

RESEARCH

Open Access



Explaining decisions of graph convolutional neural networks: patient-specific molecular subnetworks responsible for metastasis prediction in breast cancer

Hryhorii Chereda¹, Annalen Bleckmann², Kerstin Menck², Júlia Perera-Bel³, Philip Stegmaier⁴, Florian Auer⁵, Frank Kramer⁵, Andreas Leha⁶ and Tim Beißbarth^{1,7*} 

Abstract

Background: Contemporary deep learning approaches show cutting-edge performance in a variety of complex prediction tasks. Nonetheless, the application of deep learning in healthcare remains limited since deep learning methods are often considered as non-interpretable black-box models. However, the machine learning community made recent elaborations on interpretability methods explaining data point-specific decisions of deep learning techniques. We believe that such explanations can assist the need in personalized precision medicine decisions via explaining patient-specific predictions.

Methods: Layer-wise Relevance Propagation (LRP) is a technique to explain decisions of deep learning methods. It is widely used to interpret Convolutional Neural Networks (CNNs) applied on image data. Recently, CNNs started to extend towards non-Euclidean domains like graphs. Molecular networks are commonly represented as graphs detailing interactions between molecules. Gene expression data can be assigned to the vertices of these graphs. In other words, gene expression data can be structured by utilizing molecular network information as prior knowledge. Graph-CNNs can be applied to structured gene expression data, for example, to predict metastatic events in breast cancer. Therefore, there is a need for explanations showing which part of a molecular network is relevant for predicting an event, e.g., distant metastasis in cancer, for each individual patient.

Results: We extended the procedure of LRP to make it available for Graph-CNN and tested its applicability on a large breast cancer dataset. We present Graph Layer-wise Relevance Propagation (GLRP) as a new method to explain the decisions made by Graph-CNNs. We demonstrate a sanity check of the developed GLRP on a hand-written digits dataset and then apply the method on gene expression data. We show that GLRP provides patient-specific molecular subnetworks that largely agree with clinical knowledge and identify common as well as novel, and potentially druggable, drivers of tumor progression.

Conclusions: The developed method could be potentially highly useful on interpreting classification results in the context of different omics data and prior knowledge molecular networks on the individual patient level, as for example in precision medicine approaches or a molecular tumor board.

Keywords: Gene expression data, Explainable AI, Personalized medicine, Precision medicine, Classification of cancer, Deep learning, Prior knowledge, Molecular networks

*Correspondence: tim.beissbarth@bioinf.med.uni-goettingen.de

¹Medical Bioinformatics, University Medical Center Göttingen, Göttingen, Germany

⁷Campus-Institute Data Science (CIDAS), University of Göttingen, Göttingen, Germany

Full list of author information is available at the end of the article



© The Author(s). 2021 **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>. The Creative Commons Public Domain Dedication waiver (<http://creativecommons.org/publicdomain/zero/1.0/>) applies to the data made available in this article, unless otherwise stated in a credit line to the data.

Background

Gene expression profiling by microarrays or next-generation sequencing has played a significant role in identifying predictive gene signatures and discovering individual biomarkers in cancer prognosis [1]. High-throughput sequencing produces huge amounts of gene expression data that can potentially be used for deriving clinical predictors (e.g., predicting occurrence of metastases) and identifying novel drug targets. Breast cancer is one of the paradigmatic examples of the utility of high-throughput data to derive prognostic molecular signatures (PAM50, MammaPrint, OncotypeDX) [2, 3] that predict clinical outcome. Based on the expression of 50 genes, the PAM50 classifier is widely used to divide breast cancers into four main molecular subtypes: luminal A, luminal B, triple-negative/basal-like, and HER2-enriched [4]. While the two luminal subtypes are characterized by high hormone receptor expression and generally have a better prognosis, the basal-like breast cancers are a heterogeneous group of hormone receptor- and HER2-negative breast cancers that are highly proliferative and often metastasize early. MammaPrint and OncotypeDX are 70- and 21-gene expression signatures that stratify patients according to the likelihood of metastasis. Although molecular signatures have prognostic impact, a more complete analysis of the molecular characteristics in the individual patient is required for personalized breast cancer therapy [2]. We hypothesize that molecular signatures can differ from one patient to another due to the heterogeneity of breast cancers. Such molecular signatures can be depicted as patient-specific subnetworks that are parts of a molecular network representing background knowledge about biological mechanisms. Presenting interpretable patient-specific subnetworks to clinicians and researchers enables better interpretability of the data for further medical and pharmaceutical insights, and possibly, for extended treatment options.

From a machine learning (ML) perspective, the prediction of a clinical outcome is a classification task, and molecular signatures can be identified as discriminative features. One drawback is that the search for molecular signatures is based on high-dimensional gene expression datasets, where the number of genes is much higher than the number of patients. The “curse of dimensionality” leads to instability in the feature selection process across different datasets. Stability can be improved including prior knowledge of molecular networks (e.g., pathways) into ML approaches [5]. ML methods benefit from pathway knowledge since neighboring genes are not treated as independent but instead similarities among adjacent genes, which should have similar expression profiles, are captured [6].

The essence of our classification task is to predict an occurrence of distant metastasis based on gene

expression data structured by a molecular network (encoded as a graph) representing connections between genes. The patients are represented as graph signals (gene expression data) on a single graph. Since each vertex of a molecular network has a corresponding gene expression value as an attribute, we perform a graph signal classification task. Patients’ gene expression profiles create different graph signal patterns that can be learned by the means of deep learning.

In recent years, deep learning has been widely applied on image data using convolutional neural networks (CNNs). The CNNs exploit the grid-like structure of images and cannot directly process data structured in non-Euclidean domains. Examples of non-Euclidean data domains include networks in social sciences and molecular networks in biology. Recently, deep learning methods extended to domains like graphs and manifolds [7]. Graph-CNN [8] learns graph signal patterns and can be applied to our graph signal classification task.

Deep neural networks are able to model complex interactions between the input and output variables. This complexity does not allow to track what role a particular input feature plays in the output; thus, a neural network itself as a black-box ML model does not provide interpretable insights.

On the other hand, decisions proposed by neural networks have to be explained before they can be taken into account in the clinical domain [9]. The European Union’s recent General Data Protection Regulation (GDPR) restricted automated decision making produced by algorithms [10]. Article 13 of [10] specifies that clinics should provide patients with “meaningful information about the logic involved”. Article 22 of [10] states that a patient shall have the right not to be subject to an automated decision unless the patient gives a consent with it (paragraph 2.c). Therefore, the explainability of deep neural networks becomes an imperative for clinical applications.

Explanation methods aim at making classification decisions of complex ML models interpretable in terms of input variables. These methods use one of two available approaches [11]: functional or message passing. The first group of methods produces explanations out of local analysis of a prediction. It includes the sensitivity analysis, Taylor series expansion, and the model agnostic approaches LIME [12] and SHAP [13]. The second group [14, 15] provides explanations by running a backward pass in a computational graph, which generates a prediction as its output. The Layer-Wise Relevance Propagation (LRP) method [15] combines through the framework of deep Taylor decomposition [11] functional and message passing approaches to generate relevances of each input feature. For a fixed input feature, the relevance shows how much this feature influences the classifier’s decision.

The relevances are generated for each data point (in our application each patient) individually.

In image data, LRP exhibited promising results and has been applied in cancer research to identify prognostic biomarkers: Klauschen et al. [16] applied LRP for visual scoring of tumor-infiltrating lymphocytes (TIL) on hematoxylin and eosin breast cancer images. Binder et al. [17] used LRP to identify spatial regions (cancer cell, stroma, TILs) on morphological tumor images that explained predictions of molecular tumor properties (like protein expression).

There are also some interpretation methods specialized for Graph Neural Networks (GNN). In [18–20], the authors provided explanation methods that are exclusively based on and crafted only for Graph Convolutional Network [21] utilizing a convolutional architecture which is a simplified version of that of Graph-CNN [8] we use. Ying et al. [22] suggested the model-agnostic GNNExplainer that is suitable for node classification, link prediction, and graph classification, but the authors did not consider an application of their approach to graph signal classification [23, 24], which is the problem at hand. The GNN-LRP method [25] proposes explanations in the form of scored sequences of edges on the input graph (i.e., relevant walks). Such a sequence represents a path extracted from the input to the output of GNN that brings insights for GNN's decision strategy. This is useful especially for graph classification tasks, where each data point is represented as an individual graph. In our task, patients are represented as graph signals on a single graph, so that this method is not applicable.

Hence, there is still a lack of methods explaining individualized predictions in the context of graph signal classification task. Here, we adapted an existing LRP technique to graph convolutional layers of Graph-CNN [8] incorporating prior knowledge of a molecular network. Our approach generates explanations in the form of relevant subgraphs for each data point and allows to provide interpretable molecular subnetworks that are individual for each patient. According to the knowledge of the authors, an explanation method that benefits from prior knowledge and provides patient-specific subnetworks has not been shown before. The novelty of our work consists of two parts. First, we present the Graph Layer-wise Relevance Propagation (GLRP) method delivering data point-specific explanations for Graph-CNN [8]. Second, we train Graph-CNN on a large breast cancer dataset to predict an occurrence of distant metastasis and show how patient-specific molecular subnetworks assist in personalized precision medicine decisions: We interpret the classifier's predictions by patient-specific subnetworks that explain the differential clinical outcome and identify therapeutic vulnerabilities.

Methods

Gene expression data and molecular network

Protein-protein interaction network

We used the Human Protein Reference Database (HPRD) protein-protein interaction (PPI) network [26] as the molecular network to structure the gene expression data. The database contains protein-protein interaction information based on yeast two-hybrid analysis, in vitro and in vivo methods. The PPI network is an undirected graph with binary interactions between pairs of proteins. The graph is not connected.

Breast cancer data

We applied our methods to a large breast cancer patient dataset that we previously studied and preprocessed [27]. That data is compiled out of 10 public microarray datasets measured on Affymetrix Human Genome HG-U133 Plus 2.0 and HG-U133A arrays. The datasets are available from the Gene Expression Omnibus (GEO) [28] data repository (accession numbers GSE25066, GSE20685, GSE19615, GSE17907, GSE16446, GSE17705, GSE2603, GSE11121, GSE7390, GSE6532). The RMA probe-summary algorithm [29] was used to process each of the datasets, and only samples with metadata on metastasis-free survival were selected and combined together on the basis of HG-U133A array probe names. Quantile normalization was applied over all datasets. In the case of several probes mapping to one gene, only the probe with the highest average value was considered. After pre-processing the dataset contained 12,179 genes in 969 patients. The patients were assigned to one of two classes: 393 patients with distant metastasis within the first 5 years and 576 patients without metastasis having the last follow-up between 5 and 10 years. Breast cancer molecular subtypes for the patient samples were predicted in [27] utilizing *genefu* R-package [30].

