



OPEN

Breast cancer histopathology image-based gene expression prediction using spatial transcriptomics data and deep learning

Md Mamunur Rahaman¹, Ewan K. A. Millar^{2,3,4} & Erik Meijering¹✉

Tumour heterogeneity in breast cancer poses challenges in predicting outcome and response to therapy. Spatial transcriptomics technologies may address these challenges, as they provide a wealth of information about gene expression at the cell level, but they are expensive, hindering their use in large-scale clinical oncology studies. Predicting gene expression from hematoxylin and eosin stained histology images provides a more affordable alternative for such studies. Here we present BrST-Net, a deep learning framework for predicting gene expression from histopathology images using spatial transcriptomics data. Using this framework, we trained and evaluated four distinct state-of-the-art deep learning architectures, which include ResNet101, Inception-v3, EfficientNet (with six different variants), and vision transformer (with two different variants), all without utilizing pretrained weights for the prediction of 250 genes. To enhance the generalisation performance of the main network, we introduce an auxiliary network into the framework. Our methodology outperforms previous studies, with 237 genes identified with positive correlation, including 24 genes with a median correlation coefficient greater than 0.50. This is a notable improvement over previous studies, which could predict only 102 genes with positive correlation, with the highest correlation values ranging from 0.29 to 0.34.

Breast cancer is the most prevalent malignancy in women and is a molecularly diverse disease. In 2018, 2.1 million women were diagnosed with breast cancer, one every 18 seconds, and 626,679 women died, a 3.1% yearly rise¹. Early breast cancer, where the disease is limited to the breast with or without the involvement of the axillary lymph nodes, has a good prognosis with a 5-year survival of close to 90%. However, advanced metastatic disease is often incurable with existing therapies, which aim to delay progression and treat symptoms. According to the World Health Organization (WHO), early identification of cancer significantly improves the likelihood of appropriate decision making and successful treatment^{2,3}.

Pathological analysis of a tissue biopsy removed from the breast tumour is acknowledged as the gold standard method for the diagnosis, subtype classification and staging⁴. Heterogeneity is evident at the morphological and molecular levels of breast cancer, where divergent patterns of gene expression and interaction with host immune and stromal populations can vary widely within and between similar tumour types. A better understanding of spatial gene expression is a current area of very active cancer research to better capture spatial differences in gene expression which may more accurately reflect patient outcome and response to treatment, which are currently difficult to gauge.

Emerging spatial transcriptomics (ST) technologies enable profiling gene expression at a single-cell resolution while preserving spatial orientation and cellular composition within a tissue^{5,6}. ST is quickly becoming an extension of single-cell RNA sequencing (scRNAseq). To understand the complex transcriptional architecture of biological systems, it is necessary to determine how cells are arranged in space and how gene expression changes in different foci of a targeted tumour or tissue. ST methods based on next generation sequencing (NGS), such as

¹School of Computer Science and Engineering, University of New South Wales, Kensington, Sydney NSW 2052, Australia. ²Department of Anatomical Pathology, NSW Health Pathology, St. George Hospital, Kogarah, Sydney NSW 2217, Australia. ³St. George and Sutherland Clinical School, University of New South Wales, Kensington, Sydney NSW 2052, Australia. ⁴Faculty of Medicine & Health Sciences, Western Sydney University, Campbelltown, Sydney NSW 2560, Australia. ✉email: erik.meijering@unsw.edu.au

10× Genomics' Visium, Slide-Seq⁷, Slide-Seq2⁸, and HDST⁹, barcode whole transcriptomes but have restricted capture rates and resolutions greater than a single cell (around 100 μm for Visium and 10 μm for Slide-Seq). Moreover, NGS-based approaches offer unbiased profiling of large tissue sections without requiring a list of target genes, unlike image-based technologies such as in-situ hybridization (ISH) and in-situ sequencing (ISS)^{10–12}. In order to accurately and robustly analyse the data produced by ST technologies, specialised computational methods are required because of the data's intrinsic noise, high dimensionality, sparseness, and multimodality (containing histology pictures, count matrices, etc.).

Although ST provides a plethora of data, its generation is prohibitively costly, precluding its large-scale application. In addition, with commercial equipment like the 10× Genomics Visium system, significant expertise is required to create high-quality expression patterns for whole tissue samples¹³. On the other hand, compared to ST, hematoxylin & eosin (H & E)-stained histology images are not only easier and less expensive to acquire, but they are also routinely used in clinical practice. Therefore, the recent development of predicting spatial gene expression data from routine histology images offers advantages in cost, speed, and availability of such data. These projections have the potential to provide virtual ST data, which will make it possible to investigate the regional differences in gene expression on a large scale.

While existing analysis pipelines mostly use gene expression profiles, not image pixel values, the integration of imaging data with gene expression data is a new area of research¹⁴. Moreover, whole-slide images (WSIs) have been utilised to predict global gene expression patterns via HE2RNA¹⁵, demonstrating their strong correlation with transcription. HE2RNA is trained to predict gene expression patterns from The Cancer Genome Atlas (TCGA) WSIs without specialist interpretation. Similarly, ST-Net uses ST data in conjunction with DenseNet to make predictions on the spatially varying gene expression of each spot over WSIs¹⁶. Prediction of the spatially resolved transcriptome makes it possible to use images to look for biomarkers that change in various tumour location spots.

Although these techniques have performed well, they do have certain drawbacks. HE2RNA lacks the capacity to learn from ST data since it was designed for bulk RNA sequencing. Despite the fact that ST-Net is purpose-built for ST, its convolutional neural network (CNN) model is pretrained on the ImageNet dataset, which fails to take into account the observed individual spatial locations of spots. Expression of genes generally has local patterns, thus it is vital to consider spatial location when training networks to make accurate predictions.

