technologies

MDPI

Article

Analysis of Multidimensional Clinical and Physiological Data with Synolitical Graph Neural Networks

Mikhail Krivonosov [1], Tatiana Nazarenko [2], Vadim Ushakov [3], Daniil Vlasenko [3], Denis Zakharov [3], Shangbin Chen [4], Oleg Blyus [5] and Alexey Zaikin [2,*]

[1] Laboratory of Systems Medicine of Ageing, Centre for Artificial Intelligence, Department of Applied Mathematics, Lobachevsky University, Nizhny Novgorod 603022, Russia; krivonosov@itmm.unn.ru
[2] Department of Mathematics, Institute for Women's Health, University College London, London WC1H 0AY, UK; t.nazarenko@ucl.ac.uk
[3] Institute for Cognitive Neuroscience, University Higher School of Economics, 20 Myasnitskaya, Moscow 101000, Russia; uwadim@gmail.com (V.U.); vlasenko.daniil.vl@gmail.com (D.V.); dzakh76@gmail.com (D.Z.)
[4] Britton Chance Center for Biomedical Photonics, Wuhan National Laboratory for Optoelectronics-Huazhong University of Science and Technology, Wuhan 430074, China; sbchen@hust.edu.cn
[5] Wolfson Institute of Population Health, Queen Mary University of London, London EC1M 6BQ, UK; o.blyuss@qmul.ac.uk
* Correspondence: alexey.zaikin@ucl.ac.uk

Abstract: This paper introduces a novel approach for classifying multidimensional physiological and clinical data using Synolitic Graph Neural Networks (SGNNs). SGNNs are particularly good for addressing the challenges posed by high-dimensional datasets, particularly in healthcare, where traditional machine learning and Artificial Intelligence methods often struggle to find global optima due to the "curse of dimensionality". To apply Geometric Deep Learning we propose a synolitic or ensemble graph representation of the data, a universal method that transforms any multidimensional dataset into a network, utilising only class labels from training data. The paper demonstrates the effectiveness of this approach through two classification tasks: synthetic and fMRI data from cognitive tasks. Convolutional Graph Neural Network architecture is then applied, and the results are compared with established machine learning algorithms. The findings highlight the robustness and interpretability of SGNNs in solving complex, high-dimensional classification problems.

Keywords: networks; data analysis; graph neural networks

check for updates

Academic Editors: Mohammed Mahmoud and Pedro Antonio Gutiérrez

Received: 7 October 2024
Revised: 11 December 2024
Accepted: 22 December 2024
Published: 28 December 2024

Citation: Krivonosov, M.; Nazarenko, T.; Ushakov, V.; Vlasenko, D.; Zakharov, D.; Chen, S.; Blyus, O.; Zaikin, A. Analysis of Multidimensional Clinical and Physiological Data with Synolitical Graph Neural Networks. *Technologies* **2025**, *13*, 13. https://doi.org/10.3390/technologies13010013

Copyright: © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

The human body is a super-network. In the human brain alone, we have 86 billion neurons, accompanied by a similar number of glial cells, all packed into a highly complex multi-layered network structure [1]. Moreover, every organ also functions as a network, often embodying a form of local intelligence. These networks interact in intricate ways, leading to a new approach known as Network Physiology [2]. Fortunately, recent advancements have given scientists access to vast amounts of physiological and clinical data that characterise the current state of the human body. In fact, the amount of accumulated and partially available data has surpassed our ability to analyse them effectively. These data include genomic, epigenetic, chromatin assembly structure, metabolomic and proteomic datasets, EEG and ECG, as well as imaging data, such as MRI, fMRI or MEG data, to mention just a few. In all these cases, the datasets are highly multidimensional and often

serial, describing both the current state of a human body and its evolution over time, along with the dependencies between parameters of these multidimensional data.

Recently, we have gained access to the power of Artificial Intelligence (AI) algorithms, including Deep Neural Networks, Convolutional Neural Networks, Generative Adversarial Networks, Autoencoders, and, of course, Transformers, along with ChatGPT and similar algorithms from other providers. However, when it comes to patient-related data, even with the full capabilities of AI, we struggle to find the global optimum in optimisation tasks due to the "curse of dimensionality" resulting from the very high dimensionality of the data. In clinical practice and healthcare, obtaining a sufficient amount of data is challenging due to its high cost and organisational issues. Consequently, in the complex landscape of high-dimensional parameter spaces, the application of AI algorithms may lead to finding local optima instead of global ones, which can also be unstable as more data become available.

To address this problem, we observe the emergence of algorithms that lie at the intersection of AI and classical mathematics. In particular, algorithms using the internal structure of the data, ranging from manifold learning to Graph Neural Networks, have been developed and classified under the broader category of Geometric Deep Learning (GDL) [3]. There is a significant potential in utilising the internal structure of the data in order to overcome the curse of dimensionality and optimise the search for solutions or information processing. Notable advances in the application of GDL and, particularly, its subclass, Graph Neural Networks (GNNs) have been observed. GNNs are specialised AI algorithms, designed to work with data that can be represented in the form of a network or, mathematically, a graph. These algorithms are especially effective for processing data represented in the form of a graph because they utilise the structure of this graph in the architecture of the learning model. Various architectures of GNNs, including Message Passing Neural Networks [4], Graph Neural Networks with Attention [5], and its simplification, Convolutional Graph Networks [6], have been employed in multiple applications [7], demonstrating remarkable abilities to detect hidden topological changes and interdependencies in multi-dimensional parameter space. To achieve this, Deep Learning algorithms utilise the internal network structure of the data. However, the challenge remains: how can we embed multi-dimensional data into a graph representation, if a priori no links are known? Several approaches have been used, such as graphs of functional connectivity [8], correlation graphs [9] or the identification of correlation network markers using internal structure of nodes [10]. A more universal approach has been developed by us with synolitic or ensemble graphs, named after the Greek word "σύνολο" standing for "ensemble" [11]. These graphs can accommodate any kind of multi-dimensional data, relying solely on class labels in the training dataset [12–15].

Network-based approaches have become increasingly popular in the fields of biology and medicine, particularly in the emerging discipline of Network Physiology. This field seeks to understand how physiological systems interact as networks of networks, revealing dynamic interdependencies across different organ systems. Pioneering works such as those by Bashan et al. [16] and Ivanov et al. [17,18] have demonstrated how network topology can provide insights into physiological function, health and disease states. For instance, the mapping of dynamic interactions between networks of physiological systems has provided critical knowledge about homeostasis, resilience and systemic coordination in the human body. Inspired by these advances, we propose Synolitical Graph Neural Networks (SGNNs) as a novel approach to analysing multidimensional physiological data. Unlike conventional methods, SGNNs construct ensemble graphs based solely on class labels, offering a universal and interpretable framework for studying complex physiological interactions. This method aligns closely with the goals of Network Physiology,

as it emphasises relationships within high-dimensional data and leverages graph-based representations to capture latent interactions that might otherwise remain obscured. The ability of SGNNs to generalise across diverse datasets positions this methodology as a promising tool for investigating dynamic physiological networks and their implications for health and disease.