After mapping of 12,179 genes to the vertices of the PPI, the resulting PPI graph consisted of 7168 vertices (mapped genes) in 207 connected components. The main connected component had 6888 vertices, and each of the other 206 components had from 1 to 4 vertices. For further analyses, we utilized only the main connected component since the Graph-CNN requires the graph to be connected. The preprocessed data is provided in [31].

Expression data of HUVECs before and after TNF α stimulation

For validation purposes, we analyzed gene expression data from human umbilical vein endothelial cells (HUVECs) treated or not treated with tumor necrosis factor alpha TNF α [32]. The data, provided by the same authors (GEO database series: GSE144803), containing 39 sample pairs (treated and untreated), were suitable for a binary classification task and balanced. The expression data were

quantile normalized and mapped to vertices of HPRD PPIs resulting in 7798 genes in the main connected component.

Problem formulation

We focus on explaining classifier decisions of Graph-CNN adapting existing LRP approaches for graph convolutional layers. LRP should be applied as a postprocessing step to a model already trained for the ML task. The task is formulated as a binary classification of gene expression data $X \in \mathbb{R}^{n \times m}$ to a target variable $Y \in \{0, 1\}^n$. n is the number of data points (patients) and m is the number of features (genes). The information of the molecular network is presented as an undirected weighted graph $G = (V, E, A)$, where V and E denote the sets of vertices and edges respectively and A denotes the adjacency matrix. The Graph-CNN was designed to work with weighted graphs. We define weighted adjacency matrix A of dimensionality $m \times m$ since in general molecular networks can be weighted. For the unweighted HPRD PPI network, the matrix A has only “0s” and “1s” as its elements. A row x of the gene expression matrix X contains data from one data point (patient) and can be mapped to the vertices of the graph G . In such a way, values of x are interpreted as a graph signal.

A trained neural network can be represented as a function $f : \mathbb{R}_+^m \rightarrow [0, 1]$ mapping the input to the probability of the output class. The input x is a set of gene expression values $x = \{x_g\}$ where g denotes a particular gene. The function $f(x)$ computes the probability that a certain pattern of gene expression values is present w.r.t to the output class. LRP methods apply propagation rules from the output of the neural network to the input in order to quantify the relevance score $R_g(x)$ for each gene g . These relevances show how much gene g influences the prediction $f(x)$:

$$\forall x : f(x) = \sum_g R_g(x). \quad (1)$$

Equation (1) [11] demonstrates that the relevance scores are calculated w.r.t every input data point x .

Graph Convolutional Neural Network and Layer-wise Relevance propagation

Usual CNNs learn data representations on grid-like structures. The Graph-CNN [8] as a deep learning technique is designed to learn features on weighted graphs. The convolution on graphs is used to capture localized patterns of a graph signal. This operation is based on spectral graph theory. The main operator to investigate the spectrum of a graph is the graph Laplacian $L = D - A$, where D is a weighted degree matrix, and A is a weighted adjacency matrix. L is a real symmetric positive semidefinite matrix that can be diagonalized such that $L = U\Lambda U^T$, where $\Lambda = \text{diag}([\lambda_1, \dots, \lambda_m])$ is a diagonal non-negative

real valued matrix of eigenvalues, matrix U is composed of eigenvectors. Matrices U and U^T define the Fourier and the inverse Fourier transform respectively. According to the convolution theorem, the operation of graph convolution can be viewed as a filtering operation:

$$y = h_\theta(L)x = h_\theta(U\Lambda U^T)x = U h_\theta(\Lambda) U^T x, \quad (2)$$

where $x, y \in \mathbb{R}^m$, and the filter $h_\theta(\Lambda)$ is a function of eigenvalues (graph frequencies). To localize filters in space, the authors in [8] decided to use a polynomial parametrization

$$h_\theta(\Lambda) = \sum_{k=0}^{K-1} \theta_k \Lambda^k, \quad (3)$$

where $\theta \in \mathbb{R}^K$ is a vector of parameters. The order of the polynomial, which is equal to $K - 1$, specifies the local $K - 1$ hop neighborhood. The neighborhood is determined by the shortest path distance. The polynomial filter can be computed recursively, as a Chebyshev expansion, which is commonly used in graph signal processing to approximate kernels [33]. The Chebyshev polynomial $T_k(x)$ of order k is calculated as $T_k(x) = 2xT_{k-1}(x) - T_{k-2}(x)$ with $T_0 = 1$ and $T_1 = x$. The Chebyshev expansion applies for values that lie in $[-1, 1]$; therefore, the diagonal matrix of eigenvalues Λ has to be derived from a rescaled Laplacian $L = (D - A)/\lambda_{max} - I_n$. Thus, the filtering operation can be rewritten as

$$y = h_\theta(\Lambda)x = \sum_{k=0}^{K-1} \theta_k T_k(L)x = [\bar{x}_0, \dots, \bar{x}_{K-1}] \theta, \quad (4)$$

where $\bar{x}_k = 2L\bar{x}_{k-1} - \bar{x}_{k-2}$ with $\bar{x}_0 = x$ and $x_1 = Lx$. The transition in Eq. 4 is done according to the observation $(U\Lambda U^T)^k = U\Lambda^k U^T$. The filtering at the convolutional layer boils down to an efficient sequence of $K - 1$ sparse matrix-vector multiplications and one dense matrix-vector multiplication [8].

LRP is based on the theoretical framework of deep Taylor decomposition. The function $f(x)$ from Eq. (1) can be decomposed in terms of the Taylor expansion at some chosen root point x^* so that $f(x^*) = 0$. The first order Taylor expansion of $f(x)$ is:

$$\begin{aligned} f(x) &= f(x^*) + \sum_{g=1}^m \frac{\partial f}{\partial x} \Big|_{x=x^*} \cdot (x_g - x_g^*) + \epsilon \\ &= 0 + \sum_{g=1}^m R_g(x) + \epsilon \end{aligned} \quad (5)$$

where the relevances $R_g(x)$ are the partial differentials of the function $f(x)$. The details of how to choose a good root point are described in [11]. The $f(x)$ represents an output neuron of a neural network which consists of multiple layers and each layer consists of several neurons. A

neuron receives a weighted sum of its inputs and applies a nonlinear activation function. The idea of the deep Taylor decomposition is to perform a first order Taylor expansion at each neuron of the neural network. These expansions allow to produce relevance propagation rules that compute relevances at each layer in a backward pass. The rules redistribute the relevance from layer to layer starting from output until the input is reached. The value of the output represents the model's decision which is equal to the total relevance detected by the model.

LRP is commonly applied to deep neural networks consisting of layers with rectified linear units (ReLU) nonlinearities. In our experiments, we use only this activation function. Let i and j be single neurons at two consecutive layers at which the relevance should be propagated from j to i . The activation function has this form:

$$a_j = \max\left(0, \sum_i a_i w_{ij} + b_j\right) \quad (6)$$

where a_i , a_j are neurons' values, w_{ij} are weights, and b_j is bias. Noticeably, the layers of this type always have non-negative activations. The relevance propagation rule is the following:

$$R_i = \sum_j \frac{a_i w_{ij}^+}{\sum_i a_i w_{ij}^+ + \epsilon} R_j, \quad (7)$$

where w_{ij}^+ corresponds to the positive weights w_{ij} and ϵ stabilizes numerical computations [9]. We set ϵ to 1^{-10} . Equation (7) depicts the z^+ rule coming from deep Taylor decomposition [11]. The z^+ rule is commonly applied to the convolutional and fully connected layers. It favors the effect of only positive contributions to the model decisions. The first input layer can have other propagation rules that are specific to the domain [34]. In our work, we used the rule (7) for the input layer as well since the gene expression data has positive values.

In order to propagate relevance through the filtering (4), we rewrite it as follows:

$$y = \sum_{k=0}^{K-1} \theta_k T_k(L)x = [\bar{L}_0, \dots, \bar{L}_{K-1}] \theta x = Wx, \quad (8)$$

where matrix $W \in R^{m \times m}$ connects nodes y and x . The computation of matrix W is done as: $W = [\bar{L}_0, \dots, \bar{L}_{K-1}] \theta$, where $\bar{L}_k = 2L\bar{L}_{k-1} - \bar{L}_{k-2}$ with $\bar{L}_0 = I$ and $\bar{L}_1 = L$ are the Chebyshev polynomials of the Laplacian matrix.

Each convolutional layer has F_{in} channels

$$[x_1, \dots, x_{F_{in}}] \in R_+^{m \times F_{in}} \quad (9)$$

in the input feature map and F_{out} channels

$$[y_1, \dots, y_{F_{out}}] \in R^{m \times F_{out}} \quad (10)$$

of the output feature map. We consider the values of output feature maps before applying ReLU non-linearities on them. The $F_{in} \times F_{out}$ vectors of the Chebyshev coefficients $\theta_{i,j} \in R^k$ are the layer's trainable parameters. The input feature map can be transformed into a vector $\hat{x} = [x_1^T, \dots, x_{F_{in}}^T]^T \in R_+^{m \cdot F_{in}}$. We adapt Eq. (8) to compute the j^{th} channel of the output feature map based on the input feature map:

$$\begin{aligned} y_j &= [\bar{L}_0, \dots, \bar{L}_{K-1}] \cdot [\theta_{1,j}, \dots, \theta_{F_{in},j}] \cdot [x_1^T, \dots, x_{F_{in}}^T]^T \\ &= [\hat{L}_{1,j}, \dots, \hat{L}_{F_{in},j}] \cdot [x_1^T, \dots, x_{F_{in}}^T]^T \\ &= \hat{W}_j \times \hat{x} \in R^m \end{aligned} \quad (11)$$

where $\hat{L}_{i,j} = [\bar{L}_0, \dots, \bar{L}_{K-1}] \theta_{i,j} \in R^{m \times m}$, $\hat{W}_j = [\hat{L}_{1,j}, \dots, \hat{L}_{F_{in},j}] \in R^{m \times m \cdot F_{in}}$

Since the j^{th} channel of the output feature map is connected through the matrix-vector multiplication with the input feature map, \hat{W}_j can be treated as a matrix of weights joining two fully connected layers. Therefore, the relevance $R_y^j \in R_+^m$ from the j^{th} output channel can be propagated to the input feature map relevance $R_x^j \in R_+^{m \cdot F_{in}}$ according to the rule (7). Overall, the relevance propagated from the output feature map to the input feature map is:

$$R_{\hat{x}} = \sum_{j=1}^{F_{out}} R_x^j \in R_+^{m \cdot F_{in}}. \quad (12)$$

For running LRP on graph convolutional layers, one needs to compute huge and dense matrices \hat{W}_j . It requires $K - 2$ sparse matrix-matrix multiplications and one sparse to dense matrix-matrix multiplication. The computations for relevance propagation are heavier and much more memory demanding compared to the filtering (4). The code implementing our GLRP approach is available in [35].

