



## From base data to knowledge discovery – *A life cycle approach* – Using multilayer networks



Abhishek Santra <sup>a,\*</sup>, Kanthi Komar <sup>a</sup>, Sanjukta Bhowmick <sup>b</sup>, Sharma Chakravarthy <sup>a</sup>

<sup>a</sup> Information Technology Laboratory (IT Lab), The University of Texas at Arlington, USA

<sup>b</sup> Computer Science and Engineering Department, The University of North Texas, TX, USA

### ARTICLE INFO

#### Keywords:

Knowledge discovery life cycle  
Multilayer networks  
Conceptual modeling  
Objectives to computable expressions  
Decoupling-based analysis  
Drill-down & visualization  
Data and knowledge visualization

### ABSTRACT

Analysis of complex data sets to infer/discover meaningful information/knowledge involves (after data collection and cleaning): (i) Modeling the data — an approach for deriving a suitable representation of data for analysis, (ii) translating analysis objectives into computations on the generated model instance; these computations can be as simple as a query or a complex computation (e.g., community detection over multiple layers), (iii) computation of expressions generated — considering efficiency and scalability, and (iv) drill-down of results to understand them clearly. Beyond this, it is also useful to visualize results for easier understanding. Covid-19 visualization dashboard presented in this paper is an example of this.

This paper covers the above steps of data analysis life cycle using a representation (or model) that is gaining importance. With complex data sets containing multiple entity types and relationships, an appropriate model to represent the data is important. For these data sets, we first establish the advantages of Multilayer Networks (or MLNs) as a data model. Then we use an entity-relationship based approach to convert the data set into MLNs for a precise representation of the data set. After that, we outline how expected analysis objectives can be translated using keyword-mapping to aggregate analysis expressions. Finally, we demonstrate, through a set of example data sets and objectives, how the expressions corresponding to objectives are evaluated using an efficient decoupling-based approach. Results are further drilled down to obtain actionable knowledge from the data set.

Using the widely popular Enhanced Entity Relationship (EER) approach for requirements representation, we demonstrate how to generate EER diagrams for data sets and further generate, algorithmically, MLNs as well as Relational schema for analysis and drill down, respectively. Using communities and centrality for aggregate analysis, we demonstrate the flexibility of the chosen model to support diverse set of objectives. We also show that compared to current analysis approaches, a “decoupling-based” approach using MLNs is more appropriate as it preserves structure as well as semantics of the results and is very efficient. For this computation, we need to derive expressions for each analysis objective using the MLN model. We provide guidelines to translate English queries into analysis expressions based on keywords.

Finally, we use several data sets to establish the effectiveness of modeling using MLNs and their analysis using the *decoupling approach* that has been proposed recently. For coverage, we use different types of MLNs for modeling, and community and centrality computations for analysis. The data sets used are from US commercial airlines, IMDb (a large international movie data set), the familiar DBLP (or bibliography database), and the Covid-19 data set. Our experimental analyses using the identified steps validate modeling, breadth of objectives that can be computed, and overall versatility of the life cycle approach. Correctness of

\* Corresponding author.

E-mail addresses: [abhishek.santra@mavs.uta.edu](mailto:abhishek.santra@mavs.uta.edu) (A. Santra), [kanthisannappa.komar@mavs.uta.edu](mailto:kanthisannappa.komar@mavs.uta.edu) (K. Komar), [sanjukta.bhowmick@unt.edu](mailto:sanjukta.bhowmick@unt.edu) (S. Bhowmick), [sharmac@cse.uta.edu](mailto:sharmac@cse.uta.edu) (S. Chakravarthy).

results is verified, where possible, using independently available ground truth. Furthermore, we demonstrate drill-down that is afforded by this approach (due to structure and semantics preservation) for a better understanding and visualization of results.

## 1. Introduction

Big data analytics is predicated upon our ability to model and analyze disparate and complex data sets. For managing databases using relational database management systems (or RDBMSs) and data warehouses, Relational and Snowflake data models have served well. Newer data models are used in NoSQL systems.

On the mining and knowledge discovery side, data models have not received much attention. However, as the complexity of data increases, data models are needed and useful for modeling the data — both for understanding and analysis. Although graph models have been used for this, there is not much work in a systematic approach to model the data. We see the utility of several graph data models for this purpose and hence interested in providing a general framework.

*We describe the complete life cycle of data analysis along with the flow and steps that need to be automated. In this paper, we focus on a few steps, such as: (i) need for MLN for expressive modeling, (ii) keyword-based algorithmic translation of objectives, (iii) drill-down, visualization, and validation. In addition, we are using an efficient decoupling-based approach and taking some data sets through the entire life cycle. We do not address data pre-processing as there are many commercial ETL (Extract, Transform, and Load) tools available [1–5]. We are also not considering, in this paper, efficiency issues, benchmarks, or other approaches, such as neural networks. This helps us focus on the few steps and how a data set can be taken through the complete life cycle. We are also delimiting our approach, in this paper, to data sets with multiple types of entities that have multiple features and interact through different relationships. Although graphs are used as the basis, the analysis and computations performed in knowledge discovery are quite different from the ones addressed in NoSQL systems, such as Neo4J [6]. As an example, consider a data set about a group of actors and directors (as entities). Each person has some attributes associated with them, such as who they co-act with, which genre they direct, etc. (termed features.) Actors and directors can also be connected (termed relationships) if an actor is directed by a director. Implicit relationships can also be inferred between two entities if they share similar features or transitivity holds. For this data set, finding a strong group of actors and directors, where actors co-act and directors direct similar genres cannot be expressed as a query. It is a computation that requires finding communities that satisfy certain properties.*

### 1.1. Multi-entity and multi-feature data analysis challenges

A critical question is how to model multi-entity and -feature data sets that also involve explicit relationships among the entities as well as how to express & analyze objectives more precisely. Data set characteristics have grown from a single entity and feature and/or relationship type to multiple entities, features and relationships, leading to the following new challenges:

- *Model Expressiveness.* With the presence of multiple types of entities and different relationships among the same- and different-types of entities, it is important to use a model that is expressive, i.e., *preserves clarity of semantics as well as structure of the data*<sup>1</sup> being modeled.
- *Flexible Analysis Alternatives.* Analysis objectives on multi-feature data may be specified on a subset of entities/features. For example, given a data set about *actors*, *movies* in which they have acted, and *directors* who have directed those movies (see IMDb data set used in this paper), the analysis can involve actors and their movie rating, or other actors with whom they work, or any combination thereof. Also, in the IMDb data set, although both actors and directors are people, they are considered separate types of entities, as they perform different roles. The *challenge* is to allow for flexibility of mixing and matching of entities/features for analysis, while avoiding loss of information or redundancy of computations.
- *Analysis Objectives to Expressions and Their Efficient Computation:* Beyond modeling and flexibility of analysis, it is important to have a methodology for translating or mapping analysis objectives into computations on the model and further compute them efficiently. In this paper, we present an algorithmic approach to translate objectives to analysis expressions using computations available on the data model. For efficiency of computation, we use an approach that has been proposed for MLNs [7,8]. Here data structures, use of parallel/distributed approaches, and the ability to scale become important.
- *Drill-down and Visualization of Results.* Finally, understanding the results (knowledge discovered) is extremely important. The model and computation using the model should not mask the structure and semantics of the results. This is where we also think the MLN approach has an edge as compared to earlier approaches to this problem. Visualization is likely to become easier once structure and semantics of results are preserved. We have developed an architecture for visualizing both base data and analyzed results [9]. We show the results from that dashboard in this paper.

<sup>1</sup> By structure and semantics preservation, we imply that additional book keeping information and mapping is not needed before and after analysis. This is needed for drill-down, but is different from expansion of other attributes during drill-down for understanding the results better.

## 1.2. Problem statement

For a given data set  $D$  with  $T$  entity types,  $F$  features, and a set of analysis objectives  $\mathcal{O}$ , **apply** the proposed life cycle framework to: (i) generate an expressive model for the data that preserves structure and semantics, (ii) flexibly analyze combinations of entities and relationships, (iii) generate algorithmically expressions for  $\mathcal{O}$  and compute them efficiently, and (iv) drill-down and visualize knowledge discovered for easier understanding.

Recently, a generalized framework (a model and computation using that model) has been proposed that addresses some of these challenges. Earlier work in multi-feature data analysis (see Section 2) were focused either on a specific application or a specific analysis technique. In this paper, we focus on the knowledge discovery life cycle using the decoupling-approach for the MLN data model.

**Overview of the Paper.** Our *main contribution* in this paper is to develop a complete data analysis (or knowledge discovery) life cycle using a *generalized framework for modeling & computation on multi-feature, multi-entity data sets and show its versatility and effectiveness*. Based on the survey and comparison of the currently used techniques for modeling and analyzing multi-feature data, we have chosen the MLN approach for modeling the data and its analysis using the *decoupling approach*. We have chosen the EER approach for modeling data and generating the MLN model. We also demonstrate that by combining these techniques the chosen framework can handle diverse data sets with multiple features and entity types as well as varied analysis objectives i.e., the entire life cycle.

## 1.3. Contributions and roadmap

- *Expressive Modeling:* In Section 4 we compare the advantages and disadvantages of several approaches for modeling of multi-feature, multi-entity data, and show why multilayer networks (MLNs) are a better alternative with many advantages.
- *Analysis Life Cycle:* In Section 5, we present the steps of the analysis life cycle from data set description and objectives to discovering knowledge corresponding to those objectives. Further, we demonstrate how drill-down of results for visualization can be accomplished using the same framework.
- *Decoupling-based Analysis:* In Section 6 we summarize the *decoupling approach* using which information about each feature/entity type (essentially a layer) is analyzed separately and then composed efficiently to obtain results for the objectives.
- *Algorithmic Translation of Objectives:* In Section 7 we present an approach to translate analysis objectives into computation expressions using the generated MLN model characteristics and available computations.
- *Validation, Drill-down, and Visualization:* In Section 8, we present drill-down and visualization of results. Where possible, we validate our experimental results with *independently* available ground truth. Further, several visualization approaches have been used in our dashboard and is capable of visualizing both base data as well as analysis results. The goal is to facilitate better understanding of the results. We also demonstrate the computational efficiency of the decoupling approach by comparing the time with the standard network aggregation approach.

We present related work in Section 2. We present data sets used along with expected analysis objectives in Section 3. And conclusions in Section 9.

## 2. Related work

We present related work corresponding to the steps of the life cycle in this section.

### 2.1. Modeling

**EER modeling:** Since the Seventies, *EER model* [10] has served as a methodology, for database design, for representing important semantic information about the real-world applications. Relational database modeling has clearly benefited from this body of work and has motivated UML (Unified Modeling Language) for OO (Object-Oriented) design. A good EER diagram based on the data and queries to be supported, is critical and goes a long way for an error-free relational database schema. Numerous tools [11–17] have been developed for creating the EER diagram and algorithmically mapping it into relations for different commercial DBMSs.

However, with the emergence of data sets with relationships among entities and complex application requirements, such as shortest paths, important neighborhoods, dominant nodes (or groups of nodes), etc., [18,19], the relational data model was not the best choice for modeling as well as analyzing them [20]. This led to the evolution of NoSQL data models including the graph data model [21]. In many applications, such as Facebook (friendship relationship), Movies (collaborations relationship) and Twitter (follower-followee relationship), relationships needed to be modeled explicitly using the graph model. This gave access to a wide variety of analyses that were available for these data models. Recently, there has been some work in the area of graph modeling from EER diagrams, but is limited to simple attribute graphs only [10,22–24]. Only recently, an approach [25] has been proposed for using the EER approach for generating MLNs using data sets and analysis objectives. This approach has been adapted in this paper.

**MLN generation:** There are a number of tools developed for creating the EER diagram and these tools also translate the EER diagram developed and generate system-specific Relational schema for Oracle, DB2, etc. This makes the development of a model a bit easier

although the mapping from requirements to an EER diagram is still subjective, relies on experience, and hence not unique. When it comes to mining, this modeling approach is not typically adopted. A decision to use a specific mining algorithm is based on the context and experience as well as desired objectives and if needed the data is transformed into a different format. For example, association rule algorithms use different representations of data than many other mining algorithms. However, when a graph is used as a data model, choice of nodes, edges and their labels (if needed), becomes important and there may be multiple alternative ways of creating them depending on the analysis objectives. Further, creating edges may need a similarity/proximity criteria which needs to be identified or specified.

Our approach for data analysis stems from: (i) the need for analyzing the same data in multiple ways, (ii) a number of aggregate analysis (e.g., community, centrality, substructure, to name a few) alternatives that can be used, and (iii) need for generating expressions for analysis rather than a single computation as is typically done in mining. Hence, this approach and framework, although new to knowledge discovery, is effective as we illustrate it in Section 7. This problem requires further research as we consider our contribution in this paper to be a starting point.

## 2.2. Graph and MLN analysis

MLNs can be classified into Homogeneous (HoMLNs), Heterogeneous (HeMLNs) and Hybrid (HyMLNs) depending upon the characteristics of entities/nodes in each layer and their connectivity to other layers. If the entities and their types are of the same type across all layers, it is a HoMLN where same entities across layers are assumed to be connected implicitly. If the entities and their types are different from layer to layer, explicit edges are used, as needed, to connect entities between two layers and these correspond to HeMLNs. Hybrid MLNs (or HyMLNs) are a combination of these two.

**Analysis of HoMLNs:** *Community detection* algorithms have been extended to HoMLNs for identifying tightly knit groups of nodes based on different feature combinations (review: [26–29].) Algorithms based on matrix factorization [30], cluster expansion [31], Bayesian probabilistic models [32], regression [33] and spectral optimization of the modularity function based on the supra-adjacency representation [34] have been developed. Further, methods have been developed to determine *centrality measures* to identify highly influential entities [35–37]. However, all these approaches *analyze a MLN by reducing it to a simple graph* either by aggregating all (or a subset of) layers that is likely lead to loss of semantics as the entity and feature type information is lost. Other approaches that consider the entire MLN as a whole result in increased complexities due to repeated traversals of individual as well as connected layers.

Recently developed decoupling-based approaches combine partial analysis results from individual layers systematically *in a loss-less manner* to compute communities [7] or centrality hubs [8] for layer combinations. There is no aggregation of layers in this approach. Due to the “divide and conquer” approach of decoupling, this method has been shown to be more efficient as it avoids re-computation of layer communities/hubs, and also provides flexibility of analysis.

**Analysis of HeMLNs:** Majority of HeMLN work (reviews in [38,39]) focuses on developing meta-path based methods for similarity of objects [40], object classification [41], missing link prediction [42], ranking/co-ranking [43], and recommendations [44]. Few existing works propose methods to generate clusters of entities [45,46]. Most of them concentrate mainly on inter-layer edges and not the networks themselves. Moreover, existing approaches (type-independent [47] and projection [48]) do not preserve the structure or types and labels of nodes/edges without extensive mapping and unmapping before and after computation. The type independent approach collapses all layers into a single graph keeping *all* nodes and edges (including inter-layer edges) sans their types and labels. Similarly, the projection-based approach projects nodes of one layer onto another layer and uses layer neighbor and inter-layer edges to collapse two layers into a single graph with one entity type instead of two.

## 2.3. Drill-down and visualization

Drill-down of analysis results is critical especially for complex data which has both structure and semantics. For example, it is not sufficient to know the identities of objects in a community, but also additional details of the objects. Similarly, for a centrality hub. As we are using the MLNs as the data model, we also need to know the objects across layers and their inter-connections, if it is a HeMLN. From a computation/efficiency perspective, minimal information is used for analysis and the drill-down phase is used to expand upon, to the desired extent. Our algorithms, especially the decoupling-based, make it easier to perform drill-down without any additional mappings back and forth for recreating the structure. Our schema generation also separates information needed for drill-down (Relations) and information needed for analysis (MLNs) from the same EER diagram. It also generates needed information for the translation of objectives into expressions.

Visualization is not new either and there exists a wide variety of tools for visualizing both base data, results, and drilled-down information in multiple ways [49–51]. Our focus, in this paper, is to make use of available tools in the best way possible and not propose new ones. For example, we have experimented with a wide variety of tools including, maps, individual graph and community visualization, animation of features in different ways, hovering to highlight data, and real-time data fetching and display, based on user input as menu. Perhaps the main contribution of visualization is our architecture with clearly defined modules for analysis, visual output generation, and user-interaction. In addition to the efficiency aspects of the analysis module, We have also paid attention to efficiency of visualization creation and its access by caching pre-generated results (to avoid re-generation) and use of a hash lookup [9].