In this paper we show how a synolitic network representation can be coupled with Graph Neural Network analysis to classify multi-dimensional data, and we illustrate this approach by two examples of classification task solutions, to classify specially designed synthetic data fMRI data collected during a cognitive task. The paper is structured as follows. First, we describe the data and methodology, two different algorithms to apply a synolitic network approach and the architecture of the Convolutional Graph Neural Network applied (CGNN). These two methods are combined together to organise Synolitic Graph Neural Networks (SGNNs). Then we present the results of the classification, comparing it with more established machine learning algorithms and summarise the results in a Discussion chapter. Finally, we highlight why this approach excels in ensuring result interpretability and offers a notably robust methodology.

2. Material and Methods

Synthetic Data (as generated in [11]). We have chosen spherical boundaries for two classes to make classification difficult for traditional ML methods. For all modelled spheres, we considered a variety of configurations, including different dimensions (2, 3, 10, 30, 60, 90, 120, 150), as well as varying sample sizes for the Cases and Controls. Specifically, we used the following sample counts: for Cases in the training set, we considered 15, 65, 115, 165, 215 and 265 samples; for Controls in the training set, we used the same counts. The number of Case and Control samples for the test set were calculated as 25% of their corresponding training counts. The data include the following datasets:

Ideal Spheres Model. In the ideal spheres model, the data are represented within the bounds of an N-dimensional sphere with a radius of 1. Each sample corresponds to a vector that describes its position in space. Controls are defined as points with a radius between 0.01 and 0.5, while Cases are represented by points with a radius from 0.5 to 1.

Noisy Spheres Model. To create the noisy spheres model, we introduced 50 noise variables to each ideal sphere. Each sample now consists of its original coordinates, along with these additional random variables. The noise variables are drawn from a uniform distribution within the range of -1 to 1, which simulates real-world data imperfections and variability.

Broken Spheres Model. In the broken spheres model, we modified the ideal sphere by retaining only half of the original variables, while replacing the other half with random values. Each sample is thus represented by a combination of the retained coordinates and new random variables, again sourced from a uniform distribution between -1 and 1. This approach allows us to examine how the absence of key parameters impacts model performance and classification efficacy.

fMRI data from cognitive experiment. The method was tested on the research data [19], where fMRI data were recorded while subjects viewed images of objects (observation experiment) or imagined objects with their eyes closed (imagination experiment). In the observation experiment, 1200 images from 150 object categories were used (8 images per category). Each image was shown to the subject once. Each subject underwent 24 fMRI scanning cycles. All images were taken from ImageNet (http://www.image-net.org, Fall 2011 release), a large-scale image database. During the observation experiment, subjects performed the task of recalling the images in reverse order (5 attempts per cycle), with 55 images displayed per cycle. In the imagination experiment, subjects were asked to

visually imagine a sequence of 25 objects. Each object belonged to one of 50 categories. Each subject underwent 20 fMRI scanning cycles. The original voxel size in the fMRI data was $3 \times 3 \times 3$ mm^3. To reduce computational costs and smooth individual brain structure differences, the fMRI data resolution was reduced to $10 \times 10 \times 10$ mm^3.

Synolitic representation. We have utilised two approaches to represent the data in the form of a synolitic graph (SG), solely based on its labelling. The first one is more appropriate for unstructured high-dimensional data and the second one for imaging data. A synolitic network is a network where both groups contribute to defining normal and abnormal states. An application of the first approach to construct an SG works as follows. In the feature space formed by any two dimensions, a radial SVM is employed to establish the optimal boundary that separates the classes (see Figure 1). It is important that instead of a radial SVM kernel, any ML algorithm can be used to find a boundary between two classes on the plane of two features. Consequently, each point in this model is assigned a probability value indicating its likelihood of belonging to each class. For every new sample, the edge weight is calculated based on the probability of it being part of the Cases group.

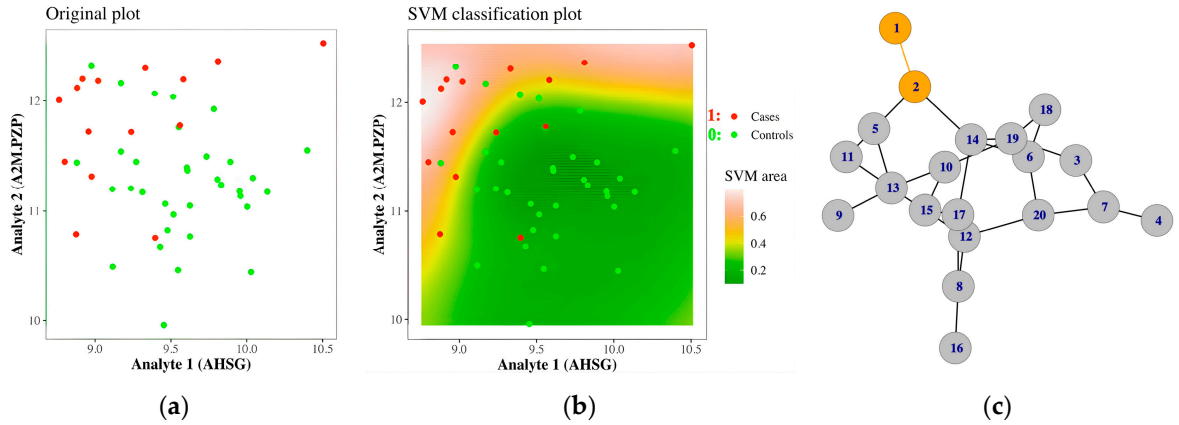


Figure 1. Representation of high-dimensional data as synolitic, or ensemble, graph. (**a**) For all pairs of the parameters (here we demonstrate just two analytes—AHSG and A2M.PZP), we plot two classes, red and green. (**b**) Using a machine learning kernel (here, radial SVM), we construct a non-linear boundary between the two classes. This boundary is used to compute the probability that a sample belongs to the red class at any point on the plane formed by the two analytes. This probability serves as the weight of the edge connecting analytes 1 (AHSG) and 2 (A2M.PZP). (**c**) The synolitic graph is constructed where features (analytes) form the nodes, and the edge weights are derived from pairwise classifiers' probabilities of belonging to one of the classes. All features are connected by pairwise classifiers, and the graph's topology depends solely on class labelling. Edges can be filtered using a threshold to binarise the connections.