GLRP on gene expression data

To demonstrate the utility of GLRP, the Graph-CNNs were trained on two gene expression datasets described in the "Gene expression data and molecular network" section. In our previous study [23], the gene expression data were standardized for the training. But in this paper, we did not standardize the data. The argument for it is the following. For the non-image data, to standardize the input features is the usual practice. However, in case of standardization, the input features are treated independently. For an image, the neighboring pixels are highly correlated. If the pixels as features are standardized across the dataset, then this can distort the pattern of the image quite

significantly and lead to misinterpretation. Analogically, feature wise standardization of microarray data changes expression patterns of genes located in the same neighborhood of a molecular network (HPRD PPI in our case). This might affect the explainability of the Graph-CNN that we aim at. Therefore, we trained the Graph-CNN directly on the quantile normalized data avoiding the additional standardization step. Instead, we subtracted the minimal value (5.84847) of the data from each cell of the gene expression matrix to keep the gene expression values non-negative. If initially, GE data was lying in [5.84847, 14.2014], now it is in the interval [0.0, 8.3529]. This transformation allows Graph-CNN to converge faster, to apply the LRP propagation rule (7) suitable for non-negative input values, and to preserve original gene expression patterns in local neighborhoods of the PPI network.

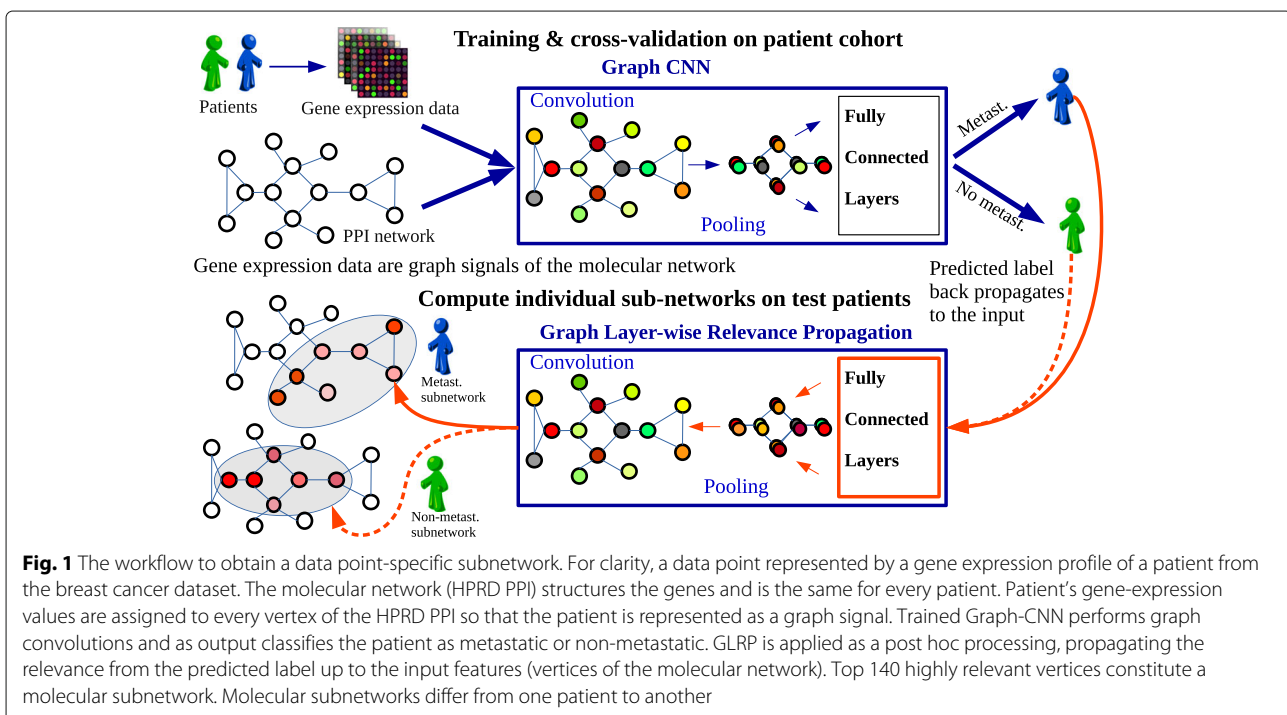
For each of the two gene expression datasets structured by the same prior knowledge (HPRD PPI), we used a 10-fold cross validation over a whole dataset to estimate the predictive performance of Graph-CNN. The hyperparameters such as the number of filters, the presence of pooling, the learning rate, and decay were tweaked manually on this 10-fold cross validation.

The architecture of the Graph-CNN trained on the HUVECs dataset and its performance are given in the “[Comparison of subnetworks derived by GLRP to gene-coexpression networks identified by WGCNA](#)” section.

For the breast cancer dataset, the Graph-CNN architecture consisted of two graph convolutional layers following

maximum pooling of size 2, and two hidden fully connected layers with 512 and 128 units respectively. Each graph convolutional layer contained 32 filters covering the vertex’ neighborhood of size 7. For the performance comparison, we trained a “glmgraph” method [36] implementing network-constrained sparse regression model using HPRD PPI network, and Random Forest without any prior knowledge as baselines. The results on 10-fold cross validations are presented in the “[GLRP to deliver patient-specific subnetworks](#)” section.

Further we generated the patient-specific (data point specific) subnetworks via GLRP. For that, each of the gene expression datasets was randomly split again: 90% training and 10% test. We retrained the Graph-CNN on 90% of data using manually selected hyperparameters from 10-fold cross validation, and propagated relevances on test data which was not “seen” by the model during training to make it more challenging. Since the LRP rule (7) propagates only positive contributions, our Graph-CNN had two output neurons for binary classification tasks that showed the probability of these two classes. For each patient in the test set, relevance was propagated by GLRP from the predicted output neuron to the input neurons representing genes (vertices) of the underlying molecular network. The workflow to deliver the patient-specific subnetworks is depicted on Fig. 1. A patient-specific subnetwork explaining the prediction was constructed from the 140 most relevant genes. Selecting more than 140 top relevant vertices entailed visualization issues. The single-



tions were deleted so that the subnetwork consisted mainly of around 130 vertices. The same workflow was applied to generate data-point-specific subnetworks for the data described in the “[Expression data of HUVECs before and after TNF \$\alpha\$ stimulation](#)” section.

Pathway analysis

Enrichment of signal transduction pathways annotated in the TRANSPATH[®] database version 2020.1 [37] in genes prioritized by GLRP were analyzed using the geneXplain platform version 6.1 [38]. The analysis based on the Fisher’s exact test [39] was carried out for gene sets obtained for individual patients from the breast cancer dataset as well as for their combination into subtype gene sets.

The following calculations were applied to investigate differences in pathway hits. Let P denote a set of pathway genes and S_i and S_k two subnetwork gene sets, so that $P_i = P \cap S_i$ and $P_k = P \cap S_k$ are the sets of pathway genes matched by the two subnetworks. The difference $\Delta P_{i,k}$ in matched pathway genes was then calculated as $|(P_i \cup P_k) \setminus (P_i \cap P_k)| / |P_i \cup P_k|$ with $|P_i \cup P_k| > 0$. For each selected pathway, we calculated $\Delta P_{i,k}$ for each pair of subnetworks and reported the median of examined pairs.

Comparison of subnetworks derived by GLRP to gene-coexpression networks identified by WGCNA

To further examine the biological relevance of subnetwork genes prioritized by GLRP and for the purpose of comparison to an already available method that uses expression and network information to prioritize gene sets, we analyzed the gene expression data described in “[Expression data of HUVECs before and after TNF \$\alpha\$ stimulation](#)” section. We compared gene sets identified in our subnetworks to gene modules and differentially expressed genes in response to TNF α identified by Rhead et al. [32]. Rhead et al. [32] reported gene modules obtained by weighted gene co-expression network analysis (WGCNA). The method has been applied in many studies and constructs a gene network based on expression measurements from which it can derive modules of co-expressed genes [40]. We trained a Graph-CNN on the gene expression data to classify the TNF α treatment status of HUVECs. The Graph-CNN architecture consisted of 2 convolutional layers with 4 and 8 filters respectively followed by one hidden fully connected layer with 128 nodes. The vertex’s neighborhood covered by graph convolutions was of size 7. No pooling was used. The performance of the Graph-CNN in 10-fold cross validation: mean $100 \cdot \text{AUC}$, accuracy, and F1-weighted were 99.49, 96.25% and 96.06%, respectively. A random forest achieved the same performance. We generated the subnetworks according to the “[GLRP on gene expression data](#)” section, retrained the Graph-CNN on 70 randomly

selected samples, and applied GLRP on 8 test samples (4 treated and 4 not treated). The test samples were predicted correctly. For each of the 8 test samples, we constructed a subnetwork. Associations between subnetwork genes sets and 16 gene modules defined by Rhead et al. [32] as well as 589 upregulated genes (log-fold change > 0.5 , FDR < 0.01), 425 downregulated genes (log-fold change < -0.5 , FDR < 0.01), and the combined set of 1014 DE genes were analyzed using the *Functional classification* tool of the geneXplain platform [41]. Fisher test calculations were carried out with a total contingency table count corresponding to the number of genes in [32, file S1 of] after mapping to Ensembl [42] gene ids (10022 genes). Rhead et al. [32] assigned a color code to the 16 gene co-expression modules and denoted them as *black, blue, brown, cyan, green, greenyellow, grey, magenta, midnight-blue, pink, purple, red, salmon, tan, turquoise, and yellow* which is maintained in results reported here.

Results

Sanity check of the implemented graph LRP

To initially validate our implemented LRP, we applied Graph-CNN on the MNIST dataset [43] in the same way as described in the paper [8]. The MNIST dataset contains 70,000 images of hand-written digits each having a size of 28 by 28 pixels. To apply Graph-CNN on the image data, we constructed an 8 nearest-neighbors graph similarly to the schema proposed in [8], with the exception that all the weights are equal to 1. The weight 1 is more natural for the graph connecting neighboring image pixels. Thus, each image is a graph signal represented by node attributes—pixel values. We achieved high classification accuracy (99.02%) on the test set for the Graph-CNN, which is comparable to the performance of classical CNN (99.33%) reported in [8]. The number of parameters was the same for both methods.

Usually, to manage box-constrained pixel values, the special pixel-specific LRP rule is applied for the input layer. This pixel-specific rule highlights not only the digits itself, but also the contours of the digits [34, Fig. 13 of]. In contrast, the rule (7) highlights only those positively relevant parts of the image where the signal of the digit is present. We kept the propagation rule (7) for the input and all other layers in all our experiments. Further, we visually compared on the same digits how the heatmaps generated by implemented GLRP correspond to the heatmaps generated by usual LRP procedure applied on classical CNN (Fig. 2).