To overcome these challenges, we have developed BrST-Net (short for Breast ST-Net), a framework based on deep learning that uses spatial transcriptomics data to predict gene expression directly from breast histopathology images (Fig. 1). This framework includes stain normalisation of tissue sections, filtering of gene and ST spots, and the generation of image patches depending on the locations of the spots. The image patches are then utilised to train CNN and transformer models, with the primary network devoted to predicting 250 highly expressed genes. We introduce an auxiliary network before the end of the main network to forecast the remaining genes, which ultimately enhances the generalisation performance of the main network by functioning as an alternative

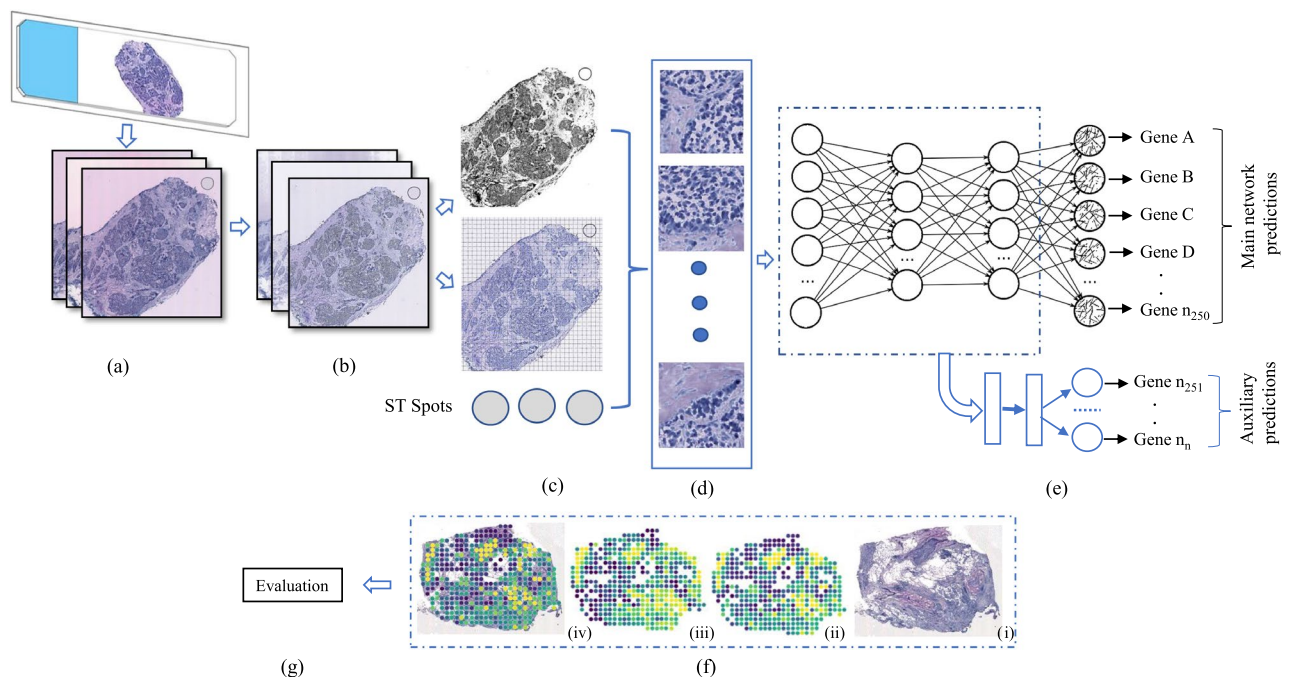


Figure 1. Overview of the proposed BrST-Net framework. (a) Original whole slide image. (b) Stain normalised image. (c) Mask image and filtered ST spot location information. (d) Extracted image patches using the spot location information. (e) Main network with an auxiliary network for gene expression prediction. (f) Prediction of an individual gene on the tissue sample, showing (i) test sample, (ii) ACTB gene prediction, (iii) ground-truth expression of gene ACTB, and (iv) gene ACTB overlaid on the tissue sample. (g) Performance evaluation.

kind of regularisation¹⁷. We have thoroughly trained and evaluated four distinct state-of-the-art deep learning architectures, which include ResNet101, Inception-v3, EfficientNet (with six different variants), and vision transformer (with two different variants), to assess their prediction capabilities on the ST dataset. Our suggested approach outperforms existing work: the framework using EfficientNet-b0 with an auxiliary network can predict 237 genes with positive correlation, whereas ST-Net can only predict 102 genes with positive correlation. ST-Net identified the top-5 predicted genes in their framework with a median correlation coefficient value of 0.34, 0.33, 0.31, 0.30, and 0.29, while our framework using EfficientNet-b0 with an auxiliary network can predict 24 genes with a median correlation coefficient greater than 0.50, which is a substantial performance boost.

Materials and methods

Dataset description. We used the fifth edition of the publicly accessible spatial transcriptomics dataset¹⁸ as also used by ST-Net¹⁶. It comprises a total of 68 WSIs of H &E-stained frozen tissue sections, from 23 breast cancer patients having various subtypes of breast cancer, including luminal A, luminal B, triple-negative, human epidermal growth factor receptor 2 (HER2) luminal, and HER2 non-luminal. Each section is scanned at 20× magnification, with image size approximately 9500 × 9500 pixels in JPEG format, and contains around 450 spots, each having a diameter of 100 μm and approximately 3000 gene expression counts. In total, the dataset comprises 30,612 spots, and includes spot coordinates, count matrices, and coordinate files.

Preprocessing of the dataset. *Stain normalization.* Histology samples stained with H &E often vary in colour between and within laboratories from one batch to the next. Stain normalisation is employed as a pre-processing step in computational pathology algorithms to regulate variances and has a significant impact on the outcome of automated image analysis methods¹⁹. In our work, we employed StainTools²⁰, using the Vahadane technique to normalize the stain variance in the target image and the Luminosity Standardizer function to standardize the brightness, both of which facilitate consistent interpretation of the images²¹.

Gene and spots filtering. To eliminate background noise from gene expression data per spot, we removed genes with a mean expression of zero across all samples, and selected spots with at least 1000 total read counts for analysis. The resulting filtered genes were stored in a ZipFile format (.npz) comprising multiple arrays, including count, pixel, patient, and index arrays. This resulted in approximately 28,792 spot files from all samples, which were then utilized to generate corresponding image patches. After filtering, approximately 6000 genes remained. The complete preprocessing pipeline is described in detail in Ref.²².

Generating image patches. Since the original images are too large for direct input to a CNN, we extracted N patches from the WSIs centered on the ST spots, using their size and position. The patches were then flattened into an $N \times (3 \times m \times n)$ matrix, where m and n are the patch's width and height, respectively, and each patch has 3 colour channels. In our experiment, a value of $m = n = 224$ pixels was selected. To ensure that the training dataset only contained informative data, we excluded patches with more than 50% white pixels, as these are unlikely to contain relevant information and may introduce noise in the data. In total, 27,652 image patches were acquired for training.