For imaging data, an SG can be constructed as follows. Each node in the graph represents a voxel from the fMRI data. The edges and their weights represent the relationships between voxels. Based on the array a, we construct a graph $g = (V, E, R, W)$, where

- $V = \{v_i\}_i$ is the set of vertices;
- $E = \{e_{ij}\}_{ij}$ is the set of undirected edges;
- $R = \{r_i\}_i$ is the set of node values;
- $W = \{w_{ij}\}_{ij}$ is the set of edge weights.

Each node v_i corresponds to voxel i, and the edge e_{ij} represents a connection between voxels i and j. The value r_i is assigned to node v_i, and the weight w_{ij} is assigned to edge e_{ij}. Each voxel i corresponds to a time series with multiple values. To assign a value r_i to each

node, we use a statistical function T that transforms the time series of voxel i into a single scalar value. This allows us to introduce a new 3D array $a^T = T(a)$, where for all x, y, z

$$a^T_{xyz} = T(a_{xyz}).$$

The values of the array a^T are used as the values of the vertices R. The specific choice of the statistic T, such as the mean of the time series or the difference between specific quantiles, depends on the performance of the method during testing.

Since edge weights represent the relationships between voxels in different brain activity modes, calculating these weights W is a critical task. We define the weight w_{ij} of an edge between two voxels as

$$w_{ij} = P(\sigma = II \mid r_i, r_j) - P(\sigma = I \mid r_i, r_j) \quad (1)$$

In other words, the weight w_{ij} is the difference between the probabilities of two brain activity modes (mode II and mode I), given the values of the incident vertices r_i and r_j. These edge weights w_{ij} can take values from -1 to 1.

- If $w_{ij} < 0$, the edge e_{ij} indicates that brain mode I is more likely;
- If $w_{ij} > 0$, the edge suggests that brain mode II is more likely.

The larger the absolute value $|w_{ij}|$, the more informative the edge is for classification. Conversely, when $|w_{ij}|$, is close to zero, the edge carries little information. In practice, to compute these probabilities, we use probabilistic classifiers.

$C_{lij} : \{\sigma|(r_i, r_j), \{\left(r_i^n, r_j^n\right)\}_n, \{\sigma_n\}_n\} \to [0, 1]$, trained on the available dataset. Formula (1) can then be rewritten as

$$w_{ij} = C_{lij}(\sigma = II \mid (r_i, r_j), \{(r_{in}, r_{jn})\}_n, \{\sigma_n\}_n) - C_{lij}(\sigma = I \mid (r_i, r_j), \{(r_{in}, r_{jn})\}_n, \{\sigma_n\}_n) \quad (2)$$

Thus, for each edge e_{ij}, a unique probabilistic classifier C_{lij} is trained to compute the weights W. As classifiers $\{C_{lij}\}_{ij}$ or calculating the edge weights W, probabilistic classifiers from the scikit-learn library were used [20]. They are based on the support vector machine method with a radial basis function kernel and standard parameters for this method.

Having determined the method for calculating node values R and edge weights W, we now describe the graph topology. Typically, synolitic networks are represented as complete graphs, allowing for interactions between all elements. However, fMRI data often involve too many vertices for this approach to be computationally feasible. Constructing a complete graph leads to a vast number of edges, requiring significant computational resources. For example, with fMRI data at a resolution of $100 \times 100 \times 100$ voxels, a complete graph would contain 1,000,000 vertices and approximately 500 billion edges. Moreover, for each edge e_{ij}, a classifier C_{lij} would need to be trained, significantly increasing computational demands.

Instead of a complete graph, we propose constructing a grid graph, where edges connect only neighbouring voxels. Two voxels are considered neighbours if they share a face, edge or corner. Specifically, an internal voxel (x, y, z) is connected to the voxels in the set:

$$\{(\mathrm{x}', \mathrm{y}', \mathrm{z}') : \mathrm{x}' \in \{\mathrm{x}-1, \mathrm{x}, \mathrm{x}+1\}, \mathrm{y}' \in \{\mathrm{y}-1, \mathrm{y}, \mathrm{y}+1\}, \mathrm{z}' \in \{\mathrm{z}-1, \mathrm{z}, \mathrm{z}+1\}, (\mathrm{x}', \mathrm{y}', \mathrm{z}') \neq (\mathrm{x}, \mathrm{y}, \mathrm{z})\}.$$

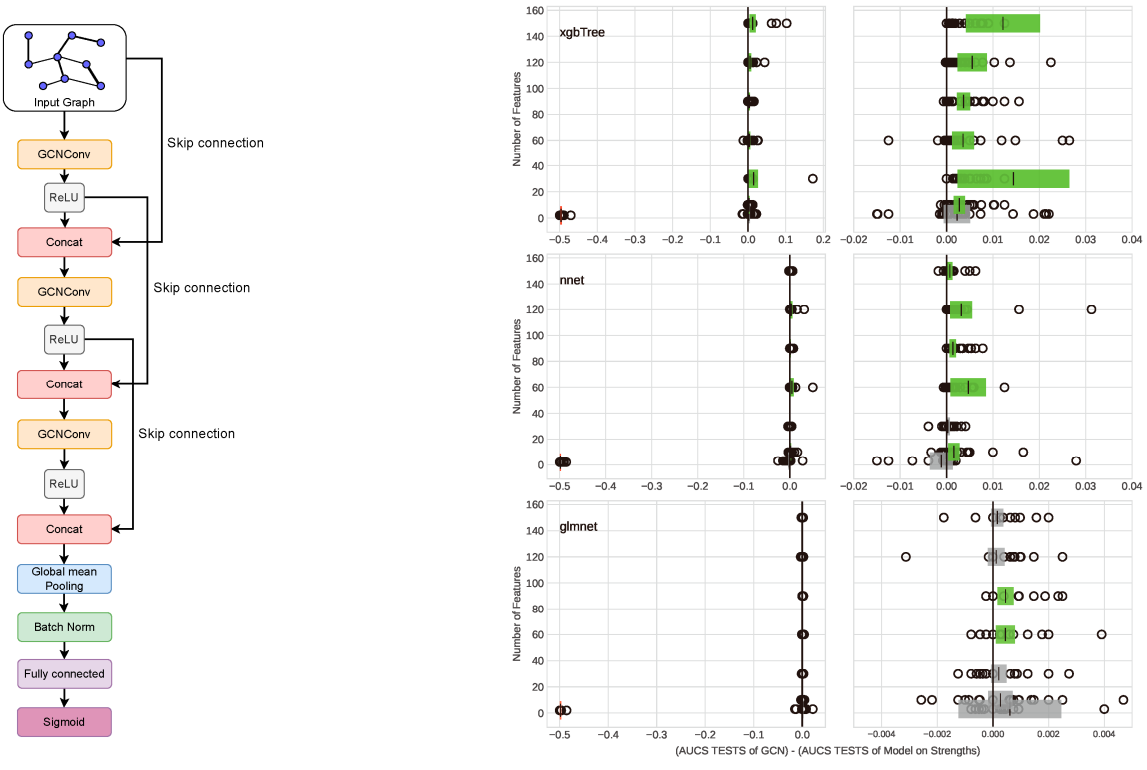
With this topology, the computational complexity is reduced from $O(n^2)$ to $O(n)$, where n is the number of voxels. Since fMRI scans include data from areas surrounding the brain, we remove edges that are incident to vertices with values below a threshold r. These edges do not carry useful information for classification, as they correspond to voxels outside the brain. We also remove edges whose absolute weight is below a threshold w. Edges with $|w| \approx 0$ may arise in two scenarios: (1) they are associated with inactive voxels,