The heatmaps were rendered only for the classes predicted by classical CNN and Graph-CNN. In this case, the classes are “6” and “3”. For the Graph-CNN, a bigger part of the digit is relevant for the classification since the covered neighborhood can be expanded up to 24 hops. Graph-CNN’s filters are isotropic; thus, they tend to cover




Fig. 2 From left to right: initial image, LRP on classical CNN and GLRP on Graph-CNN

roundish areas that concern rounded patterns (curves) of the digit (Additional file 1: Fig. S1).

Genes selected by GLRP correlate with modules identified by gene co-expression network analysis

In the analysis of TNF-induced gene expression changes in HUVECs, our procedure prioritized in total 168 genes of which 105 genes were found in subnetworks of all eight test samples (Additional file 2). Remarkably, the *green* gene module, which was the most strongly correlated one with TNF α upregulation [32], showed significant association (adjusted p value < 0.05) with the combined set of subnetwork genes, with genes found in the majority of subnetworks and also with 5 of the 8 subnetworks (Additional file 2). At the same significance level, the *turquoise* gene module described in [32] was strongly associated with 2 of 8 subnetworks and with genes found in all 8 subnetworks. In addition, both the *green* and the *turquoise* modules showed moderate association (adjusted p value < 0.1) with the majority of gene sets defined on the basis of the test subnetworks. Furthermore, we found strong (adjusted p value < 0.05) or moderately (adjusted p value < 0.1) significant overlap between upregulated genes and some subnetwork gene sets. The gene modules *cyan*, *greenyellow*, and *midnightblue* did not overlap with GLRP-derived subnetworks. These results demonstrate partial agreement between gene sets suggested by GLRP, another gene network analysis and classical differential expression analysis. Hence, the GLRP-based subnetworks gathered biologically meaningful genes and may even complement other approaches in revealing important properties of the underlying biological systems. Additionally, another two gene sets were compared with WGCNA modules: the intersection of subnetworks genes and genes that occurred in more than in 4 test samples subnetworks. Notably, the individual subnetworks shared more genes with the *green* and *turquoise* WGCNA modules than

those described gene sets, pointing out the ability of GLRP to identify sample-specific genes.

GLRP to deliver patient-specific subnetworks

We applied the GLRP to the Graph-CNN trained on gene expression data from the “Breast cancer data” section. The gene expression data was structured by a protein-protein interaction network. The standardization of features was not performed as described in the “GLRP on gene expression data” section. The prediction task performed by the Graph-CNN was to classify patients into 2 groups, metastatic and non-metastatic. The results of a 10-fold cross validation are depicted in Table 1. While Graph-CNN and glmgraph utilized the HPRD PPI network topology, a random forest did not use any prior knowledge. glmgraph was not evaluated on non-standardized data, since it had convergence issues in this case. The metrics were averaged over folds and the standard errors of their means were calculated.

The GLRP was applied as described in the “GLRP on gene expression data” section. We retrained the Graph-CNN on 872 patients and generated relevances for 97 test patients. The relevances were propagated from the Graph-CNN’s output node corresponding to the correctly predicted class. The most frequently selected features are summarized in Additional file 1: Table S1. The eukaryotic translation elongation factor EEF1A1, which is overexpressed in the majority of breast cancers and protects tumor cells from proteotoxic stress [44], was the sole factor that was selected in all of the 97 test set patients. Other frequently selected features in both non-metastatic as well as metastatic patients included genes such as the epithelial-to-mesenchymal-transition (EMT)-related gene VIM (46/58 non-metastatic, 30/39 metastatic patients), the extracellular matrix protein FN1 (43/58 non-metastatic, 22/39 metastatic patients), the actin cytoskeleton regulator CFL1 (7/58 non-metastatic, 7/39 metastatic

Table 1 Performance of Graph-CNN on metastatic event prediction, depending on normalization

Method	Std	100*AUC	Accuracy, %	F1-weighted, %
Graph-CNN	-	82.57±1.25	76.07±1.30	75.82±1.33
Random Forest	-	81.27±1.66	74.23±1.73	73.47±1.84
Graph-CNN	+	82.16±1.25	76.18±1.36	75.86±1.35
Random Forest	+	81.40±1.76	74.74±1.67	74.00±1.82
glmgraph	+	80.88±1.37	75.14±1.30	74.73±1.39

Std stands for standardization of features (genes)

patients), and the estrogen receptor ESR1 (28/58 non-metastatic, 10/39 metastatic patients) that are all known to be linked with breast cancer development and progression [45–48]. This indicates that our method successfully identified relevant key players with a general role in breast tumorigenesis.

Additionally, we show individualized PPI subnetworks delivered for four correctly predicted breast cancer patients (Table 2) from the microarray data set. Two of them had been assigned with the most common subtype luminal A (LumA), while the other two suffered from the highly aggressive basal-like subtype. In each group, one patient with early metastasis was picked and one who did not develop any within at least 5 years of follow-up.

The generated PPI subnetworks are displayed in Fig. 3. The sequence of pictures in order ABCD is the same as in the table.

Interestingly, the networks of both LumA patients contained ESR1 which fits well since this subtype is considered as estrogen receptor positive [49]. In contrast, genes often associated with the basal-like subtype and a poor prognosis such as MCL1, CTNNA1, EGFR, or SOX4 were found in the basal-like patient GSM519217 suggesting that the generated networks are capable of extracting breast cancer subtype-specific features. The comparison of the subnetworks of the non-metastatic and the metastatic patients furthermore revealed some patient-specific genes which might give valuable information about specific mechanisms of tumorigenesis and therapeutic vulnerabilities in the respective patient. In general, it seemed that the subnetworks of the non-metastatic patients contained more genes that have been

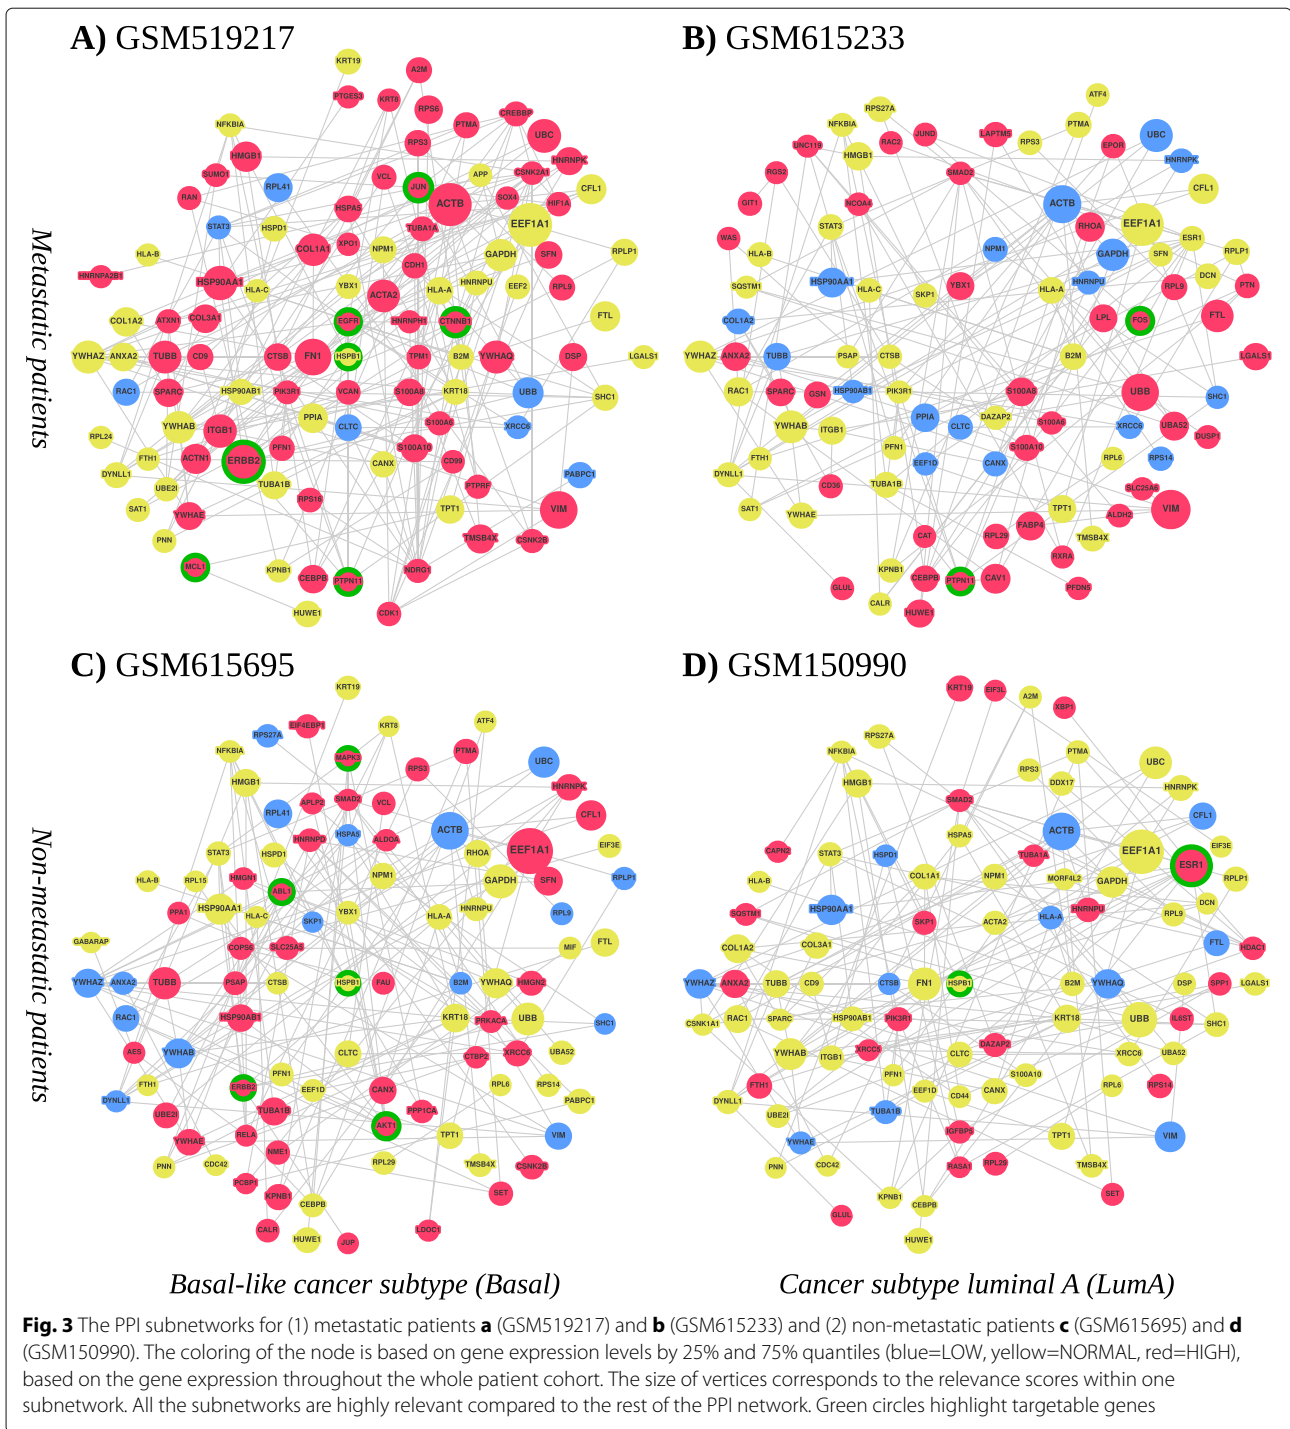
Table 2 Patients that the PPI subnetworks are generated for

