





Network alignment in psychometrics: Comparing symptom networks in eating disorders across age and diagnosis

Riccardo De Luca^{a,b,c},¹, Laura Girelli^d, Pietro H. Guzzi^e, Paolo Tieri^c,²,
Manuela Petti^a,^{*,2}

^a DIAG Department of Computer, Control and Management Engineering, Sapienza University of Rome, Via Ariosto 25, Rome, 00185, Italy

^b Unit of Artificial Intelligence and Computer Systems, Department of Engineering, Università Campus Bio-Medico di Roma, Via Alvaro del Portillo 21, Rome, 00128, Italy

^c CNR Consiglio Nazionale delle Ricerche, IAC Istituto per le Applicazioni del Calcolo, Via dei Taurini 19, Rome, 00185, Italy

^d University of Salerno, Department of Humanities, Philosophy and Education, Via Giovanni Paolo II 132, Fisciano (SA), 84084, Italy

^e Dept Surgical and Medical Sciences, University of Catanzaro "Magna Graecia", Campus Universitario "S. Venuta", Viale Europa, loc. Germaneto, Catanzaro, 88100, Italy

ARTICLE INFO

Keywords:

Network science
Network alignment
Psychometric data
Geometric deep learning
Eating disorders

ABSTRACT

Network alignment (NA) is an emerging tool in network science, enabling knowledge transferring and structural comparison of networks. While widely used in biology, social networks and computer vision, its application to network psychometrics has not been explored. In this study we present a comparative analysis of several NA algorithms, spanning four methodological categories: spectral methods, representation learning, deep matching, and genetic algorithms. These algorithms were evaluated using psychometric networks derived from the Eating Disorder Inventory-3 (EDI-3) dataset: in these networks, nodes represent symptoms/personality traits related to eating disorders, while edges encode for statistical dependencies between nodes. Initially, we tested each algorithm's robustness to structural noise by aligning a psychometric network with permuted and noise-corrupted versions of itself. The best-performing algorithm from each category was then used to align networks representing different age groups (i.e., early vs. late adolescence; early adolescence vs. adulthood) and diagnostic categories (anorexia nervosa vs. bulimia nervosa). The overall best-performing algorithm was subsequently run 100 times to identify the most consistently aligned node pairs, which were interpreted from a psychological perspective. Our results show that PALE, a representation learning-based algorithm, demonstrated the highest robustness and stability under noisy conditions and consistently outperformed others in real-world applications. Notably, PALE identified symptom/personality trait nodes with consistent functional roles across developmental stages, such as interpersonal difficulties and body-related concerns, highlighting their persistent influence within the symptom networks of patients with eating disorders.

1. Introduction

Psychometric networks are graph-based models that capture the interactions and correlations among symptoms of mental disorders, reflecting the view that such disorders may emerge not from a single latent cause but from mutually reinforcing symptoms and traits [1,2]. Such interplay can be effectively rendered and rigorously analyzed by means of graph theory and network science. Recently, network comparison approaches such as the Network Comparison Test (NCT), have been proposed in network psychometrics to compare psychological

constructs across groups by testing statistical differences in network topology, edge strength, and global strength [3].

Another possible approach to compare networks is Network Alignment (NA), the computational problem of leveraging similarities between two or more networks to create a mapping between their nodes. Based on the nature of the resulting mapping, NA can be categorized as either Local Network Alignment, to align sub-graphs, or Global Network Alignment, to create a one-to-one mapping between the network nodes [4]. In the biomedical field, NA has been successfully used in biology and neuroscience. In particular, in biology, NA has

* Corresponding author.

E-mail addresses: riccardo.deluca@unicampus.it (R. De Luca), lgirelli@unisa.it (L. Girelli), hguzzi@unicz.it (P.H. Guzzi), paolo.tieri@cnr.it (P. Tieri), manuela.petti@uniroma1.it (M. Petti).

¹ The research was conducted while the author was at DIAG, Sapienza University of Rome.

² These authors contributed equally to this work and share last authorship.

been used to compare molecular networks (e.g. PPIs) across species or conditions, helping to identify conserved pathways, functional analogs, or evolutionary relationships [4]. In neuroscience, it has been applied to compare brain connectomes derived from diffusion MRI data across individuals and developmental stages [5].

In psychometrics, NA could be utilized to compare psychological symptom networks (or other psychometric constructs) across different groups, time points, or conditions. In particular, the mapping can be used to compare the properties and roles that the symptoms assume in the different networks, e.g. to detect conceptual equivalence or differences in how symptoms relate. In fact, a symptom (or personality trait) observed in an adolescent may be replaced by a different symptom in adulthood that serves a similar psychological function or reflects the same underlying issue. Thus, while NCT provides formal hypothesis testing for group-level network differences, NA enables fine-grained symptom roles comparison.

The primary goal of this work is to propose and test the use of global network alignment (which we will refer to simply as Network Alignment for brevity) in network psychometrics. These networks are not created from observed interactions, but rather deduced through a statistical method: this peculiarity makes it an interesting novelty in the landscape of alignment algorithms.

For this aim we developed NetAlign, a framework for the development and testing of alignment algorithms, and we focused on psychometric networks associated with eating disorders, using those constructed by Punzi et al. (2023) [6] as a case study. Such networks have been built by exploiting data from the open access project Eating Disorder Inventory-3, American Clinical Cases [7], which collects responses to the EDI-3 questionnaire from 1206 female patients, aged between 11 and 75 years, who received treatment at an Ohio eating disorders' clinic between 1996 and 2015. The EDI-3 questionnaire is a standardized self-report instrument used to assess eating disorders [8]. The investigation proceeds in two main stages. First, we perform a comprehensive evaluation of multiple network alignment algorithms by testing their ability to withstand noise in a controlled, semi-synthetic environment. Second, we apply the best-performing algorithm to align psychometric networks of different age groups and diagnostic categories, comparing aligned symptoms to uncover psychologically meaningful patterns.

In Section 2, we provide an overview of the main network alignment algorithm categories and discuss two existing frameworks for the comparison of alignment algorithms, which served as inspiration for this work. In Section 3, we provide a detailed examination of the psychometric networks utilized, outline the algorithms applied, explain our analytical methodology, and present the framework developed for the study. In Section 4, we present quantitative results of the conducted experiments, and in Section 5 we discuss them. Finally, in Section 6, we draw the conclusions of this analysis and outline a potential direction for future investigations.

2. Background

2.1. State of the art of network alignment algorithms

Over the years, several algorithms have been proposed to solve the problem of global network alignment. We can group them into four main approaches: spectral methods (SMs), representation learning methods (RLMs), deep matching methods (DMMs), and genetic algorithms-based methods (GAMs). All these algorithms aim to build a one to one mapping between graphs by optimizing the associated cost function.

Spectral methods (SMs) address the problem directly through the manipulation of adjacency and/or feature matrices. The algorithms in this category define an objective function that takes as input the adjacency matrices (e.g. IsoRank [9]), the feature matrices (e.g. Bi-Align [10]), or both (e.g. FINAL [11]) and optimize it based on a similarity matrix between the vertices of the networks.

Representation learning methods (RLMs), on the other hand, exploit the concept of node embedding to encode the topological information and attributes of a node into a vector. By using similarity measures between vectors (usually, cosine similarity), it is possible to create a similarity matrix among the nodes of two or more networks. From this similarity matrix, a one-to-one mapping can then be derived [12]. Typically, they are composed of two phases: the representation phase and the mapping phase.