Gene symbol conversion. The dataset provides a comprehensive list of genes, including their Ensembl identifiers (IDs), which are widely used as a standard for gene annotation and identification. The IDs were translated from the given format to gene symbols according to the HUGO Gene Nomenclature Committee (HGNC) database (www.genenames.org) and converted into a Pickle file, retaining just the Ensembl ID and its related symbol. Pickle is a Python library to store complex data as binary files.

Data augmentation. Histology image analysis is supposed to be orientationally invariant, just like a pathologist can analyse a microscopic image from any angle. Therefore, to facilitate the generalisability of the trained networks, we used PyTorch's "torchvision.transforms" module to randomly transform image patches during training. It performs random horizontal and vertical flipping, and random rotation of 90 degrees to increase the diversity of the training data. All patches were converted to a PyTorch tensor and normalised to zero mean and unit variance.

Gene expression prediction models. We experimented with four distinct deep learning architectures, which include ResNet, InceptionNet, six variants of EfficientNet, and two vision transformer models, as the basis for our BrST-Net framework (Fig. 1), briefly described below. These main networks were trained to learn spatial features from histology images and predict gene expression levels. However, due to the high dimensionality of ST data, CNN architectures may be limited in their ability to predict gene expression levels accurately for the remaining genes. To address this limitation, we modified the main network by introducing an auxiliary network (AuxNet) just before the end of the network. The AuxNet is a simple feed-forward neural network with a single fully connected layer that predicts the remaining genes. This approach has been shown to enhance the performance of various models in handling high-dimensional data. By doing so, the AuxNet enhances the generalisation performance of the main network and helps prevent overfitting, providing an alternative form of regularization^{23–27}.

ResNet model. ResNet (Residual Network)²⁸ is a deep learning architecture widely used in computer vision tasks such as image classification, object detection, and semantic segmentation^{29,30}. It uses residual connections

to improve information propagation and learning, expressed as $F(x) = H(x) + x$, where x is the input, $H(x)$ is the learned residual mapping, and $F(x)$ is the residual block output. This helps mitigate the vanishing gradient problem in deep neural networks. In this study, ResNet101 is used as the main network and fully trained on the ST dataset, making this the first study to use ResNet for gene expression prediction.

Inception model. Inception²⁷ is a deep convolutional neural network designed for image classification. It uses a combination of convolutional and pooling layers, along with auxiliary classifiers, to capture both local and global features in the input image. The final prediction is made by combining the outputs of the main and auxiliary classifiers, represented mathematically as $y = Wx + b$, where x is the input feature vector, W is the weight matrix and b the bias term, and y is the output. Different versions exist, with Inception-v3 achieving state-of-the-art performance on benchmark datasets and now being widely used in real-world applications^{31,32}. In this study, we fully trained it on the ST dataset and evaluated its performance for gene expression prediction.

EfficientNet model. The conventional approach to increasing the learning capacity of CNN models involves tweaking the network's depth d and width w and the image resolution r . While this enhances accuracy, it often requires significant manual tuning and can result in suboptimal performance. This has been addressed in literature by examining various scaling strategies and proposing a systematic network architecture scaling method, leading to the development of EfficientNet³³. The scaling method employs a user-defined coefficient, φ , to scale up networks in a more systematic manner, according to the following equations:

$$\begin{aligned} d &= \alpha^\varphi, & w &= \beta^\varphi, & r &= \gamma^\varphi, \\ \text{s.t. } \alpha \cdot \beta^2 \cdot \gamma^2 &\approx 2, & \alpha &\geq 1, & \beta &\geq 1, & \gamma &\geq 1, \end{aligned} \quad (1)$$

where α , β , and γ are calculated through automatic hyperparameter optimisation using grid search. The scaling of the model is regulated by φ , which controls the allocation of computational resources. In pathological image analysis, multiple studies have utilized the EfficientNet model and achieved superior accuracy^{34–38}. To our knowledge, however, our study is the first to use EfficientNet for gene expression prediction using the ST dataset. Specifically, we experimented with versions b0, b1, b2, b3, b4, and b5.

Vision transformer model. Recent editions of the ImageNet Large-Scale Visual Recognition Competition have witnessed the dominance of the vision transformer (ViT) over other state-of-the-art approaches. Inspired by the challenge of natural language processing (NLP) to deal with varying sentence lengths, ViT aims to deal with dependencies at various spatial distances, and partitions an image into a fixed number of patches. In the feature extraction stage of a ViT model, a 2D image $x \in \mathbb{R}^{X \times Y \times C}$ of size $X \times Y$ pixels and C channels, is transformed into a 1D sequence $x_p \in \mathbb{R}^{N \times P^2 \times C}$ of N patches of size P^2 , where N is equal to the input sequence length of the transformer encoder and is calculated as $N = XY/P^2$. In the transformer, all patches are flattened and then projected linearly to D dimensions, referred to as patch embeddings. These patch embeddings are the input to the transformer encoder, preserving positional information³⁹. The transformer encoders consist of multiple multihead self-attention (MHSA) and multilayer perceptron (MLP) blocks. The MHSA layer in the transformer is capable of learning the attention for spots or image patches and involves a linear combination of several attention heads:

$$\text{MultiHead}(Q, K, V) = [\text{head}_1, \dots, \text{head}_c] \times W_0, \quad (2)$$

where c denotes the number of heads, W_0 are the weights used to aggregate the attention heads, and Q , K , and V represent the query, key, and value, respectively. Each head is calculated as:

$$\text{head}_i = \text{Attention}\left(QW_i^Q, KW_i^K, VW_i^V\right), \quad (3)$$

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V, \quad (4)$$

where W_i^Q , W_i^K , W_i^V are weight matrices. Here, $QK^T/\sqrt{d_k}$ refers to the attention map, whose pattern is $N \times N$, and V represents the self-attention mechanism's value, where $V = Q = K$. The attention mechanisms produces an $N \times 1024$ matrix as output. In our experiment, we evaluated the performance of the ViT-B16 and ViT-B32 models, where B16 and B32 refer to the partitioning of an image into 16 and 32 patches, respectively³⁹.

Model training and evaluation. *Loss function.* To train a network, its weights are iteratively updated to minimise a given loss function. In our framework, we use the following loss:

$$\mathcal{L} = \mathcal{L}_{\text{main}} + \lambda \cdot \mathcal{L}_{\text{aux}}, \quad (5)$$

where \mathcal{L} is the overall loss, $\mathcal{L}_{\text{main}}$ is the loss of the main network, \mathcal{L}_{aux} is the loss of the auxiliary network, with both losses being calculated using cross-entropy, and λ is a hyperparameter used to balance the contribution of the two losses. The λ value effectively controls the influence of the auxiliary network on the overall training of the model, acting as a form of regularization by adding an additional constraint that the model needs to satisfy. We experimented with various values of λ and found that a value of 40 gives good performance. Stochastic gradient descent (SGD) optimization was used to minimise the loss.