Patient's ID	Subtype	Metastatic event	Time of metastases, years	Last follow-up, years
GSM519217	Basal	1	0.9	-
GSM615233	LumA	1	0.79	-
GSM615695	Basal	0	-	5.38
GSM150990	LumA	0	-	9.93

linked to better prognostic outcomes such as JUP, PCBP1, and HMG2 in GSM615695 [50–52] or RASA1, IL6ST, KRT19, and RPS14 in GSM150990 [53–56], while the networks of both metastatic patients harbored genes that are known to be involved in aggressive tumor growth or therapy resistance which might explain the early metastatic spread in these patients. Some examples are CDK1, SFN, and XPO1 in GSM519217 [57–59] or CAV1, PTPN11, and FTL in GSM615233 [60–62].

However, not only the presence of specific genes might be important, but also their overall expression level. Our analyses identified, e.g., the EMT-related gene VIM as one of the most relevant nodes in the subnetworks of both metastatic patients in which the gene was highly expressed (> 75% quantile based on the gene expression throughout the whole patient cohort). In contrast, VIM was also present in the subnetworks of the two non-metastatic patients, however, with a lower relevance and a particularly low expression (< 25% quantile). VIM is an important marker for EMT and high expression levels correlate with a motile, mesenchymal-like cancer cell state, thus making VIM an essential effector of metastasis [45].

A comparison of subnetwork genes of 79 correctly predicted test set patients to a database of signal transduction pathways confirmed significant enrichment of pathways that have previously been associated with cancer disease mechanisms such as the EGF, ER-alpha, p53, and TGFbeta pathways as well as Caspase and beta-catenin networks. Comparisons were performed for each patient as well as for subtype gene sets formed by combining subnetwork genes of patients associated with a breast cancer subtype. Results for the 238 signaling pathways from the TRANSPATH® database that were significantly enriched with subtype genes are visualized in Fig. 4. Differences in enrichment significance may suggest that the importance of some signaling pathways detected this way is subtype-specific, e.g., for YAP ubiquitination or the VE-cadherin network (orange heatmap, Fig. 4, see also Additional file 1: Table S2 for details). The pattern of enrichment found on the level of cancer subtypes coincided well with the findings for subnetwork genes of individual patients revealing several molecular networks with elevated significance in both subtype and patient gene sets such as the EGF pathway, although the patient-level visualization did not suggest subtype-specific enrichment (green heatmap, Fig. 4). One source of these observations can be that patient subnetworks tend to be associated with certain pathways but cover different pathway components (genes). We therefore compared pathway genes in pairs of patient subnetworks for the 33 largest pathways. In 18 pathways, the median pair of patient subnetworks differed in 33% or more of the genes matched within a pathway (see also Additional file 1: Table S3 for details). These results demonstrate that the subnetworks obtained by




Graph-CNN were enriched with common signaling pathways relevant for the respective disease and can assign patient-specific priorities to pathway components.

Finally, we tested whether the subnetworks can also be used for finding potentially targetable genetic vulnerabilities that could open new options for personalized treatment decisions. We applied the “MTB report”

methodology described in [63] to identify actionable genes present in the subnetworks. For that, we extended the algorithm to match high expression with gain of function alterations, and low expression with loss of function alterations. The results are summarized in Table 3.

Although information about the presence of actionable genetic variants is missing from our patient microarray



data, the information generated by the PPI subnetworks could be used to define specific panels for subsequent sequencing. Indeed, the MTB reports highlighted specific genes that could be targeted therapeutically in each of the four patients: In the non-metastatic LumA patient GSM150990 ESR1 was proposed as therapeutic target

which is in line with current treatment regimens that use hormone therapy as the main first-line treatment of choice for this patient subgroup. In contrast, in the metastatic LumA patient GSM615233 FOS and PTPN11 were identified as novel actionable alterations. In the often rapidly relapsing basal-like patients HSPB1 and ERBB2 were

Table 3 Actionable genes identified by the MTB report workflow

Patient	Gene	Expression	Known Var	Predicts
615695	HSPB1	Normal	expression	Response to gemcitabine
	ABL1	High	GoF	Response to ABL TK inhibitors (imatinib, dasatininb, ponatinib, regorafenib. . .)
	AKT1	High	GoF	Response to PI3K, AKT, MTOR inhibitors; resistance to BRAF inhibitors
	ERBB2	High	GoF	Response to ERBB2, EGFR, MTOR, AKT inhibitors
	MAPK3	High	GoF	Resistance to EGFR inhibition
519217	HSPB1	Normal	expression	Response to gemcitabine
	CTNNB1	High	GoF	Response to everolimus + letrozole; resistance to Tankyrase inhibitors
	EGFR	High	GoF	Response to EGFR, ERBB2, HSP90 and MEK inhibitors
	ERBB2	High	GoF	Response to ERBB2, EGFR, MTOR, AKT inhibitors
	JUN	High	overexpr	Response to irbesartan (angiotensin II antagonist)
	MCL1	High	GoF	Resistance to anti-tubulin agents
615233	PTPN11	High	GoF	Response to MEK inhibitors
	FOS	High	overexpr	Response to irbesartan (angiotensin II antagonist)
150990	PTPN11	High	GoF	Response to MEK inhibitors
	HSPB1	Normal	expression	Response to gemcitabine
	ESR1	High	GoF	Response to novel ER degraders, fulvestrant, tamoxifen

Genes from the PPI subnetworks were matched to known genomic alterations (Known Var) that predict either response or resistance to drugs (Predicts). High and low gene expression were matched to gain of function (GoF) and loss of function (LoF) genomic variants, respectively

identified as common targets as well as MAPK3, AKT1, and ABL1 for the non-metastatic patient GSM615695 or EGFR, MCL1, CTNNB1, PTPN11, and JUN for the metastatic patient GSM519217, thereby suggesting novel possibilities for combinatory or alternative treatments. Taken together, GLRP provides subnetworks centered around known oncogenic drivers that seem reasonable in the context of cancer biology and can help to identify patient-specific cancer dependencies and therapeutic vulnerabilities in the context of precision oncology.

Discussion

In our work, we focused on the interpretability of a deep learning method utilizing molecular networks as prior knowledge. We implemented LRP for Graph-CNN and provided the sanity check of the developed approach on the MNIST dataset. Essentially, the main aim of the paper was to explain the prediction of metastasis for breast cancer patients by providing an individual molecular subnetwork specific for each patient. The patient-specific subnetworks provided interpretability of the deep learning method and demonstrated clinically relevant results on the breast cancer dataset.

Supposedly, the performance of Graph-CNN can be improved. The batch normalization technique [64] that is used to accelerate the training of deep neural networks is not seen to be available for the Graph-CNN, so this can be the way to enhance its performance. The LRP rule for batch normalization layers is yet another procedure to be adapted for Graph-CNN.

Another possibility to identify genes (and construct subnetworks out of them) influencing classifier decisions is to apply model-agnostic SHAP and LIME explanation methods. LIME method provides explanations of a data point based on feature perturbations. The method samples perturbations from a Gaussian distribution, ignoring correlations between features. It leads to the instability of explanations that is not favorable for personalized medicine. SHAP provides Shapley values for each feature of a data point as well but does not have such an issue, so we attempted to derive patients-specific subnetworks applying TreeExplainer and KernelExplainer from SHAP python module on Random Forest and Graph-CNN respectively. The subnetworks were built on the basis of HPRD PPI utilizing positive Shapley values, which were pushing prediction to a higher probability of corresponding class (metastatic or non-metastatic). The subnetworks obtained were mostly consisting from single vertices. In contrast, the subnetworks from GLRP and Graph-CNN were mostly connected. The SHAP's DeepExplainer approach suitable for convenient deep learning models is not applicable for Graph-CNN. The model-agnostic KernelExplainer computes SHAP values out of a debiased lasso regression. Reevaluating the model

happens several thousands numbers of times specified by a user as well as a small background dataset is needed for integrating out features. Hence, the KernelExplainer is not scalable and application of it on Graph-CNN resulted in not connected subnetworks as well.

Furthermore, the sensitivity of Graph-CNN to the changes of prior knowledge is still to be investigated. Authors in [8] showed that for the MNIST images a random graph connecting pixels significantly decreases the performance destroying local connectivity. In our case, the permutation of the vertices of the PPI network does not influence the classifier performance on standardized gene expression data. Yet, PPI network is a small world network and its degree distribution fits to the power law with the exponent $\alpha = 2.70$. It implies great connectivity between proteins and means that any two nodes are separated by less than six hops. The filters of convolutional layers cover a 7-hop neighborhood of each vertex, so we assume it still might be enough to capture the gene expression patterns. In our future work, we will investigate how the properties of the prior knowledge influence the performance and explainability of Graph-CNN.

The subnetworks generated by GLRP contained common potential oncogenic drivers which indicates that they can extract the essential cancer pathways. Indeed, our analyses identified genes associated with hormone receptor-positive breast cancer (e.g. ESR1, IL6ST, CD36, GLUL, RASA1) in the networks from the patients with estrogen receptor positive, Luma breast cancer and genes associated with the basal-like subtype (e.g., EGFR, SOX4, AKT1 as well as high levels of HNRNP1K) in the basal-like patients, underlining the biological relevance of the networks. Next to subtype-specific genes, the networks contained several oncogenes that were found in all four patients and could thus represent common drivers of breast cancer initiation and progression. One example is the actin-binding protein cofilin (CFL1) that regulates cancer cell motility and invasiveness [46]. Another interesting candidate is STAT3 which is activated in more than 40% of breast cancers and can cause deregulated cell proliferation and epithelial-to-mesenchymal transition (EMT) [65]. Our graphs not only displayed patient-specific PPI subnetworks, but also concisely visualized the relevance of each node and its expression levels. This information is potentially relevant to judge the biological significance of the gene in a patient-specific context.

