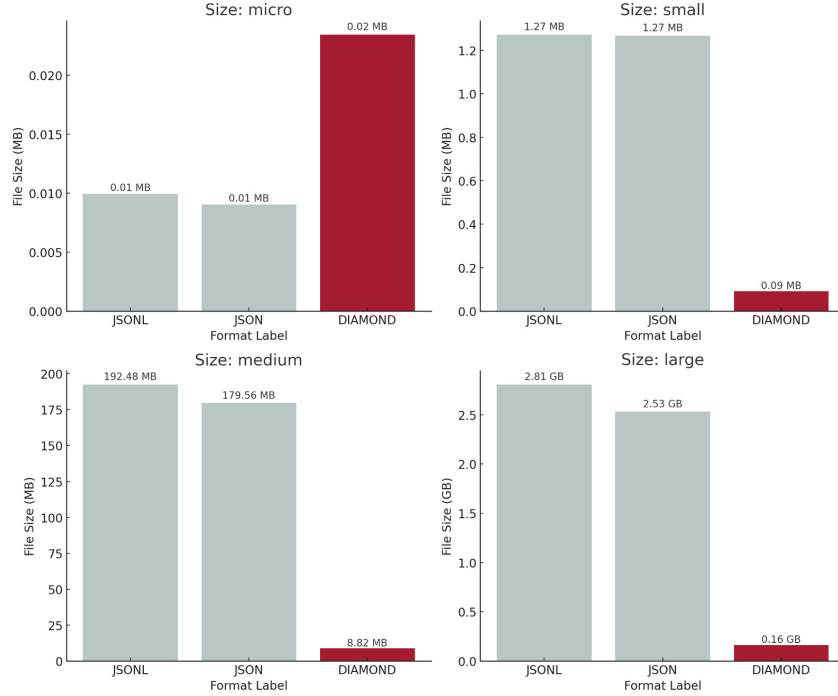


Figure 6: File size vs. graph format across graph sizes. The .diamond file size is shown in red, as compared to JSON and JSON Lines.

File size versus graph format (Figure 6). This analysis aimed to establish a baseline comparison of the on-disk storage requirements for different graph formats across representative graph sizes (micro, small, medium, and large). By fixing the graph structure and size, we directly compared the inherent storage overhead and compression effectiveness of each format without the influence of scaling property densities. As shown in Figure 6, at the smallest graph sizes with only 10 nodes and 100 edges, DIAMOND was outperformed by JSON and JSON lines. However, at even small graph sizes with 1,000 nodes and 10,000 edges, the .diamond file is significantly smaller than its JSON and JSON lines counterparts. Once the graph size scales to 10 million edges, the DIAMOND-compressed file is only 5.69% the size of JSON Lines and 6.32% the size of JSON.

File size across graph sizes and property densities (Figure 7). This analysis investigated how the file size of each graph format scales as the total number of graph elements increases, while maintaining consistent property distributions defined by μ , σ and p . We sought to understand the scalability of each format and evaluate how efficiently they handle increasing graph size under different scenarios of property density and sparsity. As depicted in Figure 7, the compression ratio achieved by DIAMOND, relative to JSON or JSON Lines, improves as properties become more dense (e.g., comparing results for $p = 0.3$ versus $p = 0.9$, note that the y -axis is shared across the panels). This observation aligns with our hypothesis that DIAMOND achieves greater compression efficiency on graphs with higher property density.

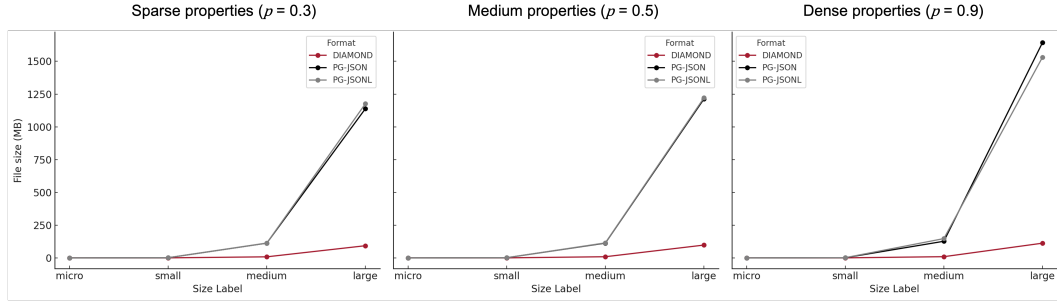


Figure 7: File size across graph sizes and property densities. The .diamond file size is shown in red, as compared to JSON and JSON Lines.

635 A.3.2 Benchmarking on real-world KGs

636 Beyond our synthetic benchmarks, we also sought to test the performance of DIAMOND on a real-
 637 world KG. When applied on PrimeKG [64], a popular biomedical KG with over 55,000 downloads on
 638 Harvard Dataverse at the time of writing, DIAMOND achieves up to $34.1\times$ compression as compared
 639 to other prevalent LPG graph formats, including CSV header, PG, YARS-PG, DOT, Cypher, and
 640 JSONL (Figure 8). It consumes only 8.9% the size of the next smallest format and uses 2.9% of
 641 the space required by the JSONL representation of the KG. Therefore, we successfully designed a
 642 compressed format for LPG graphs that outperforms state-of-the-art graph representations.