During the *representation phase*, the node embeddings are generated such that nodes that are *close* (connected by a relatively short number of steps) and/or with similar attributes correspond to similar embeddings. While in *mapping phase*, a set of anchor nodes (approximately 20% of the total number of nodes) is used to map the embeddings from the vector space of the source network to the vector space of the target network. The anchor nodes are pairs of source-target nodes that are known to be aligned [13–15]. Unlike common learning algorithms, the percentage of elements (nodes) used for training in the mapping phase is much smaller than the total number. This is because the encoding of embeddings primarily occurs during the representation phase, where the networks are analyzed separately. The mapping phase does not add new information to the embeddings but primarily serves as a translator between different vector spaces.

Deep matching methods (DMMs), similarly to RLMs, also leverage the concept of node embedding; however, they do so differently. While RLMs generate node embeddings in a preceding phase and only subsequently align them between two or more networks through a translation process, DMMs generate node embeddings by directly aiming for the alignment task [16,17].

To achieve this, DMMs exploit Graph Neural Networks (GNNs), which combine network topology with node and edge features to produce node embeddings. These embeddings are then used to calculate a correspondence between the nodes of the networks. Through a learning function, the weights of the GNN are updated, which will produce more refined embeddings in the next iteration. In this sense, we state that the embeddings are generated by directly targeting the alignment task.

Finally, the **genetic algorithm-based methods** adapt the paradigm of genetic algorithms for the network alignment problem. They consider as *population* a set of possible alignments (represented by permutation matrices). Through a cross-over function the *child alignments* are generated by inheriting and combining the characteristics of the *parent alignments*, while optimizing a metric for the conservation of edges and/or nodes [18].

2.2. Related work

In recent years, the main frameworks that have been proposed for developing and testing network alignments algorithms are the framework developed for the comparative analysis by Trung et al. (2020) [12], and ThinkMatch, developed by ThinkLab at Shanghai Jiao Tong University [19].

The framework by Trung et al. was developed for the implementation and testing of spectral-based and representation learning-based algorithms. It provides a testbed capable of handling real, synthetic, and semi-synthetic datasets. Here, a dataset is defined as a pair of networks (source and target) to be aligned, together with the optional ground truth expressed as a permutation matrix. A semi-synthetic network is generated artificially from a real network by permuting its node indices and injecting noise. The framework manages the creation of these networks, and both the noise level and its type are configurable via hyperparameters. It also includes functions for computing matching accuracy and other alignment metrics.

However, the framework relies on the node-link data format for network management. This entails converting any network to that specific format using a script supplied with the framework, which makes the incorporation of new data less straightforward. Moreover, when defining a new algorithm, one must explicitly declare its input

arguments in the main file so they can be propagated throughout the rest of the framework [12].

ThinkMatch is a more recent framework that provides modules for the development of deep graph-matching algorithms, as well as implementations of several current state-of-the-art methods. It also includes a suite of benchmark tests on the major keypoint-feature-matching datasets [19]. The framework is built upon the PyTorch Geometric library, a PyTorch spin-off tailored for graph processing [20], and the pygmtools library, developed by the same authors, which supplies graph-matching solvers in Python [21]. Furthermore, ThinkMatch streamlines the management of algorithms and experiments through the use of configuration files (in YAML format), making it straightforward to modify experiment settings and add new algorithms.

However, the framework is specifically focused on the *keypoint-feature-matching problem*, in which graphs are extracted from the image keypoints that one wishes to align. Consequently, deep matching methods are typically implemented to operate on datasets comprising multiple graph pairs, each of which has an annotated ground truth alignment. Notably, in this type of problem, it is relatively straightforward to hand-label the ground truth, facilitating the creation of accurately annotated datasets.

Split-on-nodes training. In the context of learning-based algorithms (RLMs and DMMs), we define *split-on-nodes training* as the procedure whereby, given one or more graphs to be aligned, a percentage of the nodes in each network is used for training and the remainder for validation. This is the approach typically followed by RLMs: during the mapping phase, a subset of nodes with known correct alignments (usually around 20%) is used to teach the model how to translate embeddings from one vector space to another. The model is then evaluated on the remaining nodes.

Split-on-graphs training. By contrast, split-on-graphs training refers to the training regime in which one has access to a large number of graph pairs, each with its corresponding ground-truth alignment matrix. Training proceeds in the standard fashion: most graph pairs are allocated to the training set, with the remainder reserved for validation and testing. In this scenario, 100% of the nodes in each graph pair are used, because the dataset split is performed at the graph level rather than at the level of individual nodes. This approach is typical of DMMs which, originating in computer vision, benefit from large collections of small graphs extracted from images.

Drawing inspiration from the frameworks described above, in this article we present *NetAlign*, a framework for the development and evaluation of alignment algorithms that combines the strengths of both prior works. It provides a foundation upon which alignment algorithms of any category can be implemented and is compatible with various types of datasets. Moreover, it is designed to support both *split-on-nodes* and *split-on-graphs* training regimes. Full details are provided in Section 3.

3. Materials and methods

We selected a series of network alignment algorithms from the different approaches described, and then we tested them on eating disorder-related psychometric networks. We first compared the algorithms in terms of noise resistance, a robustness test in which we try to align a network with its *permuted* and *noisy* version. Then, we selected for each category (SMs, RLMs, DMMs and GAMs) the algorithm with the higher noise tolerance, and we applied them on a real application in which we tried to align psychometric networks of different age groups and disorders. Finally, we selected the algorithm that achieved the best overall performance across all categories in the real-world application and executed an additional 100 experiments using the same pairs of networks, with the objective to analyze and interpret its most frequent predictions.

Table 1
EDI-3 network metrics. From top to bottom: Early adolescents (EA), late adolescents (LA), adults (AD), anorexia nervosa (AN), bulimia nervosa (BN) and global EDI-3 network (GL).

Network	No. patients	No. nodes	No. edges	Avg degree
EA	307	57	409	14.35
LA	298	57	411	14.42
AD	300	57	416	14.59
AN	600	55	353	12.83
BN	404	61	350	11.47
GL	1206	58	426	14.69

3.1. Psychometric networks

The psychometric networks that we employed are weighted undirected graphs that establish correlations among the symptoms and personality traits exhibited by patients with eating disorders. These networks were created by Punzi et al. [6] based on data from the open access project Eating Disorder Inventory-3, available on openICPSR [22]. The project collects responses to the EDI-3 questionnaire from 1206 female patients between the ages of 11 and 75 who received treatment at an Ohio clinic between 1996 and 2015 [7]. The raw dataset did not contain any missing values, noting that the expectation-maximization algorithm was used to impute data in 407 cases where one to five responses were missing [7].

The EDI-3 is a standardized self-report questionnaire consisting of 91 questions rated on a five-point Likert scale (0–4) and organized into 12 subscales, which assess both specific eating-related factors (body dissatisfaction, bulimia, drive for thinness) and general psychological factors (self-esteem, interpersonal security, etc.).

Redundant components that depend on a common underlying attribute were identified and combined using the *Unique Variable Analysis* method, resulting in a series of unique causal components. Subsequently, a precision matrix (the inverse of a covariance matrix) among the components was estimated using GLASSO, a Gaussian Graphical Model with LASSO regularization [23]. The precision matrix was then used as a weighted adjacency matrix to construct the graph [6]. We refer to the networks thus obtained as *EDI-3 networks*.