Next to the common genes found in all four networks, each network was characterized by several special, cancer-associated genes which are of high interest because they might represent patient-specific central signaling nodes and therapeutic vulnerabilities. Some examples are PTPN11 that is known to activate a transcriptional program associated with cancer stem cells or the EMT-related genes SOX4 or VIM that might be responsible for the high

invasive capacity of the tumors and their early metastasis formation [45, 61, 66, 67]. Interestingly, the network of the metastatic patient GSM615233 harbored the genes FABP4 and LPL which both have been shown to interact with CD36, another highly expressed node in the network, to support cell proliferation and counteract apoptosis [68–70]. In contrast, in the non-metastatic patient GSM150990 especially the interleukin receptor IL6ST and the Ras GTPase-activating protein 1 (RASA1) seem to be interesting because for both high expression levels have been linked with a favorable prognosis [53, 54]. In the other non-metastatic patient GSM615695 high levels of HMG2 and PCBP1 were identified which both have been shown to be able to inhibit cell proliferation [51, 52]. Although the experimental validation for the networks is still missing, it is tempting to speculate that these genes might contribute to the benign phenotype of the tumor in these patients.

All patient-specific subnetworks contained relevant drug targets that have been largely studied in breast cancer (e.g., ERBB2, ESR1, EGFR, AKT1). Yet, resistance mechanisms in breast cancer targeted therapies represent a big challenge; many of the identified therapeutic approaches have failed [71] due to the highly interconnected nature of signaling pathways and potential circumvents. A promising way forward could involve the molecular characterization of the tumor with transcriptomics and a parallel culture of patient-derived organoids. PPI networks could elucidate the right combination strategy by identifying central signaling nodes. Different therapeutic strategies could be tested on organoids and confirm the best strategy that synergistically blocks cancer cell escape routes and minimizes the emergence of survival mechanisms. Only the identification of relevant mechanisms of action for cell survival as well as of the factors involved in resistance for each patient, together with a more precise and personalized characterization of each cancer phenotype, may provide useful improvements in current therapeutic approaches.

Conclusions

We present a novel Graph-CNN-based feature selection method that benefits from prior knowledge and provides patient-specific subnetworks. We adapted the existing Layer-wise Relevance Propagation technique to the Graph-CNN, demonstrated it on MNIST data, and showed its applicability on a large breast cancer dataset. Our new approach generated individual patient-specific molecular subnetworks that influenced the model's decision in the given context of a classification problem. The subnetworks selected by the developed method utilizing general prior knowledge are relevant for prediction of metastasis in breast cancer. They contain common as well as subtype-specific cancer genes that match the

clinical subtype of the patients, together with patient-specific genes that could potentially be linked to aggressive/benign phenotypes. In the context of a breast cancer dataset GLRP provides patient-specific explanations for the Graph-CNN that largely agree with clinical knowledge, include oncogenic drivers of tumor progression, and can help to identify therapeutic vulnerabilities. We therefore conclude that our method GLRP in combination with Graph-CNN is a new, useful, and interpretable ML approach for high-dimensional genomic data-sets. Generated classifiers rely on prior knowledge of molecular networks and can be interpreted by patient-specific subnetworks driving the individual classification result. These subnetworks can be visualized and interpreted in a biomedical context on the individual patient level. This approach could thus be useful for precision medicine approaches such as for example the molecular tumor-board.

Supplementary Information

The online version contains supplementary material available at <https://doi.org/10.1186/s13073-021-00845-7>.

Additional file 1: Contains Supplementary Tables S1–S3 and Supplementary Figures S1, S2.

Additional file 2: Subnetwork genes obtained for 8 test samples and analysis of their association with gene modules reported by [32] as well as differentially expressed (DE) genes.

Worksheet *Subnetwork genes 8 samples* provides identifiers and gene symbols of 167 subnetwork genes, in how many and in which samples they were selected. Worksheet *Gene module enrichment* presents results of Fisher test calculations comparing subnetwork gene sets to gene modules and DE gene sets. Each row contains data for a DE gene set or a gene module consisting of the total group size and column triplets with *p*-value, adjusted *p*-value as well as the number of hits, respectively, observed in comparisons to the union of genes from 8 subnetworks, the set of genes occurring in the majority, the set of genes found in all of the subnetworks and each of the 8 samples. Highlighted are rows corresponding to *green* and *turquoise* gene modules, which were most often significantly associated with subnetwork gene sets (grey), adjusted *p*-values below 0.05 (red) and between 0.05 and 0.1 (yellow).

Abbreviations

ML: Machine learning; LRP: Layer-wise Relevance Propagation; GNN: Graph Neural Network; CNN: Convolutional Neural Network; GLRP: Graph Layer-wise Relevance Propagation; WGCNA: Weighted gene co-expression network analysis; GDPR: General Data Protection Regulation; GEO: Gene Expression Omnibus; HPRD: Human Protein Reference Database; PPI: Protein-protein interaction; ReLU: Rectified linear unit; HUVEC: Human umbilical vein endothelial cells; EMT: Epithelial-to-mesenchymal transition

Acknowledgements

We would like to acknowledge Michaela Bayerlová, Mark Gluzman, and Vladyslav Yushchenko for fruitful discussions. HC is a member of the International Max Planck Research School for Genome Science, part of the Göttingen Graduate Center for Neurosciences, Biophysics, and Molecular Biosciences. TB is a member of the Göttingen Campus Institute Data Science.

Authors' contributions

HC and TB designed the study. HC developed and implemented the approach and performed the computational experiments. AB, FK, and PS provided major contributions to the study design. AB and KM provided clinical insights as well as JPB and PS provided biological insights, performing analyses of patient-specific subnetworks. FA developed the web-site to visualize the subnetworks.

TB and AL provided machine learning insights. HC, TB, AL, KM, and PS wrote the manuscript. All authors read and approved the final manuscript.

Funding

This work was funded by the German Ministry of Education and Research (BMBF) e:Med project *MyPathSem* (031L0024) and the project *MTB-Report* by the big data initiative of the Volkswagenstiftung. KM was supported by German Research Foundation (DFG) project 424252458. We acknowledge support by the Open Access Publication Funds of the Göttingen University.

Availability of data and materials

The utilized breast cancer datasets are accessible from Gene Expression Omnibus (GEO) [28] data repository (accession numbers GSE25066, GSE20685, GSE19615, GSE17907, GSE16446, GSE17705, GSE2603, GSE11121, GSE7390, GSE6532). The HUVECs gene expression data [32] is available in GEO database (GSE144803). The HPRD PPI network can be found in [26]. The preprocessed breast cancer data, the adjacency matrix of the HPRD PPI network, and the code of the GLRP method are provided in <http://mypathsem.bioinf.med.uni-goettingen.de/resources/qlrp> [31] and <https://gitlab.gwdg.de/UKEBpublic/graph-qlrp> [35]. The web-site to explore patient-specific subnetworks is in <http://mypathsem.bioinf.med.uni-goettingen.de/MetaRelSubNetVis/> [72].

Ethics approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Competing interests

PS is an employee of geneXplain GmbH, Germany. The remaining authors declare that they have no competing interests.

Author details

¹Medical Bioinformatics, University Medical Center Göttingen, Göttingen, Germany. ²Dept. of Medicine A (Hematology, Oncology, Hemostaseology and Pulmonology), University Hospital Münster, Münster, Germany. ³Hospital del Mar Medical Research Institute (IMIM), Barcelona, Spain. ⁴geneXplain GmbH, Wolfenbüttel, Germany. ⁵IT Infrastructure for Translational Medical Research, University of Augsburg, Augsburg, Germany. ⁶Medical Statistics, University Medical Center Göttingen, Göttingen, Germany. ⁷Campus-Institute Data Science (CIDAS), University of Göttingen, Göttingen, Germany.

