

OPEN

# DeePathology: Deep Multi-Task Learning for Inferring Molecular Pathology from Cancer Transcriptome

Behrooz Azarkhalili<sup>1,4</sup>, Ali Saberi<sup>2</sup>, Hamidreza Chitsaz<sup>3</sup> & Ali Sharifi-Zarchi<sup>1,2\*</sup>

Despite great advances, molecular cancer pathology is often limited to the use of a small number of biomarkers rather than the whole transcriptome, partly due to computational challenges. Here, we introduce a novel architecture of Deep Neural Networks (DNNs) that is capable of simultaneous inference of various properties of biological samples, through multi-task and transfer learning. It encodes the whole transcription profile into a strikingly low-dimensional latent vector of size 8, and then recovers mRNA and miRNA expression profiles, tissue and disease type from this vector. This latent space is significantly better than the original gene expression profiles for discriminating samples based on their tissue and disease. We employed this architecture on mRNA transcription profiles of 10750 clinical samples from 34 classes (one healthy and 33 different types of cancer) from 27 tissues. Our method significantly outperforms prior works and classical machine learning approaches in predicting tissue-of-origin, normal or disease state and cancer type of each sample. For tissues with more than one type of cancer, it reaches 99.4% accuracy in identifying the correct cancer subtype. We also show this system is very robust against noise and missing values. Collectively, our results highlight applications of artificial intelligence in molecular cancer pathology and oncological research. DeePathology is freely available at <https://github.com/SharifBioinf/DeePathology>.

Improving the accuracy of cancer diagnosis is extremely important for millions of patients and for far more non-patients who are tested worldwide every year. Despite great advances in oncologic pathology, there has been a significant ratio of errors that potentially affect the diagnostic results and/or treatment strategies<sup>1</sup>. Our systematic search through COREMINE ([www.coremine.com](http://www.coremine.com)) revealed 7652 article abstracts containing both *neoplasms* and *diagnostic errors* (or synonymous terms), with an increasing trend over time. An M.D. Anderson Cancer Center study of 500 brain or spinal cord biopsies that were submitted to their neuropathology consultation service for a second opinion revealed 42.8% disagreement between the original and the review diagnoses, including 8.8% serious cases<sup>2</sup>. A study of 340 breast cancer patients identified differences between the first and the second pathology opinions in 80% of the cases, including major changes that altered surgical therapy occurred in 7.8% of cases<sup>3</sup>. A review of 66 thyroid cancer patients revealed a different pathological diagnosis of 18% of the cases<sup>4</sup>. A recent study verified the accuracy and reproducibility of pathologists' diagnoses of melanocytic skin lesions for 240 skin biopsy cases from 10 US states and revealed 8–75% error rates in different interpretation classes and an estimated 17.8% whole-population error rate<sup>5</sup>. Another recent study of 263 Australian Lichenoid keratosis patients revealed a diagnosis failure rate higher than 70%, including 47% of the cases misdiagnosed as basal cell carcinoma<sup>6</sup>. This situation is even worse in rare types of cancer. A study of 26 patients revealed 30.8% misdiagnosis ratio in discriminating common gastric adenocarcinoma from hepatoid adenocarcinoma of the stomach, a rare subtype of gastric cancer<sup>7</sup>. Accurate diagnosis has been also challenging for a number of cancer types, including soft tissue sarcomas that are often misdiagnosed as other types of cancer<sup>8</sup>.

One limitation of the current molecular pathology methods such as Immunohistochemistry (IHC) is the limited number of genes or proteins monitored for diagnosis. Staining biopsies using antibodies against one or

<sup>1</sup>Department of Stem Cell Biology and Technology, Royan Institute, Tehran, Iran. <sup>2</sup>Department of Computer Engineering, Sharif University of Technology, Tehran, Iran. <sup>3</sup>Department of Computer Science, Colorado State University, Fort Collins, CO, USA. <sup>4</sup>Department of Mathematics and Computer Science, Sharif University of Technology, Tehran, Iran. \*email: [asharifi@sharif.edu](mailto:asharifi@sharif.edu)

two proteins cannot discriminate between different cancer types if they have similar expression patterns of the target proteins. One possible solution is to use the whole-transcriptome of tissue biopsies<sup>9</sup>. But this approach is computationally challenging and different algorithmic and machine learning approaches have been employed so far to address this problem. A subset of research is focused on *binary classification*, e.g. discriminating between normal vs. tumor samples<sup>10–13</sup>. Stacked autoencoders are used for binary classification between glioma grades III vs. IV, and evaluated on 185 samples<sup>14</sup>. These methods, however, can have limited clinical applications since most of the molecular pathology problems are *multiclass*, e.g. assigning each sample to one of the different cancer types. Reaching high accuracy in classification problems usually becomes harder as the number of classes increases. Even a random assignment of samples to two classes will achieve 50% accuracy if the classes are *balanced* (i.e. there are an equal number of samples in each class), but a random classification will be around 3% accurate if there are 33 balanced classes. Hence, it is important to consider the number of classes for comparing the accuracy of different techniques.

Optimal Feature Weighting (OFW) is one of the earliest multiclass algorithms employed for cancer sample classification based on Microarray transcriptomes. This algorithm selects an optimal discriminative subset of genes and uses Support Vector Machines (SVM) or Classification And Regression Trees (CART). In previous work, it has been applied to five different problems, each consisting of 3 to 11 classes, without explicitly mentioning the obtained accuracy<sup>15</sup>. A combination of SVM with Recursive Feature Elimination (RFE) is used to classify Microarray data of three cancer-related problems consisting of 3 to 8 classes, with accuracy between 95% (for 8-class) to 100% (for 3-class)<sup>16</sup>. Greedy search over top-scoring gene-sets has achieved an average 88% accuracy, ranging from 48% to 100%, on seven different cancer datasets, each consisting of 3 or 4 classes with 40 to 96 samples per dataset<sup>17</sup>.