For this work, we utilized five EDI-3 networks obtained considering different age groups and the type of diagnosed disorder: early adolescents (EA, 11–17 years), late adolescents (LA, 18–20 years), adults (AD, 25+ years), anorexia nervosa (AN), and bulimia nervosa (BN). Finally, we also used a global EDI-3 (GL) network obtained using all available data. Table 1 displays the metrics of the six networks.

Fig. 1 shows the global EDI-3 networks and its degree distribution: the employed networks do not follow a scale-free distribution. EDI-3 networks have not been created from observed interactions, but rather deduced through a statistical method. This peculiarity makes it an interesting novelty in the landscape of alignment algorithms.

3.2. NetAlign: Framework for developing and testing network alignment algorithms

NetAlign draws inspiration from the framework proposed by Trung et al. [12] for the implementation of spectral-based algorithms (IsoRank, BigAlign, FINAL) and representation learning-based algorithms (PALE, DeepLink, and IONE), as well as the management of semi-synthetic datasets. From ThinkMatch, it adopts the use of PyTorch Geometric for graph handling and neural network training, the modular design, and the use of configuration files (in YAML format) for experiments and models management. Moreover, NetAlign implements both training mode: *split-on-nodes* and *split-on-graphs*.

An overview of the framework is provided in Fig. 2. It is composed of four main modules and operates based on user-supplied input,




Fig. 1. On the left the view of the global EDI-3 network with nodes representing symptoms/personality traits, edges encoding for statistical dependencies between nodes, and node color informing about the node's degree; on the right its degree distribution.




Fig. 2. NetAlign framework overview.

including the networks to be aligned, ground truth data, and a configuration file specifying algorithm settings and parameters. These inputs are first processed by a data handling module, which prepares the data for use by the alignment algorithm. The core alignment module then uses this prepared data and the configuration details to initialize and, if applicable, train the selected model. Once alignment is complete, the resulting predictions are passed to an evaluation module, which measures performance using the provided ground truth.

Like the framework by Trung et al. NetAlign is also capable of handling semi-synthetic datasets, but unlike the former, it does so dynamically. In this case, the user selects only the source network, from which a semi-synthetic target network is generated on the fly for each experiment using the specified level and type of noise. Upon completion of the experiment, the target network is discarded, and the next experiment is performed. The permutation matrix used to permute the source network thus becomes the ground truth, which is subsequently divided into training, validation, and test sets according to the ratio specified in the configuration file (Fig. 3). The reproducibility of experiments is ensured by utilizing a seed-based approach.

Within NetAlign, we also introduce SHELLEY, a tool we specifically developed to simplify the creation of deep matching algorithms and representation learning-based methods. SHELLEY provides a modular framework, allowing the independent implementation of individual algorithm components, which can then be seamlessly integrated [24].

3.3. Selected NA algorithms

Following the approach of Trung et al. we selected IsoRank, BigAlign, and FINAL for the spectral-methods category, and PALE, DeepLink, and IONE for the representation learning methods. As a genetic algorithm-based method, we employed MAGNA++. Finally,

with regard to deep matching methods, we reimplemented StableGM and SIGMA via SHELLEY, since they are among the few DMMs that have also been evaluated on large-scale biological networks [16,17].

3.3.1. Spectral methods

IsoRank was the first global network alignment algorithm to be proposed. It is based on the hypothesis that two nodes, belonging to two different networks, are similar if their neighbors are similar as well. This hypothesis has been formulated as an eigenvalue problem and the solution is extracted using k-partite matching [9].

BigAlign converts the two networks to be aligned into bipartite graphs using the nodes' own features and hand-crafted features computed from the networks' local metrics (e.g. degree, clustering coefficient). The bipartite graph thus consists of two sets of vertices: the set of original graph nodes and the set of *feature nodes*, i.e. the nodes representing the graph's features. Edges connect vertices in the first set to those in the second set. The set of feature nodes is shared between the two networks being aligned. Consequently, BigAlign's objective is to reorder the vertex set of the first bipartite graph so that it matches the vertex set of the second bipartite graph, exploiting the shared set of feature nodes [10].

FINAL combines both topological information and attribute information based on the assumption that, given a pair of nodes u_A, v_A in network A and a pair of nodes u_B, v_B in network B, u_A is similar to u_B and v_A is similar to v_B if both u_A, v_A and u_B, v_B are close neighbors in their respective networks (*topology consistency*); and moreover u_A, u_B and v_A, v_B have similar attributes (*node attribute consistency*); and finally the edges (u_A, v_A) and (u_B, v_B) have similar attributes (*edge attribute consistency*) [11].




Fig. 3. NetAlign's semi-synthetic data loader.

3.3.2. Representation learning methods

Representation learning methods exploit the concept of *node embedding*, associating each node in the network with a vector that encodes that node's topological and attribute information.

PALE is divided into two phases: the representation phase and the mapping phase. In the *representation phase*, node embeddings for the networks to be aligned are generated separately based on the assumption that embeddings of connected nodes should be similar, while embeddings of unconnected nodes should be dissimilar. Negative sampling is performed using a skip-gram model. In the *mapping phase*, a small set of known alignments is used to map the embeddings from network A into the embedding space of network B [13].

DeepLink is similar to PALE, but extends the concept of connected nodes in the representation phase with that of *random walks*; moreover, it defines a dual mapping function to account for both directions (source-to-target and target-to-source). These mapping functions are then refined using the *Linkage Dual Learning* [14].

IONE has the same mapping phase as PALE, but extends the embedding phase by also considering the embeddings of neighboring nodes. Originally designed for directed graphs, IONE extends the concept of node embeddings by concatenating two additional vectors: the input context and the output context. The idea is that, given a node *u*, its parent nodes influence its input context, while its child nodes influence its output context [15].

3.3.3. Deep matching methods

Like representation learning methods (RLMs), deep matching methods (DMMs) exploit the concept of node embeddings. However, unlike RLMs, DMMs combine *graph neural networks* with matching solver algorithms to generate embeddings directly targeting the alignment task [25]. We focus on SIGMA and StableGM, because, at the best of our investigation, they are among the few algorithms in this category that have been tested on large-scale biological networks [16,17].

StableGM introduces a hardness attention mechanism to select negative samples for a novel contrastive loss function that is both positively and negatively informative. It employs the Sinkhorn algorithm to produce a ranking of all possible alignments, and then selects as negatives those samples that achieve a high rank relative to the true alignments. In this way, StableGM is forced to better discriminate the most ambiguous pairs by jointly leveraging positive and negative examples [16].

SIGMA, instead of learning a single prediction matrix, learns a distribution of matchings using the stochastic softmax trick. In test phase, the actual matching is then obtained through a MAP estimate of the distribution. This forces the model to explore a wide range of possible solutions [17].

3.3.4. Genetic algorithm-based methods

MAGNA++ is a genetic algorithm that defines the *population* as a set of permutation matrices (alignments) connected to form a graph in which two permutations are adjacent by one transposition. The members of the population are then paired (*parents*) using a random wheel selection process, and the *child* permutation is obtained using a *crossover function* that find the middle point node between the parents in the permutation graph. The process is guided by the *survival of the fittest* principle that maximize a measure of node and edge conservation [18].

3.4. Robustness test

The robustness test evaluates the matching accuracy of the selected algorithms in aligning the global EDI-3 network with a permuted, noise-perturbed version of itself. Noise is introduced both by dropping edges randomly (*edge removal noise*) and by adding dummy edges (*edge addition noise*). The noise level is defined as the percentage of edges that are either added or removed. The noise levels tested for both modalities are (0%, 5%, 10%, 20%, and 25%), where 0% indicates that the target network was only permuted. Since the target network is generated in a semi-synthetic fashion, the ground truth corresponds to the permutation matrix used to permute the node indices. For the purposes of this study, the global EDI-3 network was selected as the source network, and for each experiment, its semi-synthetic version was generated.