Error measures. To quantify errors between predicted and actual values, we used the mean absolute error (MAE) and root mean squared error (RMSE):

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i|, \quad \text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2}, \quad (6)$$

where n is the total number of observations, y_i is the true value of observation i , and \hat{y}_i is its predicted value. Both error measures yield an average prediction error of a model, ranging from 0 to ∞ , with smaller values indicating more accurate predictions.

Correlation measure. To assess the reliability of gene expression predictions from histopathology images, we used the Pearson correlation coefficient (PCC)^{40–42}:

$$\text{PCC} = \frac{\sum_{i=1}^n (a_i - \bar{a})(b_i - \bar{b})}{\sqrt{\sum_{i=1}^n (a_i - \bar{a})^2 \sum_{i=1}^n (b_i - \bar{b})^2}}, \quad (7)$$

where a_i and b_i are the true and predicted genes, respectively, n is the total number of genes, \bar{a} is the mean of the a_i , and \bar{b} is the mean of the b_i . The PCC evaluates the linear relationship between two variables and assigns a score ranging from -1 to 1 . A score of 1 means there is a perfect positive linear correlation, a score of -1 indicates a perfect negative linear relationship, and a score of 0 implies there is no linear relationship. In practice, a PCC score between 0.5 to 1 indicates strong correlation, a score between 0.3 to 0.5 indicates medium correlation, and a score between 0.1 to 0.3 indicates weak correlation.

Measures taken to prevent overfitting. Despite the limited size of our dataset, we have implemented several strategies to prevent overfitting and ensure model robustness. A 5-fold cross-validation strategy was employed on 22 out of 23 patients, ensuring model performance isn't overly reliant on a single data subset. To evaluate model generalizability, data from one patient was reserved as an independent test set, providing a stringent test of the model's extrapolation capabilities.

Overfitting was further mitigated through various techniques. Batch gradient descent was utilized during training, and gradients were zeroed at each loop's start. Data augmentation was applied during the training process to effectively increase the training dataset size. The Stochastic Gradient Descent (SGD) optimizer was set with a weight decay parameter, introducing L2 regularization. Early stopping was implemented, halting training when validation set performance ceased to improve. A ReduceLROnPlateau scheduler was used for learning rate scheduling, reducing the learning rate when validation loss plateaued. Lastly, an auxiliary loss was introduced during training, acting as a form of regularization and adding an additional constraint to the model, thereby enhancing its generalization ability.

Experimental results

Implementation and setup. BrST-Net was implemented in PyTorch and Python 3.7. For training, the batch size was set to 32, the number of epochs to 200, the gene filter to 250, and the learning rate for SGD was set to 0.001. The models were trained on 2 GPUs with 24 CPU cores and 50 GB of RAM on the Gadi cluster of the National Computational Infrastructure in Australia. A 5-fold cross-validation was performed on 22 out of 23 patients to evaluate the performance of the trained models. The remaining patient was held out as a final test set to assess the generalisability of the models.

Quantitative results. To assess the efficacy of the different models for our framework, we fully trained and evaluated a total of 10 models on the ST dataset: ResNet101, Inception-v3, EfficientNet-b0, EfficientNet-b1, EfficientNet-b2, EfficientNet-b3, EfficientNet-b4, EfficientNet-b5, ViT-B-16 and ViT-B-32, with and without the use of AuxNet, and compared their performances in predicting the top 250 genes. From the results with AuxNet on the held-out test case (Table 1 and Fig. 2), we observe that the EfficientNet architecture is dominant, with EfficientNet-b0 performing most favourably, followed by EfficientNet-b4. The results without using AuxNet for the test case (Supplementary Table S1) clearly show the improvements brought by the proposed auxiliary network. Specifically, for each gene, the best median PCC value is higher with the use of AuxNet than without, though this higher number may be produced by a different main network model. Of the other models, ResNet101, Inception-v3, and ViT-B16 performed somewhat comparably (no consistent winner), and better than ViT-B32 and EfficientNet-b5, the latter of which generally performed worst. EfficientNet-b0 with AuxNet was able to predict 237 genes with positive correlation (Fig. 3), including 24 genes with median PCC values greater than 0.5 (strong correlation), 123 genes with PCC values between 0.3 and 0.5 (medium correlation), and 78 genes with PCC values between 0.1 and 0.3 (weak correlation).

The average MAE and average RMSE of all models with and without AuxNet (Table 2) show that for the best-performing model according to the PCC metric (Table 1) the errors on the testing data are smaller with the use of AuxNet than without. More broadly, Inception-v3, the lower-version EfficientNet models, and ViT-B16 all produce smaller errors with the use of AuxNet. Interestingly enough, the errors of these models on the training data (Table 2) as well as on the validation data (Supplementary Tables S2 and S3) can be larger with the use of AuxNet. This may be explained by the fact that AuxNet introduces additional outputs in the network, which can make the optimization process more difficult. During training, the network learns to optimize the main objective

Genes	ResNet101	Incep-v3	ENet-b0	ENet-b1	ENet-b2	ENet-b3	ENet-b4	ENet-b5	ViT-B16	ViT-B32
B2M	0.4729	0.5720	0.6325	0.5659	0.5722	0.5199	0.6221	0.1800	0.5475	0.3909
ACTG1	0.5457	0.5391	0.6233	0.5637	0.6054	0.5888	0.6220	0.2259	0.6075	0.5231
ACTB	0.6066	0.4985	0.6204	0.3723	0.4843	0.5919	0.6147	0.2718	0.6359	0.5452
TMSB10	0.3896	0.4678	0.6197	0.0634	0.4597	0.5621	0.5717	0.2038	0.5576	0.1359
GNAS	0.5019	0.5896	0.6139	0.5864	0.5886	0.5976	0.6170	0.2588	0.5872	0.4148
PTMA	0.5521	0.5072	0.5902	0.5641	0.5654	0.5423	0.5956	0.3685	0.5639	0.4762
PTPRF	0.4927	0.5269	0.5926	0.5667	0.5524	0.5359	0.5568	0.0733	0.5126	0.4643
ERBB2	0.3156	0.5033	0.5897	0.3600	0.3499	0.5930	0.5750	0.1684	0.5343	0.4626
PRDX1	0.4561	0.5329	0.5702	0.4478	0.5246	0.5405	0.5466	0.1213	0.5262	0.4789
TMSB4X	0.4238	0.4884	0.5665	0.3876	0.4203	0.5876	0.5872	0.2086	0.5717	0.5345

Table 1. Top 10 predicted genes by our framework using different models with AuxNet. The numbers are median PCC values on the held-out test patient ID BC23903. Here, to save space, “EfficientNet” is abbreviated to “ENet”, and “Inception-v3” to “Incep-v3”. Bold indicates best performance (highest PCC) per gene and italics indicates second-best.