One of the largest databases of cancer transcriptome, genome and epigenome profiles is Genomic Data Commons (GDC) that includes The Cancer Genome Atlas (TCGA) and Therapeutically Applicable Research to Generate Effective Treatments (TARGET) programs<sup>18</sup>. There have been comprehensive works to analyze GDC data from different perspectives including identification of cancer driver somatic and pathogenic germline variations<sup>19</sup>, oncogenic signaling pathways<sup>20</sup>, the role of cell-of-origin<sup>21</sup> and cancer stem cells<sup>22</sup>, relationships between tumor genomes, epigenomes and microenvironments<sup>19</sup>. However, there has been little effort directed towards developing a molecular cancer pathology framework out of this valuable data. Semi-supervised stacked sparse autoencoders are used for binary classification of three types of cancer, and evaluated on 1311 samples<sup>23</sup>. DeepGene, a feed-forward artificial neural network (ANN) based classifier, uses somatic point mutations profiles in order to assign each sample to one of 12 different cancer types. It achieved a mean 58% and maximum 64% accuracy over 3122 TCGA samples from 12 cancer types<sup>24</sup>. Another study used genetic algorithms and k-Nearest Neighbors (KNN) to classify TCGA samples based on RNA-seq transcriptome profiles. It reached about 90% accuracy for classification of 602 normal samples and 9096 samples from 31 tumor types<sup>25</sup>.

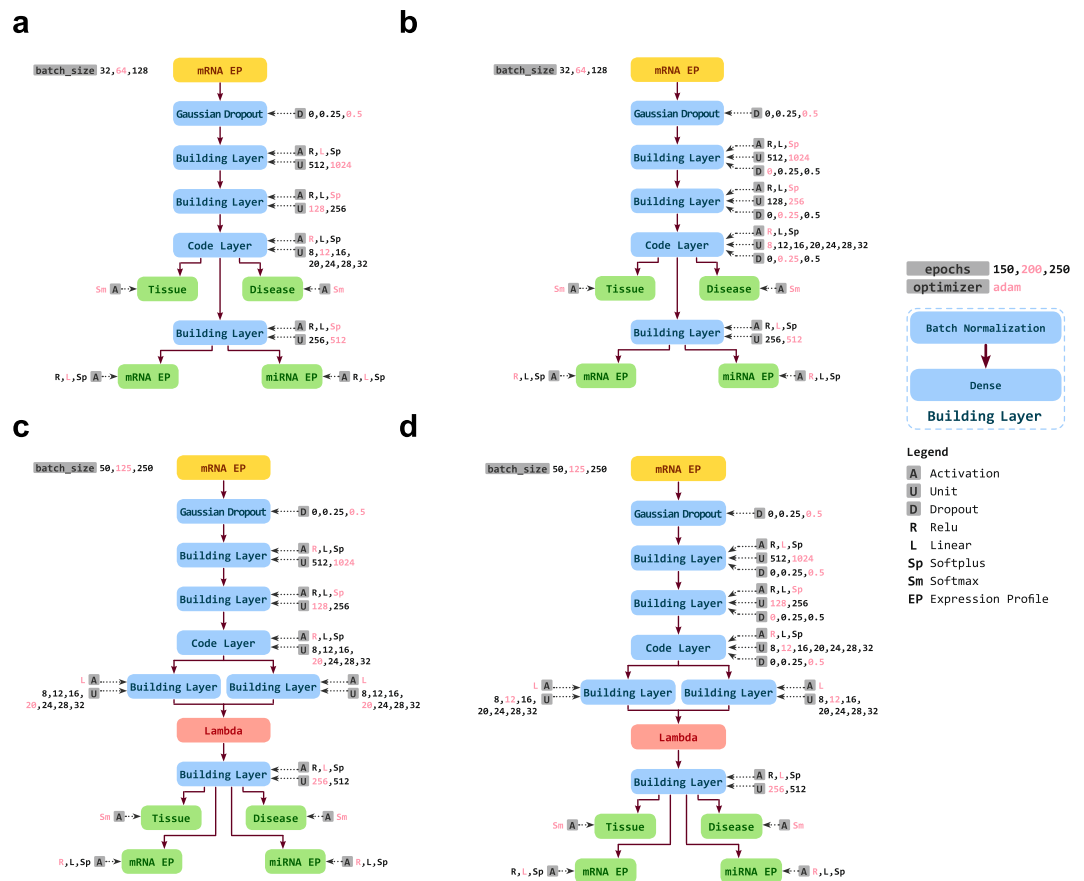
## Methods

Here we used Deep Neural Networks (DNNs) in a multi-task learning approach to infer different biological and clinical information from transcriptome profiles. We designed four different architectures as shown in Fig. 1 based on two different methods, including Contractive Autoencoder (CAE) and Variational Autoencoder (VAE). Each of our DNNs consists of two parts that are serially connected: the *encoder* part, that learns to convert a given mRNA expression profile (mRNA EP) of a clinical sample at the input layer to a latent representation, which we call Cell Identity Code (CIC), and the *decoder* part that infers multiple outputs from the CIC. While the CIC is a simple vector of numerical values in the CAEs, the VAEs encode the input into two equal-sized vectors, which represent means and standard deviations of multiple Gaussian distributions. For each method, we designed two architectures, one simpler and another consisting of dropout layer after each layer of the encoder. To make our architectures resistant to missing values and noisy data, we added a *Gaussian Dropout* layer after the input, which perturbs the data with Gaussian noise and randomly sets some input values to zero. More details are provided in the online methods.