We also evaluated the DMMs using the *split-on-graphs* training mode. The semi-synthetic dataset was dynamically generated using NetAlign, with 100 semi-synthetic networks created for each noise level. Of these, 70% were used for training and 15% for validation and 15% for testing. The entire networks were used during training, and no seed nodes were provided during validation and testing. This setup was designed to assess whether the DMMs could generalize the knowledge acquired during the training phase.

3.5. Application on eating disorder networks

The algorithms, which for each category achieved the best performance in the previous test, were then used for a real-world application. In this application, we selected five psychometric networks: age11–17 (11–17 years), age18–20 (18–20 years), age25+ (25+ years), an (anorexia nervosa), and bn (bulimia nervosa). Using the selected alignment algorithms, we aligned the networks from the different age groups and from the distinct diagnostic categories. In particular:

- Two age groups with similar ages (age11–17, age18–20).
- Two age groups with differing ages (age11–17, age25+).
- Two groups with different diagnoses (an, bn).

Table 2

Best SHELLEY configurations for SIGMA and SGM on the EDI3 dataset. **Matching** refers to the alignment algorithm used; **Feat** denotes the input node features, where *Degree* indicates that each node's degree was used as its feature, while *Shared* means that all nodes were initialized with the same vector of ones. **Embedding** specifies the graph neural network employed for embedding generation.

Matching	Feat	Embedding	Num layers	Emb dim
SIGMA	Degree	GINE	2	256
SGM	Share	GINE	2	256

The learning-based algorithms require a set of seed nodes for the training phase. Due to the nature of the networks employed, we used the principle of *node name homology*: nodes with the same name are considered to be aligned. The same principle was used to measure the matching accuracy of the predicted alignments.

We emphasize that alignments based solely on the node name homology principle may not accurately reflect reality. This is because symptoms can assume different roles within the networks across different diagnoses or developmental stages. To further investigate this phenomenon, we conducted the additional analysis.

3.6. Most recurrent predictions

The algorithm that achieved the best overall performance in the previous test was selected for an additional experiment in which we run the algorithm 100 times on the same pairs of networks to analyze the consistency of its predictions by counting how many times each pair of nodes was predicted to be aligned.

3.7. Setup

The experiments were conducted using a machine equipped with an 8-core Intel i7-7700HQ 2.80 GHz processor, 16 GiB of RAM, and an Nvidia 1050 m GPU with 2 GiB of RAM.

Regarding the RLMs and the SMs, the hyperparameters have been set using the same values as Trung et al. (2020) [12]. For MAGNA++, on the other hand, we used the recommended number of generations (2000 generations) but reduced the initial population from the recommended value to accommodate psychometric networks that have fewer than 100 nodes (initial population of 2500).

Given the use of SHELLEY for the implementation of DMMs, we were also able to test various combinations of modules. To select the best configurations of modules and hyperparameters, we performed a 5-fold cross validation using the global EDI-3 network. A permuted and noised version was created (with a 10% removal probability and a 10% addition probability), and at each iteration, 1/5 of the network was used to generate the seed node set, while the remaining 4/5 were used for validation. In this manner, we preserved the typical ratio used in network alignment algorithms. The results were then averaged. The best configurations for the DMMs can be seen in Table 2.

4. Results

In this section, we present the results of the robustness tests and the real-world application in which we aligned symptom networks across different age groups and diagnostic categories. The findings reported here focus on the computational perspective and the performance of the algorithms. We have reserved the medical interpretation of the predicted alignments for Section 5.

The algorithms' performance is expressed in terms of matching accuracy, defined in Eq. (1) [12]:

$$\text{match_acc} = \frac{\#\{\text{correctly identified node pairs}\}}{\#\{\text{groundtruth node pairs}\}} \quad (1)$$

Table 3

Matching accuracy comparison of most noise resistant algorithms on different EDI3 network pairs. The ground truth has been constructed using the node name homology. The results are the average of 10 experiments, the standard deviation is in parentheses. In bold, the best results for each network pair.

	PALE	StableGM	FINAL
<i>age11-17_age18-20</i>	0.26 (0.05)	0.060 (0.023)	0.09 (0.00)
<i>age11-17_age25+</i>	0.13 (0.04)	0.052 (0.014)	0.02 (0.00)
<i>an_bn</i>	0.11 (0.03)	0.032 (0.028)	0.06 (0.00)

4.1. Robustness test

Fig. 4 presents the results of the robustness evaluation. The **representation learning methods** (PALE, DeepLink, and IONE) exhibit the highest resilience to noise in both settings, following similar accuracy trends. Among them, PALE achieves the best overall performance. In contrast, the **spectral methods** (IsoRank, BigAlign, and FINAL) maintain acceptable accuracy only in the absence of noise or under low-noise conditions, but their performance deteriorates significantly as the noise level increases. The **deep matching methods**, implemented using SHELLEY, perform well in noise-free conditions achieving 100% matching accuracy, but their robustness declines more sharply than that of the representation learning methods. Lastly, MAGNA++, the **genetic algorithm-based method**, achieves moderately high accuracy on average, but it is affected by a large standard deviation.

Fig. 5 presents the robustness test results of DMMs trained using the *split-on-graphs* training mode. Overall, the accuracy is higher compared to the *split-on-nodes* approach shown in the previous figure. Additionally, StableGM demonstrates greater resistance to edge removal noise. However, despite this improvement, its accuracy still declines substantially when compared to the leading representation learning method, PALE.

4.2. Real world application

The algorithms that, for each category, obtained the best performance in the robustness test, have been applied to the real world application. They are: PALE (RLM), StableGM (DMM) and FINAL (SM). MAGNA++ has been excluded due to its high standard deviation. Each experiment has been repeated 10 times and the results averaged.

Table 3 shows the results. In this case, PALE is the only algorithm that achieves a matching accuracy above 10%, while the other algorithms struggle to reach this level. Moreover, compared to the previous test, the performances are drastically lower. These results confirm what was already highlighted by Trung et al. namely that global network alignment algorithms require the principles of *attribute homology* and *structural homology* to be satisfied in order to function correctly. Moreover, it should be emphasized that the ground truths used to calculate the matching accuracy for each pair of networks were artificially created using the principle of *node homology*, i.e., nodes with the same names are considered aligned.

In our case study, Network Alignment is employed to examine and compare the roles that individual symptoms assume across distinct age groups and diagnostic categories within eating disorders. Specifically, we aim to identify conceptual equivalences or divergences in symptom interrelations, for example, determining whether symptoms that are central in adolescence are supplanted by different symptoms in adulthood. Consequently, in this real-world application, a matching accuracy below 100% is expected, reflecting genuine variations in symptom profiles rather than misalignment.




Fig. 4. Robustness evaluation results: matching accuracy under edge addition noise (left) and edge removal noise (right). The shaded area behind the lines represents the standard deviation computed on 10 experiments.




Fig. 5. Robustness evaluation of DMMs trained using the *split-on-graphs* training mode. Results are shown for noise introduced through edge addition (left) and edge removal (right).

4.3. Most recurrent predictions

PALE was then employed for a further investigation, in which we analyzed the most recurrent predictions for each pair of networks. We chose PALE for its performance in noise robustness tests and in real-world application.