Finally, this paper is not about efficiency of MLN analysis as they are discussed in other places [7,8]. Also, we are not comparing graph-based analysis with other approaches, such as neural-network based knowledge discovery to keep the paper focused on the life cycle. It is possible that expression generation and analysis could be replaced with other approaches that fit this framework.

### 3. Data sets and terminology

We present a subset of the data sets we have analyzed using the approach presented in this paper with their description along with analysis objectives. We have chosen data sets from different application domains to illustrate the versatility of our framework. While much larger data sets can be generated, we selected these because reliable ground truth from orthogonal sources were available for some. Due to space constraints, we are discussing a *subset of analysis objectives* driven by coverage. The data sets chosen cover all types of MLNs and illustrate the generality of the framework.

**1. US-based Airlines:** This is a data set of six US-based airlines and their flight connections (*active in February 2018*) among US cities. This information has been collected from corresponding airline sources ([52–57].) Here all the entities are of the same type: *cities*. Two cities are related if there is a *direct flight* between them. The data set is characterized by single entity type (city). The multiple features are due to the presence of multiple airlines. A similar data set for European carriers has been analyzed in [58] in a different way.

*Analysis Objectives.* Our analysis objectives are: (i) For American, Southwest, Spirit, Delta, Frontier and Allegiant Airlines rank the top five cities, that provide the maximum coverage (A1), and (ii) predict which city (taking its population into consideration) could be selected as the next hub(s) for Allegiant Airlines to expand its coverage and avoid competition with other airlines (A2).

**2. Bibliography Database (DBLP):** As most researchers are familiar, the DBLP data set is a publicly available information repository about computer science publications in various conferences and journals. It contains author names, their affiliation, year of publication of papers, conference/journal names, and links to the papers [59]. Clearly, there are multiple entities that can be related based on different types of relationships.

*Analysis Objectives.* Our objectives for this data set are: (i) For each 3-year interval group, find the most actively publishing strong author collaboration groups (A3), (ii) for each conference-based paper group, find the most popular author collaboration group and further for each of them identify their most active 3-year interval group(s) (A4), and (iii) identify author collaboration groups who have published in conferences VLDB and SIGMOD, but have never published in conferences DASFAA and DaWaK (A5).

**3. Internet Movie Database (IMDb):** The IMDb data set is publicly available and has information about movies, TV episodes, actor, directors, ratings and genres of the movies, etc. [60]. Here the entities are of different types, such as actors, directors, movies, etc. The features can also differ since actors can be connected based on co-acting or if they have worked in movies of the same genre. This data set can also be enriched by involving additional information about actors and directors from their social media presence, such as Facebook and Twitter. This is not elaborated due to space constraints in this paper.

*Analysis Objectives.* Some of the analysis objectives for this data set are: (i) Cluster actors who have acted together and have a similar average rating (A6), (ii) find the groups of actors who have never acted together, but are highly rated on an average and have worked in similar genres (A7), (iii) identify genre-based groups of actors and directors having strong collaborations (A8); and (iv) identify, for each movie rating group the genre-based most popular actor and most popular director groups. From this result, find the actor and director groups having strong collaborations (A9).

**4. Covid-19 Data Set:** Covid-19 data reported by New York Times<sup>2</sup> [61] along with census data [62], and data from other trusted sources have been used to compile a data set for the 3141 counties in the US starting from February 2020 that includes features, such as number of daily new cases, number of daily new deaths, latitude–longitude of the county seat, the mean per capita income, population density by land area, total land area, educational qualifications, traffic movement and so on.

*Analysis Objectives.* The inclusion of this data set is for two purposes: (i) to demonstrate how MLN analysis can be applied to data sets such as Covid and (ii) to visualize *analyzed results* instead of base data as is typically done. A number of similarity analysis can be done on different desired features. In this paper, we leverage the MLN aggregate analysis approach to *visualize* how Covid has *spread geographically* in two arbitrarily selected periods. This can be used for understanding the Covid situation corresponding to pre and post a major event, such as Vaccination drive, Spring break, lockdowns, long holiday weekends, etc. Specifically, we want to visualize how counties with *similar percentage increase in covid cases/deaths* has changed across any two user-defined disjoint periods starting from February 2020, and combine that with other demographics information. Specific objective used in the visualization is given in Section 7 as (A10) and visualization results are shown in Section 8.4.

Our selected data sets and analysis objectives are quite varied to illustrate the versatility of the approach. They range from analysis of finding coverage of individual airlines, clusters of co-actors/co-authors to more complicated predictions like the next planned hub of an airline, future potential teaming of actors, high quality actor–director collaborations and active periods of most popular co-authors. *In addition, some of the data sets, based on the objectives, may come out either as homogeneous or heterogeneous or hybrid further depicting the capability and completeness of the modeling approach.*

#### 3.1. Terminology

A *Network* (or graph),  $G$  is an ordered pair  $(V, E)$ , where  $V$  is a set of vertices and  $E$  is a set of edges. An edge  $(v, u)$  is a 2-element subset of the set  $V$ . The two vertices that form an edge are said to be neighbors of each other. Here we consider graphs that are undirected (the vertices in the edge are unordered.)

<sup>2</sup> This data is collected from countries, states and health officials and verified. This includes data from The Covid Tracking Project and Johns Hopkins University. Refer to <https://www.nytimes.com/interactive/2020/us/about-coronavirus-data-maps.html> for how this data is curated.

A *Community* in a graph consists of groups of vertices that are more connected to each other than to other vertices in the graph. Several algorithms have been proposed for community detection in a simple graph. This objective is achieved by optimizing network parameters such as modularity [63] or conductance [64].

*Centrality Metrics* are used for measuring the importance of vertices. They include degree centrality (number of neighbors), closeness centrality (mean distance of the vertex from other vertices), betweenness centrality (fraction of shortest paths passing through the vertex), and eigenvector centrality (the number of important neighbors of the vertex) [65]. Choice of the metric is derived from analysis objectives.

#### 4. Modeling alternatives for multi-entity and multi-feature data

We discuss the different models using which multi-feature data can be represented as a graph and argue why using multilayer networks is a better alternative.

##### 4.1. Graph alternatives

Graphs are widely used for modeling data as rich collections of computations have been developed over the years. Their usage has become even more pronounced and important due to Internet and social networking platforms, such as Twitter, Facebook, LinkedIn. Newer systems such as Neo4J are a result of this trend. We only discuss graph alternatives in this paper due to its appropriateness for these data sets. However, we use the traditional relational model which is better suited for drill-down analysis.

1. **Single Graph or Monoplex:** Here the data set is represented by a single network or graph. The vertices represent the entities and the edges represent the similarity of end points based on a feature or the dyadic relationships between them. At most one edge is assumed between nodes and labels may not be supported.

*Advantages.* This way of modeling data as networks is very popular due to extensive research in this area and availability of several algorithms, such as detecting cliques, communities, centrality metrics, mining subgraphs, motifs, search, etc.

*Disadvantages.* Single graph, however, is not best-suited for representing *multiple entities and features*. Although labels can be used for different entities and features, it is difficult to combine features of different categories (e.g., numerical and categorical), in a meaningful way as *one labeled edge*. The problem compounds when the entities are also of different types. Moreover, as discussed in Section 6, when analyzing a subset of entities and/or associated feature types, separate graphs have to be generated for each such combination and analysis.

2. **Attribute or Knowledge Graphs:** Here additional features of the data sets can be represented by including node types in terms of labels (even multiple labels) and multiple edges, even self-loops, corresponding to relationships for different features.

*Advantages.* Attribute graphs have been successfully used in subgraph mining [66–68], querying [18,19,69] and searching [70] over multi-entity types and multi-feature data sets. They capture more semantic information than simple graphs, and can handle both multiple types of features and entities.

*Disadvantages.* Algorithms for some key analysis functions, such as community and centrality detection are not yet available for general attribute graphs. Hence, these graphs need to be converted to a monoplex for analysis. Although different features can be stored in the graph, for every subset of features, the analysis has to be done separately. If a subset of entities/features are used for analysis, elaborate book keeping is needed before and after the analysis to identify node/edge semantics. In other words, structure is not well-captured in this representation.

3. **Multilayer Networks:** Given the pros and cons of the above options, we propose modeling multi-feature, multi-entity type data sets, as multilayer networks (MLNs). Informally, MLNs are layers of single graphs (or monoplexes). Each layer, typically, captures the semantics of one particular feature along with associated entities. As in a monoplex, the graph vertices represent the entities of the data set and the edges represent similarity between the feature values or the dyadic relationship between the end point vertices. The vertices of two layers can also be connected. To differentiate, we term the edges within a layer as *intra-layer edges* and the edges across the layers as *inter-layer edges*.

There exist, primarily, two distinct types of multilayer networks – *homogeneous* and *heterogeneous*. If each layer of a MLN has the **same set of entities of the same type**, it is termed a Homogeneous MLN (or HoMLN.) For a HoMLN, intra-layer edges are shown explicitly and inter-layer edges are not shown, as they are implicit. As an example, the US-Airlines data set can be modeled using HoMLN. The nodes in each layer are the same (i.e. cities) and edges correspond to the flights between cities. Each layer captures a different airline. Within a layer, two nodes (cities) are connected if there is a direct flight between them for that airline. It is also possible to capture additional information, such as distance, number of flights per week using edge labels in this model. Modeling of this data set using the EER diagram and the generation of MLN for the US-Airlines is discussed in Section 5.

When the **set and types of entities are different across layers**, then the MLN is termed as a heterogeneous multilayer network (HeMLN). IMDb data set is an example which generates a HeMLN. Each layer has a different entity type as its nodes (e.g., actors, directors, and movies). The graph of a layer is defined with respect to the chosen features and entity types. In this case of HeMLNs, the inter-layer links are defined explicitly based on feature semantics that correspond to an edge (e.g., directs-actor, directs-movie, acts-in-a-movie). It is also possible that a data set generates a combination of the above two, termed a hybrid multilayer network (or HyMLN.)

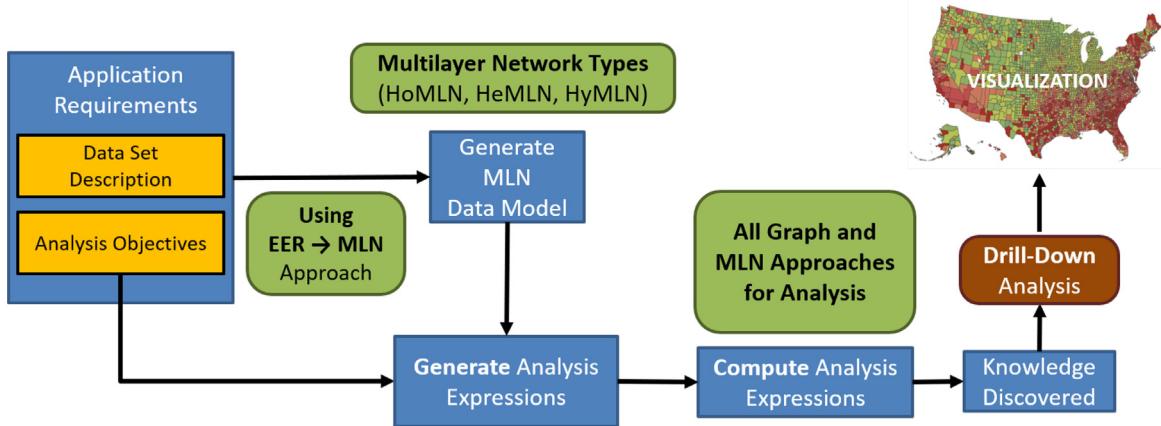


Fig. 1. Data analysis or knowledge discovery life cycle.

Note that whether a data set generates a HoMLN or HeMLN or HyMLN depends on the data set description and objectives being analyzed. Our choice of four different data sets is to showcase the effectiveness of the approach for generating appropriate MLN type needed for the analysis.

**Advantages.** Compared to the other options, multilayer networks are more expressive and elegant for modeling data sets with multiple entities, features, and relationships. In MLNs each *chosen* feature (or combination) is modeled in a separate layer and thus this model can support both heterogeneous and homogeneous data sets. MLNs are also better suited from an information representation (i.e., *structure and semantics*) viewpoint and its visualization. Instead of cluttering all the entities and relationships in a single graph (or layer), they are logically separated and hence are easy to understand. The intra- and inter-layer relationships are also separated semantically. Each incremental change to each feature or relationship, as modeled by addition/deletion of vertices and edges can be easily included without extensive re-modeling of the already created MLN. Unlike most currently used approaches there is no need to convert a MLN representation to another one (simple or attributed) for analysis when the decoupling approach, discussed in Section 6, is used. Finally, a subset of the layers can be analyzed making this model flexible from a selective analysis perspective.

**Challenges:** Having argued for MLNs for modeling, the primary challenge is to develop new algorithms for MLNs for performing analysis. This needs to be done preserving the MLN structure and semantics, as much as possible, during analysis (i.e., without collapsing them as is done by current approaches). The difficulty with the alternative approaches that collapse is the reconstruction of final analysis results to understand them clearly. This requires two mappings one before collapsing and one for reconstruction. This adds additional computation which is separate from drill-down. Semantics preservation is certainly needed for drill-down of results as shown in Section 8. Decoupling-based algorithms used in this paper, by definition, preserve both structure and semantics (removing mapping of entities and features back and forth) making it easier for drill-down analysis. Preservation of structure and semantics also facilitates visualization clarity of results directly. For computing both community and centrality *efficiently*, computation algorithms used in this paper keep the MLN structure intact.

## 5. Data analysis life cycle

Fig. 1 shows the steps of the data analysis life cycle from gathering base data and analysis objectives to discovery of knowledge, its drill-down, and visualization. Given a data set along with its description and the desired analysis objectives, the first step is to choose a model for representing the data. For this paper, we have chosen the multilayer network as the data model, based on the discussion in Section 4. This will be generated from the data set and the analysis objectives using the approach presented in Section 5.1.

Once the model is generated, as shown in the figure, expressions for analysis objectives are generated using the MLN model characteristics (along with available computations for the chosen model) and only the analysis objectives. This is an important step and is currently done using keywords and description of the analysis objectives based on some heuristics. Details of the model, understanding of the objectives as well as the computations available on this model are important for this step. Eventually, this step needs to be automated, as much as possible, using natural language processing of the objectives along with the model characteristics and heuristics. Following this step is the actual computation of the generated expressions to discover knowledge. Ideally, any available algorithms (for graph and/or MLN in our case) can be used for this purpose. Use of model characteristics and algorithms as input, the approach used in this paper to generate the expressions are described in Section 7. The results obtained need to be drilled down further using additional attributes that were not retained as part of the MLN model. This is done using additional attributes specified in the EER diagram that are mapped to a relational model from the same EER diagram<sup>3</sup>. For example, if a director id is

<sup>3</sup> It is not necessary to use a relational model for this as other alternatives are possible. We have used Oracle for drill-down and hence we generate the relational schema.

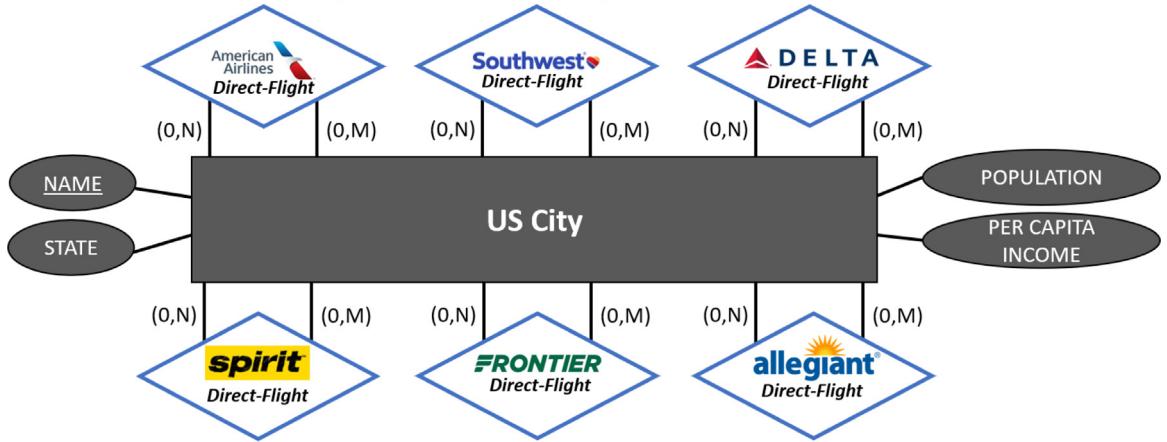


Fig. 2. US-airlines EER diagram.

used in the MLN models, details of the director in terms of location, number of movies directed etc. may be needed for drill-down. Hence, capturing all information needed for analysis and drill-down is extremely important as part of the EER diagram creation step. We will illustrate the results of detailed drill-down in Section 8. Even a better way is to visualize the drilled down results to make it is easier for a non-technical person to understand and interpret the analysis results. This is illustrated in Section 8.4.