Received: 26 August 2020 Accepted: 5 February 2021

Published online: 11 March 2021

References

- Perera-Bel J, Leha A, Beißbarth T. In: Badve S, Kumar GL, editors. Bioinformatic methods and resources for biomarker discovery, validation, development, and integration. Cham: Springer; 2019, pp. 149–64. https://doi.org/10.1007/978-3-319-95228-4_11.
- Rivenbark AG, O'Connor SM, Coleman WB. Molecular and cellular heterogeneity in breast cancer: challenges for personalized medicine. *Am J Pathol.* 2013;183(4):1113–24. <https://doi.org/10.1016/j.ajpath.2013.08.002>.
- Sørli T. Molecular classification of breast tumors: toward improved diagnostics and treatments. In: Target Discovery and Validation Reviews and Protocols. Totowa: Humana Press; 2007. p. 91–114. <https://doi.org/10.1385/1-59745-165-7-91>.
- Fragomeni SM, Sciallis A, Jeruss JS. Molecular subtypes and local-regional control of breast cancer. *Surg Oncol Clin N Am.* 2018;27(1):95–120. <https://doi.org/10.1016/j.soc.2017.08.005>.
- Porzelius C, Johannes M, Binder H, Beißbarth T. Leveraging external knowledge on molecular interactions in classification methods for risk prediction of patients. *Biom J.* 2011;53(2):190–201. <https://doi.org/10.1002/bimj.201000155>, Accessed 01 Dec 2020.
- Johannes M, Brase JC, Fröhlich H, Gade S, Gehrmann M, Fälth M, Sülthmann H, Beißbarth T. Integration of pathway knowledge into a reweighted recursive feature elimination approach for risk stratification of cancer patients. *Bioinformatics.* 2010;26(17):2136–44. <https://doi.org/10.1093/bioinformatics/btq345>.
- Monti F, Boscaini D, Masci J, Rodola E, Svoboda J, Bronstein MM. Geometric deep learning on graphs and manifolds using mixture model cnns. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR); 2017. p. 5115–24.
- Defferrard M, Bresson X, Vandergheynst P. Convolutional neural networks on graphs with fast localized spectral filtering. In: Proceedings of the 30th International Conference on Neural Information Processing Systems (NIPS); 2016. p. 3844–52.
- Yang Y, Tresp V, Wunderle M, Fasching PA. Explaining therapy predictions with layer-wise relevance propagation in neural networks. In: 2018 IEEE International Conference on Healthcare Informatics (ICHI); 2018. p. 152–62. <https://doi.org/10.1109/ICHI.2018.00025>.
- Parliament and C. of the European Union. General data protection regulation. 2016. <https://gdpr-info.eu/>.
- Montavon G, Lapuschkin S, Binder A, Samek W, Müller K-R. Explaining nonlinear classification decisions with deep Taylor decomposition. *Pattern Recogn.* 2017;65:211–22. <https://doi.org/10.1016/j.patcog.2016.11.008>.
- Ribeiro MT, Singh S, Guestrin C. “Why Should I Trust You?”: explaining the predictions of any classifier. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD); 2016. p. 1135–44.
- Lundberg S, Lee S-I. A unified approach to interpreting model predictions. In: Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS); 2017. p. 4768–77.
- Zeiler MD, Fergus R. Visualizing and understanding convolutional networks. In: Fleet D, Pajdla T, Schiele B, Tuytelaars T, editors. Computer Vision – ECCV. Cham: Springer; 2014. p. 818–33. https://doi.org/10.1007/978-3-319-10590-1_53.
- Bach S, Binder A, Montavon G, Klauschen F, Müller K-R, Samek W. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLoS ONE.* 2015;10(7):0130140. <https://doi.org/10.1371/journal.pone.0130140>.
- Klauschen F, Müller K-R, Binder A, Bockmayr M, Hägele M, Seegerer P, Wienert S, Pruner G, de Maria S, Badve S, Michiels S, Nielsen TO, Adams S, Savas P, Symmans F, Willis S, Gruosso T, Park M, Haiße-Kains B, Gallas B, Thompson AM, Cree I, Sotiriou C, Solinas C, Preusser M, Hewitt SM, Rimm D, Viale G, Loi S, Loibl S, Salgado R, Denkert C. Scoring of tumor-infiltrating lymphocytes: from visual estimation to machine learning. *Semin Cancer Biol.* 2018;52:151–7. <https://doi.org/10.1016/j.semcancer.2018.07.001>. Immuno-oncological biomarkers.
- Binder A, Bockmayr M, Hägele M, Wienert S, Heim D, Hellweg K, Stenzinger A, Parlow L, Budczies J, Goeppert B, Treue D, Kotani M, Ishii M, Dietel M, Hocke A, Denkert C, Müller K-R, Klauschen F. Towards computational fluorescence microscopy: machine learning-based integrated prediction of morphological and molecular tumor profiles. *arXiv:1805.11178 [cs]*. 2018.
- Xie S, Lu M. Interpreting and understanding graph convolutional neural network using gradient-based attribution method. *arXiv:1903.03768 [cs]*. 2019. Accessed 12 July 2020.
- Schwarzenberg R, Hübner M, Harbecke D, Alt C, Hennig L. Layerwise relevance visualization in convolutional text graph classifiers. *arXiv:1909.10911 [cs]*. 2019. Accessed 06 Nov 2020.
- Pope PE, Kolouri S, Rostami M, Martin CE, Hoffmann H. Explainability methods for graph convolutional neural networks. In: 2019 Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR); 2019. p. 10764–73. <https://doi.org/10.1109/CVPR.2019.01103>. ISSN: 2575-7075.
- Kipf TN, Welling M. Semi-supervised classification with graph convolutional networks. *arXiv:1609.02907 [cs, stat]*. 2016. Accessed 09-01-2017.
- Ying R, Bourgeois D, You J, Zitnik M, Leskovec J. GNNExplainer: generating explanations for graph neural networks. *Adv Neural Inf Process Syst.* 2019;32:9240–51.
- Chereda H, Bleckmann A, Kramer F, Leha A, Beißbarth T. Utilizing molecular network information via graph convolutional neural networks to predict metastatic event in breast cancer. *Stud Health Technol Inform.* 2019;267:181–6. <https://doi.org/10.3233/SHT1190824>.
- Rhee S, Seo S, Kim S. Hybrid approach of relation network and localized graph convolutional filtering for breast cancer subtype classification. In: Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence. Stockholm, Sweden: International Joint Conferences on Artificial Intelligence Organization; 2018. p. 3527–34. <https://doi.org/10.24963/ijcai.2018/490>. <https://www.ijcai.org/proceedings/2018/490>.

25. Schnake T, Eberle O, Lederer J, Nakajima S, Schütt KT, Müller K-R, Montavon G. XAI for graphs: explaining graph neural network predictions by identifying relevant walks. arXiv:2006.03589 [cs, stat]. 2020. Accessed 29 Oct 2020.
26. Keshava Prasad TS, Goel R, Kandasamy K, Keerthikumar S, Kumar S, Mathivanan S, Telikicherla D, Raju R, Shafreen B, Venugopal A, Balakrishnan L, Marimuthu A, Banerjee S, Somanathan DS, Sebastian A, Rani S, Ray S, Harys Kishore CJ, Kanth S, Ahmed M, Kashyap MK, Mohmood R, Ramachandra YL, Krishna V, Rahiman BA, Mohan S, Ranganathan P, Ramabadrans S, Chaerkady R, Pandey A. Human protein reference database?2009 update. *Nucleic Acids Res.* 2009;37:767–72. <https://doi.org/10.1093/nar/gkn892>.
27. Bayerlová M, Menck K, Klemm F, Wolff A, Pukrop T, Binder C, Reißbarth T, Bleckmann A. Ror2 signaling and its relevance in breast cancer progression. *Front Oncol.* 2017;7:135. <https://doi.org/10.3389/fonc.2017.00135>.
28. Barrett T, Wilhite SE, Ledoux P, Evangelista C, Kim IF, Tomashevsky M, Marshall KA, Phillippy KH, Sherman PM, Holko M, Yefanov A, Lee H, Zhang N, Robertson CL, Serova N, Davis S, Soboleva A. NCBI GEO: archive for functional genomics data sets—update. *Nucleic Acids Res.* 2013;41(Database issue):991–5. <https://doi.org/10.1093/nar/gks1193>.
29. Irizarry RA, Hobbs B, Collin F, Beazer-Barclay YD, Antonellis KJ, Scherf U, Speed TP. Exploration, normalization, and summaries of high density oligonucleotide array probe level data. *Biostatistics.* 2003;4(2):249–64. <https://doi.org/10.1093/biostatistics/4.2.249>.
30. Gendoo DMA, Ratanasirigulchai N, Schröder MS, Paré L, Parker JS, Prat A, Haibe-Kains B. Genefu: an R/Bioconductor package for computation of gene expression-based signatures in breast cancer. *Bioinformatics.* 2016;32(7):1097–9. <https://doi.org/10.1093/bioinformatics/btv693>.
31. Bayerlová M, Chereda H. Preprocessed breast cancer data. 2020. <http://mypathsem.bioinf.med.uni-goettingen.de/resources/glrp>.
32. Rhead B, Shao X, Quach H, Ghai P, Barcellos LF, Bowcock AM. Global expression and CpG methylation analysis of primary endothelial cells before and after TNF α stimulation reveals gene modules enriched in inflammatory and infectious diseases and associated DMRs. *PLoS ONE.* 2020;15(3):0230884. <https://doi.org/10.1371/journal.pone.0230884>.
33. Hammond DK, Vandergheynst P, Gribonval R. Wavelets on graphs via spectral graph theory. *Appl Comput Harmon Anal.* 2011;30(2):129–50. <https://doi.org/10.1016/j.acha.2010.04.005>.
34. Montavon G, Samek W, Müller K-R. Methods for interpreting and understanding deep neural networks. *Digit Signal Process.* 2018;73:1–15. <https://doi.org/10.1016/j.dsp.2017.10.011>.
35. Chereda H. Graph layer-wise relevance propagation (GLRP). Gitlab. 2020. <https://gitlab.gwdg.de/UKEBpublic/graph-lrp>.
36. Chen L, Liu H, Kocher J-PA, Li H, Chen J. glmgraph: an R package for variable selection and predictive modeling of structured genomic data. *Bioinformatics.* 2015;31(24):3991–3. <https://doi.org/10.1093/bioinformatics/btv497>.
37. Krull M, Voss N, Choi C, Pistor S, Potapov A, Wingender E. TRANSPATH $\text{\textcircled{R}}$: an integrated database on signal transduction and a tool for array analysis. *Nucleic Acids Res.* 2003;31(1):97–100. <http://dx.doi.org/10.1093/nar/gkg089>. <https://academic.oup.com/nar/article-pdf/31/1/97/7127458/gkg089.pdf>.
38. Koschmann J, Bhar A, Stegmaier P, Kel A, Wingender E. “Upstream analysis”: an integrated promoter-pathway analysis approach to causal interpretation of microarray data. *Microarrays.* 2015;4(2):270–86. <https://doi.org/10.3390/microarrays4020270>.
39. Fisher RA. On the interpretation of χ^2 from contingency tables, and the calculation of P. *J R Stat Soc.* 1922;85(1):87–94. <https://doi.org/10.2307/2340521>.
40. Zhang B, Horvath S. A general framework for weighted gene co-expression network analysis. *Stat Appl Genet Mol Biol.* 2005;4:17. <https://doi.org/10.2202/1544-6115.1128>.
41. Kolpakov F, Poroikov V, Selivanova G, Kel A. GeneXplain—identification of causal biomarkers and drug targets in personalized cancer pathways. *J Biomol Tech.* 2011;22(Suppl):16.
42. Yates AD, Achuthan P, Akanni W, Allen J, Allen J, Alvarez-Jarreta J, Amode MR, Armean IM, Azov AG, Bennett R, Bhai J, Billis K, Boddu S, Marugán JC, Cummins C, Davidson C, Dodiya K, Fatima R, Gall A, Giron CG, Gil L, Grego T, Haggerty L, Haskell E, Hourlier T, Izuogu OG, Janacek SH, Juettemann T, Kay M, Lavidas I, Le T, Lemos D, Martinez JG, Maurel T, McDowall M, McMahon A, Mohanan S, Moore B, Nuhn M, Ohed DN, Parker A, Parton A, Patricio M, Sakhivel MP, Abdul Salam AI, Schmitt BM, Schuilenburg H, Sheppard D, Sycheva M, Szuba M, Taylor K, Thormann A, Threadgold G, Vullo A, Walts B, Winterbottom A, Zadissa A, Chakiachvili M, Flint B, Frankish A, Hunt SE, Ilesley G, Kostadima M, Langridge N, Loveland JE, Martin FJ, Morales J, Mudge JM, Muffato M, Perry E, Ruffier M, Trevanion SJ, Cunningham F, Howe KL, Zerbino DR, Flicek P. Ensembl 2020. *Nucleic Acids Res.* 2020;48(D1):682–8. <https://doi.org/10.1093/nar/gkz2966>.
43. Lecun Y, Bottou L, Bengio Y, Haffner P. Gradient-based learning applied to document recognition. *Proc IEEE.* 1998;86(11):2278–324. <https://doi.org/10.1109/5.726791>.
44. Lin C-Y, Beattie A, Baradaran B, Dray E, Duijf PHG. Contradictory mRNA and protein misexpression of EEF1A1 in ductal breast carcinoma due to cell cycle regulation and cellular stress. *Sci Rep.* 2018;8(1):13904. <https://doi.org/10.1038/s41598-018-32272-x>.
45. Sharma P, Alsharif S, Fallatah A, Chung BM. Intermediate filaments as effectors of cancer development and metastasis: a focus on keratins, vimentin, and nestin. *Cells.* 2019;8(5):497. <https://doi.org/10.3390/cells8050497>.
46. Wang W, Eddy R, Condeelis J. The cofilin pathway in breast cancer invasion and metastasis. *Nat Rev Cancer.* 2007;7(6):429–40. <https://doi.org/10.1038/nrc2148>.
47. Lin T-C, Yang C-H, Cheng L-H, Chang W-T, Lin Y-R, Cheng H-C. Fibronectin in cancer: Friend or foe. *Cells.* 2019;9(1):27. <https://doi.org/10.3390/cells9010027>.
48. Feng Y, Spezia M, Huang S, Yuan C, Zeng Z, Zhang L, Ji X, Liu W, Huang B, Luo W, Liu B, Lei Y, Du S, Vuppalapati A, Luu HH, Haydon RC, He T-C, Ren G. Breast cancer development and progression: Risk factors, cancer stem cells, signaling pathways, genomics, and molecular pathogenesis. *Genes Dis.* 2018;5(2):77–106. <https://doi.org/10.1016/j.gendis.2018.05.001>.
49. Perou CM, Sørlie T, Eisen MB, van de Rijn M, Jeffrey SS, Rees CA, Pollack JR, Ross DT, Johnsen H, Akslen LA, Fluge Ø, Pergamenschikov A, Williams C, Zhu SX, Lønning PE, Børresen-Dale A-L, Brown PO, Botstein D. Molecular portraits of human breast tumours. *Nature.* 2000;406(6797):747–52. <https://doi.org/10.1038/35021093>.
50. Bailey CK, Mittal MK, Misra S, Chaudhuri G. High motility of triple-negative breast cancer cells is due to repression of plakoglobin gene by metastasis modulator protein SLUG. *J Biol Chem.* 2012;287(23):19472–86. <https://doi.org/10.1074/jbc.m112.345728>.
51. Shi H, Li H, Yuan R, Guan W, Zhang X, Zhang S, Zhang W, Tong F, Li L, Song Z, Wang C, Yang S, Wang H. PCBP1 depletion promotes tumorigenesis through attenuation of p27 Kip1 mRNA stability and translation. *J Exp Clin Cancer Res.* 2018;37(1):187. <https://doi.org/10.1186/s13046-018-0840-1>.
52. Fan B, Shi S, Shen X, Yang X, Liu N, Wu G, Guo X, Huang N. Effect of HMGN2 on proliferation and apoptosis of MCF-7 breast cancer cells. *Oncol Lett.* 2018;17(1):1160–6. <https://doi.org/10.3892/ol.2018.9668>.
53. Liu Y, Liu T, Sun Q, Niu M, Jiang Y, Pang D. Downregulation of Ras GTPase-activating protein 1 is associated with poor survival of breast invasive ductal carcinoma patients. *Oncol Rep.* 2014;33(1):119–24. <https://doi.org/10.3892/or.2014.3604>.
54. Mathe A, Wong-Brown M, Morten B, Forbes JF, Braye SG, Avery-Kiejda KA, Scott RJ. Novel genes associated with lymph node metastasis in triple negative breast cancer. *Sci Rep.* 2015;5(1):15832. <https://doi.org/10.1038/srep15832>.
55. Saha S, Kim K, Yang G-M, Choi H, Cho S-G. Cytokeratin 19 (KRT19) has a role in the reprogramming of cancer stem cell-like cells to less aggressive and more drug-sensitive cells. *Int J Mol Sci.* 2018;19(5):1423. <https://doi.org/10.3390/ijms19051423>.
56. Zhou X, Hao Q, Liao J-M, Liao P, Lu H. Ribosomal protein S14 negatively regulates c-Myc activity. *J Biol Chem.* 2013;288(30):21793–801. <https://doi.org/10.1074/jbc.m112.445122>.
57. Alexandrou S, George S, Ormandy C, Lim E, Oakes S, Caldon C. The proliferative and apoptotic landscape of basal-like breast cancer. *Int J Mol Sci.* 2019;20(3):667. <https://doi.org/10.3390/ijms20030667>.
58. Neve RM, Chin K, Fridlyand J, Yeh J, Baehner FL, Fevr T, Clark L, Bayani N, Coppe J-P, Tong F, Spellman PT, DeVries S, Lapuk A, Wang NJ, Kuo W-L, Stilwell JL, Pinkel D, Albertson DG, Waldman FM, McCormick F, Dickson RB, Johnson MD, Lippman M, Ethier S, Gazdar A, Gray JW. A collection of breast cancer cell lines for the study of functionally distinct cancer subtypes. *Cancer Cell.* 2006;10(6):515–27. <https://doi.org/10.1016/j.ccr.2006.10.008>.