or (2) they connect voxels that are equally involved in both brain activity modes, rendering the edges uninformative for classification. As a result, edges in the set

$$\{e_{ij} : r_i < r \mid r_j < r \mid |w_{ij}| < w\}_{ij}$$

are removed from the graph g. The parameter r is chosen based on the fMRI machine's voxel values for areas outside the brain, typically a small positive number near zero. The parameter w determines the significance of the remaining edges in the graph. The larger the w, the more edges will be pruned. As a result, each internal voxel retains no more than 26 neighbouring voxels, and the degree of each node in the graph is limited to 26.

Graph Neural Network. For the classification of graphs obtained at the previous stage, a simple Graph Neural Network was used, with the architecture shown in Figure 2, left. The graph to be classified was fed into a graph convolutional layer [5], followed by the ReLU non-linearity [21]. After the non-linearity, Batch Normalisation [22] and DropOut [23] layers were applied to prevent the overfitting of the neural network. Then, "skip connections" [24] were used, where the output of the convolutional layer after the non-linearity was connected to the input data of the layer. Similarly, two more convolutional layers were applied. As a result, node embeddings were obtained, taking into account the influence of neighbouring vertices of the graph. To transition from node embeddings to an embedding of the entire graph, Global Max Pooling [25,26] was used. Next, after another Batch Normalisation layer, the data were fed into a fully connected layer, after which class membership coefficients were calculated using a sigmoid function. Cross-entropy was used as the loss function.



(**a**) (**b**)

Figure 2. *Cont.*

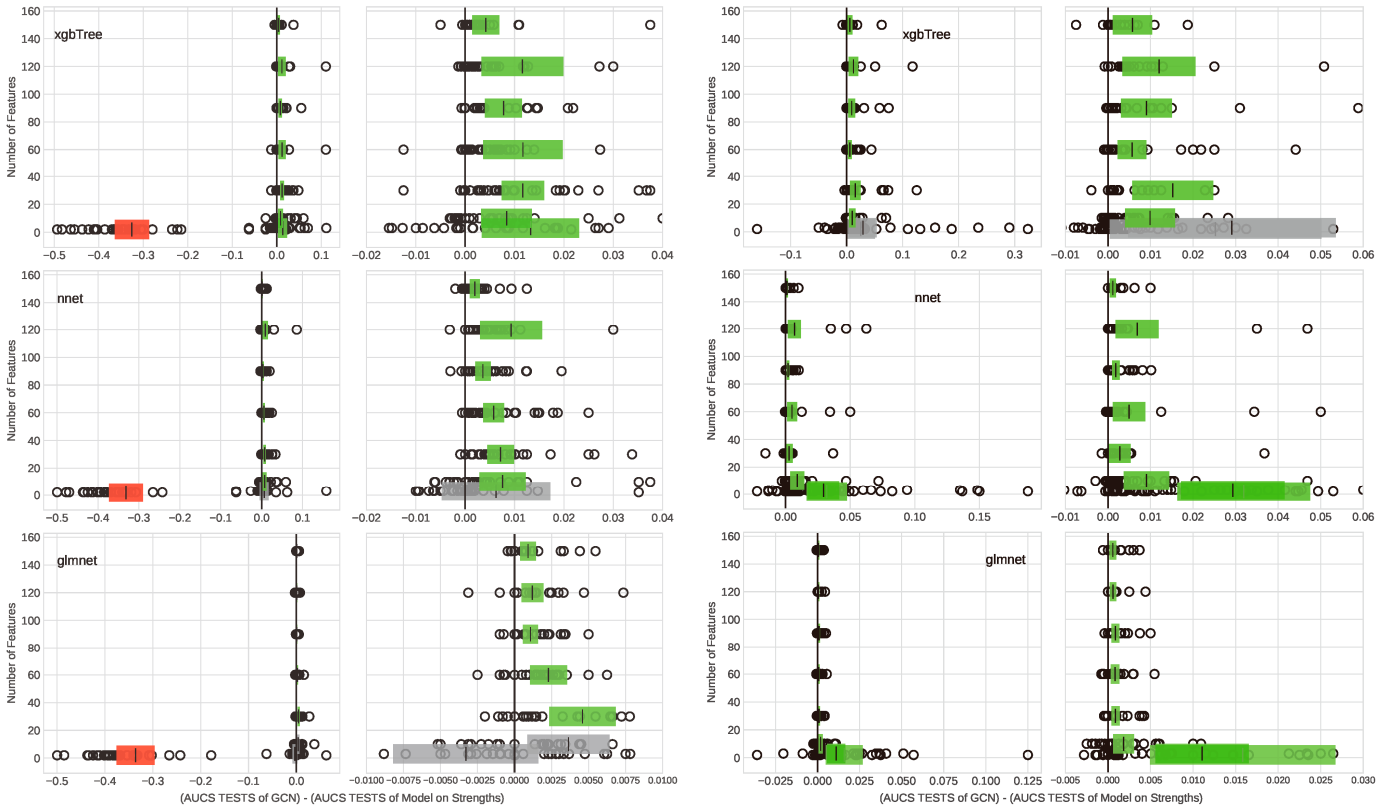


Figure 2. (**a**) Architecture of the Convolutional Graph Neural Network. Here "GCNConv" stands for Graph Convolution Network convolution operation and "Concat" for concatenation, i.e., when the feature of a node is concatenated with features from other layers. (**b**–**d**) Comparison of SGNN classification performance against traditional methods for different synthetic datasets: (**b**) ideal spheres, (**c**) broken spheres and (**d**) noisy spheres. The vertical axis represents the number of features used in the analysis. The horizontal axis represents the difference in AUC (Area Under the Curve) for the classification ROC (Receiver Operating Characteristic) curves between the SGNN classifiers and traditional ones. Colour bars (green, red or grey) shows the confidence interval for the mean difference. Green bars indicate scenarios where SGNNs outperformed traditional methods with statistically significant differences (two-sided paired Wilcoxon signed rank test). Grey bars represent no significant difference, while red bars indicate cases where the traditional ML model performed significantly better than SGNNs. The hollow circles shows individual differences. The models compared are (from top to bottom in each subfigure): xgbTree, neural network (nnet) and generalised linear model (glmnet). The superior performance of SGNNs in (**b**–**d**) can be attributed to their ability to effectively capture relationships between analytes in high-dimensional data, unlike traditional methods that often struggle with this task.