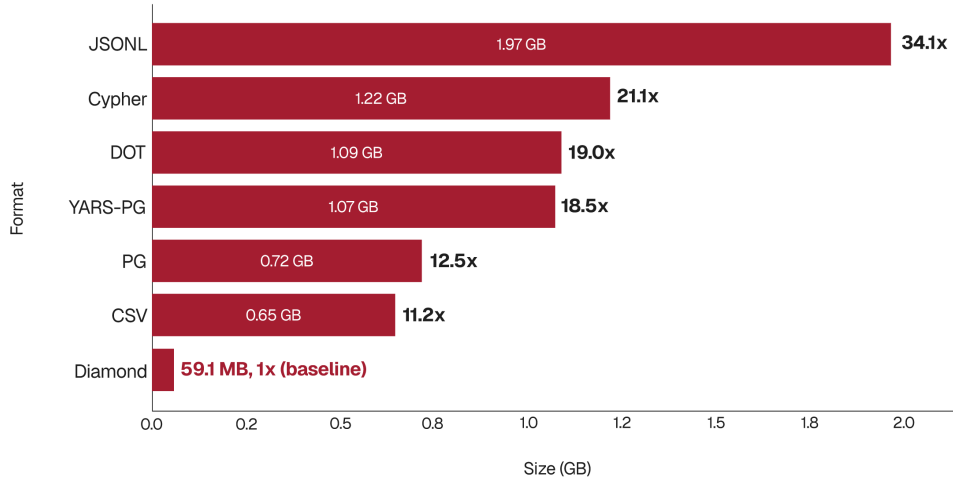


Figure 8: DIAMOND performance on a real-world KG. As compared to other popular graph-encoding formats, DIAMOND achieves up to a $34.1\times$ compression ratio when used to compress PrimeKG [11].

643 A.4 Production-readiness of DIAMOND

644 The DIAMOND library employs a continuous integration and continuous deployment (CI/CD) pipeline
 645 based on GitHub actions to automate testing and streamline code release. CI/CD generates releases
 646 based on the code modified in a pull request. The CI/CD pipeline is as follows:

647 **Python continuous integration.** If Python binding code is altered, a workflow is triggered that
 648 performs linting, uses Bandit for security analysis, uses interrogate to evaluate docstring coverage,
 649 performs static analysis and type checking with mypy, and performs unit testing with pytest. All
 650 these checks are orchestrated via a Makefile.

651 **Rust continuous integration.** Similarly, modifications to the Rust binding code trigger a separate
652 workflow to lint and format the Rust code. A scheduled workflow is also in place to clear CI/CD
653 caches.

654 **Continuous deployment.** Upon pushing Python binding code to the trunk, a release workflow
655 is executed. This comprehensive workflow generates Python wheels for a range of platforms and
656 architectures, including various Linux targets (x86_64, x86, aarch64, armv7, s390x, ppc64le),
657 Windows (x64, x86), and macOS (x86_64, aarch64), creates a release following semantic versioning;
658 and publishes the package to PyPI.

659 A.5 Multi-language support in DIAMOND

660 The following code snippets illustrate the core graph loading and saving operations using the
661 PropertyGraph class across different language bindings for the DIAMOND library. First, the
662 user can instantiate a PropertyGraph object by reading data from a source file in the JSON Lines
663 format using methods like read_pg_jsonl or language-specific equivalents. This process parses the
664 input data and constructs the in-memory graph representation. Subsequently, the write_diamond
665 (or equivalent) method allows the user to serialize the PropertyGraph object into the .diamond
666 binary format. This can be accomplished as follows in the Rust core:

```
1      use diamond_core::PropertyGraph;
2
3      fn main() -> Result<(), Box<dyn std::error::Error>> {
4          let pg = PropertyGraph::read_pg_jsonl("./data/my_graph.jsonl")
5              .expect("Error reading graph");
6          println!("Successfully read graph");
7
8          pg.write_diamond("./data/my_graph.diamond")
9              .expect("Error writing diamond file");
10         println!("Successfully wrote graph");
11     }
```

667 The TypeScript binding allows the same operations to be performed in TypeScript:

```
1      import { PropertyGraph } from "diamond-graph";
2
3      const pg = PropertyGraph.readPgJsonl(inputJsonlPath);
4      console.log(`Successfully read graph from ${inputJsonlPath}`);
5      pg.writeDiamond(outputDiamondPath);
```

668 Finally, the Python binding allows the same operations to be performed in Python:

```
1      from diamond_graph import PropertyGraph
2
3      pg = PropertyGraph.read_pg_jsonl("./data/my_graph.jsonl")
4      pg.write_diamond("./data/my_graph.diamond")
```

669 A.6 Identifying graph differences through recursion

670 To compute the differences between two graphs, we recursively apply two generic primitives,
671 computeSetChanges and computeListChanges. These graph deltas are then useful in the merging
672 process.

673 A.6.1 Computing set changes

674 As described in Algorithm 15, we first define a primitive to determine the changes necessary to
 675 transform an initial “source” set of elements, denoted as S_{source} , into a “target” set, S_{target} , by
 676 identifying elements that must be added, removed, or modified. We define a comparison function,
 677 $f_{\text{compare}}(e_1, e_2)$, which evaluates any two elements e_1 and e_2 (one from the source and one from the
 678 target) and determines if they are identical (Equal), if the target element is a modified version of the
 679 source element (Modified), or if they are otherwise distinct.

680 Next, we attempt to establish correspondences between elements in S_{source} and S_{target} . The algorithm
 681 iterates through each element $s_i \in S_{\text{source}}$. For each s_i that has not yet been matched, it then scans
 682 through elements $t_j \in S_{\text{target}}$ that also remain unmatched. Upon comparing s_i and t_j using f_{compare} ,
 683 if an Equal relationship is found, s_i is recorded as unchanged, and both s_i and t_j are marked as
 684 accounted for, preventing their re-evaluation. The search for a match for s_i then concludes. Similarly,
 685 if f_{compare} indicates a Modified relationship, the pair (s_i, t_j) is stored to signify that s_i transforms
 686 into t_j , both elements are marked as accounted for, and the algorithm moves to the next source
 687 element. This systematic pairing ensures that each element from either set is part of at most one such
 688 Equal or Modified relationship.