Given a pair of networks to align, while maintaining the same set of seed nodes, we ran PALE 100 times and collected information on the alignments that were predicted most frequently.

We used the same network pairs from the real-world application, and the most recurrent alignments for each pair are presented in Table 4 (age11–17 – age18–20), Table 5 (age11–17 – age25+), and Table 6

(an – bn). In each table, in addition to the node pairs and the number of times each alignment was predicted, the reference scale according to the EDI-3 questionnaire and the description of the corresponding symptoms/personality traits are also provided. In the case of nodes with redundant symptoms (e.g., 16_25_49), we report the description only for the first (e.g., 16).

The visualization of the aligned networks, as well as the corresponding alignment matrices on 100 PALE iterations, can be seen in Fig. 6, Fig. 7 and Fig. 8. In the first row of each figure are represented the networks with the node colored by their degree, while in the second row the alignment matrix obtained by running PALE 100 times. The greener the squares, the more times the pair of nodes was predicted to align. The red squares represent the seed nodes.

Table 4
Most recurrent PALE prediction for age11–17 and age18–20.

Pair ID	age11–17 node	age18–20 node	Num. pre- dictions	EDI-3 scale	Description item 1	Description item 2
1	23_69_73	23_69_73	100/100	II	<i>I can communicate with others easily (reverse)</i>	<i>I can communicate with others easily (reverse)</i>
2	70_79_83	61	100/100	ED/B	<i>I say things impulsively that I regret having said.</i>	<i>I eat or drink in secret.</i>
3	68_86	68_86	97/100	A	<i>I would like to be in total control of my bodily urges.</i>	<i>I would like to be in total control of my bodily urges.</i>
4	34	34	96/100	II	<i>I have trouble expressing my emotions to others.</i>	<i>I have trouble expressing my emotions to others.</i>
5	67	70_79_83	95/100	ED/ED	<i>Other people would say that I am emotionally unstable.</i>	<i>I say things impulsively that I regret having said.</i>
6	66	56	94/100	A/PA	<i>I am ashamed of my human weaknesses.</i>	<i>I feel empty inside (emotionally).</i>
7	84	21_26_51_60	92/100	PA/ID	<i>I feel like I am losing out everywhere.</i>	<i>I get confused about what emotion I am feeling.</i>
8	77	40	90/100	ID/ID	<i>I can't get rid of strange ideas in my head.</i>	<i>I get confused as to whether or not I am hungry.</i>
9	56	35	89/100	PA/MF	<i>I feel empty inside (emotionally).</i>	<i>The demands of adulthood are too great.</i>
10	37	37	88/100	LSE	<i>I feel safe with myself.</i>	<i>I feel safe with myself.</i>
11	21_26_51_60	41	87/100	ID/LSE	<i>I get confused about what emotion I am feeling.</i>	<i>I have a low opinion of myself.</i>
12	11	82	86/100	DT/A	<i>I feel extremely guilty after overeating.</i>	<i>I think relaxing is just a waste of time.</i>
13	24	67	86/100	PA/ED	<i>I wish I were someone else.</i>	<i>Other people would say that I am emotionally unstable.</i>
14	40	88	86/100	ID/A	<i>I get confused as to whether or not I am hungry.</i>	<i>Suffering makes you a better person.</i>

Table 5
Most recurrent PALE prediction for age11–17 and age25+.

Pair ID	age11–17 node	age25+ node	Num. pre- dictions	EDI-3 scale	Description item 1	Description item 2
1	68_86	42	100/100	A/LSE	<i>I would like to be in total control of my bodily urges.</i>	<i>I feel that I can achieve my standards.</i>
2	76	13_43	100/100	IA/P	<i>People understand my real problems.</i>	<i>Only outstanding performance is acceptable in my family.</i>
3	24	61	99/100	PA/B	<i>I wish I were someone else.</i>	<i>I eat or drink in secret.</i>
4	2_12_19	2_12_19	99/100	BD/BD	<i>I think that my stomach is too big.</i>	<i>I think that my stomach is too big.</i>
5	61	50	97/100	B/LSE	<i>I eat or drink in secret.</i>	<i>I feel that I am a valuable person.</i>
6	17	35	96/100	IA/MF	<i>I trust others.</i>	<i>The demands of adulthood are too great.</i>
7	23_69_73	37	95/100	II/LSE	<i>I can communicate with others easily.</i>	<i>I feel safe with myself.</i>
8	70_79_83	67	92/100	ED/ED	<i>I say things impulsively that I regret having said.</i>	<i>Other people would say that I am emotionally unstable.</i>
9	11	71	87/100	DT/-	<i>I feel extremely guilty after overeating.</i>	<i>I go out of my way to experience pleasure.</i>
10	35	21_26_51_60	83/100	MF/ID	<i>The demands of adulthood are too great.</i>	<i>I get confused about what emotion I am feeling.</i>
11	29_36_52_63	84	81/100	P/PA	<i>As a child, I tried very hard to avoid disappointing my parents and teachers.</i>	<i>I feel like I am losing out everywhere.</i>
12	21_26_51_60	33	81/100	ID/ID	<i>get confused about what emotion I am feeling.</i>	<i>I don't know what's going on inside me.</i>
13	10_27	1	78/100	LSE/DT	<i>I feel ineffective as a person.</i>	<i>I eat sweets and carbohydrates without feeling nervous.</i>
14	71	89	75/100	-/IA	<i>I go out of my way to experience pleasure.</i>	<i>I know there are people who love me.</i>

Table 6
Most recurrent PALE prediction for anorexia nervosa and bulimia nervosa.

Pair ID	AN node	BN node	Num. pre- dictions	EDI-3 scale	Description item 1	Description item 2
1	3_6_14_48	82	100/100	MF/A	<i>I wish that I could return to the security of childhood.</i>	<i>I think relaxing is just a waste of time.</i>
2	80	65	100/100	PA/IA	<i>I feel that people give me the credit I deserve.</i>	<i>People I really like end up disappointing me.</i>
3	53	37	99/100	B/LSE	<i>I have the thought of trying to vomit in order to lose weight.</i>	<i>I feel safe with myself.</i>
4	8_44	74_75	98/100	ID/IA,A	<i>I get frightened when my feelings are too strong.</i>	<i>I feel trapped in relationships.</i>
5	17	80	93/100	IA/PA	<i>I trust others.</i>	<i>I feel that people give me the credit I deserve.</i>
6	65	35	90/100	IA/MF	<i>People I really like end up disappointing me.</i>	<i>The demands of adulthood are too great.</i>
7	29_36_52_63	87	88/100	P/II	<i>As a child, I tried very hard to avoid disappointing my parents and teachers.</i>	<i>I would rather spend time by myself than with others.</i>
8	54	24	88/100	IA/PA	<i>I need to keep people at a certain distance (feel uncomfortable if someone tries to get too close).</i>	<i>I wish I were someone else.</i>
9	20	30	84/100	PA/IA	<i>I feel generally in control of things in my life.</i>	<i>I have close relationships.</i>
10	42	50	84/100	LSE/LSE	<i>I feel that I can achieve my standards.</i>	<i>I feel that I am a valuable person.</i>
11	47	3_6_14_48	84/100	BD/MF	<i>I feel bloated after eating a normal meal.</i>	<i>I wish that I could return to the security of childhood.</i>
12	69_73	17	83/100	II/IA	<i>I feel relaxed in most group situations.</i>	<i>I trust others.</i>
13	23	47	81/100	II/BD	<i>I can communicate with others easily.</i>	<i>I feel bloated after eating a normal meal.</i>
14	90	90	80/100	ED/ED	<i>I feel like I must hurt myself or others.</i>	<i>I feel like I must hurt myself or others.</i>




Fig. 6. PALE most recurrent alignments between age11–17 and age18–20. The red squares in the alignment matrix represent the seed nodes.