3. Results

We compare the SGNN approach with other traditional ML models by applying it to the synthetically generated data chosen. We consider three ML models (xgbTree, nnet, glmnet) from the caret package in R. We chose them because the principles of their training are based on different static principles and they all produce a selection of features. We train the model with function train (caret package R) using scaling and centring for data preprocessing; with a selection of hyperparameters that are set by default; and using cross-validation (number of folds is 5). A comparison of SGNN and more traditional ML classification methods applied to synolitic graphs is shown in Figure 2, right, and it shows that SGNN works better in almost all cases. Taking account of the fact that we have shown that synolitic network classification works better than the direct application of

these ML techniques in [11], we can conclude that SGNN provides the best classification results in comparison to classification based on synolitic graphs or the xgbTree, nnet or glmnet algorithms.

For the classification of fMRI data from a cognitive experiment [19] (see Figure 3a), a method for representing fMRI data in the form of synolitic graphs with further classification with GNNs was implemented and tested (Figure 3b). In fact, within this experiment, two modes can be distinguished in which the brain of the subject functions. The first mode occurs when the subject is sequentially shown 55 blocks, 50 of which are different images, and 5 repeat the previous image. If the subject sees a repeated image, they must press a button. In the second mode, the subject's brain is active when they are asked to sequentially imagine 25 objects. After imagining each object, the subject is asked to rate the clarity of the image they imagined on a five-point scale by pressing the corresponding buttons. In this work, we assessed the effectiveness of our method for classifying the two brain modes using fMRI: during the observation experiment and the imagination experiment.

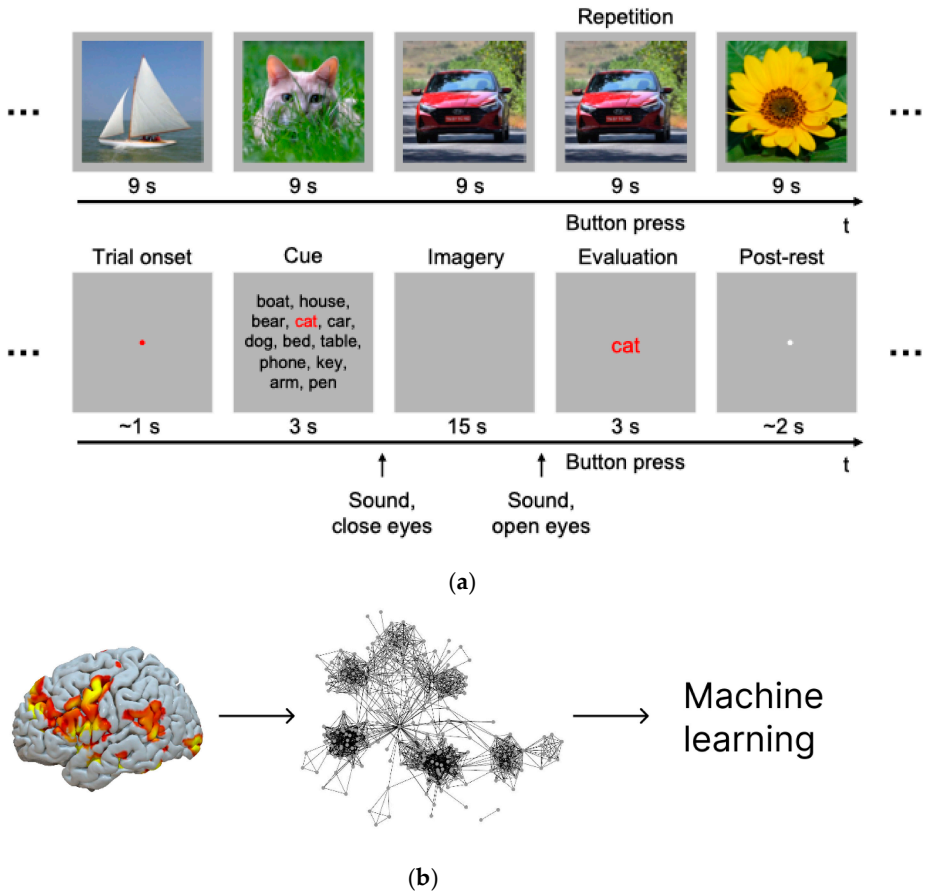


Figure 3. (**a**) A scheme of the cognitive experiment, a comparison of the seen object and imagined object. Right up: during the experiment, serial fMRI data have been recorded. (**b**) An fMRI image has been represented as a synolitic graph and then classified either by traditional ML methods or Convolutional Graph Neural Network.

Five subjects participated in the data collection, each undergoing 44 fMRI scans in both modes on different days, with 24 scans in the visual perception mode and 20 in the imagination mode. The sequences of images and objects for the same subject varied. The sample was divided, so that 30% of each mode's data for each individual was placed in the test set and 70% in the training set. The division into test and training parts was performed as follows: the seen were split as 17 training and 7 test images, the imagined were split as 14 training and 6 test images. A feature of the data is that they were obtained for five

subjects, and for each of them, the data were split into test and training sets. This approach does not allow us to claim that the method's effectiveness was tested on independent data.

However, it allows us to check whether it is possible to predict the behaviour of the person on whom the method was trained. SGNNs have shown 100% accuracy in distinguishing between the two regimes, better than the accuracy achieved only with synolitic graphs and an analysis of graph parameters. Note that while we acknowledge that cross-validation could further improve the evaluation, in the current review of the analysis, we leveraged the fact that we initially calculated the edges for all voxels. This approach, in effect, averages the model performance over the voxel space, which provides a form of generalisation that we believe compensates, to some extent, for the lack of traditional cross-validation.