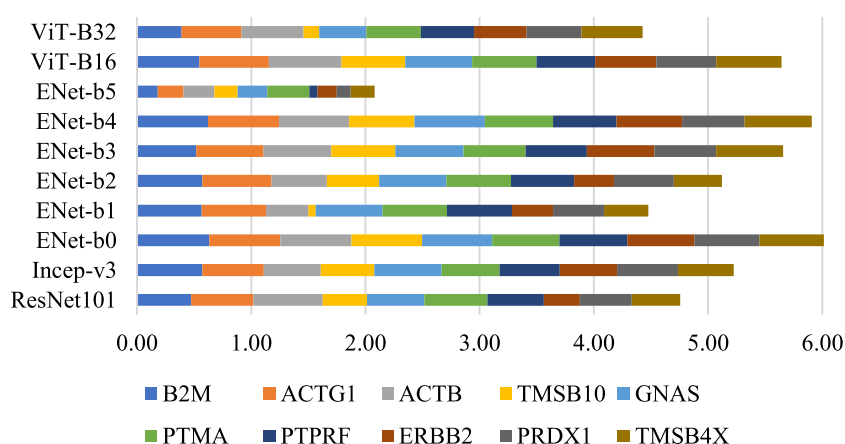


Figure 2. Stacked bar chart summarising the predictive performances of the different models.

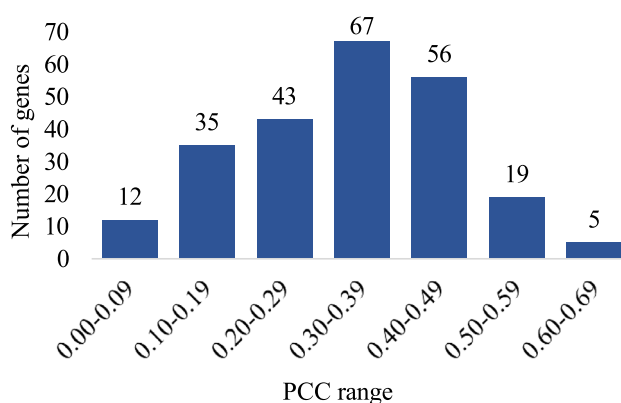


Figure 3. Gene prediction PCC value histogram of the 237 positive correlations produced by EfficientNet-b0 + AuxNet.

while also optimizing the auxiliary objectives. Consequently, this may result in a more complex optimization problem, and the network may not converge as quickly or as well as without the use of AuxNet. However, the use of AuxNet can still enhance the generalisation performance of the models on the unseen testing data, which is crucial in practical applications. Thus, the slight increase in training error is a reasonable trade-off for the improved generalisation performance on the testing data.

Model	Training data		Testing data	
	aMAE	aRMSE	aMAE	aRMSE
ResNet101 + AuxNet	0.7823	0.9654	0.7609	0.9344
Inception-v3 + AuxNet	0.7831	0.9650	0.7403	0.9036
EfficientNet-b0 + AuxNet	0.7730	0.9528	0.7490	0.9144
EfficientNet-b1 + AuxNet	0.7489	0.9259	0.7345	0.8976
EfficientNet-b2 + AuxNet	0.7618	0.9399	0.7277	0.8898
EfficientNet-b3 + AuxNet	0.7858	0.9677	0.7481	0.9135
EfficientNet-b4 + AuxNet	0.8037	0.9899	0.7718	0.9409
EfficientNet-b5 + AuxNet	0.8042	0.9904	0.7905	0.9716
ViT-B16 + AuxNet	0.7909	0.9754	0.7262	0.8893
ViT-B32 + AuxNet	0.8039	0.9909	0.7699	0.9403
ResNet101	0.7830	0.9675	0.7484	0.9158
Inception-v3	0.7801	0.9657	0.7470	0.9148
EfficientNet-b0	0.7456	0.9242	0.7499	0.9148
EfficientNet-b1	0.7529	0.9324	0.7462	0.9122
EfficientNet-b2	0.7610	0.9419	0.7397	0.9042
EfficientNet-b3	0.8008	0.9878	0.7465	0.9120
EfficientNet-b4	0.7635	0.9446	0.7445	0.9100
EfficientNet-b5	0.7761	0.9586	0.7374	0.9017
ViT-B16	0.7878	0.9730	0.7621	0.9324
ViT-B32	0.7961	0.9829	0.7607	0.9311

Table 2. Average mean absolute error (aMAE) and average root mean squared error (aRMSE) of different models with and without AuxNet on the training and testing data.

Visualisation of gene expression. To better understand the gene expression predictions, the top predicted genes with high PCC values were visualised on the corresponding histology image (Fig. 4 and Supplementary Figs. S1 and S2). The visual comparison of the true and predicted expressions of gene B2M and their PCC values for the test dataset (Fig. 4) shows a strong correlation between the two. The yellow portions correspond to regions where B2M is strongly expressed, while the blue portions correspond to regions where that gene is poorly expressed. We observe that the yellow portion of the tissue is highly correlated with the presence of the black tumour annotation, indicating that this gene is colocalized with cancer cells in the image. This demonstrates the effectiveness of our proposed framework BrST-Net in predicting local gene expression from tissue images for a selected panel of genes.