What makes our architectures different from conventional autoencoders is its particular design to learn *useful* latent representations of the input data. Altered weights of the encoder part result in different latent representations of the same input, which can be decoded to the same output by different decoder weights. Hence, there is an infinitely large number of latent representation sets for a given set of data, however, the question is whether all of these representations are useful. Some latent representations might be extremely cryptic, while the others might be very useful for classification of biological samples, or obtaining other information. For instance, classification would be much easier if all cells of each type cluster together, distantly from the other clusters, in the latent representation space. The challenge is how to train the network in order to learn useful latent representations of the mRNA EPs.

To address this challenge, we designed the decoder part of each network to simultaneously learn four different classification and regression problems using the CIC, in a multi-task learning scheme: (i) reproducing mRNA EP as one of the outputs that is as close to the original mRNA EP in the input as possible, (ii) predicting a miRNA expression profile (miRNA EP) that is as close as possible to the experimentally measured miRNA EP of the same sample, (iii) predicting the sample tissue of origin, among 27 different tissues, and (iv) predicting the sample disease state, which can be either normal or one of 33 different cancer types. All parts of each network, including encoder, decoder and classifiers were trained simultaneously.

From the above tasks, (i) and (ii) can be viewed as non-linear regression, and (iii) and (iv) are classification. Importantly, the multi-task learning part of the DNN is aimed to accomplish all of these tasks only by getting the CIC as the input. Task (i) is to ensure the CIC stores much of the information in the original mRNA EP. Due to task (ii), we selected from GDC a subset of samples having both mRNA and miRNA EPs available. Furthermore, we removed 11 miRNA-encoding genes from the mRNA profiles, to make this task non-trivial.



**Figure 1.** Different architectures of DNN used for this study. (a) Contractive Autoencoder (CAE), (b) Dropout Contractive Autoencoder (Dropout-CAE), (c) Variational Autoencoder (VAE), (d) Dropout Variational Autoencoder (Dropout-VAE). For each network, the layers are shown as boxes and the connections between them as arrows. The evaluated hyperparameters of each layer are shown next to each layer. The hyperparameter values in red show the optimal parameters identified through hyperparameter optimization. See extended methods for more details.

The other key advantage of our model is *hyperparameter optimization*. In addition to the internal network parameters (i.e. synapse weights and neuron bias values), each network has a set of hyperparameters, including the number of neurons in each layer, activation functions, size of mini-batches, standard deviations of the Gaussian noises, ratios of dropout layers, and the number of training epochs. Altered values of hyperparameters greatly affect the network results. Due to their cryptic inter-dependencies, all hyperparameters are required to be optimized simultaneously. For this purpose, we performed a comprehensive hyperparameter optimization through a Bayesian approach that runs in a number of iterations. In each iteration, it tries to find a set of hyperparameters that has the maximum likelihood of optimizing the network training outcome by integrating the results of all previous iterations in a Bayesian model. This process has superior advantages over grid search or random search of the hyperparameters in reducing the search space and more direct approaching the optimal hyperparameters using much fewer iterations.

**Data pre-processing.** Transcriptome profiles of 11,500 samples were obtained from the Genomic Data Commons (GDC) consortium<sup>18</sup>, including The Cancer Genome Atlas (TCGA)<sup>19</sup> and the Therapeutically Applicable Research to Generate Effective Treatments (TARGET) databases. For each mRNA profile, its corresponding miRNA profile was needed to train the model. So we kept the samples that both mRNA and miRNA expression profiles were available. This resulted in a total number of 10787 samples. For training a VAE, the total number of samples should be a multiple of batch size. Since a batch size of 250 was chosen for the model training process, total number of samples should be a multiple of 250. Thus we kept 10750 samples for the analysis, including 10124 tumor and 626 normal samples. The mRNA and miRNA profiles of each sample were matched by the ID values of the patients, using the TCGAbiolinks R/Bioconductor package<sup>26</sup>. We randomly selected about 10% of the whole data as the test dataset, and the remaining samples were assigned to the training dataset.

One of our tasks was to predict expression level of miRNAs from mRNA profile. Hence we removed all miRNA genes from mRNA profiles, to ensure the answer to this task is not provided to the network as input. As the result, each mRNA and miRNA expression profile was of sizes 19671 and 2588, respectively. The precise number of samples from each type of tissue and tumor is provided in Tables S1 and S2.

**Deep autoencoder.** An interesting Neural Network (NN) architecture is autoencoder (AE), which compresses the input into a latent space representation, and then reconstructs the input back from this representation. In other words, an AE seeks to learn an identity-like mapping function  $f$  such that  $f(x) \approx x$ . AEs can reduce the dimensionality of data without losing significant information and can be trained using unlabeled data, hence they are widely used in different problems including data compression, dimensionality reduction, manifold learning, and feature learning<sup>27</sup>.