### 5.1. EER $\rightarrow$ MLN mapping

Enhanced/Extended Entity Relationship (EER) approach is widely used for modeling data (and functionality desired) from which, typically, a database schema is derived. This has been used for all three types of traditional databases — hierarchical, network, and relational. This modeling technique predates the UML (Unified Modeling Language) widely used today for modeling object-oriented applications as well as algorithmic flow and activity diagrams. The purpose of the EER approach is to convert data requirements gathered during the knowledge acquisition phase to develop a more precise and unambiguous model/representation that can be used for algorithmically generating the MLN schema in our case. We have used the same underlying principles for creating an EER diagram of the data set and analysis objectives (queries are used, instead, in the database context) from which a MLN as well as a relational schema are generated. MLN schema/model is used for deriving expressions for knowledge discovery corresponding to the objectives while the relational schema/model is used for drill-down of the knowledge discovered. Once the EER diagram is generated, generation of MLN is done algorithmically. The details of motivation and an algorithm to translate an EER diagram to a MLN can be found in [25]. Below, a brief overview is provided through examples using the data sets and analysis objectives described in Section 3. For generating EER diagrams, typical heuristics used are: nouns as entities, verbs as relationships, and adjectives as attributes. Other considerations based on objectives (e.g., coverage of objectives in our case) are also taken into account.

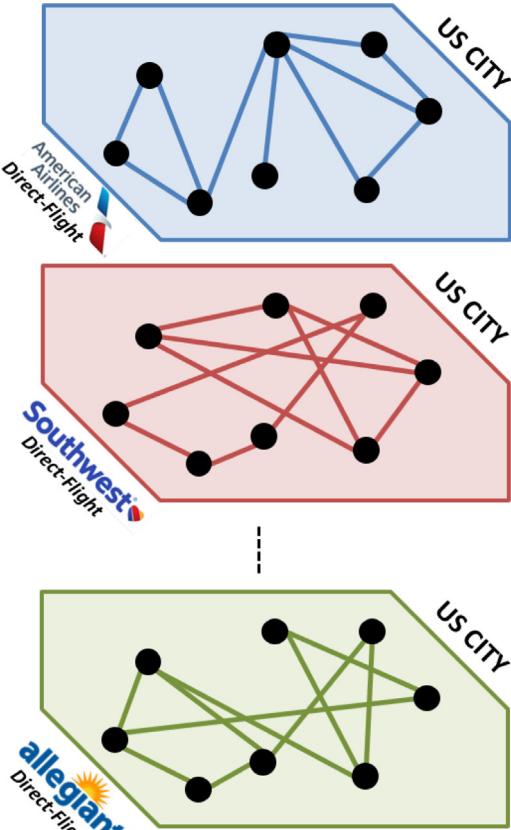
#### 5.1.1. EER modeling of the US-airlines data set

Creating an EER diagram for the US airline data set is relatively straight forward. Although the data set contains airports and their unique codes along with flight information between airports, for simplicity of analysis,<sup>4</sup> we have used cities instead of airports. Other information such as flight number, number of flights per week are also be available. The data set contains *direct* flight information for American, Southwest, Spirit, Delta, Allegiant, and Frontier, that were active in February 2018. The analysis objectives are:

- (A1) For American, Southwest, Spirit, Delta, Frontier and Allegiant Airlines rank the top five cities, that provide the maximum coverage.
- (A2) Predict which city (taking its population into consideration) could be selected as the next hub(s) for Allegiant Airlines to expand its coverage and avoid competition with other airlines.

Since the objectives are to analyze *each airline* for maximum coverage, it is clear that each airline needs to be modeled separately in the EER diagram. For that, *US City* can be modeled as an entity. The direct airline flights are modeled as self-relationship between cities connecting them. A relationship is used for each airline in the EER diagram - *American Direct-Flight*, *Southwest Direct-Flight*, etc. The resulting EER diagram is shown in Fig. 2. Since the analysis objectives are for these airlines, only the cities that are served by all the airlines are considered. If an individual airline is analyzed, all cities served by that airline can be included. The objectives also indicate the need for additional information for each hub for objective (A2) as only the hub information is not sufficient. Additional

<sup>4</sup> There may be multiple airports associated with a city. For example, DAL and DFW are two major airports associated with the Dallas-Fort Worth metroplex.



(a) US Airlines Homogeneous Multilayer Network  
(6 Layers)

| US City                                 |       |            |                   |
|---|-------|------------|-------------------|
| Name                                    | State | Population | Per Capita Income |
| <b>American Airlines Direct-Flight</b>  |       |            |                   |
| City1Name                               |       | City2Name  |                   |
| <b>Southwest Airlines Direct-Flight</b> |       |            |                   |
| City1Name                               |       | City2Name  |                   |
| <b>Delta Airlines Direct-Flight</b>     |       |            |                   |
| City1Name                               |       | City2Name  |                   |
| <b>Spirit Airlines Direct-Flight</b>    |       |            |                   |
| City1Name                               |       | City2Name  |                   |
| <b>Frontier Airlines Direct-Flight</b>  |       |            |                   |
| City1Name                               |       | City2Name  |                   |
| <b>Allegiant Airlines Direct-Flight</b> |       |            |                   |
| City1Name                               |       | City2Name  |                   |

(b) US Airlines Relations  
(obtained as by product used for Drill-Down Analysis)

Fig. 3. MLN and relational schema generated for the US-airlines EER diagram shown in Fig. 2.

information about entities (e.g., population, per capita income etc.) is modeled as attributes of relation US City (see Fig. 3(b)) that will be used for drill-down analysis and ranking of cities as will be shown in Section 8.

When the algorithm given in [25] is applied, the 6 layer homogeneous MLN shown in Fig. 3(a) is generated. In addition, the relations shown in Fig. 3(b) are also generated for drill-down analysis and additional computations, if any.

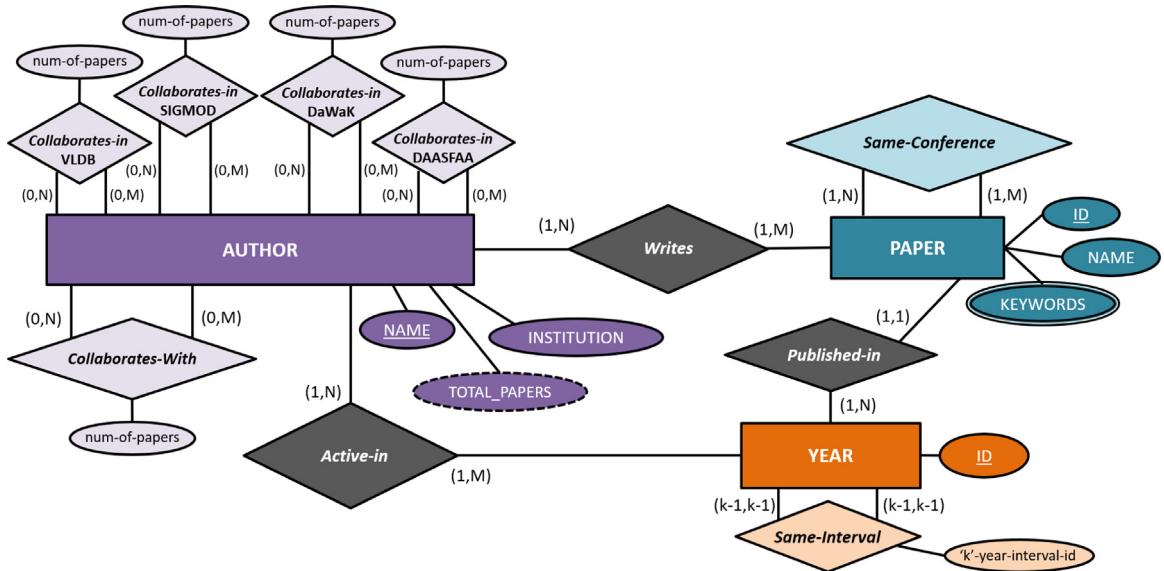
#### 5.1.2. EER modeling of the DBLP data set

Let us consider the DBLP data set along with analysis objectives described in Section 3. An EER diagram for that is constructed as follows. For DBLP, the data set consists of all publications from VLDB, SIGMOD, ICDM, KDD, DaWaK and DASFAA from 2001 to 2018. The analysis objectives are as follows<sup>5</sup>:

- (A3) For each 3-year interval group, find the most actively publishing strong author collaboration groups.
- (A4) For each conference-based paper group, find the most popular author collaboration group and further for each of them identify their most active 3-year interval group(s).
- (A5) Identify author collaboration groups who have published in conferences VLDB and SIGMOD, but have never published in conferences DASFAA and DaWaK.

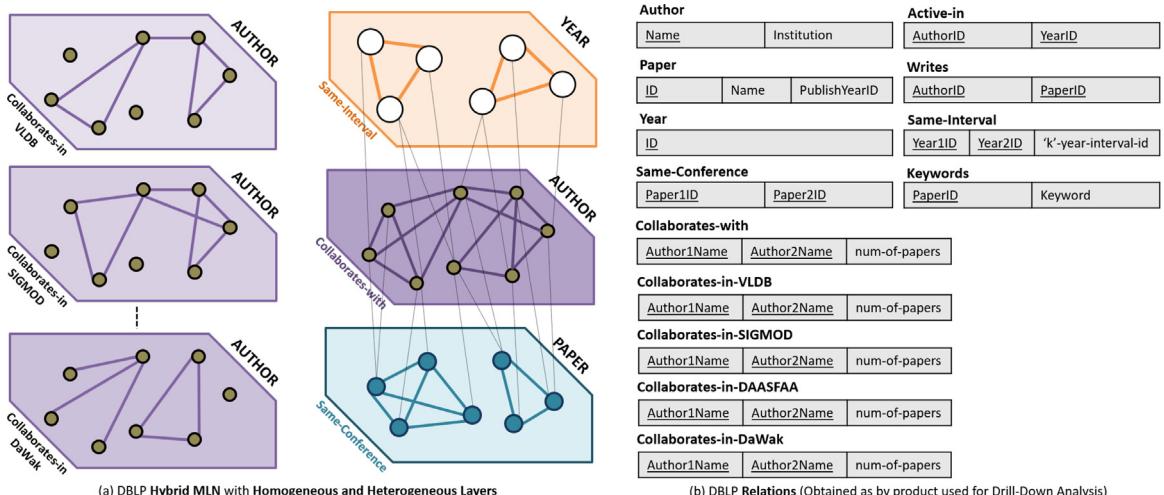
Based on data set description and analysis objectives (A3)–(A5), the EER diagram shown in Fig. 4 has been developed by following the same steps used in the previous example: (i) identifying Entities, (ii) identifying Relationships, including self and binary

<sup>5</sup> Note that some of the objectives indicate specific analysis requirements, such as 3-year periods and others are stated in general terms, such as collaboration groups. These essentially are parameters associated with the objectives that need to be resolved for generating layers prior to analysis. Hence, these are modeled as parameters of the relationships whose values are needed for creating the layer graph prior to actual analysis, but are not needed for generating the MLN schema. This is important as these parameters provide a way to perform new analysis on the same data set without changing the model and expressions generated. Layers that are affected by the parameters need to be re-generated. This adds significant advantage for the breadth of analysis supported by the approach presented in this paper.



**Fig. 4.** DBLP EER diagram for objectives (A3) to (A5).

relationships, and (iii) Cardinality information. *Author*, *Paper* and *Year* come out as three entities, where some entity characteristics, such as *institution* and *keywords* are modeled as attributes of *Author* and *Paper* entities, respectively. *Total\_Papers* attribute of (*Author*) entity is shown as a derived attribute as it can be calculated using *writes* binary relationship. *Collaborates-with*, *Same-Conference* and *Same-Interval* are the three self relationships that relate two *authors* if they have worked together on papers, two *papers* if they are published in same conference, and two instances of *years* if they are in the same disjoint interval, respectively.



**Fig. 5.** Result of EER  $\rightarrow$  MLN algorithm on the DBLP EER diagram shown in [Fig. 4](#).

The *Collaborates-with* and *Same-Interval* self relationships are associated with the attributes *num-of-papers* and '*k*'-*year-interval-id*, respectively, to capture the parameters implicitly specified in the objectives. The value of these relationship attribute parameters become the basis for relating two entities and connecting two nodes in the MLN layer graph.

In addition to the *Collaborates-with* relationship, the author entity is associated to 4 other self relationships - *Collaborates-in-VLDB*, *Collaborates-in-SIGMOD*, *Collaborates-in-DASFAA* and *Collaborates-in-DaWak*, that capture collaboration relationship between two authors for specific conferences as required by objective (A5). Even these have the attribute parameter, *num-of-papers*, that defines the implicit notion of collaboration in the objective. Since *Author* and *Paper* are distinct entities by definition, a binary relationships is needed to capture paper authorship which is a many-to-many relationship. Hence, *Writes*, and similarly *Active-in* and *Published-in* binary relationships capture the information to indicate if an author has written a paper, whether an author was actively publishing in a year and in which year a paper was published, respectively. Finally, the data characteristics and intuitive assumptions have been used to deduce the min-max cardinality.

Once the EER diagram is developed using the data set and analysis objectives, the algorithm in [25] is used to generate the MLN schema shown in Fig. 5(a) which happens to be Hybrid. The expression generation phase will demonstrate how this hybrid MLN will be used appropriately for expression generation of the objectives. The five self-relationships with the Author entity (*Collaborates-With*, *Collaborates-in-VLDB*, *Collaborates-in-SIGMOD*, *Collaborates-in-DASFAA* and *Collaborates-in-DaWaK*) are mapped to five *homogeneous* AUTHOR layers. For generating these layers, two authors are connected if they have collaborated on at least 3 papers, that is num-of-papers parameter value is at least 3. The other two self-relationships (*Same-Interval*, and *Same-Conference*) are mapped to two layers - YEAR-Same-Interval and PAPER-Same-Conference. Based on the requirement of analyzing 3-year periods, two year nodes are connected in the layer graph if they have same value of ‘ $k$ ’-year-interval-id attribute parameter, where  $k = 3$ . For instance, for the years 2001 to 2018 present in the data set, [2001–2003] is interval 1 and [2016–2018] is interval 6. The binary relationships *Writes*, *Active-in* and *Published-in* correspond to inter-layer edges between the corresponding layers representing the entities. Few inter-layer edges are not illustrated in the figure to maintain the visual clarity of the figure. Additionally, using the same EER diagram, a Relational schema is also generated as shown in Fig. 5(b). These are used for drill-down analysis.

For the rest of the data sets used in this paper, we will only show the MLN schema generated for them in Section 7 and not go into the details of the EER modeling due to space constraints. A similar process as illustrated above is used.

## 6. Analysis alternatives for multilayer networks (MLNs)

Since there are not many algorithms available for MLNs and there are a number of widely-used algorithms for single graphs for community, centrality, and substructure discovery, current approaches to MLN analysis take advantage of this by converting a MLN into a single graph.

### 6.1. Single graph approaches

The basic idea is to map the multilayer networks to an equivalent single graph in various ways [71,72]. However, through this process, many of the information in the multilayer graphs can be lost, if appropriate mappings are not created and used. In some cases mappings can become fairly complicated. There are mainly two approaches for converting a MLN into a single layer network. The first, used for homogeneous MLNs, is to *aggregate the edges of the multilayer network*. Specifically, given two vertices  $v$  and  $u$ , the edges between them from each layer are aggregated to form a single aggregated edge. This process is repeated for all the vertex pairs. Some typical aggregation functions are Boolean AND (intersection), OR (union) or linear functions when the edges are weighted. An example, from homogeneous MLNs, would be aggregating routes of different airlines [58] by applying OR (union). This will give rise to multiple edges between nodes or a single edge (if desired from analysis perspective). Mapping has to capture this information clearly and used before and after analysis.

For heterogeneous MLNs, aggregation is performed in many ways. The first is *type independent* [47], that is ignore the different types of the entities (and labels) present, and essentially treat it as a homogeneous MLN with a subset of vertices in each layer. The second method is *projection-based* [48]. Here, if two vertices in a layer are connected to a common vertex in another layer, then an edge is inferred between them. Such “projections” of one layer onto another layer produce inferred edges and then these edges are aggregated.

A third approach, used for HeMLNs, is to *transform the multilayer network into an attribute graph*, where the vertices and edges are labeled based on their types. This graph is analyzed to find specified subgraphs, such as patterns of authors, papers and venues [39] or vulnerabilities in infrastructure networks [73].