689 Following this matching phase, the algorithm identifies elements for removal by examining the source
 690 set. Any element $s_i \in S_{\text{source}}$ that was not marked as matched during the previous step is considered
 691 to be absent from the target set (either directly or as a modified version) and is thus designated for
 692 removal. Conversely, elements for addition are identified by examining the target set. Any element
 693 $t_j \in S_{\text{target}}$ that remains unmarked is interpreted as a new element not present in the source set and is
 694 designated for addition.

695 Ultimately, the algorithm outputs four collections: a list of elements to be added (L_{add}), a list of
 696 elements to be removed (L_{remove}), a list of pairs representing modifications (L_{modify}), and a list of
 697 elements that were found in both sets and remained unchanged ($L_{\text{no_change}}$). These collections
 698 collectively define the delta transforming S_{source} into S_{target} . The primary computational load arises
 699 from the nested iterative search for matches, leading to a worst-case time complexity of $O(|S_{\text{source}}| \cdot$
 700 $|S_{\text{target}}|)$ comparisons.

701 A.6.2 Computing list changes

702 Next, as shown in Algorithm 14, we define a primitive to determine the most efficient sequence of
 703 changes necessary to transform an initial “source” list of elements into a “target” list using dynamic
 704 programming. Note that, unlike Algorithm 15, this primitive deals with ordered lists rather than
 705 unordered sets.

706 The first step involves constructing the cost matrix, where each cell represents the minimum number of
 707 operations required to convert a prefix of the source sequence into a prefix of the target sequence.
 708 The matrix edges are initialized, which corresponds to transforming a sequence into an empty one
 709 (requiring deletions) or an empty sequence into a target sequence (requiring insertions). Then, the
 710 rest of the matrix is iteratively computed. For any given pair of prefixes, the last elements of these
 711 prefixes are considered to calculate the transformation cost. If these elements are deemed identical by
 712 a provided comparison function $f_{\text{compare}}(e_1, e_2)$, no new cost is incurred, and the value is carried over
 713 from a previous state. If they differ, the algorithm explores the costs of three potential operations:
 714 deleting the element from the source, inserting the element into the target, or modifying the source
 715 element to match the target element. The comparison function can assign different costs based on
 716 whether differing elements are considered “modified” versions or entirely distinct. The algorithm
 717 always chooses the operation that produces the minimum cumulative cost for that particular cell.

718 Once this cost matrix is computed, the value in the cell corresponding to the full source and target
 719 sequences represents the total minimum cost for the entire transformation. Starting from this final
 720 cell, the algorithm traces a path back to the beginning of the matrix. The path taken is determined
 721 by reversing the equality, modification, addition, or removal decisions made during the matrix
 722 construction. This path directly translates into the sequence of operations that optimally transform
 723 the source sequence into the target sequence. The final output is this ordered list of changes. This
 724 algorithm is quadratic in terms of the lengths of the two sequences, both in time and memory.

725 A.6.3 Leveraging change primitives to compute graph differences

726 We recursively utilize the following primitives to compare two graphs, starting at the highest level
 727 and traversing all the way down to single property lists. This level of detail enables us to write very
 728 flexible logic for the graph merging strategies.

Algorithm 8 COMPARE_GRAPHS

Input: source, target ▷ full graph objects
Output: graphChanges ▷ {nodeChanges, edgeChanges}
 1: nodeChanges \leftarrow COMPARENODES(source.nodes, target.nodes)
 2: edgeChanges \leftarrow COMPAREEDGES(source.edges, target.edges)
 3: **return** { nodeChanges, edgeChanges }

Algorithm 9 COMPARE_NODES

Input: source, target ▷ lists of nodes
Output: nodeChanges ▷ add / remove / modify / unchanged summary
 1: **procedure** CMPNODE(n_a, n_b)
 2: **if** NODEKEY(n_a) \neq NODEKEY(n_b) **then**
 3: **return** Different
 4: **end if**
 5: labelDiffs \leftarrow COMPARELABELS(n_a.labels, n_b.labels)
 6: propertyDiffs \leftarrow COMPAREPROPERTIES(n_a.properties, n_b.properties)
 7: labelsEqual \leftarrow every change in labelDiffs is Equal
 8: propsEqual \leftarrow propertyDiffs.add = remove = modify = 0
 9: **return if** labelsEqual \wedge propsEqual **then** Equal **else** Modified
 10: **end procedure**
 11: nodeChanges \leftarrow COMPUTESETCHANGES(Set(source), Set(target), cmpNode)
 12: **return** nodeChanges

Algorithm 10 COMPARE_EDGES

Input: source ▷ list of edges in graph A
Input: target ▷ list of edges in graph B
Output: edgeChanges ▷ add / remove / modify / unchanged summary
 1: **procedure** CMPEDGE(e_a, e_b) ▷ returns Equal, Modified, Different
 2: **if** EDGEKEY(e_a) \neq EDGEKEY(e_b) **then**
 3: **return** Different ▷ different IDs \rightarrow cannot match
 4: **end if**
 5: labelDiffs \leftarrow COMPARELABELS(e_a.labels, e_b.labels)
 6: propertyDiffs \leftarrow COMPAREPROPERTIES(e_a.properties, e_b.properties)
 7: sameEnds \leftarrow (e_a.from = e_b.from) \wedge (e_a.to = e_b.to)
 8: sameDir \leftarrow (e_a.undirected = e_b.undirected)
 9: labelsEqual \leftarrow every change in labelDiffs is Equal
 10: propsEqual \leftarrow propertyDiffs.add = remove = modify = 0
 11: **return if** sameEnds \wedge sameDir \wedge labelsEqual \wedge propsEqual **then** Equal **else**
 Modified
 12: **end procedure**
 13: edgeChanges \leftarrow COMPUTESETCHANGES(Set(source), Set(target), cmpEdge)
 14: **return** edgeChanges