The 3D CNN network is also known to be suitable for analysing fMRI data (see e.g., [27], hence, it would be interesting to discuss SGNNs vs. 3D CNNs. SGNNs offer several advantages, including transforming data into a graph structure that reduces dimensionality and focuses on key relationships within the dataset. This is particularly beneficial for fMRI data, which often contains noisy or irrelevant features. The graph is also class-driven, enhancing class separability, and provides interpretability through insights into the data's structure, such as community patterns and node importance. These characteristics make SGNNs valuable in healthcare contexts, where understanding the model's decisions is crucial. On the other hand, 3D CNNs directly operate on raw fMRI data, bypassing the need for graph construction and simplifying the workflow. They have demonstrated high accuracy in tasks like schizophrenia detection due to their ability to learn hierarchical features from spatially structured data.

However, there are trade-offs between the two approaches. SGNNs are more generalisable across various high-dimensional datasets, making them suitable for diverse applications. In contrast, 3D CNNs excel with spatially structured data but may not perform as well with high-dimensional features that lack spatial organisation. Additionally, SGNNs inherently perform feature selection through the graph construction process, while 3D CNNs rely on raw data that may require extra preprocessing to achieve a similar interpretability. It is also worth noting that SGNNs and 3D CNNs can complement each other. For example, 3D CNNs can be used to preprocess fMRI data by identifying meaningful spatial features, which can then be incorporated into a graph-based framework, combining the strengths of both methods. In conclusion, while 3D CNNs are effective for specific tasks like schizophrenia classification, SGNNs provide a more adaptable and interpretable framework for analysing a variety of high-dimensional datasets. The choice between SGNNs and 3D CNNs depends on the goals of the analysis, whether the focus is on achieving a high accuracy with spatially structured data or deriving generalisable insights and interpretability across different types of data.

It is worthwhile also to discuss the relevance of previous GNN models like BrainGNN in fMRI data abalysis [28]. BrainGNN relies on functional networks from predefined regions of interest (ROIs), which are represented as nodes with edges based on their functional connectivity. This approach is specialised for fMRI data but requires prior knowledge of the ROIs functional connectivity. In contrast, SGNNs construct graphs directly from the dataset, guided by class labels, making them more flexible and generalisable to various domains beyond fMRI, though without relying on spatial–functional priors. While BrainGNN provides interpretability by identifying important ROIs or connections, SGNNs offer a broader interpretability through graph properties, such as community structures and key nodes, applicable to various datasets. In terms of predictive performance, BrainGNN excels with fMRI data's network structure, but SGNNs achieve competitive results by transforming high-dimensional data into graph representations, reducing overfitting risks.

BrainGNN is specifically designed for fMRI data, while SGNNs are domain-agnostic and can handle various types of high-dimensional datasets. SGNNs' generalisability makes them suitable for situations where predefined network structures are unavailable, but they may be less specialised for domain-specific tasks like ROI-based analysis. However, SGNNs and BrainGNN complement each other. BrainGNN is more suited for tasks with domain-specific knowledge, while SGNNs offer a versatile framework for diverse datasets. The choice between them depends on the application and the available prior knowledge. SGNNs are an excellent alternative when functional connectivity information is limited or when a more general approach is needed.

4. Discussion

The integration of synolitic graph representation with Graph Neural Networks (GNNs) in the form of SGNNs offers a promising framework for addressing complex challenges in the analysis of multidimensional data. By leveraging the strengths of both methodologies, we can enhance our understanding and processing of data structures that are often intricate and high-dimensional. One of the most significant advantages of employing a synolitic graph representation lies in its applicability to any form of multidimensional data. More traditional data representation methods often struggle with the high-dimensional space of parameters, leading to issues such as the curse of dimensionality and machine learning overfitting. In contrast, Deep Geometric Learning and GNNs capture the intrinsic and difficult to detect relationships and structures within the data, allowing for a more coherent representation and analysis. This capability is particularly beneficial when working with datasets that do not conform to linear relationships or where features are interdependent. By using a synolitic graph-based approach we enable Deep Geometric Learning (DGL) for any kind of data, reconstructing the topological properties of the data solely from its labelling and assignment to different classification outputs.

Synolitic graph representation is, in fact, an ensemble of independent pairwise classifiers, building, however, a connected graph. The combination of ensemble learning with the advantages of DGL enhances the robustness of our models. Ensemble learning, aggregating the predictions from multiple classifiers, can significantly reduce variance and improve performance, especially in high-dimensional settings. When integrated with GNNs, this approach enables the models to learn diverse representations and capture various aspects of the data structure. The GDL further enriches this combination by allowing the model to learn from geometric properties of the data. An essential feature of our approach is its ability to exploit the internal, often hidden, structure of the data, relying solely on assignments to different classes. This method allows for a more detailed understanding of the data, as it is focused on relationships and interactions between classes rather than focusing solely on individual features. Furthermore, GNNs can effectively propagate information through the graph, leading to improved classification and clustering outcomes. This class-based assignment strategy not only streamlines the learning process but also helps the model discern subtle patterns that may be overlooked by traditional machine learning methods.

The synolitic graph representation of multidimensional data goes beyond merely transforming datasets into networks; it uncovers hidden relationships and structures that are often not visible in the original feature space. By examining the characteristics of these networks, we can gain valuable insights into the data's underlying structure and understand how Synolitical Graph Neural Networks (SGNNs) utilise these structures for classification. One key aspect of this interpretability is the emergence of community-like structures during the transformation process. These communities represent clusters of data points with a high similarity, indicating cohesion within a class or separation between classes, which can significantly impact the classification accuracy of SGNNs. In our experi-

ments with both synthetic and fMRI data, the presence of high-modularity communities might be linked to better class separability and a more robust SGNN performance. Another important factor is the identification of key nodes in the network. Nodes with high-degree or high-betweenness centrality often play critical roles. High-degree nodes act as hubs that consolidate connections within a class, thereby strengthening intra-class relationships, while high-betweenness nodes tend to lie at decision boundaries, helping to delineate class boundaries. These nodes exert a disproportionate influence over the learning process, affecting how feature embeddings are learned and propagated during graph convolution. The heterogeneity of the network structure, defined by the distribution of node degrees and edge weights, reflects the dataset's complexity. Dense sub-networks or highly connected hubs correspond to stable intra-class patterns, while sparsely connected regions may highlight outliers or transitional points between classes. SGNNs take advantage of this heterogeneity by focusing on the most informative regions of the network for classification.

By structuring the data in a graph format, we reduce the noise inherent in high-dimensional datasets, enabling the model to focus on the most relevant features. The graph-based representation allows for effective regularisation techniques that can combat overfitting by promoting smoothness and coherence with the learned representations. The hierarchical nature of GNNs permits multi-scale learning, where information is aggregated from various layers of the graph. This mechanism provides a natural way to deal with high-dimensional data by focusing on meaningful relationships rather than individual dimensions. As a result, we observe improved performance across a range of classification tasks, while maintaining interpretability due to the graph representation of the high-dimensional data. The combination of synolitic graph representation with Graph Neural Networks offers a versatile and powerful approach for analysing multidimensional data and its classification. Combining ensemble learning and Deep Geometric Learning effectively addresses the challenges of overfitting, the identification of the data's internal structure and the curse of dimensionality, providing us with an effective solution for a wide array of applications. Future research should focus on further exploring the implications of this approach in various domains, enhancing its applicability and efficiency and continuing to refine the integration of these methodologies to unlock their full potential.