In Section 5 we provide an interpretation of these alignments.

5. Discussion

In light of the results of the current study, our most notable finding is that network alignment has proven to be a powerful tool for

applications to psychometric networks, in order to compare symptom structures across different developmental stages or clinical populations. In this context, PALE, a representation learning–based algorithm, stands out as the most robust and reliable algorithm, showing superior performance across varying noise conditions. Unlike spectral and deep matching methods, whose accuracy drops significantly with increasing




Fig. 7. PALE most recurrent alignments between age11–17 and age25+. The red squares in the alignment matrix represent the seed nodes.




Fig. 8. PALE most recurrent alignments between the anorexia and bulimia networks. The red squares in the alignment matrix represent the seed nodes. Note that since the alignment matrix is not square, the main diagonal does not perfectly bisect the grid cells.

noise, PALE maintains consistent and high alignment accuracy. It is notably the only method to achieve over 10% accuracy under high-noise scenarios, highlighting its resilience where others fail. Furthermore, our results show that PALE consistently outperformed others in real-world applications. Indeed, in our application to psychometric networks derived from the Eating Disorder Inventory-3 (EDI-3) dataset, among the noteworthy findings, PALE produced consistent alignment results for specific nodes – particularly those related to interpersonal difficulties and rigid control over bodily impulses – across early and late adolescent patients with eating disorders. This may indicate developmental stability in the role these symptoms play within the network structure. These results are consistent with previous research highlighting the key role of interpersonal problems and ascetic tendencies in the development and maintenance of eating disorders across adolescent stages [26–28]. PALE also revealed consistent alignment results between early

adolescence and adulthood for nodes related to body-related concerns (2_12_19: “I think that my stomach is too big”). This pattern may reflect the persistence in the role of these symptoms in the psycho-pathological functioning of individuals with eating disorders across developmental stages. This is consistent with longitudinal research showing that body image concerns – particularly dissatisfaction with specific body parts – tend to persist from adolescence into adulthood [29].

However, not all symptoms demonstrated such stability in their role in the network across developmental stages. For example, the role fulfilled by the node 70_79_83 (“I say things impulsively that I regret having said”) in the early adolescence network was instead taken on by 61 (“I eat or drink in secret”) in the late adolescence network. These patterns may reflect age-related differences in how specific symptoms are integrated within the broader network of functioning in adolescents with eating disorders. In particular, this finding may suggest that,

while in early adolescence this role within the network is expressed through overt verbal behavior, in older adolescents it is assumed by more covert manifestations, such as secretive eating. This interpretation aligns with the idea that impulsive traits can manifest differently across developmental stages [30].

About differences between early adolescents and adults, the node 68_86 (*"I would like to be in total control of my bodily urges"*) in early adolescence and the node 42 (*"I feel that I can achieve my standards"*) in adulthood, for instance, occupied divergent positions within their respective networks. This pattern suggests that although both nodes are thematically linked to control, they do not play the same structural role across developmental stages. Rather than expressing the same symptomatology in identical ways, individuals at different ages may organize and experience control-related concerns through distinct symptom pathways. Specifically, in early adolescence, the role of control within the network appears to be articulated through the management of bodily impulses, while in adulthood, it is represented in relation to personal standards and achievement.

Similarly, the nodes 76 (*"People understand my real problems"*, reverse-coded) and 13_43 (*"Only outstanding performance is acceptable in my family"*) assumed different roles within the symptom networks of early adolescents and adults. In early adolescence, the functional role within the network appears to be fulfilled by experiences of being misunderstood, reflecting a heightened salience emotional validation during this developmental period. In contrast, in adulthood, a similar structural role is occupied by perceptions of familial pressure for perfection, pointing to an increased internalization of external expectations. These differences are consistent with developmental models emphasizing the evolving nature of social needs and self-definition processes across adolescence and adulthood [31,32].

It is worth noting that only a very limited number of nodes demonstrated consistent alignment across diagnostic groups. In fact, only one item exhibited consistent alignment over 80% of the time, whereas all others fell below the 40% threshold. This may suggest that, although anorexia nervosa and bulimia nervosa share several clinical features, the underlying psychological mechanisms differ substantially between the two disorders [33]. The only node that held a comparable role in the symptom networks of both anorexia nervosa (AN) and bulimia nervosa (BN) was 90 (*"I feel like I must hurt myself or others"*). This alignment may suggest that self-aggressive impulses represent a transdiagnostic feature cutting across different eating disorder presentations.

Among the nodes that assumed divergent roles across the diagnostic networks, for instance, the item 3_6_14_48 (*"I wish that I could return to the security of childhood"*) in the AN network corresponded functionally to 82 (*"I think relaxing is just a waste of time"*) in the BN network. Although both items may reflect difficulties in managing emotional vulnerability, they seem to manifest through distinct pathways: in AN, through a longing for regression and protection, and in BN, through overactivation and hyperproductivity. This may suggest that patients with AN may cope with emotional stress by withdrawing and idealizing past security, whereas those with BN may channel discomfort into action and control-oriented behaviors. These distinctions align with previous observations that individuals with AN often exhibit avoidant and constricted emotional styles, while those with BN show more impulsive and dysregulated coping strategies [34,35]. Similarly, the items 80 (*"I feel that people give me the credit I deserve"*) (reverse-coded) and 65 (*"People I really like end up disappointing me"*) did not occupy equivalent positions in the networks of AN and BN patients. This finding may indicate that, while in AN this role within the network is expressed through a need for external validation and a tendency to minimize one's own worth, in BN it is assumed by feelings of relational disappointment. This pattern highlights how similar interpersonal concerns may take on distinct forms within the symptom architecture of each disorder [32].

These findings may carry practical implications for clinical intervention. The stable alignment of symptoms related to interpersonal difficulties and rigid control over bodily impulses across early and late

adolescence suggests that these symptoms may warrant particular clinical attention throughout the adolescent developmental period. Similarly, the persistence of body image concerns from adolescence into adulthood points to the potential importance of maintaining focused interventions on these symptoms across the continuum of care. The shifting role of impulsivity-related symptoms across developmental stages highlights the need for developmentally sensitive assessments, as the same underlying vulnerability may manifest differently depending on age. Finally, the transdiagnostic stability of self-aggressive impulses (item 90) across AN and BN suggests that this symptom may represent a clinically meaningful target for cross-diagnostic treatment approaches.

6. Conclusions

6.1. General considerations

In this study, we propose for the first time the use of network alignment in network psychometrics. These networks are not created from observed interactions, but rather deduced through a statistical method: this peculiarity makes it an interesting novelty in the landscape of alignment algorithms. Our results show that PALE, a representation-learning-based algorithm, demonstrated the highest robustness and stability under noisy conditions and consistently outperformed others in real-world applications. These findings support the idea that successful global network alignment relies on both structural and attribute homology, which PALE effectively captures. In our clinical case study, PALE revealed stable alignment of nodes related to interpersonal difficulties and self-control issues in early and late adolescents with eating disorders, suggesting that these symptoms consistently play comparable functional roles within the overall network structure over time. It also identified persistent alignment of body-related concerns, such as dissatisfaction with specific body parts, from adolescence into adulthood. These results align with longitudinal evidence on the stability of body image issues in disordered eating. By uncovering these stable, clinically meaningful patterns, PALE demonstrates a potential for advancing network-based psychopathology research.