Algorithm 11 COMPARE_LABELS

Input: srcLabels, tgtLabels ▷ ordered lists
Output: labelChanges ▷ add / remove / modify / unchanged summary
1: **procedure** CMPLABEL(a, b)
2: **return** if a = b **then** Equal **else** Different
3: **end procedure**
4: **return** COMPUTELISTCHANGES(srcLabels, tgtLabels, cmpLabel)

Algorithm 12 COMPARE_PROPERTIES

Input: srcProps, tgtProps ▷ maps key → list
Output: propChanges ▷ set-style diff for (key, list) entries
1: **procedure** CMPENTRY((k_a, v_a), (k_b, v_b))
2: **if** k_a ≠ k_b **then**
3: **return** Different
4: **end if**
5: listDiffs ← COMPAREPROPERTYLIST(v_a, v_b)
6: **return** if every change in listDiffs is Equal **then** Equal **else** Modified
7: **end procedure**
8: propChanges ← COMPUTESETCHANGES(Set(Entries(srcProps)),
 Set(Entries(tgtProps)), cmpEntry)
9: **return** propChanges

Algorithm 13 COMPARE_PROPERTY_LIST

Input: srcList, tgtList ▷ ordered value lists
Output: valueChanges ▷ list-style diff result
1: **procedure** CMPVALUE(a, b)
2: **return** if a = b **then** Equal **else** Different
3: **end procedure**
4: **return** COMPUTELISTCHANGES(srcList, tgtList, cmpValue)

729 A.7 Using graph differences to solve the three-way merge

730 When merging LPGs, a crucial first step is to reliably identify corresponding nodes and edges
731 across different versions of the graph: base, ours, and theirs. We achieve this using unique keys
732 generated by the nodeKey and edgeKey functions. This keying strategy allows us to distinguish
733 between structural changes (deletion or addition) and modifications. If a change causes an element's
734 key to differ from its base version, then that change is regarded as structural (*i.e.*, the old element
735 was deleted and a new element was added). A modification occurs when an element retains the *same*
736 key but its labels or properties are altered. Similarly to our diffing computation approach, we use
737 primitives based on sets and lists to perform all the necessary computations to merge graphs across
738 two branches.

739 The mergeSets algorithm implements a three-way merge for sets, designed to reconcile differences
740 between a local version (S_{ours}) and a remote version (S_{theirs}), both derived from a common ancestor
741 (S_{base}). The algorithm is generic, operating on elements of type T . It leverages a caller-provided
742 comparison function, $f_{\text{compare}} : (T, T) \rightarrow \text{ComparisonResult}$, to determine if two elements are
743 identical, modified, or distinct. A diffing function, f_{diff} (which defaults to computeSetChanges), is
744 used to calculate the changes ($\text{SetChanges}\langle T \rangle$), comprising additions, removals, and modifications)
745 between S_{base} and S_{ours} , and between S_{base} and S_{theirs} . An important aspect of this algorithm is its
746 pluggable conflict resolution mechanism, defined by a MergeStrategy $\langle T \rangle$ (denoted as Σ), which
747 dictates how disagreements are handled. The function returns a MergeResult $\langle T \rangle$ containing the
748 merged set (S_{merged}) and an array of any conflicts encountered.

749 The process begins by computing the deltas: $\Delta_{\text{ours}} = f_{\text{diff}}(S_{\text{base}}, S_{\text{ours}}, f_{\text{compare}})$ and $\Delta_{\text{theirs}} =$
750 $f_{\text{diff}}(S_{\text{base}}, S_{\text{theirs}}, f_{\text{compare}})$. These deltas itemize elements added to, removed from, or modified

751 in S_{ours} and S_{theirs} relative to S_{base} . For efficient lookup, modifications are stored in maps (M_{ours} ,
 752 M_{theirs}), mapping base elements to their modified versions, and deletions are stored in sets (D_{ours} ,
 753 D_{theirs}). The algorithm then iterates through each element $e_{\text{base}} \in S_{\text{base}}$ to determine its fate in S_{merged} .
 754 The following cases arise. Let e'_{ours} be the modified version of e_{base} .

755 First, cases 1-4 consider when the element was modified in S_{ours} (Figure 9, left panel).

756 **Case 1: modified identically in both branches.** If e_{base} was also modified in S_{theirs} to e'_{theirs} (a
 757 “modify-modify” scenario), and if $f_{\text{compare}}(e'_{\text{ours}}, e'_{\text{theirs}})$ yields `Equal`, e'_{ours} (or e'_{theirs}) is added to
 758 S_{merged} .

759 **Case 2: modified differently in both branches.** Otherwise, a `MergeConflict<T>` of type
 760 “modify-modify” is created with e_{base} , e'_{ours} , and e'_{theirs} . The strategy Σ is invoked. If it returns
 761 a resolved value, that value is added to S_{merged} ; otherwise, the conflict is recorded.

762 **Case 3: modified in ours, deleted in theirs.** If e_{base} was deleted in S_{theirs} (a “modify-delete”
 763 scenario, from ours/theirs perspective): A conflict of type “modify-delete” is created (with e_{base} and
 764 e'_{ours}). The strategy Σ is invoked. If resolved, the result is added to S_{merged} ; otherwise, the conflict is
 765 recorded.

766 **Case 4: modified in ours, unchanged in theirs.** If e_{base} was unchanged in S_{theirs} (not modified or
 767 deleted): e'_{ours} is added to S_{merged} .