*Issues.* Single graph/network approach has the advantage that many analysis algorithms for community and hub detection are available (e.g., Infomap [74], Louvain [75] being prominent ones for community detection). However, the aggregation approaches preserve neither structure nor semantics of MLNs (without explicit mapping and unmapping) as they aggregate layers. Importantly, aggregation approaches are likely to result in some information loss or distortion of properties [72] or hide the effect of different entity types and/or different intra- or inter-layer relationship combinations [76]. In cases, where the multilayer network is converted to an attribute graph, algorithms for aggregate computations (e.g., community, hub) do not exist. Again, they have to be separated into simple graphs (with at most one edge between any pair of nodes) for analyzing. This adds not only additional cost, but the purpose of modeling is defeated to a large extent. Some approaches use the multilayer network as a whole [29,77] and use inter-layer edges, but do not preserve the layer semantics completely. An alternative is to separate desired subgraphs and use single network algorithms which defeats the purpose of modeling as attribute graphs and is inefficient.

### 6.2. The decoupling approach

Network decoupling is a method by which MLNs can be analyzed *without being transformed*. The decoupling approach preserves the structure and semantics of the layers natively in the result and at the same time can take advantage of the existing algorithms. The *network decoupling* approach developed in [7,8] is summarized in Fig. 6(b) and is applied as follows, for a given analysis function  $\Psi$  and composition function,  $\Theta$ :

- (i) Use the analysis function  $\Psi$  to analyze each layer individually.
- (ii) Second, for any two chosen layers, apply a *composition function*  $\Theta$  to compose the partial results from each layer to generate intermediate results.
- (iii) Finally, apply the composition process until the expression is computed.

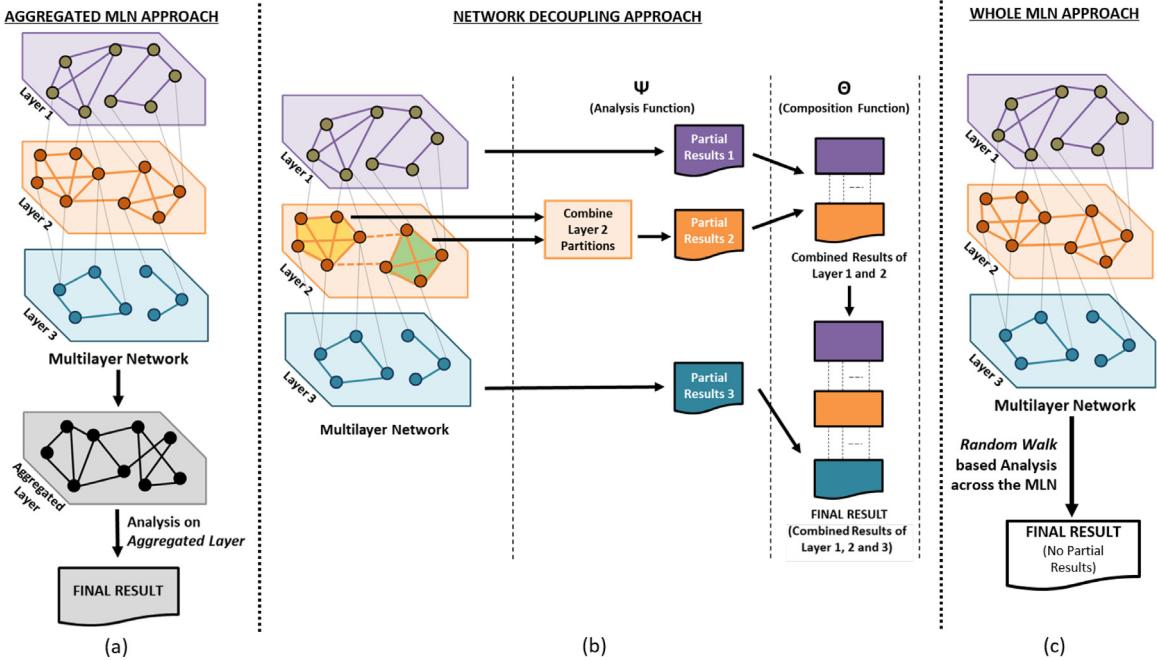


Fig. 6. Alternative approaches for MLN analysis.

This is in contrast to current approaches described earlier. Fig. 6(a) indicates aggregation-based approaches where structure and semantics are lost (sans mapping.) Fig. 6(c) illustrates MLN approaches where only inter-layer edges are used instead of all edges. **Advantages:** The decoupling approach has several advantages over the traditional methods. By using the aggregation approach, information pertaining to the individual layers is lost and it is difficult to measure their relative importance to the system as a whole. In contrast, network decoupling retains the semantic information of each layer and therefore their individual importance and contribution can be measured. This approach also facilitates the mix and match of the features and relationships. In the aggregation approach, each time a subset of features is selected, the analysis has to be recomputed, even when the subsets might have overlaps. This leads to redundant computations. Using the decoupling approach, redundant analysis are avoided, since each layer, corresponding to a particular feature is analyzed separately, and then combined. Finally and importantly, the structure and semantics are preserved in the results explicitly a there is no conversion needed in this approach.

**Challenges.** The decoupling approach can be applied to both HoMLN and HeMLN and hence to HyMLN as well. Moreover, the success of this approach is dependent on correctly matching the analysis objectives using appropriate  $\Psi$  as the aggregate function and  $\Theta$  as the composition function. In Section 7, we show how the network decoupling approach can be applied for our data sets and appropriately determine  $\Psi$  and  $\Theta$  for the diverse analysis objectives (A1) through (A10).

A number of algorithms that use the decoupling approach have been developed for community detection for both HoMLN [7] and HeMLN [78] as well as centrality detection [8] for HoMLN. An algorithm for substructure discovery on MLNs has been developed in [79]. There are also some algorithms that compute substructures [80] and community [29] directly on MLNs without collapsing or aggregating them. In this paper, we use the decoupling based algorithms as they cover the needs of all analysis objectives under consideration.

## 7. Generating analysis expressions from objectives

As indicated in Section 5.1, the modeling of data sets and the generation of a MLN (HoMLN/HeMLN/HyMLN) depends mainly on the relationships identified on the entities in the data set. Typically, self-relationships generate HoMLNs, n-ary (mostly binary as we do not use/support hyper edges in MLNs yet) relationships generate HeMLNs. The EER diagram is also converted to a relational schema for drill-down and additional processing as specified in objectives. The details of creating edges in each layer come from the attributes (deemed as parameters) from the relationships used for MLN generation. For example, for (A4), the 3-year period explicitly provided in the analysis objective is modeled in the EER diagram as an attribute of the relationship same-interval and is used for creating the layer graph. Each three consecutive years form a clique in that layer. If a threshold is needed for creating edges of the layer *Year* to capture similarity, it becomes a parameter of the analysis alternatives. As an example, for (A3) 3 publications together has been used for the parameter *num-of-papers* attribute associated with the collaborates-with relationship that generates the *Author* layer. This is a parameter that can be modified to perform a different set of analysis. The MLN model nor the expressions change, except the graph of the layer generated based on this parameter value.

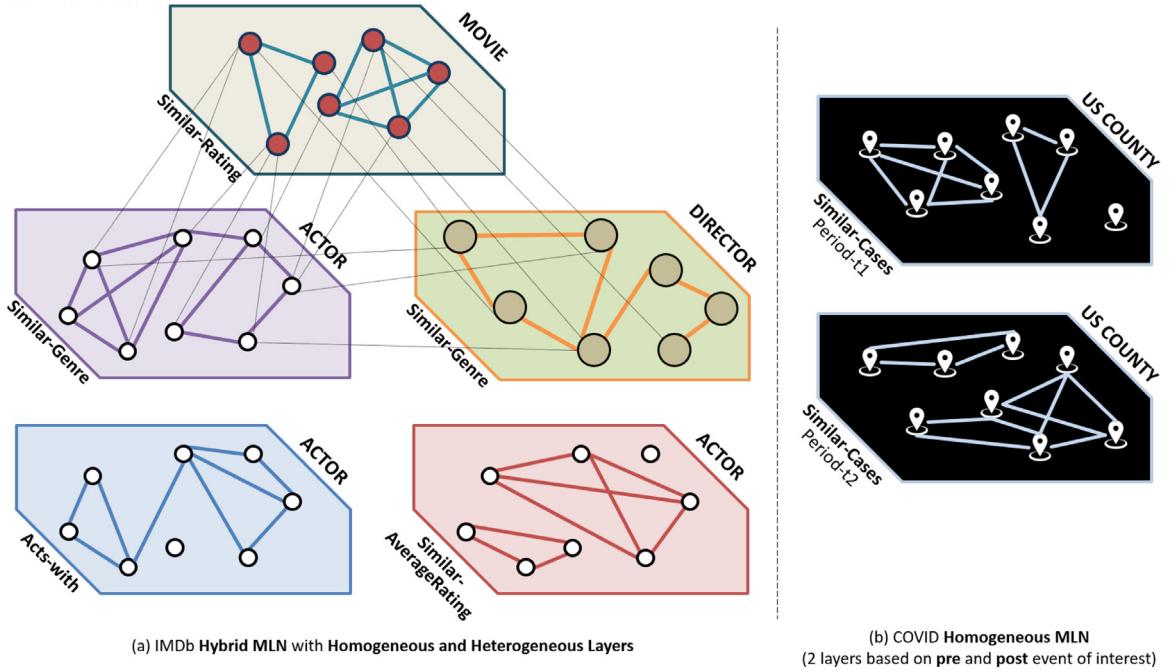


Fig. 7. Multilayer networks generated for remaining analysis objectives.

In addition to these, Table 2 is generated indicating possible  $\Theta$  for each pair of layers for the DBLP data set. Similar tables are generated for each data set. This is dependent on the outcome of modeling. Table 1 is a table that indicates available  $\Psi$  options for each layer. A similar table is generated for each layer produced during modeling of each data set. This depends on algorithms available for computing expressions as shown in Fig. 1. This is independent of the modeling step.

For this paper, we assume community and centrality (both node and closeness) for  $\Psi$ . Other possible  $\Psi$  options (shown in Table 1) can be any graph-based analysis approach like degree centrality, interesting substructure discovery and so on. The complement of a graph (unary NOT) is another  $\Psi$  option that produces a graph with complement set of edges based on the input graph. For composition of homogeneous layers, we assume binary AND and OR binary compositions. For composing heterogeneous layers, we assume Maximum Weighted Matching (MWM) bipartite approach [81] as discussed in [78].

Briefly, algorithms for computing 2-layer communities for HoMLNs use either Boolean AND or OR composition. Algorithms for these are in [7,8]. Boolean NOT operator (complement of a graph) can be used for any layer and further composed with other layers using AND or OR. Multiple layer community computation is done by applying the operators on the result of the previous step. The order of Boolean operator application can be user-specified (or generated as we show in this paper.)

Composing HeMLNs for community detection is challenging since the entities are of different types. As described in [78], each community is considered to be a meta-node. Two meta-nodes in two different layers are connected if there is at least one inter-layer edge between them. The weight of these edges (meta-edges) between the meta-nodes is given by the number of inter-layer edges between them. This construction creates a bipartite graph. These meta nodes (communities) in the bipartite graph are paired using the composition function ( $\Theta$ ) Maximum Weighted Matching (MWM) as proposed by Jack Edmonds [81]. Thus, the *paired meta-nodes* correspond to the heterogeneous MLN communities.

The MLNs derived for the US-Airlines and the DBLP data sets (described in Section 5) are shown in Figs. 3 and 5, respectively. In Figs. 7, we show the MLNs derived after modeling the IMDb (Fig. 7(a)) and Covid (Figure 7 (b)) data sets as EER diagrams and applying the algorithm in [25]. The figures do not include a few inter-layer edges for IMDb MLN to maintain clarity. The EER diagrams and other relations derived for drill-down follow the approach illustrated for the other two data sets.

We use the characteristics of the MLNs generated to create a table for each data set for looking up  $\Theta$  during translation. Another table is created for each layer in each data set to indicate what operations ( $\Psi$ ) are available. The  $\Psi$  (Table 1) and  $\Theta$  (Table 2) lookup tables have been shown for DBLP MLN. In addition, as the current approach is based on extracted keywords and their interpretation (e.g., nouns as layers, verbs as  $\Psi$ , and conjunctions as  $\Theta$ ), another table is used for lookup during translation. This is shown in Table 3 which groups keywords and their possible synonyms in each category with the corresponding choice for computation. The scope of translation depends on the available  $\Psi$  and  $\Theta$ . Objectives that cannot be mapped to available computable operations are indicated as such. Then, we convert each analysis objective for that data set to expressions that can be computed on the MLN model. That is  $\Psi$  and  $\Theta$  are inferred. Also, additional computations on the result (e.g., sorting, ranking, top-k, ...) may need to be performed based on the wording in the objectives in addition to translation.

The challenge here is to automate the generation of analysis expressions from objectives expressed in English. From our experience in analyzing many data sets for diverse objectives, these objectives can be expressed in multiple ways and can be manually

**Table 1**  
 $\Psi$  lookup table for DBLP MLN layers.

|           |  |                   |                      |                        |                  |
|-----------|--|-------------------|----------------------|------------------------|------------------|
| L1        | Community  | Degree centrality | Closeness centrality | Substructure discovery | Complement (NOT) |
| L2        | Community  | Degree centrality | Closeness centrality | Substructure discovery | Complement (NOT) |
| ...       |  |                   |                      |                        |                  |
| L7        | Community  | Degree centrality | Closeness centrality | Substructure discovery | Complement (NOT) |
| Layer IDs | L1: AUTHOR-Collaborates-With, L2: AUTHOR-Collaborates-in-VLDB, L3: AUTHOR-Collaborates-in-SIGMOD, L4: AUTHOR-Collaborates-in-DASFAA, L5: AUTHOR-Collaborates-in-DaWaK, L6: PAPER-Same-Conference, L7: YEAR-Same-Interval |                   |                      |                        |                  |

**Table 2**  
 $\Theta$  lookup table for DBLP MLN layer pairs.

|           |  |         |         |         |     |     |    |
|-----------|--|---------|---------|---------|-----|-----|----|
| L1        |  |         |         |         |     |     |    |
| L2        | AND, OR  |         |         |         |     |     |    |
| L3        | AND, OR  | AND, OR |         |         |     |     |    |
| L4        | AND, OR  | AND, OR | AND, OR |         |     |     |    |
| L5        | AND, OR  | AND, OR | AND, OR | AND, OR |     |     |    |
| L6        | MWM  | MWM     | MWM     | MWM     | MWM |     |    |
| L7        | MWM  | MWM     | MWM     | MWM     | MWM | MWM |    |
|           | L1   | L2      | L3      | L4      | L5  | L6  | L7 |
| Layer IDs | L1: AUTHOR-Collaborates-With, L2: AUTHOR-Collaborates-in-VLDB, L3: AUTHOR-Collaborates-in-SIGMOD, L4: AUTHOR-Collaborates-in-DASFAA, L5: AUTHOR-Collaborates-in-DaWaK, L6: PAPER-Same-Conference, L7: YEAR-Same-Interval |         |         |         |     |     |    |

**Table 3**  
Keyword-based lookup for operators.

| Keywords/Phrases                   | Mapped to $\Psi/\Theta$         |
|------------------------------------|---------------------------------|
| group, cluster, strong/dense group | $\Psi$ = Community              |
| coverage                           | $\Psi$ = Closeness centrality   |
| direct neighbors, hubs             | $\Psi$ = Degree centrality      |
| frequent/interesting patterns      | $\Psi$ = Substructure Discovery |
| never, not                         | $\Psi$ = Complement (NOT)       |
| and, but, yet                      | $\Theta$ = AND for HoMLN layers |
| or, either                         | $\Theta$ = OR for HoMLN layers  |
| and, but, yet, for each, for every | $\Theta$ = MWM for HeMLN layers |

translated into several alternative expressions. This is due to multiple interpretations which lead to multiple expressions for the same objective. Figuring out the order of computations of expressions is another challenge. As explained earlier, we use a keyword-based heuristics approach in this paper as there is no general purpose NLP-based approach that we are aware of. Below, we explain our approach and provide intuitive explanations of how they are generated.