In general, the performances observed in this case study suggest that NA can be considered another suitable tool among the methods available to psychologists for comparing and studying psychometric symptom networks.

6.2. Limitations and future work

Global network alignment remains an open and challenging problem. Even PALE, which in our tests demonstrated the best performance both in terms of noise reduction and real-world applications, has limitations that are important to highlight.

PALE, being a representation learning method, requires a set of known alignments to map embeddings from the vector space of one network to another. While this set is easily obtainable in robustness tests, the same is not true for real-world applications, where alignments between two networks must be manually verified.

It should be noted that in the real-world application test, we created the set of known alignments ourselves, based on the node names homology hypothesis, that is the assumption that nodes with the same name are likely to play the same role in their respective networks. However, it is possible that different symptoms play different roles depending on the age or diagnosis of the patients analyzed.

A substantive limitation of the present study is that the seed alignments used in the real-world application were defined according to a node-name homology criterion, namely, the assumption that nodes sharing the same label across networks are likely to play the same functional role. As already noted in the "Application on Eating Disorder Networks" subsection, alignments based solely on node-name homology may not accurately reflect reality, because the same symptom can occupy different structural positions depending on age, illness stage,

or diagnostic group. In this first study, we adopted this strategy as a pragmatic and reproducible way to define a minimal set of anchors for benchmarking alignment methods and constructing an approximate ground truth. Nevertheless, this choice may bias the resulting alignments and should therefore be interpreted with caution.

A natural extension of this work is to move toward weakly supervised or fully unsupervised seed selection strategies. One possible refinement, already suggested above, is to combine semantic and structural information by retaining as seeds only nodes that share the same label and exhibit similar centrality profiles across networks. More ambitiously, future work should investigate seed-free pipelines in which node embeddings are first learned independently in each network, then aligned through an unpaired space-alignment procedure, for example by adapting recent universal-geometry approaches for cross-space translation [36], and finally used to infer candidate correspondences solely from structural similarity. Such a procedure could also be iterated, re-estimating embeddings on the basis of progressively refined matches. Systematically evaluating these unsupervised or weakly supervised variants will be an important step toward reducing dependence on manually specified correspondences and improving the ecological validity of network alignment in psychopathology.

For such a reason, it is important to refine the criteria for selecting seed nodes. One possible approach could be to use node centrality measures to find the one that varies most between the two networks, and combine it with the node names hypothesis. In this way, one could select as known alignments those nodes that have the same name and, at the same time, a similar centrality measure.

A further limitation concerns the use of global network alignment, which imposes a one-to-one correspondence between nodes in the compared networks. This formulation follows the standard setting adopted in the network alignment literature, especially in biology and connectomics, where the objective is typically to identify an approximate bijection between proteins, brain regions, or other network elements [4, 5]. In the present case, this choice is also supported by previous findings suggesting that at least some symptom analogies may persist across developmental stages in eating disorder networks. In particular, Schlegl et al. reported high correlations in network structure across adolescent and adult patients with anorexia nervosa and bulimia nervosa, while Christian et al. showed that several symptoms retain a central role across multiple developmental stages, despite broader age-related differences in network organization [37,38]. Nevertheless, these findings do not imply that a strict one-to-one functional counterpart exists for every symptom. For this reason, the one-to-one constraint should not be interpreted as a claim that symptom equivalence always holds across developmental stages or diagnoses, but rather as a methodological device that enables a systematic comparison between networks. By enforcing a complete matching, the alignment process makes it possible to identify which symptoms are consistently paired across networks and which are instead associated with approximate substitutes, thereby generating hypotheses about stable and shifting symptom roles. At the same time, more flexible formulations may prove more appropriate in contexts where symptom functions reorganize substantially. Future work should therefore investigate local or partial alignment strategies, as well as many-to-one and many-to-many correspondences, in order to better capture non-bijective relationships between symptoms and provide a more realistic account of developmental and diagnostic heterogeneity.

In this perspective, it is worth noting that Jha and coworkers [36] have recently proposed a method to translate text embeddings to and from a universal latent representation, with the goal of translating embeddings from one vector space to another without any paired data, encoders, or predefined sets of matches. Adapting this method for node embeddings could replace the supervised mapping phase of RLMS, thereby eliminating the need to define a set of known alignments.

CRediT authorship contribution statement

Riccardo De Luca: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Data curation. **Laura Girelli:** Writing – review & editing, Writing – original draft, Supervision, Investigation, Conceptualization. **Pietro H. Guzzi:** Writing – review & editing, Supervision, Methodology, Investigation. **Paolo Tieri:** Writing – review & editing, Supervision, Methodology, Investigation, Conceptualization. **Manuela Petti:** Writing – review & editing, Supervision, Methodology, Investigation, Funding acquisition, Conceptualization.

Ethics statement

This study did not require ethical approval as it exclusively used publicly available datasets that are anonymized and were collected with appropriate ethical oversight by the original data providers. No new data were collected, and no interaction with human subjects occurred in the course of this research.

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

We thank Clara Punzi for providing the psychometric networks used in this study and for her valuable scientific support.

The research was supported by Sapienza University of Rome, Progetti di Ateneo 2024 (RM12419113174547).

Data availability

The dataset analyzed during the current study is available in the openICPSR repository.

References

[1] D.J. Robinaugh, R.H. Hoekstra, E.R. Toner, D. Borsboom, The network approach to psychopathology: A review of the literature 2008–2018 and an agenda for future research, *Psychol. Med.* 50 (3) (2020) 353–366.

[2] C. Ramos-Vera, V.R. Franco, J.V. Saldarriaga, A.S. Barrientos, Psychometric networks and their implications for the treatment and diagnosis of psychopathologies, in: S. Misciagna (Ed.), *Psychometrics*, IntechOpen, Rijeka, 2022, <http://dx.doi.org/10.5772/intechopen.105404>.

[3] C.D. van Borkulo, R. van Bork, L. Boschloo, J.J. Kossakowski, P. Tio, R.A. Schoevers, D. Borsboom, L.J. Waldorp, Comparing network structures on three aspects: A permutation test, *Psychol. Methods* 28 (6) (2023) 1273–1285.

[4] C.-Y. Ma, C.-S. Liao, A review of protein–protein interaction network alignment: From pathway comparison to global alignment, *Comput. Struct. Biotechnol. J.* 18 (2020) 2647–2656.

[5] M. Milano, P.H. Guzzi, O. Tymofieva, D. Xu, C. Hess, P. Veltri, M. Cannataro, An extensive assessment of network alignment algorithms for comparison of brain connectomes, *BMC Bioinformatics* 18 (Suppl 6) (2017) 235.

[6] C. Punzi, P. Tieri, L. Girelli, M. Petti, Network-based validation of the psychometric questionnaire EDI-3 for the assessment of eating disorders, *Sci. Rep.* 13 (1) (2023) 1578.

[7] J. Brookings, Eating disorder inventory-3, American clinical cases, 2020, <http://dx.doi.org/10.3886/E109443V2>.

[8] D. Garner, Eating disorder inventory-3 (EDI-3) professional manual. Lutz, FL: Psychological assessment resources, *Int. J. Eat. Disord.* 35 (2004) 478–479.

[9] R. Singh, J. Xu, B. Berger, Global alignment of multiple protein interaction networks with application to functional orthology detection, *Proc. Natl. Acad. Sci.* 105 (35) (2008) 12763–12768.