768 Next, cases 5-7 consider when the element was deleted in S_{ours} , not modified (Figure 9, left panel).

769 **Case 5: deleted in ours, modified in theirs.** If e_{base} was modified in S_{theirs} to e'_{theirs} , a conflict of
 770 type “delete-modify” is created (with e_{base} and e'_{theirs}). The strategy Σ is invoked. If resolved, the
 771 result is added to S_{merged} ; otherwise, the conflict is recorded.

772 **Case 6: deleted in ours, unchanged in theirs.** If e_{base} was unchanged in S_{theirs} , e_{base} is considered
 773 deleted and is not added to S_{merged} .

774 **Case 7: deleted in both.** If e_{base} was also deleted in S_{theirs} , e_{base} is considered deleted and is not
 775 added to S_{merged} .

776 Next, cases 8-10 consider when the element was unchanged in S_{ours} , not modified or deleted (Figure 9,
 777 left panel).

778 **Case 8: unchanged in ours, modified in theirs.** If e_{base} was modified in S_{theirs} to e'_{theirs} , e'_{theirs} is
 779 added to S_{merged} .

780 **Case 9: unchanged in ours, deleted in theirs.** If e_{base} was deleted in S_{theirs} , e_{base} is not added to
 781 S_{merged} .

782 **Case 10: unchanged in both.** If e_{base} was also unchanged in S_{theirs} , e_{base} is added to S_{merged} .

783 After processing all elements from S_{base} , the algorithm handles additions. Let A_{ours} be the set of
 784 elements added in S_{ours} and A_{theirs} be those added in S_{theirs} . The algorithm iterates through each
 785 $a_{\text{ours}} \in A_{\text{ours}}$. It attempts to find a corresponding $a_{\text{theirs}} \in A_{\text{theirs}}$ such that $f_{\text{compare}}(a_{\text{ours}}, a_{\text{theirs}})$ is
 786 either `Equal` or `Modified`. A set $S_{\text{matchedTheirs}}$ tracks elements from A_{theirs} that have already been
 787 matched. If such a match $a_{\text{match}} \in A_{\text{theirs}}$ is found, then a_{match} is added to $S_{\text{matchedTheirs}}$. The following
 788 cases then arise (Figure ??, right panel).

789 **Case 11: added identically to both.** If $f_{\text{compare}}(a_{\text{ours}}, a_{\text{match}})$ was `Equal`, so a_{ours} is added to S_{merged}
 790 (representing a common addition).

791 **Case 12: added differently to both.** If $f_{\text{compare}}(a_{\text{ours}}, a_{\text{match}})$ was **Modified** (an “add-add-different”
 792 scenario), so a conflict of type “add-add-different” is created with a_{ours} and a_{match} . The strategy Σ is
 793 invoked. If resolved, the result is added to S_{merged} ; otherwise, the conflict is recorded.

794 **Case 13: added in ours, not added in theirs.** If no such match is found for a_{ours} : a_{ours} is considered
 795 a unique addition by “ours” and is added to S_{merged} .

796 **Case 14: not added in ours, added in theirs.** Finally, any elements $a_{\text{theirs}} \in A_{\text{theirs}}$ that were not
 797 in $S_{\text{matchedTheirs}}$ are considered unique additions by “theirs” and are added to S_{merged} .

Present in base					Not present in base				
Ours	Theirs				Ours	Theirs			
	+	-	~	/		+	-	~	/
+	×	×	×	×	+	11,12	×	×	13
-	×	7	5	6	-	×	×	×	×
~	×	3	1,2	4	~	×	×	×	×
/	×	9	8	10	/	14	×	×	×

No conflict Potential conflict Impossible case

Figure 9: Cases for merging set primitives. Case 2 results in a “modify-modify-different” conflict, Case 3 in a “modify-delete” conflict, Case 5 in a “delete-modify” conflict, and Case 12 in a “add-add-different” conflict. All green cases results in no conflict, and cases 1 and 11 specifically result in no conflict because the added or modified elements in both branches are equal modulo f_{compare} . All red cases cannot occur (*e.g.*, a branch cannot add an element that already existed in base).

798 The $\text{MergeStrategy}\langle T \rangle \Sigma$ is a function type ($\text{MergeConflict}\langle T \rangle \times \mathcal{F}_{\text{compare}} \rightarrow \text{Resolution}\langle T \rangle$).
 799 A $\text{Resolution}\langle T \rangle$ can either indicate a successful resolution with a resulting value, or that the conflict
 800 remains unresolved. The default strategy, `throwAllConflictsStrategy`, leaves all conflicts
 801 unresolved. An example strategy like `timestampWins` might resolve “modify-modify” conflicts by
 802 selecting the version with a more recent `updatedAt` timestamp.

803 The algorithm concludes by returning S_{merged} and a list of all $\text{MergeConflict}\langle T \rangle$ objects for conflicts
 804 that were not resolved by the chosen strategy Σ .

805 A.7.1 Merging list primitives

806 The `mergeLists` algorithm performs a three-way merge for ordered lists (arrays), reconciling a
 807 local version (L_{ours}) and a remote version (L_{theirs}) against a common ancestor (L_{base}). This func-
 808 tion is generic for elements of type T . It relies on a comparison function, $f_{\text{compare}} : T \times T \rightarrow$
 809 ComparisonResult , to ascertain equality or difference between elements. A diffing function, f_{diff}
 810 (defaulting to `computeListChanges`), is employed to generate sequences of changes, Δ_{ours} and
 811 Δ_{theirs} , representing the transformations from L_{base} to L_{ours} and L_{base} to L_{theirs} , respectively. Each
 812 change object in these sequences specifies an operation (such as `Add`, `Remove`, `Modified`, or `Equal`),
 813 associated elements, and relevant indices from L_{base} . A pluggable $\text{ListMergeStrategy}\langle T \rangle$, de-
 814 noted as Σ_{list} , is used for resolving conflicts. The function outputs a $\text{ListMergeResult}\langle T \rangle$, which
 815 includes the merged list, L_{merged} , and an array of any unresolved $\text{ListMergeConflict}\langle T \rangle$ objects,
 816 C_{list} .