Computational costs. To evaluate the computational cost of training our framework, we recorded the training times of all considered models in our experiments. From the results (Table 3) we observe that as EfficientNet increased in scale from b0 to b5, the training time gradually increased accordingly, as expected. ResNet101 and Inception-v3 were about as costly as the lower-scale versions of EfficientNet, while ViT-B16 was about as expensive as the higher-scale versions of EfficientNet, and ViT-B32 was the fastest network. For a few models, the use of AuxNet added comparatively little to the total cost, and for most models, it even slightly reduced the cost, demonstrating that AuxNet is a lightweight network that can improve accuracy while incurring

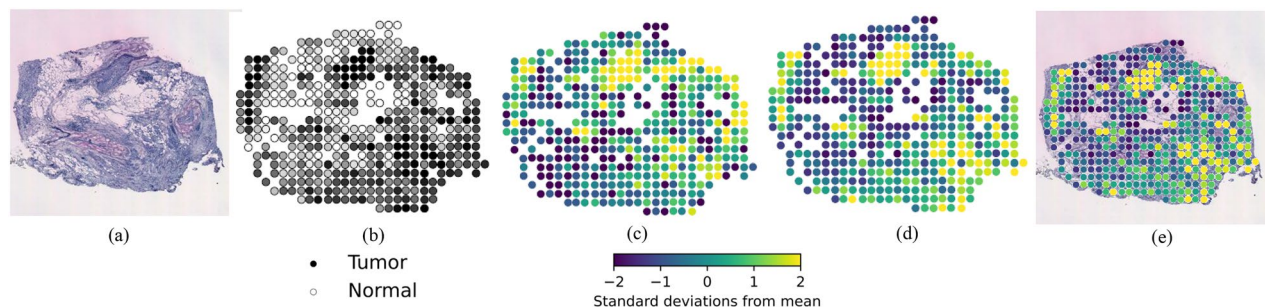


Figure 4. Visualisation of gene B2M expression for patient BC23903 slide number C2. (a) Histopathology sample from the test data. (b) Binary label of tumour (black) and normal (white) areas, with gray parts indicating areas that are not purely tumor or normal. (c) Ground truth expression of gene B2M. (d) Predicted expression of gene B2M (PCC = 0.6325). (e) Visualisation of predicted gene B2M expression overlaid on the tissue slice.

Model	Training time (hours:minutes:seconds)
ResNet101 + AuxNet	23:17:50
Inception-v3 + AuxNet	20:04:46
EfficientNet-b0 + AuxNet	19:04:28
EfficientNet-b1 + AuxNet	23:54:54
EfficientNet-b2 + AuxNet	23:30:49
EfficientNet-b3 + AuxNet	25:02:51
EfficientNet-b4 + AuxNet	29:58:24
EfficientNet-b5 + AuxNet	37:02:11
ViT-B16 + AuxNet	33:09:25
ViT-B32 + AuxNet	17:08:06
ResNet101	23:58:59
Inception-v3	20:05:58
EfficientNet-b0	19:49:15
EfficientNet-b1	23:07:10
EfficientNet-b2	22:35:53
EfficientNet-b3	25:00:06
EfficientNet-b4	30:02:30
EfficientNet-b5	37:36:46
ViT-B16	32:18:28
ViT-B32	17:18:28

Table 3. Computational cost of our framework using different models.

minimal computational overhead. For all models, while training takes many hours, their application at test time takes only a few seconds per image, making them computationally very feasible in clinical practice.

Discussion

Spatial transcriptomics technologies can profile gene expression for complete transcriptomics at almost single-cell resolution, with spatial positions matched with H & E-stained histological images (WSIs). While WSIs are inexpensive, accessible, and commonly generated in clinics, generating ST data is very costly and complicated, and currently, only a limited number of research centres are capable of doing so. Therefore, predicting gene expression directly from WSIs is useful.

In this paper, we proposed BrST-Net, a framework for predicting gene expression using ST data. We fully trained and tested four distinct deep learning architectures, which include ResNet101, Inception-v3, EfficientNet (with six different variants), and vision transformer (with two different variants), for our framework to predict 250 genes based on ST breast cancer data. Of all considered models, the combination of EfficientNet-b0 and our proposed AuxNet was able to predict 237 genes with positive median correlation, including 24 genes with strong correlation (PCC value over 0.5), 123 genes with medium correlation (PCC values between 0.3 and 0.5), and 78 genes with weak correlation (PCC values between 0.1 and 0.3). EfficientNet-b4 and EfficientNet-b3 also performed relatively well, as did Inception-v3 and ViT-B16, while EfficientNet-b5 performed worse and ViT-B32 generally did poorly. The fact that, among the transformer models, ViT-B16 outperformed ViT-B32 in gene expression prediction, may be attributed to the smaller dimensionality of ViT-B16, which enhances its ability to capture complex gene expression relationships. These findings emphasize the significance of selecting an appropriate model architecture for the given task and dataset.

In our research, we chose to evaluate ten models and diverse architectures, including convolutional neural networks (CNNs) and transformer models. This diversity allowed us to comprehensively evaluate different types of models on the task of predicting gene expression from histopathological images. We found that the EfficientNet architecture, particularly EfficientNet-b0, performed most favorably, likely due to its balanced scaling of network width, depth, and resolution, which effectively handles high-dimensional data. The Vision Transformer (ViT) models, specifically ViT-B16, also performed well due to their smaller dimensionality, which enhances their ability to capture complex gene expression relationships. We also introduced an auxiliary network (AuxNet) as a form of regularization, which improved the performance of all the models by adding an additional constraint, thereby enhancing their generalization ability. We believe that the choice of model architecture is significant for the given task and dataset, as different architectures have different strengths and weaknesses, and their suitability for a task depends on the specific characteristics of the data and the task requirements.

While the choice of model architecture is significant, it is equally important to prioritize a data-centric approach in spatial transcriptomics. This involves ensuring high-quality data through careful preprocessing and normalization, and leveraging large-scale datasets for training robust and generalizable models. Strategies to manage and learn from large-scale data, such as regularization techniques and efficient computational frameworks, are crucial in this context.

In our study, we have made several preprocessing choices that have potential implications on the results. For instance, we have implemented a filter to consider only spots with total read counts greater than or equal to 1000. This decision ensures that we focus on spots with sufficient information, reducing the impact of noise in our data. However, we acknowledge that altering this threshold could potentially affect the results, and further investigation into the optimal threshold could be beneficial. Additionally, we have applied a logarithmic transformation to the count data, a common normalization technique in gene expression analysis, to stabilize the variance and make the data more normally distributed. We also selected the top 250 genes based on their mean expression across all spots, a decision that reduces the dimensionality of the data and focuses on the most informative genes. These preprocessing choices, including the number of genes to keep, are hyperparameters of our pipeline. While these were chosen based on empirical observations and literature recommendations, we did not perform systematic hyperparameter optimization in this study. We acknowledge that different choices of these hyperparameters could potentially impact the results, and future work could explore the effects of these hyperparameters in more detail.