59. Taylor J, Sendino M, Gorelick AN, Pastore A, Chang MT, Penson AV, Gavrilu EI, Stewart C, Melnik EM, Chavez FH, Bitner L, Yoshimi A, Lee SC-W, Inoue D, Liu B, Zhang XJ, Mato AR, Dogan A, Kharas MG, Chen Y, Wang D, Soni RK, Hendrickson RC, Prieto G, Rodriguez JA, Taylor BS, Abdel-Wahab O. Altered nuclear export signal recognition as a driver of oncogenesis. *Cancer Discov.* 2019;9(10):1452–67. <https://doi.org/10.1158/2159-8290.cd-19-0298>.
60. Qian X-L, Pan Y-H, Huang Q-Y, Shi Y-B, Huang Q-Y, Hu Z-Z, Xiong L-X. Caveolin-1: a multifaceted driver of breast cancer progression and its application in clinical treatment. *OncoTargets Ther.* 2019;12:1539–52. <https://doi.org/10.2147/ott.s191317>.
61. Aceto N, Sausgruber N, Brinkhaus H, Gaidatzis D, Martiny-Baron G, Mazzarol G, Confalonieri S, Quarto M, Hu G, Balwierz PJ, Pachkov M, Elledge SJ, van Nimwegen E, Stadler MB, Bentires-Alj M. Tyrosine phosphatase SHP2 promotes breast cancer progression and maintains tumor-initiating cells via activation of key transcription factors and a positive feedback signaling loop. *Nat Med.* 2012;18(4):529–37. <https://doi.org/10.1038/nm.2645>.
62. Chekhun VF, Lukyanova NY, Burlaka AP, Bezdenezhnykh NA, Shpyleva SI, Tryndyak VP, Beland FA, Pogribny IP. Iron metabolism disturbances in the MCF-7 human breast cancer cells with acquired resistance to doxorubicin and cisplatin. *Int J Oncol.* 2013;43(5):1481–6. <https://doi.org/10.3892/ijo.2013.2063>.
63. Perera-Bel J, Hutter B, Heining C, Bleckmann A, Fröhlich M, Fröhling S, Glimm H, Brors B, Beißbarth T. From somatic variants towards precision oncology: Evidence-driven reporting of treatment options in molecular tumor boards. *Genome Med.* 2018;10(1):18. <https://doi.org/10.1186/s13073-018-0529-2>.
64. Ioffe S, Szegedy C. Batch normalization: accelerating deep network training by reducing internal covariate shift. *Proceedings of the 32nd International Conference on Machine Learning, PMLR.* 2015;37:448–456. <http://proceedings.mlr.press/v37/loff15.html>.
65. Banerjee K, Resat H. Constitutive activation of STAT 3 in breast cancer cells: a review. *Int J Cancer.* 2015;138(11):2570–8. <https://doi.org/10.1002/ijc.29923>.
66. Bentires-Alj M, Paez JG, David FS, Keilhack H, Halmos B, Naoki K, Maris JM, Richardson A, Bardelli A, Sugarbaker DJ, Richards WG, Du J, Girard L, Minna JD, Loh ML, Fisher DE, Velculescu VE, Vogelstein B, Meyerson M, Sellers WR, Neel BG. Activating mutations of the Noonan syndrome-associated SHP2/PTPN11 gene in human solid tumors and adult acute myelogenous leukemia. *Cancer Res.* 2004;64(24):8816–20. <https://doi.org/10.1158/0008-5472.can-04-1923>.
67. Zhang J, Liang Q, Lei Y, Yao M, Li L, Gao X, Feng J, Zhang Y, Gao H, Liu D-X, Lu J, Huang B. SOX4 induces epithelial to mesenchymal transition and contributes to breast cancer progression. *Cancer Res.* 2012;72(17):4597–608. <https://doi.org/10.1158/0008-5472.can-12-1045>.
68. Guaita-Esteruelas S, Bosquet A, Saavedra P, Gumà J, Girona J, Lam EW-F, Amillano K, Borràs J, Masana L. Exogenous FABP4 increases breast cancer cell proliferation and activates the expression of fatty acid transport proteins. *Mol Carcinog.* 2016;56(1):208–17. <https://doi.org/10.1002/mc.22485>.
69. Liang Y, Han H, Liu L, Duan Y, Yang X, Ma C, Zhu Y, Han J, Li X, Chen Y. CD36 plays a critical role in proliferation, migration and tamoxifen-inhibited growth of ER-positive breast cancer cells. *Oncogenesis.* 2018;7(12):98. <https://doi.org/10.1038/s41389-018-0107-x>.
70. Kuemmerle NB, Rysman E, Lombardo PS, Flanagan AJ, Lipe BC, Wells WA, Pettus JR, Froehlich HM, Memoli VA, Morganelli PM, Swinnen JV, Timmerman LA, Chaychi L, Fricano CJ, Eisenberg BL, Coleman WB, Kinlaw WB. Lipoprotein lipase links dietary fat to solid tumor cell proliferation. *Mol Cancer Ther.* 2011;10(3):427–36. <https://doi.org/10.1158/1535-7163.mct-10-0802>.
71. Nakai K, Hung MC, Yamaguchi H. A perspective on anti-EGFR therapies targeting triple-negative breast cancer. *Am J Cancer Res.* 2016;6(8):1609–23.
72. Auer F. Patient specific molecular sub-networks responsible for metastasis in breast cancer. 2020. <http://mypathsem.bioinf.med.uni-goettingen.de/MetaRelSubNetVis>.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Ready to submit your research? Choose BMC and benefit from:

- fast, convenient online submission
- thorough peer review by experienced researchers in your field
- rapid publication on acceptance
- support for research data, including large and complex data types
- gold Open Access which fosters wider collaboration and increased citations
- maximum visibility for your research: over 100M website views per year

At BMC, research is always in progress.

Learn more biomedcentral.com/submissions