817 Initially, the algorithm computes the delta sequences: $\Delta_{\text{ours}} = f_{\text{diff}}(L_{\text{base}}, L_{\text{ours}}, f_{\text{compare}})$ and $\Delta_{\text{theirs}} =$
 818 $f_{\text{diff}}(L_{\text{base}}, L_{\text{theirs}}, f_{\text{compare}})$. Pointers, p_{ours} and p_{theirs} , are initialized to traverse Δ_{ours} and Δ_{theirs} ,
 819 respectively. The core of the algorithm is a loop that continues as long as there are changes to process
 820 in either delta sequence. Inside the loop, let c_{ours} be the current change from Δ_{ours} and c_{theirs} be from
 821 Δ_{theirs} .

822 First, we consider if one delta sequence is exhausted. If c_{theirs} is undefined (all changes in Δ_{theirs}
 823 processed), the remaining changes $c_{\text{ours}} \in \Delta_{\text{ours}}$ are processed. If c_{ours} indicates an `Add` or `Modified`
 824 operation, its target element is added to L_{merged} . If it’s an `Equal` operation, its source element is

825 added. Remove operations are implicitly handled by not adding the element. p_{ours} is incremented. A
 826 symmetric process occurs if c_{ours} is undefined.

827 Next, we consider if both delta sequences have changes. Let op_{ours} and $\text{op}_{\text{theirs}}$ be the operations for
 828 c_{ours} and c_{theirs} . First, we consider addition operations.

829 **Concurrent additions.** If $\text{op}_{\text{ours}} = \text{Add}$ and $\text{op}_{\text{theirs}} = \text{Add}$, if $f_{\text{compare}}(c_{\text{ours}}.\text{targetElement},$
 830 $c_{\text{theirs}}.\text{targetElement})$ is `Equal`, $c_{\text{ours}}.\text{targetElement}$ is added to L_{merged} . Otherwise, an “add-add-
 831 concurrent” conflict is created with $c_{\text{ours}}.\text{targetElement}$ and $c_{\text{theirs}}.\text{targetElement}$. This conflict is
 832 passed to the `handleListConflict` helper, which uses Σ_{list} for resolution. In either case, both p_{ours}
 833 and p_{theirs} are incremented.

834 **Unilateral additions on ours or theirs.** If $\text{op}_{\text{ours}} = \text{Add}$ (and $\text{op}_{\text{theirs}}$ is not `Add`), $c_{\text{ours}}.\text{targetElement}$
 835 is added to L_{merged} . p_{ours} is incremented. Similarly, if $\text{op}_{\text{theirs}} = \text{Add}$ (and op_{ours} is not `Add`),
 836 $c_{\text{theirs}}.\text{targetElement}$ is added to L_{merged} . p_{theirs} is incremented.

837 Next, we consider non-add operations (referencing L_{base} elements). Both op_{ours} and $\text{op}_{\text{theirs}}$ are either
 838 `Modified`, `Remove`, or `Equal`. Let $\text{idx}_{\text{ours}} = c_{\text{ours}}.\text{sourceIndex}$ and $\text{idx}_{\text{theirs}} = c_{\text{theirs}}.\text{sourceIndex}$.

839 **Non-add operation on the same element.** If $\text{idx}_{\text{ours}} = \text{idx}_{\text{theirs}}$ (both changes refer to the same
 840 element in L_{base}), let $e_{\text{base}} = L_{\text{base}}[\text{idx}_{\text{ours}}]$. We have the following cases:

- 841 • If $\text{op}_{\text{ours}} = \text{Equal}$ and $\text{op}_{\text{theirs}} = \text{Equal}$, $c_{\text{ours}}.\text{sourceElement}$ is added to L_{merged} .
- 842 • If $\text{op}_{\text{ours}} = \text{Modified}$ and $\text{op}_{\text{theirs}} = \text{Modified}$: If $f_{\text{compare}}(c_{\text{ours}}.\text{targetElement},$
 843 $c_{\text{theirs}}.\text{targetElement})$ is `Equal`, $c_{\text{ours}}.\text{targetElement}$ is added. Otherwise, a “modify-modify”
 844 conflict (with e_{base} , idx_{ours} , $c_{\text{ours}}.\text{targetElement}$, $c_{\text{theirs}}.\text{targetElement}$) is handled via Σ_{list} .
- 845 • If $\text{op}_{\text{ours}} = \text{Modified}$ and $\text{op}_{\text{theirs}} = \text{Remove}$, a “modify-delete” conflict (with e_{base} , idx_{ours} ,
 846 $c_{\text{ours}}.\text{targetElement}$) is handled.
- 847 • If $\text{op}_{\text{ours}} = \text{Remove}$ and $\text{op}_{\text{theirs}} = \text{Modified}$, a “delete-modify” conflict (with e_{base} , idx_{ours} ,
 848 $c_{\text{theirs}}.\text{targetElement}$) is handled.
- 849 • If $\text{op}_{\text{ours}} = \text{Remove}$ and $\text{op}_{\text{theirs}} = \text{Remove}$, the element is implicitly deleted.