Similar to the heuristics used for EER modeling, nouns, verbs, and conjunctions are identified from objectives. We have highlighted parts of each objective used for expression generation. Phrases in the objective are underlined, *italicized*, and shown in **bold** below to indicate their use, respectively, for layer selection (using nouns), analysis function ( $\Psi$ ) determination (using verb forms), and composition function ( $\Theta$ ) identification (using conjunctions) for the generation of the expression. These can be isolated by an NLP keyword/phrase analysis and used for looking up  $\Psi$  and  $\Theta$  using the three tables indicated earlier. If a mapping is not possible (or keyword not properly recognized), lookup of the tables will fail. Additional computations, as needed, are also inferred from the keywords/phrases in the objective, in the form of FILTER operation. For example, predict is translated to ranking based on sort and choosing the top-k values. Listing of top k entities translates to sorting and retaining top k entities. Here, for the list of the objectives, the FILTER operation is applied (as per requirement) at the end of the analysis expression or earlier as shown with each expression.

### 7.1. Expressions generated for US-airlines objectives

(A1) For American, Southwest, Spirit, Delta, Frontier and Allegiant Airlines rank the top five cities, that provide the maximum *coverage*.

Coverage in (A1) corresponds to cities from which *one can cover most number of cities using least number of flights*. Hence, Coverage can be used to translate the objective intent to *closeness centrality* as  $\Psi$  as shown in Table 3. For airlines, since low closeness centrality value means low average distance (number of flights) to cover cities, that will provide cities which can be ranked further to fetch the top-5. The expression derived is shown below.

Expression:  $\Psi$ (Each layer), FILTER = top-5 using closeness value;

$\Psi$  = Closeness Centrality;  $\Theta$  = N/A

(A2) Predict which city (taking its population into consideration) could be selected as the next hub(s) for Allegiant Airlines to expand its *coverage* and avoid competition with other airlines.

Analysis (A2) is more complicated and perhaps beyond the current approach. However, it is clear from the objective and the mapping table that  $\Psi$  is closeness centrality. However, prediction is for cities that are not currently Allegiant hubs which requires eliminating the hubs generated from the  $\Psi$  computation. Avoiding competition indicates non-overlap with cities that have high closeness value for other airlines resulting in the expression shown below.  $\Psi$  and  $\Theta$  can be looked up. Sorting on population at the end is clear. Using this simple approach, it is difficult to derive the differences. This is a good example of the challenge we mentioned earlier that is difficult for the keyword-based heuristics translation!

Note that this objective can be specified and computed for any airline. In addition, other city attributes (than population), such as mean/median income or combinations can be used. These do not affect either the model or the expression generation. Ordering the resulting cities based on the population (or any attribute) and choosing the top one is the last step.

Expression:  $(\Psi(\text{USCITY-Allegiant-Direct-Flight}) -$

$\{\Psi(\text{USCITY-Allegiant-Direct-Flight}) \Theta \Psi(\text{Each of Remaining Airlines Separately})\}) - \{\text{Active Allegiant Hubs}\}; \text{FILTER} = \text{sort-by-population};$

$\Psi = \text{Closeness Centrality}; \Theta = \text{AND}$

## 7.2. Expressions generated for DBLP objectives

(A3) **For each 3-year interval group**, find the most actively publishing **strong author collaboration groups**.

Once the layers are identified for (A3) as AUTHOR-Collaborates-With and YEAR-Same-Interval, it can be looked up from Table 2 as a HeMLN operation. Based on the keyword mapping Table 3, the required analysis is community detection (phrases = *group*, **strong group**), which have to be combined using MWM composition due to the identification of heterogeneous layers (phrase = **for each** serves as conjunction), resulting in the expression shown below.

Expression:  $\Psi(\text{YEAR-Same-Interval}) \Theta \Psi(\text{AUTHOR-Collaborates-With});$

$\Psi = \text{Community}; \Theta = \text{MWM}$

(A4) **For each conference-based paper group**, find the most popular **author collaboration group** and further **for each** of them identify their most active 3-year interval **group(s)**.

For (A4), 3 layers are identified from the objective. The italicized phrases indicate community detection within layers (phrases = *group*). Further, MWM is used as the  $\Theta$  throughout to generate the final HeMLN communities (phrases = **for each**, **and**), based on mapping and lookup Table 3 and Table 2. Order is implied by the keyword further.

Expression:  $(\Psi(\text{PAPER-Same-Conference}) \Theta_1$

$\Psi(\text{AUTHOR-Collaborates-With})) \Theta_2 \Psi(\text{YEAR-Same-Interval});$

$\Psi = \text{Community}; \Theta_i = \text{MWM}$

(A5) Identify **author collaboration groups** who have published in conferences **VLDB and SIGMOD, but have never published in conferences DASFAA and DaWaK**.

Here 4 layers are identified. The objective clearly indicates the need for the NOT operator (phrase=*never* for the layers Author-Collaborates-in-DASFAA and Author-Collaborates-in-DaWaK). The community indication is clear (phrase = *group*). Moreover, all are homogeneous layers which support boolean composition, as looked up from Table 2. First, we compute 2-layer community using AND (between Author-Collaborates-in-VLDB and Author-Collaborates-in-SIGMOD), which is again AND composed with the NOT communities of the other two conferences (phrases = **and**, **but**). Thus, generating the final expression.

Expression:  $(\Psi(\text{AUTHOR-Collaborates-in-VLDB}) \Theta_1$

$\Psi(\text{AUTHOR-Collaborates-in-SIGMOD})) \Theta_3$

$(\Psi(\text{NOT}(\text{AUTHOR-Collaborates-in-DASFAA})) \Theta_2$

$\Psi(\text{NOT}(\text{AUTHOR-Collaborates-in-DaWaK})))$ ;

$\Psi = \text{Community}; \Theta_i = \text{AND}$

Note that the expression generated for this can be re-written to improve computation efficiency. We do not further discuss that in this paper except to indicate that expressions generated can be further optimized using any means available without affecting correctness.

## 7.3. Expressions generated for IMDb objectives

(A6) **Cluster actors who have acted together and have a similar average rating.**

Expression generation for (A6) is quite straightforward with the identification of ACTOR-Acts-with and ACTOR-Similar-AverageRating layers. The *cluster* keyword indicates community computation on the individual layers that are combined further using the AND composition (conjunction = **and**.)

Expression:  $\Psi(\text{ACTOR-Acts-with}) \Theta \Psi(\text{ACTOR-Similar-AverageRating});$

$\Psi = \text{Community}; \Theta = \text{AND}$

(A7) Find the groups of actors who have never acted together, **but** are highly rated on an average **and** have worked in similar genres. For (A7), the layers identified are ACTOR-Similar-AverageRating, ACTOR-Similar-Genre, and ACTOR-Acts-with. These are homogeneous layers as well and community detection is required (phrase = *group*). Based on the objective, NOT is applied to layer ACTOR-Acts-With (phrase = *never*) before composing it with layer Actor-Similar-AverageRating, which is then finally composed with layer Actor-Similar-Genre using AND composition (phrases = **but**, **and**) as  $\Theta$ s throughout leading to the generated expression. sort-on-AverageRating needs to be inferred from the keyword highly in the objective to output top k results. Order of composition is inferred from the objective as given.

Expression:  $(\Psi(\text{NOT}(\text{ACTOR-Acts-with})) \Theta_1 \Psi(\text{ACTOR-Similar-AverageRating})) \Theta_2 \Psi(\text{ACTOR-Similar-Genre}), \text{FILTER} = \text{sort-on-AverageRating}; \Psi = \text{Community}; \Theta_i = \text{AND}$

(A8) Identify genre-based groups of actors **and** directors having strong collaborations.

This objective is quite straightforward and is similar to others done earlier for 2 layer HeMLN community expression generation.

Expression:  $\Psi(\text{ACTOR-Similar-Genre}) \Theta \Psi(\text{DIRECTOR-Similar-Genre}); \Psi = \text{Community}; \Theta = \text{MWM}$

(A9) Identify, **for each** movie rating *group* the genre-based most popular actor **and** most popular director groups. From this result, find the actor **and** director groups having strong collaborations.

Here MOVIE-Similar-Rating, ACTOR-Similar-Genre and DIRECTOR-Similar-Genre are the 3 identified layers which needs to be processed for detecting the HeMLN communities ( $\Psi$  phrases = *group, range*;  $\Theta$  phrases = **for each, and**). Based on the objectives, the layer MOVIE-Similar-Rating becomes a common layer that is composed with ACTOR-Similar-Genre and DIRECTOR-Similar-Genre layers. This ordering is derived from the objective and translated to the generated expression by the subscripts of  $\Theta$  -  $\Theta_{M,A}$  for MOVIE-Similar-Rating and ACTOR-Similar-Genre composition, and  $\Theta_{M,D}$  for MOVIE-Similar-Rating and DIRECTOR-Similar-Genre composition. Finally, the DIRECTOR-Similar-Genre and ACTOR-Similar-Genre are composed. MWM (Maximum Weighted Matching) is used for the  $\Theta$  operation, thus completing the generated expression.

Expression:  $((\Psi(\text{MOVIE-Similar-Rating}) \Theta_{M,A} \Psi(\text{ACTOR-Similar-Genre})) \Theta_{M,D} \Psi(\text{DIRECTOR-Similar-Genre})) \Theta_{D,A} \Psi(\text{ACTOR-Similar-Genre}); \Psi = \text{Community}; \Theta = \text{MWM}$

#### 7.4. Expressions generated for covid objectives

(A10) Visualize how the geographical regions corresponding to the clusters of US counties with rise (or decline) in daily confirmed cases shift in the -

- consecutive 7-day periods pre and post 2020 spring break and,
- month-apart 3-day periods pre and post vaccination drive?

For the spring break and vaccination based identified layers in (A10) (i) and (A10) (ii), respectively, community detection (phrase = *group*) needs to be performed, as captured by the final expression.

Expression:  $\Psi(\text{Each Layer}); \Psi = \text{Community}; \Theta = \text{N/A}$

To summarize, Table 4 gives the mapping of each analysis question (A1) to (A10) to their actual computation specification (in left to right order), analysis function ( $\Psi$ ) and composition function ( $\Theta$ ).

## 8. Knowledge discovery, drill-down, and visualization

We computed the results for each analysis objective using the expressions derived and shown in Table 4 and compare it, where possible, with independently available ground truth. This helps validate both the modeling and analysis aspects of the life cycle approach proposed. We will not focus on the efficiency of the decoupling approach as it has been established elsewhere [7,8]. Structure- and semantics-preserving aspects of the decoupling approach allows us to drill-down and show details of experimental results.

### 8.1. US-airline analysis results

**MLN Details:** Based on the direct flights that were active in February 2018, US-Airline MLN layers are generated whose statistics are shown in Table 5.

(A1) **For American, Southwest, Spirit, Delta, Frontier and Allegiant Airlines rank the top five cities, that provide the maximum coverage.** Based on the expression derived, we computed the closeness centrality for each layer. We ranked the cities in each layer according to their closeness centrality value. Top 5 hubs (*higher rank, fewer flights required for coverage, more central city*) were identified for each airline. *For all 6 airlines, the ground truth obtained from [82] matched our results.* In Table 6 we have listed top 5 hubs for 4 airlines. As a drill-down byproduct, it is interesting to see common hubs (highlighted) between airlines which is also verified by the ground truth.

**Table 4**

Generated MLN expression for each analysis objective.

| Analysis   | Mapping   | $\Psi$               | $\Theta_i$ |
|--|---|----------------------|------------|
|  | Computation expression  |                      |            |
| <b>US-Airline HoMLN layers:</b> USCITY-American-DirectFlight, USCITY-Southwest-DirectFlight, USCITY-Delta-DirectFlight, USCITY-Spirit-DirectFlight, USCITY-Frontier-DirectFlight, USCITY-Allegiant-DirectFlight      |   |                      |            |
| (A1)   | Apply $\Psi$ on each layer, FILTER = top-5 using closeness value  | Closeness centrality | N/A        |
| (A2)   | $(\Psi(\text{USCITY-Allegiant-Direct-Flight}) - \{\Psi(\text{USCITY-Allegiant-Direct-Flight}) \Theta \Psi(\text{Each of Remaining Airlines Separately})\}) - \{\text{Active Allegiant Hubs}\}; \text{FILTER} = \text{sort-by-population}$ | Closeness centrality | AND        |
| <b>DBLP MLN layers:</b> AUTHOR-Collaborates-With, AUTHOR-Collaborates-in-VLDB, AUTHOR-Collaborates-in-SIGMOD, AUTHOR-Collaborates-in-DASFAA, AUTHOR-Collaborates-in-DaWaK, PAPER-Same-Conference, YEAR-Same-Interval |   |                      |            |
| (A3)   | YEAR-Same-Interval $\Theta$ AUTHOR-Collaborates-With  | Community            | MWM        |
| (A4)   | PAPER-Same-Conference $\Theta_1$ AUTHOR-Collaborates-With $\Theta_2$ YEAR-Same-Interval   | Community            | MWM        |
| (A5)   | (AUTHOR-Collaborates-in-VLDB $\Theta_1$ AUTHOR-Collaborates-in-SIGMOD) $\Theta_3$ (NOT(AUTHOR-Collaborates-in-DASFAA) $\Theta_2$ NOT(AUTHOR-Collaborates-in-DaWaK))   | Community            | AND        |
| <b>IMDb MLN layers:</b> ACTOR-Acts-with, ACTOR-Similar-Genre, ACTOR-Similar-AverageRating, DIRECTOR-Similar-Genre, MOVIE-Similar-Rating  |   |                      |            |
| (A6)   | ACTOR-Acts-with $\Theta$ ACTOR-Similar-AverageRating  | Community            | AND        |
| (A7)   | NOT(ACTOR-Acts-with) $\Theta_1$ ACTOR-Similar-AverageRating $\Theta_2$ ACTOR-Similar-Genre, Filter = sort-AverageRating   | Community            | AND        |
| (A8)   | ACTOR-Similar-Genre $\Theta$ DIRECTOR-Similar-Genre   | Community            | MWM        |
| (A9)   | MOVIE-Similar-Rating $\Theta_{1,M}$ ACTOR-Similar-Genre $\Theta_{2,M}$ DIRECTOR-Similar-Genre $\Theta_{3,D}$ ACTOR-Similar-Genre  | Community            | MWM        |
| <b>COVID HoMLN layers:</b> USCOUNTY-Similar-Cases-Period-t1, USCOUNTY-Similar-Cases-Period-t2  |   |                      |            |
| (A10)  | Apply $\Psi$ on each layer  | Community            | N/A        |

**Table 5**

US-Airlines HoMLN layer statistics.

| Layer                   | Number of nodes | Number of edges |
|-------------------------|-----------------|-----------------|
| American-Direct-Flight  | 270             | 746             |
| Southwest-Direct-Flight | 270             | 717             |
| Delta-Direct-Flight     | 270             | 688             |
| Frontier-Direct-Flight  | 270             | 346             |
| Spirit-Direct-Flight    | 270             | 189             |
| Allegiant-Direct-Flight | 270             | 379             |

**Table 6**

(A1): Cities with maximum US travel coverage.

| (a)            | (b)            | (c)            | (d)            |
|----------------|----------------|----------------|----------------|
| American       | Southwest      | Frontier       | Spirit         |
| <i>Dallas</i>  | <i>Chicago</i> | <i>Denver</i>  | Ft. Lauderdale |
| <i>Chicago</i> | <i>Denver</i>  | <i>Orlando</i> | Las Vegas      |
| Charlotte      | Baltimore      | Austin         | Orlando        |
| Philadelphia   | <i>Dallas</i>  | Las Vegas      | Detroit        |
| Phoenix        | Las Vegas      | Philadelphia   | <i>Chicago</i> |

**Table 7**

(A2): Next allegiant hub by rank.

| City                | Pop. (2019) |
|---------------------|-------------|
| <b>Grand Rapids</b> | 198,401     |
| Elko                | 20,304      |
| Montrose            | 19,238      |

(A2) Predict which city (taking its population into consideration) could be selected as the next hub(s) for Allegiant Airlines to expand its coverage and avoid competition with other airlines.

On computing the expression, from the high closeness centrality cities of the Allegiant airline, we eliminated all those cities that are also high closeness in each of the competitor airlines. From this set of cities, we ranked those cities that are not currently Allegiant hubs, based on their *population*. This information is available from the *City* relation (Fig. 3(b)) that was obtained as a by-product of the EER → MLN process. Table 7 shows the resulting set of cities where Allegiant Airline can potentially expand its operations. We validated our result by the fact that **Grand Rapids has been converted to a hub by Allegiant as of July 6, 2019** [83].

## 8.2. DBLP analysis results

**MLN Details:** Table 8 provides the statistics of the different DBLP MLN layers. Publications from 2001 to 2018 were considered for generating the first three layers - AUTHOR-Collaborates-With, PAPER-Same-Conference and YEAR-Same-Interval. For the generation of remaining 4 layers, the respective conference publications were selected from 2003–2007. Louvain algorithm was used to generate the layer-wise communities [75].