Graph Neural Networks (GNNs) offer distinct advantages for explainability and interpretability, using white [20,29] or black [30–32] box methods. GNNs and SGNNs represent data through entities (nodes) and relationships (edges), which correspond to high-level concepts or knowledge items, making them easier to interpret compared to other models like convolutional neural networks (CNNs) that rely on pixel-level data or latent space vectors. This entity-based representation simplifies understanding for humans, as it reflects familiar structures found in real-world systems like social or molecular networks. Additionally, the possibility of including attention mechanisms, particularly in Graph Attention Networks (GATs), can enhance interpretability by assigning weights to different node embeddings. These weights indicate which nodes or edges are most important for a particular prediction, allowing users to focus on the most relevant features. This attention-driven interpretability is especially valuable in tasks where understanding how different parts of the graph influence the outcome is critical. Localised explanations also play a key role in the interpretability of SGNNs [33–36]. SGNNs are particularly suited for explaining node-level predictions, which is important for providing instance-specific explanations. These localised interpretations enable a clearer understanding of individual decisions rather than requiring explanations at the whole-graph level. These characteristics make SGNNs particularly well suited for explainability and interpretability, especially when working with structured data like social networks, biological data or molecules

SGNNs exhibit robust characteristics against adversarial attacks for several key reasons [33–36]. First, SGNNs operate in highly structured data environments, such as graphs, where perturbations are not as straightforward as in other data domains like images or text. The complexity of the graph structure and the interdependence between nodes (due to message-passing mechanisms) makes it harder for adversarial changes to propagate through the network and distort predictions effectively. Adversarial attacks on SGNNs typically involve discrete changes, like adding or removing edges or altering node attributes, which require more sophisticated and targeted perturbation strategies. Moreover, SGNNs naturally aggregate information from neighbouring nodes, which means that small, localised changes to a graph (e.g., perturbing a few nodes or edges) may not have a substantial impact on the overall prediction. The robustness of SGNNs also stems from the fact that perturbations affecting one part of the graph may be mitigated by the network's ability to pool information from other, unperturbed parts of the graph, leading to more stable predictions.

The findings in this paper demonstrate the effectiveness of SGNNs for analysing high-dimensional physiological data, offering a strong classification performance and interpretability. These features make SGNNs particularly valuable in Network Physiology, where understanding interactions between physiological systems is crucial. SGNNs could be used to analyse dynamic organ system interactions under various physiological or pathological conditions. By transforming time-series data into graph representations, SGNNs capture interdependencies that align with the concept of mapping the "human physiolome", as described by [3,18]. They also have potential for disease diagnostics and monitoring, as disruptions in network topology are linked to pathological states. SGNNs provide a scalable way to detect these disruptions across diverse datasets, which could aid in the early diagnosis or monitoring of conditions such as cardiovascular disease, sleep disorders or neurological issues. Unlike traditional static network approaches, SGNNs can incorporate temporal changes in physiological interactions, aligning with the need for a dynamic network analysis to study how systems adapt to stress, exercise or disease. Additionally, SGNNs could contribute to personalised medicine by identifying individual-specific physiological patterns and deviations, helping tailor interventions based on unique network topologies. While this work highlights SGNNs' potential, future research should test their application in Network Physiology, comparing SGNN-derived networks with established physiological networks to validate their relevance. Integrating SGNNs with temporal graph architectures could further improve their ability to capture dynamic physiological interactions. We hope this work inspires future applications in understanding the complex interplay of physiological systems and advancing network-based medical diagnostics and therapeutics.

Author Contributions: Conceptualisation, A.Z.; methodology, All; formal analysis, M.K., T.N., V.U. and D.V.; writing—review and editing, All; supervision, A.Z. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the Analytical Centre for the Government of the Russian Federation (agreement identifier 000000D730324P540002, grant No 70-2023-001320 dated 27 December 2023).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data supporting the reported results can be found via links in [11] (synthetic data) and [23] (fMRI data).