Each autoencoder consists of two parts, the *encoder* and the *decoder*, which can be defined as transition functions  $f$  and  $g$  such that:

$$\begin{aligned} f &: X \rightarrow \mathcal{F} \\ g &: \mathcal{F} \rightarrow X \\ f, g &= \arg \min_{\phi, \psi} \mathcal{J}(x, \tilde{x}(\phi, \psi, x)) \end{aligned}$$

where  $\psi$  and  $\phi$  are the encoder and the decoder functions, respectively,  $\tilde{x}(\phi, \psi, x) = \phi \circ \psi(x)$  is the reconstruction of input vector  $x$  and  $\mathcal{J} = \sum_{x \in D} L(x, \tilde{x}(\phi, \psi, x))$  is the total loss, which is evaluated as the summation of reconstruction error  $L$  on training dataset  $D$ .

**Autoencoders variants.** While training an AE, we aim not only to reconstruct the input from the latent representation, but also to extract beneficial features in this representation. Therefore, it is vital to utilize some forms of regularization techniques to avoid overfitting or useless representations, even if the AE can reconstruct the input with minimal loss<sup>27,28</sup>. Different regularization techniques can veritably produce different variations of objective functions, and subsequently different features extracted from the data. The next sub-section explains the autoencoder variants of this study.

*Denoising autoencoder.* A beneficial form of regularization is used in *denoising autoencoders* (DAE), where the input vector  $x$  is slightly corrupted and the autoencoder is expected to reconstruct the clean data from the latent representation<sup>29</sup>. As a result, the DAE learns to resist against input noise and overfitting. The following objective function is used for training DAEs:

$$\mathcal{J} = \sum_{x \in D} \mathbb{E}_{\hat{x} \sim q(\hat{x}|x)} L(x, \tilde{x}(\phi, \psi, \hat{x})) \quad (1)$$

where the expectation is evaluated over the corrupted versions  $\hat{x}$  of the original data  $x$ , obtained from a corruption function  $q(\hat{x}|x)$ . This objective is optimized by stochastic gradient descent (SGD) or another iterative optimization algorithm. Additive isotropic Gaussian noise and binary masking noise are among the most frequently used corruption processes.

Dropout<sup>30,31</sup> is another regularization technique that introduces some noise to the nodes of any hidden layer, in contrast to the DAE that adds noise only to the input layer. The foremost dropout techniques utilized in deep learning are Bernoulli and Gaussian. In the former case, the output values of individual nodes are either dropped to zero with a probability  $1 - p$ , or kept unchanged with a probability  $p$ . In the latter case, a multiplicative one-centered Gaussian noise  $\mathcal{N}(1, \sigma^2)$  is applied to the output values of the nodes.

*Contractive autoencoder.* An alternative regularization technique is used in contractive autoencoder (CAE). If  $h = \psi(x)$  is the latent representation of the input data  $x$ , the regularization term in CAE is the total squares of all partial derivatives of  $h$  with respect to each dimension of the previous layer; therefore, the objective function of a CAE is expressed as:

$$\mathcal{J} = \sum_{x \in D} L(x, \tilde{x}(\phi, \psi, x)) + \lambda \left\| \left\| J_{\psi}(x) \right\|_F \right\|_F^2 \quad (2)$$

where the penalty term  $\left\| \left\| J_{\psi}(x) \right\|_F \right\|_F^2$  is the Frobenius norm of the Jacobian matrix of the encoder activations with respect to the input, and  $\lambda$  is a balancing factor. The main goal of this term is to enforce the learned representation to be robust against small variations in the input data.

*Variational autoencoder.* Variational autoencoder (VAE) is a special type of autoencoder with additional constraints on the encoded representations. It assumes that a latent, unobserved random variable  $z$  exists, which can lead to the observations  $x$  by some stochastic mapping. As a result, its objective is to approximate the distribution of the latent variable  $z$  given the observations  $x$ .

VAEs replace deterministic functions in the encoder and decoder by stochastic mappings; and compute the objective function in virtue of the density functions of the random variables:

$$\mathcal{J}(\phi, \theta, x) = D_{KL} \left( q_{\phi}(z|x) \left\| \left\| p_{\theta}(z) \right\| \right) - \mathbb{E}_{q_{\phi}((z|x))} (\log(p_{\theta}(x|z))) \quad (3)$$

where  $D_{KL}$  stands for the Kullback–Leibler divergence, and  $q$  is the distribution approximating the true latent distribution of  $z$ , and  $\theta, \phi$  are the parameters of each distribution. The prior distribution over the latent variables is generally set to standard multivariate Gaussian  $p_{\theta}(z) = \mathcal{N}(0, \sigma^2 I)$ ; however, alternative distribution have also been recently considered.

Although AEs can not generally be able to construct meaningful outputs from arbitrary encodings, VAE can learn a model of the data that can generate new samples from scratch by random sampling from the latent distribution. Therefore, VAEs are among the generative models.

**Models architecture.** Our problem consists of two regression, namely reproducing mRNA and miRNA profiles from the latent variables, and two classification tasks to determine the tissue and disease state of each sample. All mRNA and miRNA expression profiles were normalized in  $[0, 1]$  by max-norm method of Scikit-Learn package<sup>32</sup>. We subsequently constructed the multi-input and multi-output models in Keras<sup>33</sup>. Keras is a high-level Python NN library that runs on top of either TensorFlow<sup>34</sup> or Theano<sup>35</sup>.

We constructed four different NN architectures:

- Variational autoencoder (VAE)
- Dropout-VAE, an extension of VAE with Bernoulli and Gaussian dropout layers being utilized for denoising.
- Contractive autoencoder (CAE)
- Dropout-CAE, the extended CAE with Bernoulli and Gaussian dropout layers.

To make each network tolerate input noise, an optional Gaussian noise  $\mathcal{N}(0, \sigma^2)$  can be added to the input.