Table 8  
DBLP MLN statistics.

| Layer                         | Number of nodes | Number of edges |
|-------------------------------|-----------------|-----------------|
| AUTHOR-Collaborates-With      | 16,918          | 2483            |
| PAPER-Same-Conference         | 10,326          | 12,044,080      |
| YEAR-Same-Interval            | 18              | 18              |
| AUTHOR-Collaborates-in-VLDB   | 5116            | 3912            |
| AUTHOR-Collaborates-in-SIGMOD | 5116            | 3303            |
| AUTHOR-Collaborates-in-DASFAA | 5116            | 1519            |
| AUTHOR-Collaborates-in-DaWaK  | 5116            | 679             |

**(A3) For each 3-year interval group, find the most actively publishing strong author collaboration groups.** As per the expression, on applying MWM on the community bipartite graph created with all Paper and Author communities, we obtained 6 community pairings for the *co-author groups who have published most number of times in each 3-year period*, (shown in Fig. 8 with list of few prominent authors.) This visualization was accomplished by the drill-down of raw results with the help of relations obtained earlier in Fig. 5(b). The author community ids have been shown that are generated by the community detection algorithm in the  $\Psi$  phase of the decoupling approach. Quality of these results are validated by the following facts,

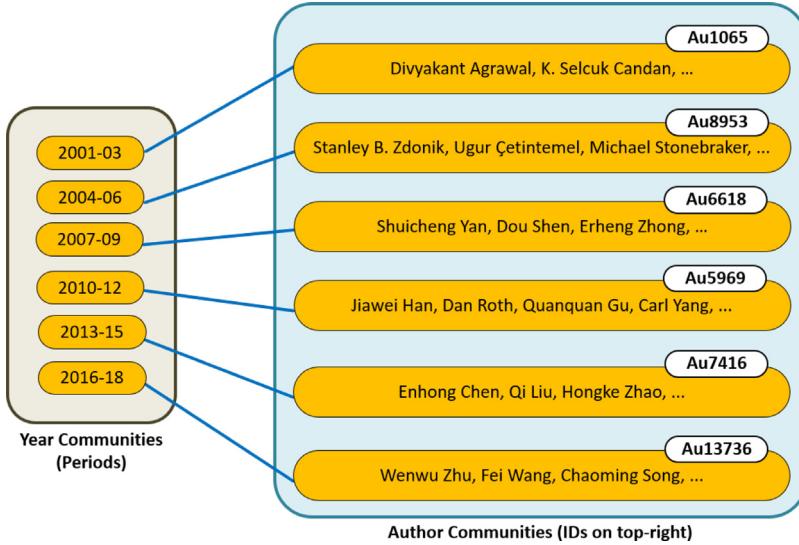


Fig. 8. (A3): Most active co-author groups.

- For the period from 2004 to 2006, the most active group included prominent researchers like **Stanley B. Zdonik (ACM Fellow from 2006)**, **Ugur Çetintemel** and **Michael Stonebraker (Turing Award Recipient, 2014)** who co-authored papers on stream processing engines in this period garnering over **2000 citations**.
- From 2010 to 2012, **Jiawei Han**, one of the most cited computer science researchers (194 752 citations), was involved in more than **200 publications**. In the beginning of this period, he had even won the McDowell Award, the highest technical award made by IEEE.

Such insightful results can be further drilled-down to find active periods of co-author subgroups, research labs and universities. **(A4) For each conference-based paper group, find the most popular author collaboration group and further for each of them identify their most active 3-year interval group(s).** In order to generate the required communities, based on the expression in Table 4, the *most popular author groups for each conference* are obtained by MWM (first composition). The matched 6 author communities are carried forward to find the *disjoint year periods* in which they were *most active* (second composition.) 6 communities are obtained (path shown by **bold blue lines** in Fig. 9.)

Few prominent names have been drilled-down and shown in Fig. 9 based on citation count (from Google Scholar profiles.) For example, for **SIGMOD**, **VLDB** and **ICDM** the most popular researchers include **Srikanth Kandula (15 188 citations)**, **Divyakant Agrawal (23727 citations)** and **Shuicheng Yan (52 294 citations)**, respectively who were active in different periods in the past 18 years.

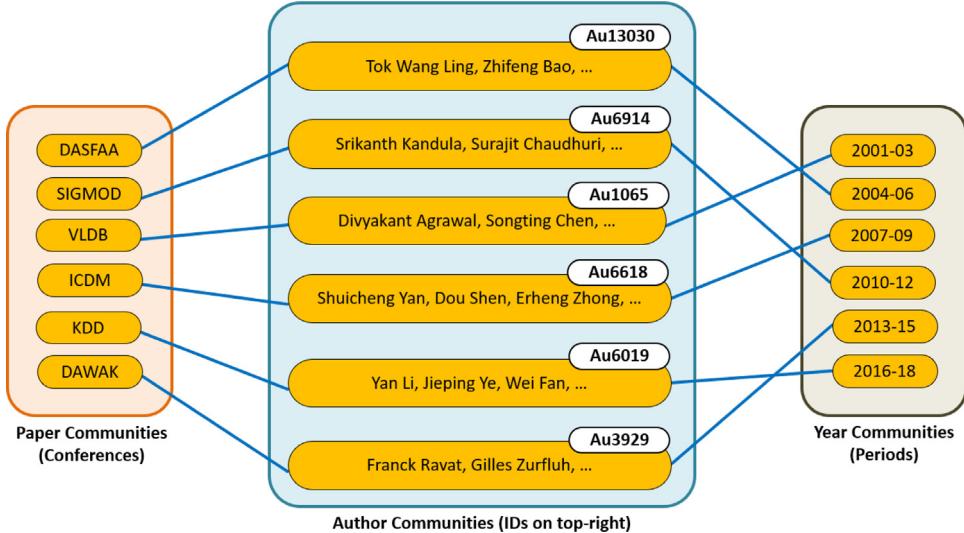


Fig. 9. (A4): Active periods for popular conference-wise co-author groups.

**(A5) Identify author collaboration groups who have published in conferences VLDB and SIGMOD, but have never published in conferences DASFAA and DaWaK.**<sup>6</sup> 102 communities are obtained after computing the expression. Drill-down results have been shown in Fig. 10 for few well-known groups most of whose members had collaborated on a paper that was published in both VLDB and SIGMOD (high ranked), but never in DASFAA or DaWaK (low to medium ranked).



Fig. 10. (A5): Author groups with high quality research.

There is a high probability that the work done by these groups is not only of good quality but also widely accepted. This claim is validated through the following facts:

- Fig. 10(a) community has researchers like Surajit Chaudhari who won the VLDB 10-Year Best Paper Award (2007) with Vivek Narasayya and VLDB Best Paper Award (2008) with Nicolas Bruno, apart from winning ACM SIGMOD Contributions Award (2004).
- Fig. 10(b) has researchers like Divyakant Agrawal who has 24 000+ citations (Google scholar).

<sup>6</sup> Note that these conferences can be changed and analyzed using the *same* expression, but with different layers. Parameterization of analysis objectives is one of the advantages of this approach!.

- **Peter A. Boncz and Stefan Manegold** from Fig. 10(c) group not only published a **highly cited paper** (350+ citations for MonetDB/XQuery) in SIGMOD 2006, but also have won the **VLDB 10-year award**.

### 8.3. IMDb analysis results

**MLN Details:** For IMDb MLN, we extracted, for the top 500 actors, the movies they have worked in (7500+ movies with 4500+ directors). The actor set was repopulated with the co-actors from these movies, giving a total of 9000+ actors. As explained for DBLP, the relationship attribute parameters in the EER model help in quantifying the similarity of actors and directors based on movie genres they have worked in. A vector was generated with the number of movies for each genre he/she has acted-in/directed. In order to consider the similarity with respect to *frequency of genres*, two actors/directors are connected if the Pearson's Correlation between their corresponding genre vectors is at least 0.9 (Other values can also be used based on similarity strength.) Widely used Louvain method [75] is used to detect layer-wise communities ( $\Psi$ ). Table 9 provides the layer statistics of the generated MLN.

**Table 9**  
IMDb MLN statistics.

| Layer                      | Number of nodes | Number of edges |
|----------------------------|-----------------|-----------------|
| ACTOR-Acts-with            | 9485            | 45,581          |
| ACTOR-SimilarAverageRating | 9485            | 13,945,912      |
| ACTOR-Similar-Genre        | 9485            | 996,527         |
| DIRECTOR-Similar-Genre     | 4510            | 250,845         |
| MOVIE-Similar-Rating       | 7951            | 8,777,618       |

**(A6) Cluster actors who have acted together and have a similar average rating.** 2430 actor groups with similar average ratings were detected in which *most of the actor pairs* have worked with each other. Few drill-down observations on the results:

- For the most popular average actor rating, [6-7), the largest co-actor groups were from Hollywood (876 actors), Indian (44 actors), Hong Kong (12 actors) and Spanish (9 actors) movies.
- Among the Hollywood movie based groups, the top group included co-actors Al Pacino, Robert De Niro, and Will Smith. Pacino and De Niro acted together in famous movies like Heat and Godfather Part II. Interestingly, co-actors from less known movies, such as Smith and De Niro, in Shark Tale also come up.
- Famous Bollywood stars like Amitabh Bachchan, Shah Rukh Khan belonged to largest top rated Indian group. They acted together in many highly rated bollywood movies.
- Jackie Chan (along with other lesser known actors) was among the prominent actors from the co-actor group from Hong Kong.

**(A7) Find the groups of actors who have never acted together, but are highly rated on an average and have worked in similar genres.** Following the expression, we detected 900 groups of actors most of whom have not worked together but have similar genre preferences and average rating. From the results, we drilled down into the communities that corresponded to *high average rating* and have listed a few *recognizable actors* along with *prominent genres* from those communities in Table 10. Out of these, as per reports in 2017, there had been talks of casting Johnny Depp and Tom Cruise in pivotal roles in Universal Studios' cinematic universe titled **Dark Universe** [84].

**Table 10**  
(A7): Highly rated genre actors who have not co-acted.

| Actors                   | Common prominent genres |
|--------------------------|-------------------------|
| Dafoe, Crowe             | Action, Crime           |
| Swank, Winslet           | Drama                   |
| Hanks, Witherspoon, Diaz | Comedy, Romance         |
| Depp, Cruise             | Adventure, Action       |
| DiCaprio, Gosling        | Crime, Romance          |
| Cage, Banderas           | Action, Thriller        |
| Grant, Hudson, Stone     | Comedy, Romance         |

**(A8) Identify genre-based groups of actors and directors having strong collaborations.** On computing the expression, 49 similar genre-based community pairs are obtained, where most actor-director pairs have interacted with each other at least once. Intuitively, a group of actors that prominently works in some genre (say, Drama, Action, Romance, ...) must pair up with the group of directors who primarily make movies in the *same genre*.

In Fig. 11 we have drilled-down and visualized the community pairings for the Action and Comedy genres with few famous actors and directors from each community. Such pairings may help production houses to sign up actors and directors for different movie genres. Recently, **Vin Diesel signed up for Avatar 2 and 3 (Action movie) which is being directed by James Cameron and this will be the first time they will be collaborating** [85]. Interestingly, even though they did not work together ever, we paired them together in the groups that corresponded to the Action genre on the basis of *high interaction among other similar actors and directors*.

**(A9): Identify, for each movie rating group the genre-based most popular actor and most popular director groups. From this result, find the actor and director groups having strong collaborations.** When finding the communities across three layers, using the expression in Table 4, we first combine results of each of two layers (ACTOR-Similar-Genre, DIRECTOR-Similar-Genre) with that of common layer (MOVIE-Similar-Rating) to find most popular genre-based group for each movie rating. Fig. 12(a) shows drill-down results of one such *intermediate* combination, where actors (community A144) and directors (community D91) are paired

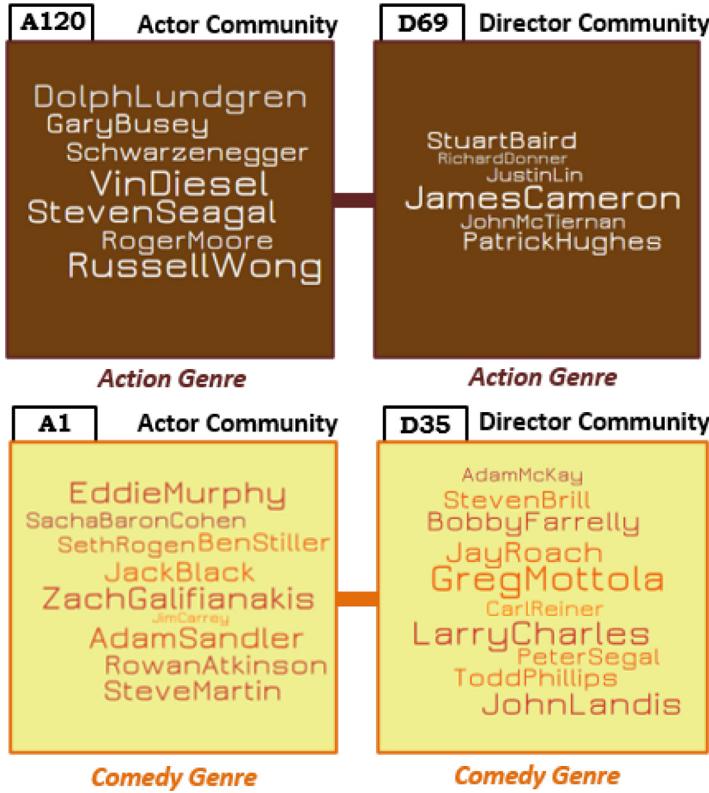


Fig. 11. (A8): Genre-based actor–director groups.

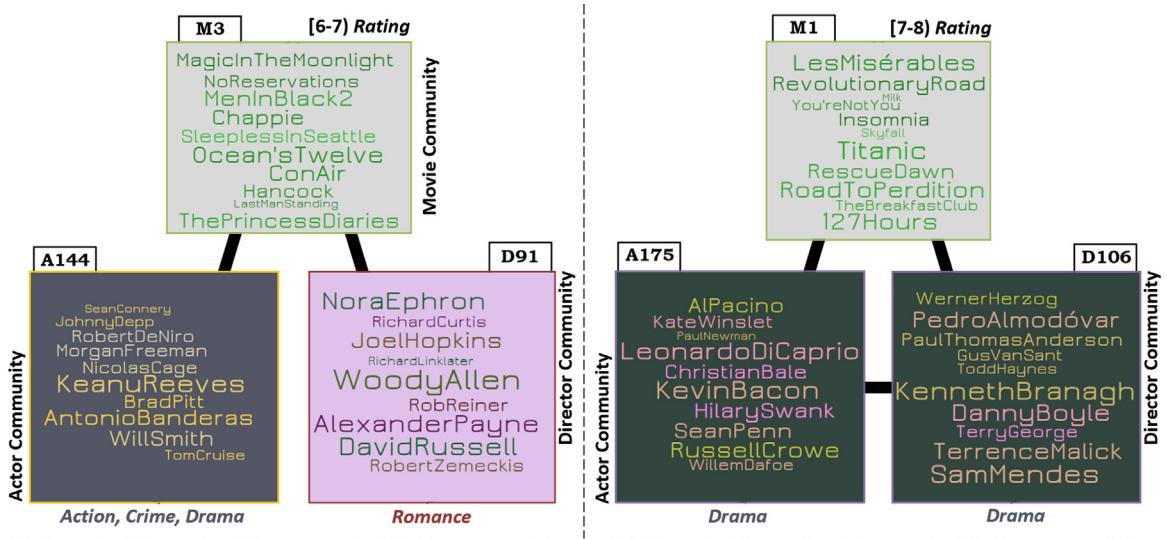


Fig. 12. (A9): Sample HeMLN community match results.

with movies (community M3). However, the most popular actor and director groups for [6-7) movie rating (represented by M3) do not have many interactions among them as they belong to different dominant genre groups.

Finally, the interactions between DIRECTOR-Similar-Genre communities and ACTOR-Similar-Genre communities are calculated to complete the analysis expression listed in Table 4. Only one HeMLN community drilled-down and visualized in Fig. 12(b) was obtained. The drill-down of Fig. 12(b) indicates, both the popular groups for [7-8) movie rating are from Drama genre and many

of these actor–director pairs have collaborated on many movies, such as **Leonardo DiCaprio**, **Kate Winslet** with **Sam Mendes** for **Revolutionary Road**, **Sean Penn** with **Gus Van Sant** for **Milk** and so on. Thus, popular groups A175 and D106 paired up with each other.

Most importantly, it is possible by drilling-down into the results to flesh out **potential actor–actor or actor–director collaborations** based on *identifying the missing links for high degree nodes in the generated HeMLN communities*. One such combination is **DiCaprio–Swank–Mendes** who *never worked together even though most of their movies belong to highly rated drama genre*.