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Whitwell, H.J.; Bacalini, M.G.; Blyuss, O.; Chen, S.; Garagnani, P.; Gordleeva, S.Y.; Jalan, S.; Ivanchenko, M.; Kanakov, O.; Kustikova, V.; et al. The Human Body as a Super Network: Digital Methods to Analyze the Propagation of Aging. *Front. Aging Neurosci.* **2020**, *12*, 136. [CrossRef] [PubMed]
2. Bartsch, R.P.; Liu, K.K.L.; Bashan, A.; Ivanov, P.C. Network physiology: How organ systems dynamically interact. *PLoS ONE* **2015**, *10*, e0142143. [CrossRef] [PubMed]
3. Bronstein, M.M.; Bruna, J.; Cohen, T.; Veličković, P. Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges. *arXiv* **2021**, arXiv:2104.13478.
4. Gilmer, J.; Schoenholz, S.S.; Riley, P.F.; Vinyals, O.; Dahl, G.E. Neural Message Passing for Quantum Chemistry. In Proceedings of the 34th International Conference on Machine Learning, Sydney, Australia, 6–11 August 2017. PMLR 70 2017.
5. Kipf, T.; Welling, M. Semi-supervised classification with graph convolutional networks. In Proceedings of the International Conference on Learning Representations, San Juan, PR, USA, 2–4 May 2016.
6. Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Lio, P.; Bengio, Y. Graph attention networks. In Proceedings of the International Conference on Learning Representations, Vancouver, BC, Canada, 30 April–3 May 2018.
7. Wu, L.; Cui, P.; Pei, J.; Zhao, L. *Graph Neural Networks: Foundations, Frontiers, and Applications*; Springer: Singapore, 2022.
8. Saeidi, M.; Karwowski, W.; Farahani, F.V.; Fiok, K.; Hancock, P.A.; Sawyer, B.D.; Christov-Moore, L.; Douglas, P.K. Decoding Task-Based fMRI Data with Graph Neural Networks, Considering Individual Differences. *Brain Sci.* **2022**, *12*, 1094. [CrossRef] [PubMed]
9. Gorban, A.N.; Tyukina, T.A.; Pokidysheva, L.I.; Smirnova, E.v. Dynamic and thermodynamic models of adaptation. In *Physics of Life Reviews*; Elsevier B.V: Amsterdam, The Netherlands, 2021; Volume 37, pp. 17–64. [CrossRef]
10. Bartlett, T.E.; Zaikin, A. Detection of epigenomic network community oncomarkers. *Ann. Appl. Stat.* **2016**, *10*, 1373–1396. [CrossRef]
11. Nazarenko, T.; Whitwell, H.J.; Blyuss, O.; Zaikin, A. Parenclitic and Synolytic Networks Revisited. *Front. Genet.* **2021**, *12*, 733783. [CrossRef] [PubMed]
12. Nazarenko, T.; Blyuss, O.; Whitwell, H.; Zaikin, A. Ensemble of correlation, parenclitic and synolitic graphs as a tool to detect universal changes in complex biological systems: Comment on "Dynamic and thermodynamic models of adaptation" by A.N. Gorban et al. *Phys. Life Rev.* **2021**, *38*. [CrossRef]
13. Demichev, V.; Tober-Lau, P.; Nazarenko, T.; Thibeault, C.; Whitwell, H.; Lemke, O.; Röhl, A.; Freiwald, A.; Szyrwiel, L.; Ludwig, D.; et al. A time-resolved proteomic and prognostic map of COVID-19 disease progression and predicts outcome. *Cell Systems* **2021**, *12*, 780–794. [CrossRef]
14. Demichev, V.; Tober-Lau, P.; Nazarenko, T.; Aulakh, S.K.; Whitwell, H.; Lemke, O.; Röhl, A.; Freiwald, A.; Mittermaier, M.; Szyrwiel, L.; et al. A proteomic survival predictor for COVID-19 patients in intensive care. *PLoS Digit. Health* **2022**, *1*, e0000007. [CrossRef]
15. Krivonosov, M.; Nazarenko, T.; Bacalini, M.G.; Vedunova, M.; Franceschi, C.; Zaikin, A.; Ivanchenko, M. Age-related trajectories of DNA methylation network markers: A parenclitic network approach to a family-based cohort of patients with Down Syndrome. *Chaos Solitons Fractals* **2022**, *165*, 112863. [CrossRef]
16. Bashan, A.; Bartsch, R.P.; Kantelhardt, J.W.; Havlin, S.; Ivanov, P.C. Network physiology reveals relations between network topology and physiological function. *Nat. Commun.* **2012**, *3*, 702. [CrossRef]
17. Ivanov, P.C.; Bartsch, R.P. Network physiology: Mapping interactions between networks of physiologic networks. In *Networks of Networks: The last Frontier of Complexity*; Springer International Publishing: Cham, Switzerland, 2014; pp. 203–222.
18. Ivanov, P.C. The new field of network physiology: Building the human physiolome. *Front. Netw. Physiol.* **2021**, *1*, 711778. [CrossRef] [PubMed]
19. Horikawa, T.; Kamitani, Y. Generic Object Decoding of Seen and Imagined Objects Using Hierarchical Visual Features. *Nat. Commun.* **2017**, *8*, 15037. [CrossRef] [PubMed]
20. Baldassarre, F.; Azizpour, H. Explainability techniques for graph convolutional networks. *arXiv* **2019**, arXiv:190513686.
21. Fukushima, K. Visual Feature Extraction by a Multilayered Network of Analog Threshold Elements. *IEEE Trans. Syst. Sci. Cybern.* **1969**, *5*, 322–333. [CrossRef]
22. Ioffe, S.; Szegedy, C. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. *arXiv* **2015**, arXiv:1502.03167.
23. Srivastava, N.; Hinton, G.; Krizhevsky, A.; Salakhutdinov, R. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *J. Mach. Learn. Res.* **2014**, *15*, 1929–1958.
24. Xu, K.; Zhang, M.; Jegelka, S.; Kawaguchi, K. Optimization of Graph Neural Networks: Implicit Acceleration by Skip Connections and More Depth. *arXiv* **2021**, arXiv:2105.04550.
25. Grattarola, D.; Zambon, D.; Bianchi, F.M.; Alippi, C. Understanding Pooling in Graph Neural Networks. *arXiv* **2021**. [CrossRef] [PubMed]

26. Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; et al. Scikit-learn: Machine Learning in Python. *J. Mach. Learn. Res.* **2011**, *12*, 2825–2830.
27. Naveed Iqbal Qureshi, M.; Oh, J.; Lee, B. 3D-CNN based discrimination of schizophrenia using resting-state fMRI. *Artif. Intell. Med.* **2019**, *98*, 10–17. [CrossRef] [PubMed]
28. Li, X.; Zhou, Y.; Dvornek, N.; Zhang, M.; Gao, S.; Zhuang, J.; Scheinost, D.; Staib, L.H.; Ventola, P.; Duncan, J.S. BrainGNN: Interpretable Brain Graph Neural Network for fMRI Analysis. *Med. Image Anal.* **2021**, *74*, 102233. [CrossRef]
29. Sanchez-Lengeling, B.; Wei, J.; Lee, B.; Reif, E.; Wang, P.; Qian, W.; McCloskey, K.; Col-well, L.; Wiltschko, A. Evaluating attribution for graph neural networks. In *Advances in Neural Information Processing Systems*; Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H., Eds.; Curran Associates, Inc.: Newry, UK, 2020; Volume 33, pp. 5898–5910.
30. Huang, Q.; Yamada, M.; Tian, Y.; Singh, D.; Yin, D.; Chang, Y. Graphlime: Local interpretable model explanations for graph neural networks. *arXiv* **2020**, arXiv:200106216. [CrossRef]
31. Zhang, Y.; Defazio, D.; Ramesh, A. Relex: A model-agnostic relational model explainer. *arXiv* **2020**, arXiv:200600305.
32. Vu, M.N.; Thai, M.T. Pgm-explainer: Probabilistic graphical model explanations for graph neural networks. *arXiv* **2020**, arXiv:201005788.
33. Zügner, D.; Günnemann, S. Adversarial attacks on graph neural networks via meta learning. In Proceedings of the International Conference on Learning Representations, ICLR, New Orleans, LA, USA, 6–9 May 2019.
34. Zügner, D.; Günnemann, S. Certifiable robustness and robust training for graph convolutional networks. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, Anchorage, AK, USA, 4–8 August 2019; pp. 246–256.
35. Zügner, D.; Günnemann, S. Certifiable robustness of graph convolutional networks under structure perturbations. In Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Association for Computing Machinery, KDD '20, Virtual Event, 6–10 July 2020; pp. 1656–1665. [CrossRef]
36. Zügner, D.; Akbarnejad, A.; Günnemann, S. Adversarial attacks on neural net- works for graph data. In Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, London, UK, 19–23 August 2018; pp. 2847–2856.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.