**Loss functions.** We utilized *cosine similarity* as the loss for the classification tasks. Another choice was *categorical cross entropy*, but we preferred to use the former function since it provides bounded results, which can be easily used in hyperparameter optimization. We also used *mean squared error (MSE)* as the loss function for the regression tasks. The total loss in our multi-task learning problem was the weighted sum of individual losses, where the weights were  $5 \times 10^{-1}$  for each classification task, and  $1 \times 10^{-3}$  for each regression task. To evaluate the model performance in the validation dataset, we used accuracy as the classification metric, and MSE and mean absolute error (MAE) for the regression tasks.

For presentation of the results, we preferred to use balanced accuracy. For imbalanced datasets, the conventional accuracy can be misleading. For instance if 95% of the data are normal, classifying all samples as normal gives 95% accuracy. Balanced accuracy solves this problem by normalizing the number of correctly predicted samples of each class by the size of the same class. More specifically, we used the python package scikit-learn implementation of balanced accuracy. If  $y_i$ ,  $\hat{y}_i$ , and  $w_i$  are the true label, predicted label, and weight of sample  $i$ , respectively, then the adjusted sample weights  $\hat{w}_i$  and balanced accuracy are computed as follows:

$$\hat{w}_i = \frac{w_i}{\sum_j 1(y_i = y_j)w_j} \quad (4)$$

$$\text{Balanced accuracy}(y, \hat{y}, w) = \frac{1}{\sum \hat{w}_i} \sum_i 1(\hat{y}_i = y_i)\hat{w}_i \quad (5)$$

where  $1(x)$  is the indicator function.

**Batch normalization.** A major issue in training the DNNs is altered distribution of each layer inputs, as the weights and other parameters of the previous layers change. As a result of this *internal covariate shift*, the learning rate should be lowered that causes reduced training speed rate, and also the initialization parameters should be assigned carefully<sup>36</sup>. A well-known approach to address these challenges is *batch normalization*.

Let's  $x_{1...m}$  be the values of an activation  $x$  during a mini-batch. Then the batch normalized values  $y_{1...m}$  are computed as follows:

$$y_i = \gamma \frac{x_i - \mu}{\sqrt{\sigma^2 + \varepsilon}} + \beta \quad (6)$$

where  $\mu$  and  $\sigma^2$  are the mean and variance of the  $x_{1...m}$  values,  $\gamma$  and  $\beta$  are the scaling parameters to be learned, and  $\varepsilon$  is a small constant value added to the mini-batch variance for numerical stability.

This method basically standardizes the inputs of each layer in such a way that they have a mean and standard deviation of zero and one, respectively. Then scales the standardized values by a linear function, with parameters that are learned. Batch normalization is analogous to how the inputs to the networks are standardized, but it can be performed for each internal layer of a DNN. It turns out that, extending this technique to hidden layers can significantly improve the training speed.

As depicted in Fig. 1, each *building layer* of our network is consisted of a batch normalization layer, followed by a dense layer. Our experience showed the batch normalization layers significantly improve the training speed of the networks.

**Hyperparameters tuning.** Many machine learning methods have a set of architectural parameters, called *hyperparameters*, which are determined prior to training the model. For example, the number of layers, the number of neurons per layer, the type of activation functions, and the type of dropout or noises are among the DNN hyperparameters. Since the values of hyperparameters affect the model architecture, hyperparameter optimization can be considered as model selection technique.

A key advantage of our work is optimizing the network hyperparameters in order to achieve the best classification and regression objectives. Besides two simple methods of hyperparameter tuning including grid search and random search, there exists a more advanced Bayesian algorithm<sup>37</sup>.

Bayesian optimization, in contrast to random or grid search, keeps track of the past evaluations and utilizes them to define a probabilistic model mapping hyperparameters to a probability distribution of the outcome of the objective function  $P(s|h)$ , where  $s$  and  $h$  are the objective function score and hyperparameters. This is called a *surrogate* model for the objective function, and is easier to be optimized in comparison with the objective function itself. In each iteration, the Bayesian optimization selects optimal hyperparameters based on the surrogate function as the next set of hyperparameters to be evaluated by the actual objective function. Briefly, it works as the following algorithm:

- Build a surrogate probability model of the objective function.
- Find the hyperparameters that perform best on the surrogate.
- Evaluate these hyperparameters by the actual objective function.
- Update the surrogate model by incorporating the new results.
- Repeat steps 2–4 for the specified number of iterations or running time.

The main advantage of the Bayesian optimization is becoming more intelligent by continuously updating the surrogate probability model after each evaluation of the objective function. By intelligently selecting hyperparameters that are more likely to optimize the objective function in each iteration, they can find better set of hyperparameters in a fewer iterations, in comparison with random and grid search<sup>38,39</sup>.

To exploit Bayesian optimization in our model selection procedure, we utilized a Python module called Hyperopt<sup>40</sup>. Hyperopt provides efficacious algorithms and parallel infrastructure. We needed to define a search space and an objective function for the hyperparameter optimization. The next sub-section explains the search space. As the objective function, we used an average of the total losses of all data, which was split to 80% training and 20% test dataset.

While we benefited from all data in hyperparameter optimization, we performed 5-fold cross-validation while measuring classification accuracy to ensure the hyperparameters are not overfitted to some specific portion of the data.

**Hyperparameters search space.** For each architecture, we selected a set of hyperparameters; and for each hyperparameter, we defined a set of discrete values as the search space. The Cartesian product of all of these search spaces was used as the total search space of the hyperparameter optimization process.