#### 8.4. Covid-19 analysis results

**MLN Details:** Each node in a layer corresponds to a county (3141 nodes.) County nodes are connected as a clique if they have the similar increase, using several bands from spike (> 100% increase) to big dip (100% decrease) and a few in between, in the number of Covid new cases/deaths/hospitalizations/... and hence varies.

Based on the analysis objective, 2 disjoint intervals (each ranging from 1 to 30 days) are selected — either arbitrarily or before and after based on an event (e.g., July 4th, Thanksgiving) to visually understand the effect of Covid between the chosen two intervals. This is translated to the generation of layers for each interval, respectively, where the 3141 US counties with similar number of new cases are connected from each interval. As per the expression, communities are detected (using Louvain [75]) for the individual layers to find the geographical regions (or counties) that have the *same* percentage of increase or decrease (using bands) and 2 maps are displayed side-by-side using different colors and counties within the same band having the same color. The colors range from **spike** to a **big dip** in the number of new cases/deaths/hospitalizations/.... Different interval selections can be made based around the event of interest to analyze and visualize the effects through the live dashboard that makes use of the multilayer network architecture underneath.<sup>7</sup> The effects of two major events visualized in this paper are discussed below. More details can be found in [9,86].

**(A10) (i) Visualize how the geographical regions corresponding to the clusters of US counties with rise in daily confirmed cases shift in the consecutive 7-day intervals pre and post 2020 Spring Break.** For most US counties, the spring break was till the third week of March in 2020. For the *pre* spring break layer, the 7-day intervals used were Feb18–Feb24 and Feb25–Mar2. For the *post* spring break layer, the 7-day intervals used were Mar20–Mar26 and Mar27–Apr2. The drilled-down results have been visualized in Fig. 13 that show how *post the spring break there was a spike in the number of daily cases in counties across the US*. Various reports attributed this massive surge due to the widespread travel to popular tourist destinations during the break leading to *crowds and non-adherence to social distancing norms* [87–89].

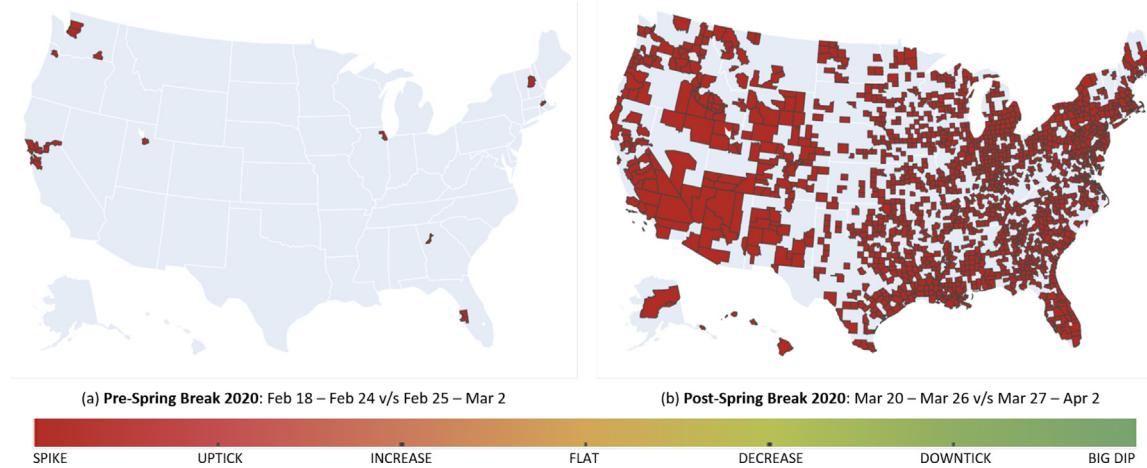


Fig. 13. (A10) (i) **Spike** in cases due to the 2020 spring break.

**(A10) (ii) Visualize how the geographical regions corresponding to the clusters of US counties with decline in daily confirmed cases shift in month-apart 3-day periods pre and post the Vaccination Drive:** The vaccination drive in the US began from December 14, 2020 [90]. For the *pre* vaccination drive layer, the 3-day intervals considered were Sep20–Sep22 and Oct21–Oct23 in 2020. For the *post* vaccination drive layer, the 3-day intervals were Jan20–Jan22 and Feb21–Feb23 in 2021. The community (groups of counties) results have been drilled-down from the individual layers and the ones displaying a downward trend have been visualized in Fig. 14. This illustration clearly shows how the **vaccination drive has become one of the reasons that has led to controlling the spread of COVID across US in the past few months**. This fact is also verified from independent sources that say how the administration of the vaccine has led to a *decline in severe cases, hospitalizations and deaths* not only in the US, but some other parts of the world as well [91,92].

More layers and decoupling-based compositions addressing the analysis objectives listed below are being developed with a revised version of visualization that provides more interaction and choices.

<sup>7</sup> Dashboard: <https://itlab.uta.edu/cowiz/>, Youtube Video: <https://www.youtube.com/watch?v=4vJ56FYBSCg>.

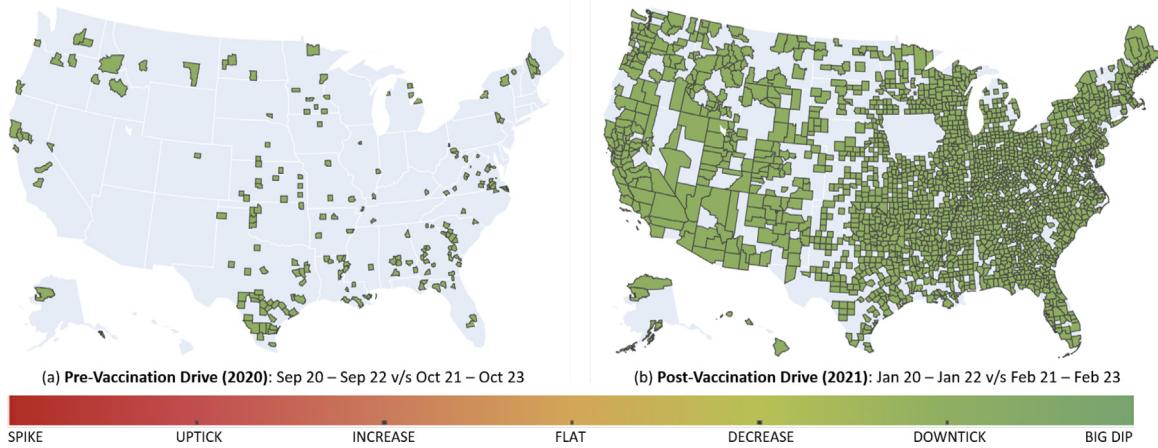


Fig. 14. (A10) (ii) Positive impact of the vaccination drive in the US.

- What is the effect of traffic movement on new cases across major (or centrally connected) counties? How to choose a county for lockdown so that it has maximum impact?
- Compare the increase/decrease in the number of new cases with respect to average education, per-capita income, mask usage and population density.

#### 8.5. Efficiency implications of the decoupling approach

Although a newly-proposed approach [7], we have used the decoupling-based algorithms for the analysis of the objectives. This is based on the efficiency benefits of this approach apart from its flexibility and ability to preserve structure and semantics. Below, we provide some efficiency results to make the point. More can be found in [7,8,78,93,94]

**Experimental Set up:** We used a quad-core 8th generation Intel i7 processor Linux machine with 8 GB memory for all of our analysis. The layer-wise results (communities or centrality nodes) are generated *once* and can be done in *parallel*. Thus, this one time cost is bounded by the layer that takes maximum time. Moreover, the cost of composing the partial results using Boolean AND (HoMLN Centrality Nodes and Communities) or Maximum Weighted Matching (HeMLN communities) is *significantly less*. For HoMLN analysis, we compare the total computational cost of the decoupling approach and the traditional single graph approach which includes the *time to generate the combined layer* followed by generating the degree/closeness hubs or communities. For HeMLN, we show how the *incremental cost* of generating the community pairings is minimal. In this section, we showcase the efficiency results of few of the analysis objectives.

**IMDb MLN Analysis Cost:** The IMDb MLN layer details, including the number of nodes and edges, can be found in Table 9. In case of analysis (A7), 2 Boolean AND compositions are required as per Table 4 to generate communities. Here the one time cost for finding the layer-wise communities is bound by the *ACTOR-Similar-AverageRating* layer (densest layer in Table 9). Overall, **64% reduction in computation time** is observed with the *decoupling approach* (228.865 s) as compared to *single graph approach* (636.45 s) as shown in Fig. 15(a).

Fig. 15(b) shows the execution time for the *one-time and iterative composition costs* for analysis (A9) - the most complex community we have computed which uses 3 iterations of pairings. Here, M, D and A represent the MOVIE-Similar-Rating, ACTOR-Similar-Genre and DIRECTOR-Similar-Genre layers, respectively. The difference in one-time cost for the 3 layers follow their edge density, that can be derived from Table 9. Iteration cost includes *creating the bipartite graph with the layer-communities that are part of the result as meta-nodes (depicted by, for example, A' instead of A in case of D'-A' iteration)*, *computing meta edge weights*, and *MWM cost*. We can also see how the iterative cost is insignificant as compared to the one time cost (by an order of magnitude.) Even the *cost of all iterations together* (0.2582 s) is still almost *an order of magnitude less than the largest one-time cost* (3.173 s for MOVIE-Similar-Rating layer.) The *additional incremental cost for computing HeMLN communities is extremely small*.

In summary, extensive experiments on the diverse data sets validate the **MLN decoupling approach from an efficiency perspective**.

#### 9. Conclusions

The success of any data analysis life cycle is predicated on our ability to: (i) appropriately model the data set and automate schema generation as much as possible, (ii) provide analysis alternatives using the model used, (iii) map user-specified objectives to analysis expressions, (iv) develop algorithms for computing expressions, preferably efficiently, and (iv) ability to drill-down and visualize results for ease of understanding and interpretation. This, as we know, is an iterative process.

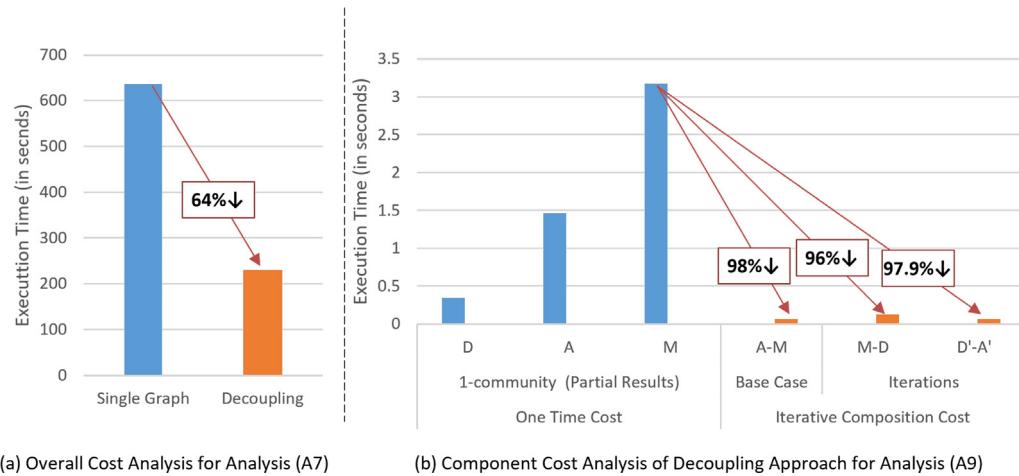


Fig. 15. Efficiency of decoupling approach for MLN analysis.

In this paper, our focus has been to address the complete life cycle for aggregate analysis using the Multilayer Networks (or MLNs) as the underlying data model. In this paper, we demonstrate how to create EER diagrams for “multi-entity, multi-feature” data sets. We use an algorithm that maps the EER diagram into MLNs of different kinds as appropriate. We have shown how user-specified objectives, in English, can be translated using heuristics based on keywords to aggregate analysis expressions for all types of MLNs. We have also demonstrated the applicability of the *decoupling approach* for efficient analysis of complex data sets using multilayer networks. Drill-down and visualization of the results is an important but sometimes ignored component. We have shown how analysis results can be visualized using a general-purpose approach with a Covid-19 dashboard.

We are working on further automating the translation of objectives into analysis expressions by using natural language processing and model characteristics. We are also working on developing decoupling-based efficient algorithms for aggregate analysis for centrality, substructure discovery, motif detection, to name a few, on MLNs. We are also broadening our analysis to include non-Boolean compositions for homogeneous MLNs, weighted graphs, and labeled graphs. This will further extend the expressive power of the MLN data model and the automation of the data discovery life cycle.

#### CRediT authorship contribution statement

**Abhishek Santra:** Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Resources, Data curation, Writing – original draft, Writing – review & editing, Visualization. **Kanthy Komar:** Software, Validation, Investigation, Resources, Data curation, Visualization. **Sanjukta Bhowmick:** Conceptualization, Methodology, Writing – original draft, Supervision, Project administration, Funding acquisition. **Sharma Chakravarthy:** Conceptualization, Methodology, Formal analysis, Writing – original draft, Writing – review & editing, Supervision, Project administration, Funding acquisition.

#### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

#### Acknowledgments

For this work, Dr. S. Chakravarthy was partially supported by NSF grants 1955798 and 2102393, Dr. A. Santra was partially supported by NSF Grant 1955798, and Dr. S. Bhowmick was partially supported by NSF grants 1916084 and 2120414.

#### References

- [1] Google cloud dataflow. <https://cloud.google.com/dataflow>.
- [2] AWS glue. <https://aws.amazon.com/glue/>.
- [3] Azure data factory. <https://azure.microsoft.com/en-in/services/data-factory/>.
- [4] IBM DataStage. <https://www.ibm.com/in-en/products/infoSphere-dataStage>.
- [5] Informatica. <https://www.informatica.com/in/products/data-integration/powercenter.html>.
- [6] Neo4J NoSQL system. [neo4j.com](https://neo4j.com).

[7] A. Santra, S. Bhowmick, S. Chakravarthy, Efficient community re-creation in multilayer networks using boolean operations, in: Int. Conf. on Computational Science, Zurich, Switzerland, 2017, pp. 58–67.

[8] A. Santra, S. Bhowmick, S. Chakravarthy, HUBify: Efficient estimation of central entities across multiplex layer compositions, in: IEEE International Conference on Data Mining Workshops, 2017.

[9] Kunal Samant, Endrit Memeti, Abhishek Santra, Enamul Karim, Sharma Chakravarthy, CoWiz: Interactive Covid-19 visualization based on multilayer network analysis, in: 37th IEEE International Conference on Data Engineering, ICDE 2021, Chania, Crete, Greece, April 19-22, 2021, 2021, <https://itlab.uta.edu/cowiz/> and <https://www.youtube.com/watch?v=4vJ56FYBSCg>.

[10] Peter Pin-Shan Chen, The entity-relationship model—toward a unified view of data, *ACM Trans. Database Syst.* 1 (1) (1976) 9–36.

[11] ERWIN data modeler, <https://erwin.com/products/erwin-data-modeler/>.

[12] ERDplus, <https://erdplus.com/>.

[13] IBM InfoSphere data architect, <https://www.ibm.com/us-en/marketplace/infosphere-data-architect>.

[14] Toad data modeler, <http://www.toadworld.com/products/toad-data-modeler>.

[15] DBeaver, <https://dbeaver.io/>.

[16] DBSchema, <http://www.dbschema.com/>.

[17] ER studio, <https://www.idera.com/er-studio-data-architect-software>.

[18] Nandish Jayaram, Arijit Khan, Chengkai Li, Xifeng Yan, Ramez Elmasri, Querying knowledge graphs by example entity tuples, in: *IEEE Trans. Knowl. Data Eng.*, 27, (10) 2015, pp. 2797–2811.

[19] Soumyava Das, Abhishek Santra, Jay Bodra, Sharma Chakravarthy, Query processing on large graphs: Approaches to scalability and response time trade offs, *Data Knowl. Eng.* 126 (2020) 101736.

[20] Sharma Chakravarthy, Ramji Beera, Ramanathan Balachandran, DB-subdue: Database approach to graph mining, in: PAKDD, 2004, pp. 341–350.

[21] Renzo Angles, Claudio Gutierrez, Survey of graph database models, *ACM Comput. Surv.* 40 (1) (2008) 1–39.

[22] Noa Roy-Hubara, Lior Rokach, Bracha Shapira, Peretz Shoval, Modeling graph database schema, *IT Prof.* 19 (6) (2017) 34–43.