[10] D. Koutra, H. Tong, D. Lubensky, Big-align: Fast bipartite graph alignment, in: 2013 IEEE 13th International Conference on Data Mining, IEEE, 2013, pp. 389–398.

[11] S. Zhang, H. Tong, Final: Fast attributed network alignment, in: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2016, pp. 1345–1354.

[12] H.T. Trung, N.T. Toan, T. Van Vinh, H.T. Dat, D.C. Thang, N.Q.V. Hung, A. Sattar, A comparative study on network alignment techniques, *Expert Syst. Appl.* 140 (2020) 112883.

[13] T. Man, H. Shen, S. Liu, X. Jin, X. Cheng, Predict anchor links across social networks via an embedding approach, in: *Ijcai*, vol. 16, 2016, pp. 1823–1829.

[14] F. Zhou, L. Liu, K. Zhang, G. Trajcevski, J. Wu, T. Zhong, Deeplink: A deep learning approach for user identity linkage, in: IEEE INFOCOM 2018-IEEE Conference on Computer Communications, IEEE, 2018, pp. 1313–1321.

[15] L. Liu, W.K. Cheung, X. Li, L. Liao, Aligning users across social networks using network embedding, in: *Ijcai*, vol. 16, 2016, pp. 1774–1780.

[16] G. Ratnayaka, Q. Wang, Y. Li, Contrastive learning for supervised graph matching, in: *Uncertainty in Artificial Intelligence*, PMLR, 2023, pp. 1718–1729.

[17] L. Liu, M.C. Hughes, S. Hassoun, L. Liu, Stochastic iterative graph matching, in: *International Conference on Machine Learning*, PMLR, 2021, pp. 6815–6825.

[18] V. Vijayan, V. Saraph, T. Milenković, Magna++: Maximizing accuracy in global network alignment via both node and edge conservation, *Bioinformatics* 31 (14) (2015) 2409–2411, <http://dx.doi.org/10.1093/bioinformatics/btv161>, [arXiv:https://academic.oup.com/bioinformatics/article-pdf/31/14/2409/49034921/bioinformatics_31_14_2409.pdf](https://academic.oup.com/bioinformatics/article-pdf/31/14/2409/49034921/bioinformatics_31_14_2409.pdf).

[19] R. Wang, Z. Guo, J. Yan, ThinkMatch, 2019, MulanPSL-2.0 license, <https://github.com/Thinklab-SJTU/ThinkMatch>.

[20] M. Fey, J.E. Lenssen, Fast graph representation learning with PyTorch geometric, 2019, arXiv preprint [arXiv:1903.02428](https://arxiv.org/abs/1903.02428).

[21] R. Wang, Z. Guo, W. Pan, J. Ma, Y. Zhang, N. Yang, Q. Liu, L. Wei, H. Zhang, C. Liu, Z. Jiang, X. Yang, J. Yan, Pygmtools: A python graph matching toolkit, *J. Mach. Learn. Res.* 25 (33) (2024) 1–7, URL <https://jmlr.org/papers/v25/23-0572.html>.

[22] J. Brookings, Eating disorder inventory-3, American clinical cases, 2020, <http://dx.doi.org/10.3886/E109443V2>, URL <https://www.openicpsr.org/openicpsr/project/109443/version/V2/view>.

[23] J. Friedman, T. Hastie, R. Tibshirani, Sparse inverse covariance estimation with the graphical lasso, *Biostatistics* 9 (3) (2007) 432–441, <http://dx.doi.org/10.1093/biostatistics/kxm045>, [arXiv:https://academic.oup.com/biostatistics/article-pdf/9/3/432/57071116/biostatistics_9_3_432.pdf](https://academic.oup.com/biostatistics/article-pdf/9/3/432/57071116/biostatistics_9_3_432.pdf).

[24] R. De Luca, M. Petti, P.H. Guzzi, P. Tieri, SHELLEY: Exploring learning-based network alignment on biological data, in: 2024 IEEE International Conference on Bioinformatics and Biomedicine, BIBM, IEEE, 2024, pp. 6943–6950.

[25] Y. Lin, M. Yang, J. Yu, P. Hu, C. Zhang, X. Peng, Graph matching with bi-level noisy correspondence, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2023, pp. 23362–23371.

[26] M. Carfagno, E. Barone, E. Arsenio, R. Bello, L. Marone, A. Volpicelli, G. Cascino, A.M. Monteleone, Mediation role of interpersonal problems between insecure attachment and eating disorder psychopathology, *Eat. Weight. Disorders-Studies Anorex. Bulim. Obes.* 29 (1) (2024) 43.

[27] S. Fassino, A. Pierò, C. Gramaglia, G.A. Daga, M. Gandione, G.G. Rovera, G. Bartocci, Clinical, psychological, and personality correlates of asceticism in anorexia nervosa: from saint anorexia to pathologic perfectionism, *Transcult. Psychiatry* 43 (4) (2006) 600–614.

[28] A. Hartmann, A. Zeeck, M.S. Barrett, Interpersonal problems in eating disorders, *Int. J. Eat. Disord.* 43 (7) (2010) 619–627.

[29] P.K. Keel, M.G. Baxter, T.F. Heatherton, T.E. Joiner Jr., A 20-year longitudinal study of body weight, dieting, and eating disorder symptoms, *J. Abnorm. Psychol.* 116 (2) (2007) 422.

[30] D. Romer, Adolescent risk taking, impulsivity, and brain development: Implications for prevention, *Dev. Psychobiol.: J. Int. Soc. Dev. Psychobiol.* 52 (3) (2010) 263–276.

[31] L. Steinberg, A.S. Morris, Adolescent development, *Annu. Rev. Psychol.* 52 (1) (2001) 83–110.

[32] S. Barakat, S.A. McLean, E. Bryant, A. Le, P. Marks, S. Touyz, S. Maguire, Risk factors for eating disorders: findings from a rapid review, *J. Eat. Disord.* 11 (1) (2023) 8.

[33] J. Treasure, T.A. Duarte, U. Schmidt, Eating disorders, *Lancet* 395 (10227) (2020) 899–911, [http://dx.doi.org/10.1016/S0140-6736\(20\)30059-3](http://dx.doi.org/10.1016/S0140-6736(20)30059-3).

[34] L. Claes, C. Nederkoorn, W. Vandereycken, R. Guerrieri, H. Vertommen, Impulsiveness and lack of inhibitory control in eating disorders, *Eat. Behav.* 7 (3) (2006) 196–203.

[35] J. Treasure, D. Stein, S. Maguire, Has the time come for a staging model to map the course of eating disorders from high risk to severe enduring illness? An examination of the evidence, *Early Interv. Psychiatry* 9 (3) (2015) 173–184.

[36] R. Jha, C. Zhang, V. Shmatikov, J.X. Morris, Harnessing the universal geometry of embeddings, 2025, arXiv preprint [arXiv:2505.12540](https://arxiv.org/abs/2505.12540).

[37] S. Schlegl, K.E. Smith, L. Vierl, R.D. Crosby, M. Moessner, C. Neumayr, U. Voderholzer, Using network analysis to compare diagnosis-specific and age-specific symptom networks in eating disorders, *Int. J. Eat. Disord.* 54 (8) (2021) 1463–1476.

[38] C. Christian, V.L. Perko, I.A. Vanzhula, J.P. Tregarthen, K.T. Forbush, C.A. Levinson, Eating disorder core symptoms and symptom pathways across developmental stages: A network analysis, *J. Abnorm. Psychol.* 129 (2) (2020) 177.