Below is a description of the hyperparameters:

- **Units:** A critical feature of each layer is the number of its neurons. Based on our experiences in a prior work, we selected a wide range of different values as the search space of each layer. These values were selected in a decreasing order for the encoder part, to gradually narrow it down from the input to the latent (code) layer, followed by an increasing order in the decoder part.
- **Activation Functions:** A widely used activation function is Rectified Linear Unit (ReLU), which is defined as  $f(x) = \max(x, 0)$ <sup>41</sup>. Since gradients can readily flow whenever the input to the ReLU function is positive, gradient descent optimization is much simpler for ReLU than sigmoid activation functions. Although a few efficacious activation functions such as Swish<sup>42</sup> and Selu<sup>43</sup> have been recently introduced; However, there has been no significant and universal enhancement to utilize them instead of ReLU function. Despite promising features, ReLU can cause some difficulties in autoencoders; including gradient explosion and saturation. Hence we additionally used linear activation functions  $f(x) = x$  and SoftPlus activation function  $f(x) = \ln(1 + \exp(x))$ , as comprehensively discussed<sup>44</sup>.
- **Dropout Rate:** Dropout layer plays an important role by acting as regularizer to prevent overfitting. It has one main tunable parameter, which is the noise rate. We chose {0, 0.25, 0.5} as the set of noise rates. By putting 0 in the options, we provided Dropout-CAE and Dropout-VAE models. By this way, these models become more flexible to either exploit or not to exploit dropout layers into their own structures to achieve the maximum performance.

For each NN architecture, we set Hyperopt to iterate over 200 different networks and select the best one among them.

**Dimension reduction analysis and comparison with other algorithms.** We observed the performance of our architecture in reducing the initial dimension of the data by visualizing Principal Components Analysis (PCA) and t-distributed Stochastic Neighbor Embedding (t-SNE) of the mRNA profiles as well as CICs. This was done by performing 5-fold cross-validation and using the CIC of each sample when it was used in the test-dataset.

We compared the classification accuracy of the proposed deep learning method against other classification methods including *KNN*, *Extra Tree Classifier*, *Random Forest Classifier*, *Stochastic Gradient Descent (SGD) Classifier* and *Support Vector Machine (SVM)*. We selected these algorithms because they were available for parallel execution in Scikit-Learn package<sup>32</sup>. Moreover, they are among the most efficient and versatile classification algorithms. The inputs of all models were the same mRNA expression profiles. Each other machine learning algorithm was trained to solve one of the classification problems (i.e. tissue or disease). To ensure the

hyperparameters of the other classifiers are selected properly, we performed 100 iterations of hyperparameter optimization on each classification algorithm.

## Results and Discussion

**Performance of the four different DNN architectures on different tasks.** We trained the hyperparameter optimized networks for 200 epochs using the training dataset. As shown in Fig. S1, network architecture and hyperparameters have a big impact on the regression and classification results. Figure S2 shows the performance of these networks on the training dataset during the training, and Fig. S3 shows their performance on the test dataset. The CAE and Dropout-CAE architectures depicted the best performance in all four tasks. For the regression tasks (i.e. reproducing the mRNA and predicting the miRNA EP) the CAE architecture had a slightly better performance over the Dropout-CAE, but for the classification tasks, both networks performed very similarly. Furthermore, there was no sign of overfitting for these architectures as the test dataset errors did not increase during the training. Also, all accuracy and MSEs reached stable levels, that showed 200 training epochs was sufficient. The red values in Fig. 1 depict the optimal hyperparameter values. A striking observation was the size of CIC, determined by hyperparameter optimization. Although there were higher values up to 32 available as options for the size of code layer, all architectures were optimized by using smaller sizes between 8 to 20 (1). Particularly, the Dropout-CAE that outperformed other architectures was optimized by a CIC of size 8 (1b). This means that the original mRNA profile of size 19671 is encoded into an 8-dimensional latent space, and then all of the mRNA and miRNA expression profiles, tissue and disease state of the samples are obtained from this CIC space.

**DNN outperforms classical algorithms in tissue and cancer type classification.** We compared the balanced accuracy of tissue and cancer classification using our DNN in comparison with five widely-used other classification algorithms including Extra Tree, k-Nearest Neighbors, Random Forest, Stochastic Gradient Descent and SVM methods. The input of each algorithm was the whole mRNA EP of one sample, and the tissue or disease label of the sample was expected as the output. To make fair comparisons, we also performed hyperparameter optimization for all of classical classifiers (Fig. S4). The results are provided in Fig. 2a,b. While our DNN reaches balanced accuracy of 98.1% for tissue and 95.2% for disease classification, the balanced accuracy of other methods is at most 95.1% and 90.9% for tissue and disease classification, respectively.

We questioned whether classical algorithms could provide better results by having our CIC rather than the original mRNA EP as the input. For this purpose, we performed a new experiment in which the 8-dimensional CIC was given to each algorithm, including our DNN, as the input and the tissue or disease state of the sample was expected as the output (Fig. 2c,d). We used 5-fold cross validation in this experiment. Again, the accuracy of our deep classifier (that was stated earlier) was significantly better than the best accuracy obtained by other algorithms (McNemar  $p$ -value  $< 1e - 15$ ).