In comparison to the first method in this area, ST-Net¹⁶, our BrST-Net framework improved the gene expression prediction performance quite substantially. We have completely trained the most recent state-of-the-art CNN models and transformers and compared their performances on the ST dataset, while ST-Net employed a DenseNet trained on natural images (such as cats, dogs, and flowers), which may not be optimal for histopathology. In our suggested approach, in addition to the core network, an auxiliary network is added to predict the remaining genes. The auxiliary loss helps reduce the fading gradients problem and stabilizes and regularizes the training. Whereas ST-Net can predict only 102 genes with positive correlation, our framework was able to predict 237. ST-Net revealed the top-5 predicted genes with a median correlation coefficient value of (0.34, 0.33, 0.31, 0.30, and 0.29), and for smoothed data (0.49, 0.50, 0.50, 0.52, and 0.43). In contrast, with our framework using EfficientNet-b0 with the proposed AuxNet, we could predict 24 genes with a median correlation coefficient greater than 0.50, which represents a considerable increase in performance.

Notwithstanding the merits of BrST-Net, there is still much room for further study and development in this field. To mitigate the risk of overfitting and to ensure the robustness and generalizability of our model, we employed a rigorous out-of-fold cross-validation strategy on 22 out of the 23 patients. This strategy ensures that our model is not overly reliant on a specific subset of the data and can generalize well to unseen data. However, we recognize that the generalizability of our model can be impacted by various factors in practical scenarios, including data acquisition techniques. Nevertheless, it is worth noting that the BrST-Net framework holds potential for implementation on a large-scale dataset, which can lead to improved performance and further enhance the model's generalizability. As is well known, deep learning methods require a large dataset to obtain good results. Thus we expect future studies using larger datasets to increase the prediction performance and robustness of the models. A richer dataset from a larger number of patients would also expand our framework's generalisability. Because each model has unique performance characteristics, combining two or more models can be beneficial. Also, it may be possible to do further gene expression filtering and model training for predicting target gene expression. There is a need to further improve the number of genes predicted with high correlation, which are of established relevance in cancer biology and treatment. In the meantime, BrST-Net could serve as an inexpensive and fast high-throughput screening tool for large numbers of patient samples to direct downstream definitive molecular analyses.

Data availability

Data used in this study is available from Mendeley Data at <https://data.mendeley.com/datasets/29ntw7sh4r>. The code for the BrST-Net framework can be found at the following GitHub repository: <https://github.com/Mamunur-20/BrSTNet>.

Received: 26 May 2023; Accepted: 7 August 2023

Published online: 21 August 2023

References

- Harbeck, N. *et al.* Breast cancer. *Nat. Rev. Dis. Primers* **5**, 66. <https://doi.org/10.1038/s41572-019-0111-2> (2019).
- Zhou, X. *et al.* A comprehensive review for breast histopathology image analysis using classical and deep neural networks. *IEEE Access* **8**, 90931–90956. <https://doi.org/10.1109/ACCESS.2020.2993788> (2020).
- Krithiga, R. & Geetha, P. Breast cancer detection, segmentation and classification on histopathology images analysis: A systematic review. *Arch. Comput. Meth. Eng.* **28**, 2607–2619. <https://doi.org/10.1007/S11831-020-09470-W/TABLES/2> (2021).
- Li, X. *et al.* A comprehensive review of computer-aided whole-slide image analysis: From datasets to feature extraction, segmentation, classification and detection approaches. *Artif. Intell. Rev.* **55**, 4809–4878. <https://doi.org/10.1007/s10462-021-10121-0> (2022).
- Ståhl, P. L. *et al.* Visualization and analysis of gene expression in tissue sections by spatial transcriptomics. *Science* **353**, 78–82. <https://doi.org/10.1126/SCIENCE.AAF2403> (2016).
- Larsson, L. *et al.* Spatially resolved transcriptomics adds a new dimension to genomics. *Nat. Meth.* **18**, 15–18. <https://doi.org/10.1038/s41592-020-01038-7> (2021).
- Rodrigues, S. G. *et al.* Slide-seq: A scalable technology for measuring genome-wide expression at high spatial resolution. *Science* **363**, 1463–1467. https://doi.org/10.1126/SCIENCE.AAW1219/SUPPL_FILE/AAW1219S1.MOV (2019).
- Stickels, R. R. *et al.* Highly sensitive spatial transcriptomics at near-cellular resolution with Slide-seqV2. *Nat. Biotechnol.* **39**, 313–319. <https://doi.org/10.1038/s41587-020-0739-1> (2020).
- Vickovic, S. *et al.* High-definition spatial transcriptomics for in situ tissue profiling. *Nat. Methods* **16**, 987–990. <https://doi.org/10.1038/s41592-019-0548-y> (2019).
- Ke, R. *et al.* In situ sequencing for RNA analysis in preserved tissue and cells. *Nat. Methods* **10**, 857–860. <https://doi.org/10.1038/nmeth.2563> (2013).
- Young, A. P. *et al.* A technical review and guide to RNA fluorescence in situ hybridization. *PeerJ* **8**, e8806. <https://doi.org/10.7717/PEERJ.8806/SUPP-1> (2020).