[23] Jaroslav Pokorný, Conceptual and database modelling of graph databases, in: Proceedings of the 20th International Database Engineering & Applications Symposium, 2016, pp. 370–377.

[24] Renzo Angles, The property graph database model, in: AMW, 2018.

[25] Kanthi Sannappa Komar, Abhishek Santra, Sanjukta Bhowmick, Sharma Chakravarthy, EER→MLN: EER approach for modeling, mapping, and analyzing complex data using multilayer networks (MLNs), in: Gillian Dobbie, Ulrich Frank, Gerti Kappel, Stephen W. Liddle, Heinrich C. Mayr (Eds.), Conceptual Modeling - 39th International Conference, ER 2020, Vienna, Austria, November 3-6, 2020, Proceedings, in: Lecture Notes in Computer Science, 12400, Springer, 2020, pp. 555–572, [http://dx.doi.org/10.1007/978-3-030-62522-1\\_41](http://dx.doi.org/10.1007/978-3-030-62522-1_41), [https://doi.org/10.1007/978-3-030-62522-1\\_41](https://doi.org/10.1007/978-3-030-62522-1_41).

[26] Jungeun Kim, Jae-Gil Lee, Community detection in multi-layer graphs: A survey, *SIGMOD Record* 44 (3) (2015) 37–48.

[27] S. Fortunato, C. Castellano, Community structure in graphs, in: *Ency. Complex. Syst.*, 2009, pp. 1141–1163.

[28] Zhige Xin, Community Detection in Social Networks (Ph.D. thesis), University of California, Davis, 2018.

[29] Matteo Magnani, Obaida Hanteer, Roberto Interdonato, Luca Rossi, Andrea Tagarelli, Community detection in multiplex networks, *ACM Comput. Surv.* 54 (3) (2021) 48:1–48:35.

[30] Xiaowen Dong, Pascal Frossard, Pierre Vandergheynst, Nikolai Nefedov, Clustering with multi-layer graphs: A spectral perspective, *IEEE Trans. Signal Process.* 60 (11) (2012) 5820–5831.

[31] Huajing Li, Zaiqing Nie, Wang-Chien Lee, Lee Giles, Ji-Rong Wen, Scalable community discovery on textual data with relations, in: Proceedings of the 17th ACM Conference on Information and Knowledge Management, ACM, 2008, pp. 1203–1212.

[32] Zhiqiang Xu, Yiping Ke, Yi Wang, Hong Cheng, James Cheng, A model-based approach to attributed graph clustering, in: Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data, ACM, 2012, pp. 505–516.

[33] Deng Cai, Zheng Shao, Xiaofei He, Xifeng Yan, Jiawei Han, Mining hidden community in heterogeneous social networks, in: Proceedings of the 3rd International Workshop on Link Discovery, ACM, 2005, pp. 58–65.

[34] Han Zhang, Chang-Dong Wang, Jian-Huang Lai, S. Yu Philip, Modularity in complex multilayer networks with multiple aspects: a static perspective, in: *Applied Informatics*, Vol. 4, Springer Berlin Heidelberg, 2017, p. 7.

[35] Manlio De Domenico, Albert Solé-Ribalta, Elisa Omodei, Sergio Gómez, Alex Arenas, Centrality in interconnected multilayer networks, 2013, arXiv preprint [arXiv:1311.2906](https://arxiv.org/abs/1311.2906).

[36] Albert Solé-Ribalta, Manlio De Domenico, Sergio Gómez, Alex Arenas, Centrality rankings in multiplex networks, in: *Procs. of 2014 ACM Conf. on Web Science*, ACM, 2014, pp. 149–155.

[37] Qianyi Zhan, Jiawei Zhang, Senzhang Wang, S Yu Philip, Junyuan Xie, Influence maximization across partially aligned heterogenous social networks, in: PAKDD (1), 2015, pp. 58–69.

[38] Chuan Shi, Yitong Li, Jiawei Zhang, Yizhou Sun, S Yu Philip, A survey of heterogeneous information network analysis, *IEEE Trans. Knowl. Data Eng.* 29 (1) (2017) 17–37.

[39] Yizhou Sun, Jiawei Han, Mining heterogeneous information networks: a structural analysis approach, *ACM SIGKDD Explor. Newsl.* 14 (2) (2013) 20–28.

[40] C Wang, Y Sun, Y Song, J Han, Y Song, L Wang, M Zhang, Relsim: relation similarity search in schema-rich heterogeneous information networks, in: Proceedings of the 2016 SIAM International Conference on Data Mining, SIAM, 2016, pp. 621–629.

[41] Chengguang Wang, Yangqiu Song, Haoran Li, Ming Zhang, Jiawei Han, Text classification with heterogeneous information network kernels, in: AAAI, 2016, pp. 2130–2136.

[42] Jiawei Zhang, Philip S. Yu, Yuanhua Lv, Organizational chart inference, in: Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, 2015, pp. 1435–1444.

[43] Chuan Shi, Yitong Li, S. Yu Philip, Bin Wu, Constrained-meta-path-based ranking in heterogeneous information network, *Knowl. Inf. Syst.* 49 (2) (2016) 719–747.

[44] Chuan Shi, Zhiqiang Zhang, Ping Luo, Philip S Yu, Yading Yue, Bin Wu, Semantic path based personalized recommendation on weighted heterogeneous information networks, in: Proceedings of the 24th ACM International Conference on Information and Knowledge Management, ACM, 2015, pp. 453–462.

[45] David Melamed, Community structures in bipartite networks: A dual-projection approach, *PLoS One* 9 (5) (2014) e97823.

[46] Yizhou Sun, Jiawei Han, Peixiang Zhao, Zhipun Yin, Hong Cheng, Tianyi Wu, Rankclus: integrating clustering with ranking for heterogeneous information network analysis, in: Proceedings of the 12th International Conference on Extending Database Technology: Advances in Database Technology, ACM, 2009, pp. 565–576.

[47] Manlio De Domenico, Vincenzo Nicosia, Alex Arenas, Vito Latora, Layer aggregation and reducibility of multilayer interconnected networks, 2014, CoRR abs/1405.0425.

[48] A.J. Berenstein, M.P. Magarinos, A. Chernomoretz, F. Aguero, A multilayer network approach for guiding drug repositioning in neglected diseases, *PLoS* (2016).

[49] The centre for disease control covid dashboard, <https://covid.cdc.gov/covid-data-tracker/>.

[50] Johns Hopkins university covid dashboard, <https://www.arcgis.com/apps/opsdashboard/index.html#/bda7594740fd40299423467b48e9ecf6>.

[51] The university of Washington covid dashboard, <https://hgis.uw.edu/virus/>.

[52] American airlines routes, <http://aa.flmaps.com/en>.

- [53] Southwest airlines routes. [https://www.southwest.com/flight/routemap\\_dyn.html](https://www.southwest.com/flight/routemap_dyn.html).
- [54] Delta airlines routes. <http://dl.fltmaps.com/en>.
- [55] Spirit airlines routes. <https://www.spirit.com/RouteMaps.aspx>.
- [56] Frontier airlines routes. <https://flights.flyfrontier.com/en/sitemap/city-to-city-flights/page-1>.
- [57] Allegiant airlines routes. <https://www.allegiantair.com/route-map>.
- [58] Alessio Cardillo, Jesús Gómez-Gardenes, Massimiliano Zanin, Miguel Romance, David Papo, Francisco Del Pozo, Stefano Boccaletti, Emergence of network features from multiplexity, *Sci. Rep.* 3 (2013).
- [59] DBLP dataset. <http://dblp.uni-trier.de/xml/>.
- [60] The internet movie database. <ftp://ftp.fu-berlin.de/pub/misc/movies/database/>.
- [61] Data from the new york times, based on reports from state and local health agencies. <https://github.com/nytimes/covid-19-data>.
- [62] US census data. <https://data.census.gov/cedsci/>.
- [63] Aaron Clauset, Mark E.J. Newman, Cristopher Moore, Finding community structure in very large networks, *Phys. Rev. E* 70 (6) (2004) 066111.
- [64] Jure Leskovec, Kevin J. Lang, Anirban Dasgupta, Michael W. Mahoney, Community structure in large networks: Natural cluster sizes and absence of large well-defined clusters, 2008.
- [65] Mark Newman, *Networks: An Introduction*, Oxford University Press, Inc., New York, NY, USA, 2010.
- [66] Soumyava Das, Sharma Chakravarthy, Duplicate reduction in graph mining: Approaches, analysis, and evaluation, *IEEE Trans. Knowl. Data Eng.* 30 (8) (2018) 1454–1466.
- [67] Michihiro Kuramochi, George Karypis, Finding frequent patterns in a large sparse graph\*, *Data Min. Knowl. Discov.* 11 (3) (2005) 243–271.
- [68] Lawrence B. Holder, Diane J. Cook, Surnjani Djoko, Substructure discovery in the SUBDUE system, in: *Knowledge Discovery and Data Mining*, 1994, pp. 169–180.
- [69] Soumyava Das, Ankur Goyal, Sharma Chakravarthy, Plan before you execute: A cost-based query optimizer for attributed graph databases, in: *DaWaK 2016*, Porto, Portugal, September 6–8, 2016, 2016, pp. 314–328.
- [70] Y. Hao, H. Cao, Y. Qi, C. Hu, S. Brahma, J. Han, Efficient keyword search on graphs using MapReduce, in: *IEEE International Conference on Big Data*, Santa Clara, CA, USA, 2015, pp. 2871–2873.
- [71] S. Boccaletti, G. Bianconi, R. Criado, C.I. del Genio, J. Gómez-Gardeñes, M. Romance, I. Sendiña-Nadal, Z. Wang, M. Zanin, The structure and dynamics of multilayer networks, *Phys. Rep.* (2014).
- [72] M. Kivelä, A. Arenas, M. Barthelemy, J.P. Gleeson, Y. Moreno, M.A. Porter, Multilayer networks, 2013, CoRR <abs/1309.7233>.
- [73] Joydeep Banerjee, Chenyang Zhou, Das, Arunabha Sen, On robustness in multilayer interdependent networks, in: Erich Rome, Marianthi Theocharidou, Stephen Wolthusen (Eds.), *Critical Information Infrastructures Security*, Springer International Publishing, Cham, 2016, pp. 247–250.
- [74] Ludvig Bohlin, Daniel Edler, Andrea Lancichinei, Martin Rosvall, Community detection and visualization of networks with the map equation framework, 2014.
- [75] Vincent D. Blondel, Jean-Loup Guillaume, Renaud Lambiotte, Etienne Lefebvre, Fast unfolding of community hierarchies in large networks, 2008, CoRR <abs/0803.0476>.
- [76] Manlio De Domenico, Albert Solé-Ribalta, Sergio Gómez, Alex Arenas, Navigability of interconnected networks under random failures, *Proc. Natl. Acad. Sci.* (2014).
- [77] James D. Wilson, John Palowitch, Shankar Bhamidi, Andrew B. Nobel, Community extraction in multilayer networks with heterogeneous community structure, *J. Mach. Learn. Res.* 18 (1) (2017) 5458–5506.
- [78] Abhishek Santra, Kanthi Sannappa Komar, Sanjukta Bhowmick, Sharma Chakravarthy, A new community definition for MultiLayer networks and a novel approach for its efficient computation, 2020, arXiv preprint <arXiv:2004.09625>.
- [79] Anish Rai, *MLN-SUBDUE: Decoupling Approach-based Substructure Discovery In Multilayer Networks (MLNs)* (Master's thesis), The University of Texas at Arlington, 2020.
- [80] Brigitte Boden, Stephan Günemann, Holger Hoffmann, Thomas Seidl, Mining coherent subgraphs in multi-layer graphs with edge labels, in: *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, in: *KDD '12*, ACM, 2012, pp. 1258–1266.
- [81] Jack Edmonds, Maximum matching and a polyhedron with 0, 1-vertices, *J. Res. Natl. Bureau Stand. B* 69 (125–130) (1965) 55–56.
- [82] Antonio Diaz, Major airline hubs, 2017, <https://Www.Travelmiles101.Com/ListOf-Major-Airline-Hubs/>.
- [83] Justin Dawes, Grand rapids is 'sweet spot' for airline base, Grand Rapids Bus. J. (GRBJ) (2019).
- [84] Jacob Stoltzworth, Dark universe: Johnny depp and Javier bardem join Tom Cruise in universal's monster movie franchise, 2017, <https://Www.Independent.Co.Uk/Us>.
- [85] Jacob Stoltzworth, Avatar 2 and 3: Vin diesel joins cast of James Cameron's long awaited sequels, 2019, <https://Www.Independent.Co.Uk/Us>.
- [86] David Rademacher, Jacob Valdez, Endrit Memeti, Kunal Samant, Abhishek Santra, Sharma Chakravarthy, Modviz: a modular and extensible architecture for drill-down and visualization of complex data, in: Mirjana Ivanovic, Marite Kirikova, Laila Niedrite (Eds.), *Digital Business and Intelligent Systems - 15th International Baltic Conference, Baltic DB&IS 2022*, Riga, Latvia, July 4–6, 2022, Proceedings, in: *Communications in Computer and Information Science*, 1598, Springer, 2022, pp. 232–250, [http://dx.doi.org/10.1007/978-3-031-09850-5\\_16](http://dx.doi.org/10.1007/978-3-031-09850-5_16), [https://doi.org/10.1007/978-3-031-09850-5\\_16](https://doi.org/10.1007/978-3-031-09850-5_16).
- [87] Michael T. Nietzel, New study: College spring break helped spread the coronavirus. <https://www.forbes.com/sites/michaelnietzel/2020/06/19/new-study-college-spring-break-helped-spread-the-coronavirus/?sh=7fdb618a4e52>.
- [88] Patricia Mazzei, Frances Robles, The costly toll of not shutting down spring break earlier. <https://www.nytimes.com/2020/04/11/us/florida-spring-break-coronavirus.html>.
- [89] Daniel Mangrum, Paul Niekamp, JUE insight: College student travel contributed to local COVID-19 spread, *J. Urban Econ.* (2020) 103311.
- [90] US administers 1st doses of Pfizer coronavirus vaccine. <https://abcnews.go.com/US/us-administer-1st-doses-pfizer-coronavirus-vaccine/story?id=74703018>.
- [91] Reis Thebault, Four reasons experts say coronavirus cases are dropping in the United States. <https://www.washingtonpost.com/health/2021/02/14/why-coronavirus-cases-are-dropping/>.
- [92] Ehud Rinott, Reduction in COVID-19 patients requiring mechanical ventilation following implementation of a national COVID-19 vaccination program—Israel, December 2020–February 2021, in: *MMWR. Morbidity and Mortality Weekly Report*, 70, 2021.
- [93] Xuan-Son Vu, Abhishek Santra, Sharma Chakravarthy, Lili Jiang, Generic multilayer network data analysis with the fusion of content and structure, in: *CiCLing 2019*, la Rochelle, France, 2019.
- [94] Hamza Reza Pavel, Abhishek Santra, Sharma Chakravarthy, Degree centrality algorithms for homogeneous multilayer networks, in: *Proceedings of the 14th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management*, IC3K 2022, 2022.



**Abhishek Santra** is a Senior Lecturer of Computer Science and Engineering Department at The University of Texas at Arlington (UT Arlington), Texas. He is also a Post-Doctoral Researcher at the Information Technology Lab (ITLab) at UT Arlington. He received his Ph. D. Degree in Computer Science from the UT Arlington in 2020. Prior to this, he completed his Master's in Computer Science from the University of Delhi, India. His interest areas include multi-layer networks, graph mining, databases, and multi-feature data analysis.



**Kanthi Komar** is a Data Scientist at Journera. Prior to this she worked as a Sr. Analyst at Discover Financial Service working on predictive analytical models. She completed her thesis in Graph Mining and received her Master's degree in Computer Science from University of Texas at Arlington in 2019. Her interests include multi-layer network modeling, data mining and machine learning.



**Sanjukta Bhowmick** is a Research Associate Professor in the Computer Science and Engineering department at the University of North Texas. She obtained her Ph.D. from the Pennsylvania State University. She had a joint postdoc at Columbia University and Argonne National Lab. Her current research is on understanding change in complex network analysis, with a focus on developing scalable algorithms for large dynamic networks and developing uncertainty quantification metrics for network analysis.



**Sharma Chakravarthy** is Professor of Computer Science and Engineering Department at The University of Texas at Arlington (UTA). He is the founder director of the Information Technology laboratory (IT Lab) at UTA. He is the co-author of the book: Stream Data Processing: A Quality of Service Perspective (2009). His current research includes multilayer network analysis, video situation analysis using stream data processing, and scalability of mining and querying algorithms. He has published over 150 papers in refereed international journals and conference proceedings.