**CIC is a better dimensionality reduction than those obtained by classical algorithms.** An important question is whether the CIC can provide more discriminative dimensionality reduction, in comparison with classical algorithms such as Principal Components Analysis (PCA), or its derivatives. To address this question, we reduced the dimensionality of each sample to 8 using the top principal components of PCA as well as its three other derivatives (Incremental, Kernel and Sparse PCA, see scikit-learn documentation for more details). Then we compared the 8-dimensional CICs with the 8-dimensional vectors obtained by the other algorithms. To make a fair comparison, we used trained an independent Ensemble-learning algorithm consisting of one SVM and another Random Forest classifier for each dimensionality reduction algorithm. The input of the Ensemble learning model was an 8-dimensional vector of one algorithm, and the output was either tissue or disease state.

As shown in Fig. 2e,f, the maximum classification accuracy for tissue and disease states from the PCA derivatives were 38.7% and 33.2%, respectively. Ensemble-learning model, however, could be trained to predict the tissue and disease state with 93.7% and 89.3% using CICs. These results were significantly higher (McNemar  $p$ -value  $< 1e - 15$ ) than those obtained by PCA derivatives.

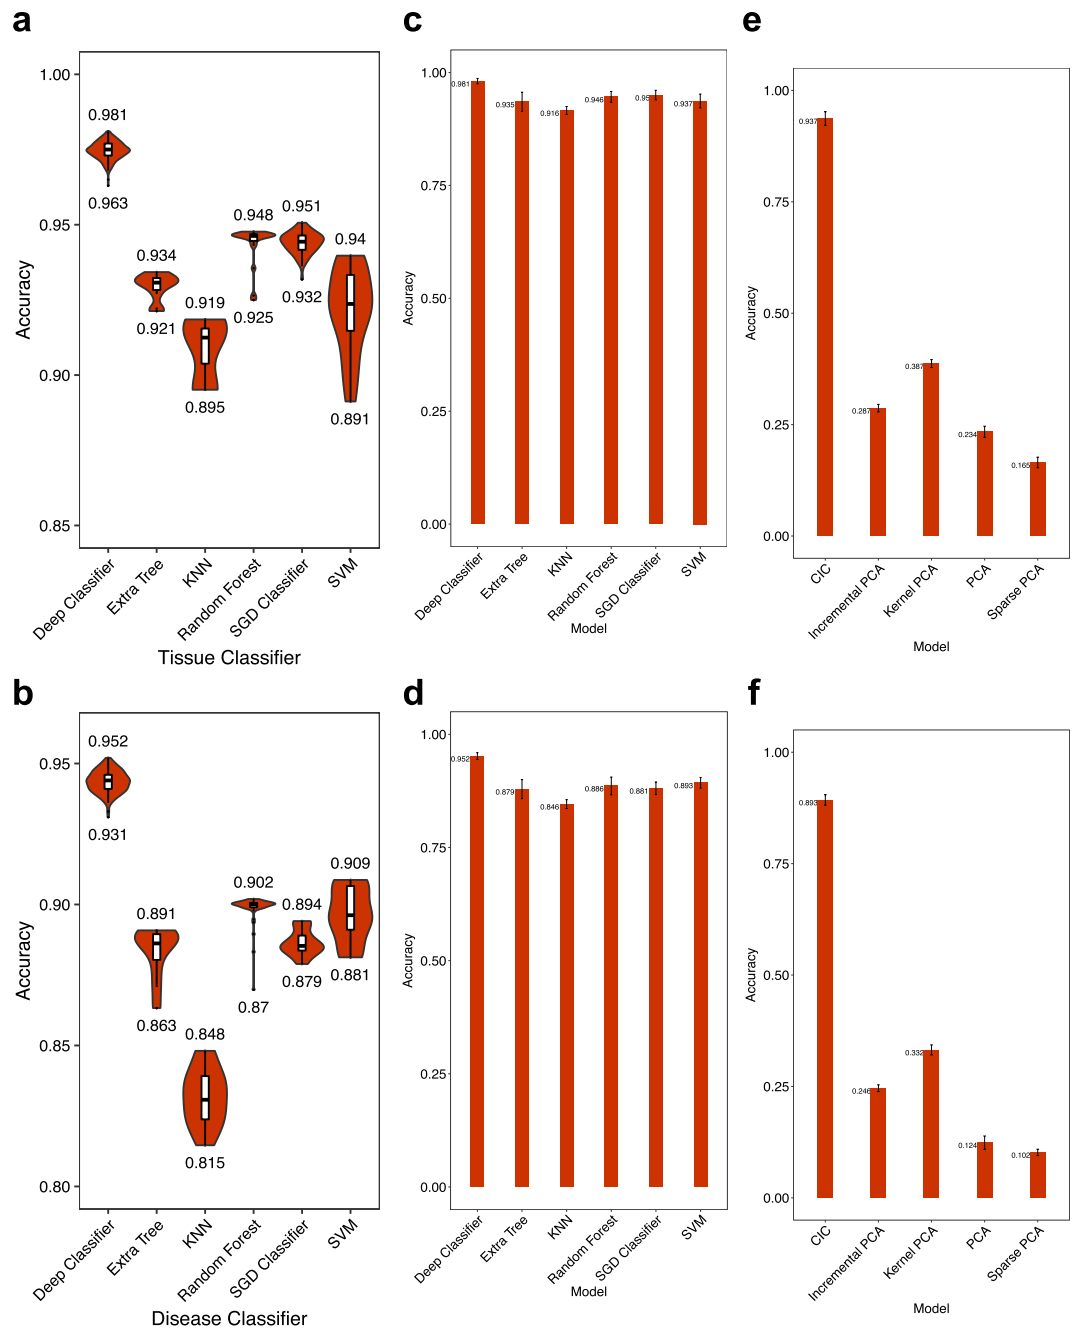
### Samples are better discriminated by CICs rather than the original mRNA expression profiles.