850 Both p_{ours} and p_{theirs} are incremented.

851 **Non-add operation on different elements.** If $\text{idx}_{\text{ours}} < \text{idx}_{\text{theirs}}$, c_{ours} is processed. If op_{ours} is
 852 `Equal` or `Modified`, the respective element ($c_{\text{ours}}.\text{sourceElement}$ or $c_{\text{ours}}.\text{targetElement}$) is added to
 853 L_{merged} . p_{ours} is incremented. If $\text{idx}_{\text{theirs}} < \text{idx}_{\text{ours}}$, c_{theirs} is processed similarly. p_{theirs} is incremented.

854 The `handleListConflict` helper function takes a `ListMergeConflict<T>`, the strategy Σ_{list} ,
 855 L_{merged} , C_{list} , and f_{compare} . It calls $\Sigma_{\text{list}}(\text{conflict}, f_{\text{compare}})$. If the returned `Resolution<T>` in-
 856 dicates `resolved` as `true` and provides a value (not `undefined`), this value is pushed to L_{merged} .
 857 Otherwise, the conflict is added to C_{list} . The default strategy marks all conflicts as unresolved. The
 858 algorithm concludes by returning an object containing L_{merged} and the list C_{list} of any unresolved
 859 conflicts.

860 A.7.2 Merging graphs

861 A similar approach can be taken to merging graphs through primitives than was taken before in
 862 Appendix A.6.3. The only caveat is that recursive calls of primitives to lower level aspects of the
 863 graph cannot be treated using a one-size-fits-all solution like for diffing. The merging strategy is
 864 in charge of make the sub-call to the appropriate primitive to resolve a merge conflict as it sees fit.
 865 For example, suppose an “add-add-different” conflict arises. One strategy might always pick the
 866 element with the latest timestamp, requiring no sub-calls, while another, that might attempt to union
 867 the properties of both version will have to use call a merge primitive on the two elements.

Algorithm 14 COMPUTE_LIST_CHANGES

Input: source ▷ original list
Input: target ▷ desired list
Input: compare ▷ function returning Equal, Modified or Different
Output: changes ▷ ordered list of edit operations

```

1: m ← length(source), n ← length(target)
2: matrix ← (m+1) × (n+1) table filled with zeros
3:
4: for i = 0 to m do ▷ initialize first row/column
5:   matrix[i][0] ← i
6: end for
7: for j = 0 to n do
8:   matrix[0][j] ← j
9: end for
10:
11: for i = 1 to m do ▷ dynamic-programming fill
12:   for j = 1 to n do
13:     cmp ← compare(source[i-1], target[j-1])
14:     if cmp = Equal then
15:       matrix[i][j] ← matrix[i-1][j-1]
16:     else
17:       removeCost ← matrix[i-1][j] + 1
18:       addCost ← matrix[i][j-1] + 1
19:       replaceCost ← matrix[i-1][j-1] + if cmp = Modified then 1 else 2
20:       matrix[i][j] ← min(removeCost, addCost, replaceCost)
21:     end if
22:   end for
23: end for
24:
25: changes ← empty list ▷ back-trace to build edit script
26: i ← m, j ← n
27: while i > 0 or j > 0 do
28:   if i > 0 ∧ j > 0 ∧ compare(source[i-1], target[j-1]) = Equal then
29:     prepend Equal(i-1, j-1) to changes
30:     i-, j-
31:   else if i > 0 ∧ j > 0 ∧ compare(source[i-1], target[j-1]) = Modified ∧
matrix[i][j] = matrix[i-1][j-1] + 1 then
32:     prepend Modified(i-1, j-1) to changes
33:     i-, j-
34:   else if j > 0 ∧ (i = 0 ∨ matrix[i][j] = matrix[i][j-1] + 1) then
35:     prepend Add(j-1) to changes
36:     j-
37:   else
38:     prepend Remove(i-1) to changes
39:     i-
40:   end if
41: end while
42: return changes

```

Algorithm 15 COMPUTE_SET_CHANGES

Input: sourceSet, targetSet ▷ sets of elements
Input: compare ▷ returns Equal, Modified or Different
Output: changes ▷ object containing add, remove, modify, unchanged

```

1: source ← list of elements in sourceSet
2: target ← list of elements in targetSet
3: toAdd ← empty list
4: toRemove ← empty list
5: toModify ← empty list of pairs (old, new)
6: toDoNothing ← empty list
7: sourceMatched ← list of false of length source
8: targetMatched ← list of false of length target
9:
10: for i = 0 to length(source) - 1 do ▷ match equal or modified pairs
11:   if sourceMatched[i] then continue
12:   end if
13:   for j = 0 to length(target) - 1 do
14:     if targetMatched[j] then continue
15:     end if
16:     cmp ← compare(source[i], target[j])
17:     if cmp = Equal then
18:       mark sourceMatched[i] and targetMatched[j] as true
19:       prepend source[i] to toDoNothing
20:       break inner loop
21:     else if cmp = Modified then
22:       prepend (oldValue: source[i], newValue: target[j]) to toModify
23:       mark sourceMatched[i] and targetMatched[j] as true
24:       break inner loop
25:     end if
26:   end for
27: end for
28:
29: for i = 0 to length(source) - 1 do ▷ collect unmatched elements
30:   if not sourceMatched[i] then prepend source[i] to toRemove
31:   end if
32: end for
33: for j = 0 to length(target) - 1 do
34:   if not targetMatched[j] then prepend target[j] to toAdd
35:   end if
36: end for
37: return SetChanges(toAdd, toRemove, toModify, toDoNothing)