12. Rao, A. *et al.* Exploring tissue architecture using spatial transcriptomics. *Nature* **596**, 211–220. <https://doi.org/10.1038/s41586-021-03634-9> (2021).
13. Monjo, T. *et al.* Efficient prediction of a spatial transcriptomics profile better characterizes breast cancer tissue sections without costly experimentation. *Sci. Rep.* **12**, 4133. <https://doi.org/10.1038/s41598-022-07685-4> (2022).
14. Tan, X. *et al.* SpaCell: Integrating tissue morphology and spatial gene expression to predict disease cells. *Bioinform.* **36**, 2293–2294. <https://doi.org/10.1093/BIOINFORMATICS/BTZ914> (2020).
15. Schmauch, B. *et al.* A deep learning model to predict RNA-Seq expression of tumours from whole slide images. *Nat. Commun.* **11**, 3877. <https://doi.org/10.1038/s41467-020-17678-4> (2020).
16. He, B. *et al.* Integrating spatial gene expression and breast tumour morphology via deep learning. *Nat. Biomed. Eng.* **4**, 827–834. <https://doi.org/10.1038/s41551-020-0578-x> (2020).
17. Palacio, S. *et al.* Contextual classification using self-supervised auxiliary models for deep neural networks. *Int. Conf. Pattern Recogn.* <https://doi.org/10.1109/ICPR48806.2021.9412175> (2020).
18. Stenbeck, L. *et al.* Human breast cancer in situ capturing transcriptomics. *Mendeley Data* <https://doi.org/10.17632/29ntw7sh4r.5> (2021).
19. Ciompi, F. *et al.* The importance of stain normalization in colorectal tissue classification with convolutional networks. *IEEE Int. Symp. Biomed. Imaging* <https://doi.org/10.1109/ISBI.2017.7950492> (2017).
20. Byfield, P. StainTools: tools for tissue image stain normalization and augmentation in Python 3 (2022). <https://github.com/Peter554/StainTools>.
21. Vahadane, A. *et al.* Structure-preserving color normalization and sparse stain separation for histological images. *IEEE Trans. Med. Imaging* **35**, 1962–1971. <https://doi.org/10.1109/TMI.2016.2529665> (2016).
22. Seal, R. L. *et al.* Genenames.org: The HGNC resources in 2023. *Nucleic Acids Res.* **51**, D1003–D1009. <https://doi.org/10.1093/NAR/GKAC888> (2023).
23. Gehlot, S., Gupta, A. & Gupta, R. Sdct-auxnet#: Dct augmented stain deconvolutional cnn with auxiliary classifier for cancer diagnosis. *Med. Image Anal.* **61**, 101661 (2020).
24. Yu, D., Duan, H., Fang, J. & Zeng, B. Predominant instrument recognition based on deep neural network with auxiliary classification. *IEEE/ACM Trans. Audio Speech Lang. Process.* **28**, 852–861 (2020).
25. Xia, Z. & Kim, J. Enhancing mask transformer with auxiliary convolution layers for semantic segmentation. *Sensors* **23**, 581 (2023).
26. Palacio, S., Engler, P., Hees, J. & Dengel, A. Contextual classification using self-supervised auxiliary models for deep neural networks. In *2020 25th International Conference on Pattern Recognition (ICPR)*, 8937–8944. <https://ieeexplore.ieee.org/abstract/document/9412175/> (IEEE, 2021).
27. Szegedy, C. *et al.* Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1–9. <https://ieeexplore.ieee.org/abstract/document/7298594/> (2015).
28. He, K. *et al.* Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778 (2016).
29. Al-Hajja, Q. A. & Adebajo, A. Breast cancer diagnosis in histopathological images using ResNet-50 convolutional neural network. *IEEE Int. IOT Electron. Mechatron. Conf.* <https://doi.org/10.1109/IEMTRONICS51293.2020.9216455> (2020).
30. Jiang, Y. *et al.* Breast cancer histopathological image classification using convolutional neural networks with small SE-ResNet module. *PLoS ONE* **14**, e0214587. <https://doi.org/10.1371/journal.pone.0214587> (2019).
31. Al Husaini, M. A. S. *et al.* Thermal-based early breast cancer detection using Inception V3, Inception V4 and modified Inception MV4. *Neural Comput. Appl.* **34**, 333–348. <https://doi.org/10.1007/s00521-021-06372-1> (2022).
32. Saini, M. & Susan, S. Data augmentation of minority class with transfer learning for classification of imbalanced breast cancer dataset using Inception-V3. *Iberian Conf. Pattern Recogn. Image Anal.* https://doi.org/10.1007/978-3-030-31332-6_36 (2019).
33. Tan, M. & Le, Q. EfficientNet: rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning (ICML)*, 6105–6114 (2019).
34. Wang, J. *et al.* Boosted EfficientNet: detection of lymph node metastases in breast cancer using convolutional neural networks. *Cancers* **13**, 661. <https://doi.org/10.3390/CANCERS13040661> (2021).
35. Kallipolitis, A. *et al.* Ensembling EfficientNets for the classification and interpretation of histopathology images. *Algorithms* **14**, 278. <https://doi.org/10.3390/A14100278> (2021).
36. Ahmad, N. *et al.* Transfer learning-assisted multi-resolution breast cancer histopathological images classification. *Vis. Comput.* **38**, 2751–2770. <https://doi.org/10.1007/S00371-021-02153-Y/FIGURES/21> (2022).
37. Byeon, S.-J. *et al.* Automated histological classification for digital pathology images of colonoscopy specimen via deep learning. *Sci. Rep.* **12**, 12804. <https://doi.org/10.1038/s41598-022-16885-x> (2022).
38. Munien, C. & Viriri, S. Classification of hematoxylin and eosin-stained breast cancer histology microscopy images using transfer learning with EfficientNets. *Comput. Intell. Neurosci.* **2021**, 5580914. <https://doi.org/10.1155/2021/5580914> (2021).
39. Dosovitskiy, A. *et al.* An image is worth 16x16 words: transformers for image recognition at scale. *arXiv:2010.11929* <https://doi.org/10.48550/arxiv.2010.11929> (2020).
40. Pearson, K. F. Liii. on lines and planes of closest fit to systems of points in space. *Lond. Edinb. Dublin Philos. Mag. J. Sci.* **2**, 559–572. <https://doi.org/10.1080/14786440109462720> (1901).
41. Benesty, J. *et al.* Pearson correlation coefficient. In *Noise Reduction in Speech* (ed. Benesty, J.) 37–40 (Springer, 2009).
42. Onik, M. M. H. *et al.* Prediction of a gene regulatory network from gene expression profiles with linear regression and Pearson correlation coefficient. *arXiv* <https://doi.org/10.48550/arxiv.1805.01506> (2018).

Acknowledgements

EKAM is supported by a Researcher Exchange and Development in Industry (REDI) Fellowship from MTPConnect and MRFF Australia. The authors thank the National Computational Infrastructure (NCI) Australia for providing access to the Gadi high-performance computing system, which contributed to the research results reported in this paper.

Author contributions

M.M.R. contributed to the design and implementation of the research, to the analysis of the results and to the writing of the manuscript. E.K.A.M. and E.M. supervised the project, contributed to the conception and design of the study, and provided critical review and revisions of the manuscript. All authors discussed the results and contributed to the final manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-023-40219-0>.

Correspondence and requests for materials should be addressed to E.M.

Reprints and permissions information is available at www.nature.com/reprints.

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



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 Author(s) 2023