```

Algorithm 16 MERGE_SETS

Input: baseSet, oursSet, theirsSet ▷ three versions of the same set
Input: compare ▷ equality / modified / different
Input: diffFn ▷ produces SetChanges
Input: strategy ▷ conflict-resolution policy
Output: merged ▷ final reconciled set
Output: conflicts ▷ unresolved merge conflicts
1: merged ← empty set ▷ initialise
2: conflicts ← empty list
3:
4: ours ← diffFn(baseSet, oursSet, compare) ▷ compute deltas
5: theirs ← diffFn(baseSet, theirsSet, compare)
6:
7: ourMods ← map base→new from ours.modify ▷ index modifications / deletions
8: theirMods ← map base→new from theirs.modify
9: ourDel ← set of ours.remove
10: theirDel ← set of theirs.remove
11:
12: **for all** item in baseSet **do** ▷ iterate through base elements
13: oursChanged ← item ∈ ourMods
14: theirsChanged ← item ∈ theirMods
15: oursRemoved ← item ∈ ourDel
16: theirsRemoved ← item ∈ theirDel
17: **if** oursChanged **then**
18: ourV ← ourMods[item]
19: **if** theirsChanged **then**
20: theirV ← theirMods[item]
21: **if** compare(ourV, theirV)=Equal **then**
22: add ourV to merged
23: **else**
24: create conflict (modify-modify, item, ourV, theirV)
25: res ← strategy(conflict, compare)
26: **if** res.resolved **then** add res.value to merged
27: **else** push conflict
28: **end if**
29: **end if**
30: **else if** theirsRemoved **then**
31: create conflict (modify-delete, item, ourV)
32: handle with strategy as above
33: **else**
34: add ourV to merged ▷ ours only
35: **end if**
36: **continue**
37: **end if**

```

38:   if oursRemoved then
39:     if theirsChanged then
40:       theirV  $\leftarrow$  theirMods[item]
41:       create conflict (delete-modify, item, theirV)
42:       handle with strategy
43:     end if
44:     continue ▷ deleted in ours
45:   end if
46:   if theirsChanged then
47:     add theirMods[item] to merged
48:   else if not theirsRemoved then
49:     add item to merged
50:   end if
51: end for
52:
53: ourAdds  $\leftarrow$  set of ours.add ▷ stage 4 : handle additions
54: theirAdds  $\leftarrow$  set of theirs.add
55: matchedT  $\leftarrow$  empty set ▷ their additions already paired
56: for all a in ourAdds do
57:   match  $\leftarrow$  nil; cmpRslt  $\leftarrow$  nil
58:   for all b in theirAdds do
59:     if b  $\in$  matchedT then continue
60:   end if
61:   c  $\leftarrow$  compare(a,b)
62:   if c = Equal  $\vee$  c = Modified then
63:     match  $\leftarrow$  b; cmpRslt  $\leftarrow$  c; break
64:   end if
65: end for
66: if match then
67:   add match to matchedT
68:   if cmpRslt = Equal then
69:     add a to merged ▷ identical add
70:   else
71:     create conflict (add-add-modified, a, match)
72:     handle with strategy
73:   end if
74: else
75:   add a to merged ▷ unique add
76: end if
77: end for
78: for all b in theirAdds do
79:   if b  $\notin$  matchedT then add b to merged
80: end if
81: end for
82: return {merged, conflicts}

```

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We have, to the best of our ability, considered each claim and we believe each to be accurate.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: They are described in Section 4.3.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (*e.g.*, independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, *e.g.*, if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: To the best of our ability, we have made the details necessary to reproduce the experimental results of the paper available in our methodological description.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (*e.g.*, in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (*e.g.*, a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (*e.g.*, with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (*e.g.*, to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027

Answer: [No]

Justification: The code for the environment will be publicly released via an open-source GitHub repository at a future date.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (*e.g.*, for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

Question: Does the paper specify all the training and test details (*e.g.*, data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: The DIAMOND compression experiments reported in the paper are deterministic given our choice of KG and hardware, both of which are reported in the paper. Therefore, error bars or confidence intervals are not provided, as there is no run-to-run variability to quantify.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (*e.g.*, Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (*e.g.* negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Yes, this information is provided for every experiment.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (*e.g.*, preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We have not identified potential harms caused by the research process and there is unlikely to be negative societal impact or harmful consequences from the adoption of new developer tooling for graph generation. It is possible that graphs get generated automatically that are incorrect, but this is independent of the method of construction, as it happens even now due to human error and faulty automations. The impact of this is also very limited.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (*e.g.*, if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: We have not identified potential harms caused by the research process, and there are unlikely that negative societal impact or harmful consequences will arise from the adoption of new developer tooling for graph generation. It is possible that graphs can be generated automatically that are factually incorrect; however, this risk is independent of the

method of construction, as it happens even now due to human error and faulty automations. The impact of this is also limited.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (*e.g.*, disinformation, generating fake profiles, surveillance), fairness considerations (*e.g.*, deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (*e.g.*, gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (*e.g.*, pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (*e.g.*, code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.

- The name of the license (*e.g.*, CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (*e.g.*, website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- 1184 • Depending on the country in which research is conducted, IRB approval (or equivalent)
1185 may be required for any human subjects research. If you obtained IRB approval, you
1186 should clearly state this in the paper.
- 1187 • We recognize that the procedures for this may vary significantly between institutions
1188 and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the
1189 guidelines for their institution.
- 1190 • For initial submissions, do not include any information that would break anonymity (if
1191 applicable), such as the institution conducting the review.

1192 16. **Declaration of LLM usage**

1193 Question: Does the paper describe the usage of LLMs if it is an important, original, or
1194 non-standard component of the core methods in this research? Note that if the LLM is used
1195 only for writing, editing, or formatting purposes and does not impact the core methodology,
1196 scientific rigorousness, or originality of the research, declaration is not required.

1197 Answer: [NA]

1198 Justification: [NA]

1199 Guidelines:

- 1200 • The answer NA means that the core method development in this research does not
1201 involve LLMs as any important, original, or non-standard components.
- 1202 • Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>)
1203 for what should or should not be described.