The TCGA data is integrated from different studies, hence one might ask whether the original mRNA expression profiles of each study are discriminated from the other studies due to batch effects. In that case, classification of different cancer types might only learn artificial patterns of batch effects that discriminate studies, rather than the real biological features of transcriptomes that discriminate tissues and cancer types.

To address this question, we compared the original 19671-dimensional mRNA expression profiles against the 8-dimensional CIC feature space that we learn using Dropout-CAE. It is noteworthy that each sample in the dataset will be used exactly once as a test sample in our cross-validation procedure; hence, all the datasets utilized for visualization and statistical analysis have been resulted in combining all corresponding samples generated by each fold of cross-validation. For visualization of the samples in each space, we used Principal Components Analysis (PCA) as a linear dimension reduction algorithm, and t-distributed Stochastic Neighbor Embedding (t-SNE)<sup>45</sup> as a non-linear manifold-learning algorithm.

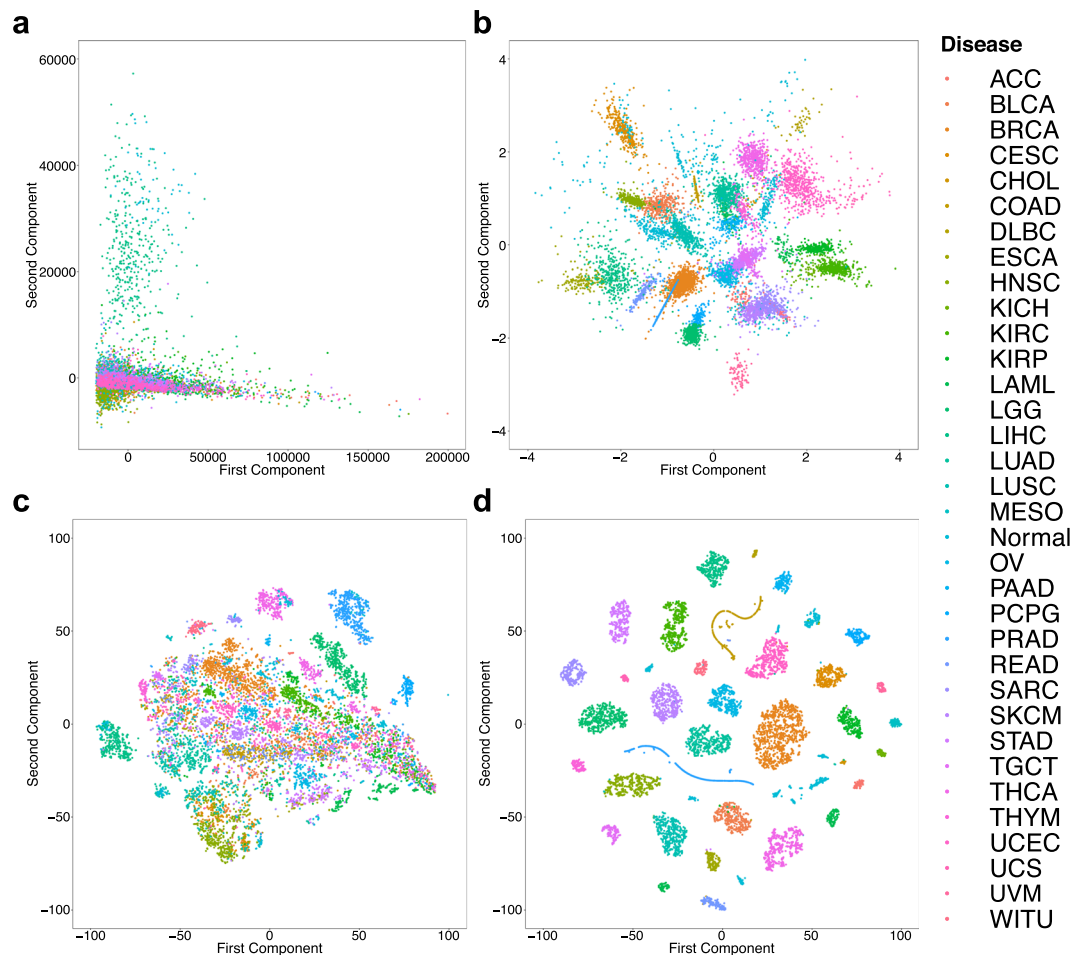
As shown in Fig. 3a, the original mRNA expression profiles from most of the studies overlap and it's impossible to discriminate disease states of the samples from 2D PCA. The same PCA projection, however, has a significantly better discrimination of disease states if applied to the 8-dimensional latent space of the CICs (Fig. 3b).

A more recent manifold learning algorithm called t-SNE is known to provide better representation of biological samples such as transcriptomes or single-cell expression profiles in low-dimensional spaces than linear methods such as PCA. We compared t-SNE visualization of original mRNA expression profiles (Fig. 3c) against latent CIC representations (Fig. 3d). Evidently, the 8-dimensional CIC space is significantly better discriminative



**Figure 2.** Accuracy of the DNN, in comparison with classical dimension reduction and classification algorithms. (Left) Each shows the distribution of tissue (a) or disease (b) classification accuracy obtained by 200 iterations of hyperparameter optimization for each algorithm. The leftmost violin shows our DNN. The input to each algorithm is the whole mRNA EP, and the tissue or disease type is expected as the output. (Middle) Accuracy of the same algorithms for tissue (c) or disease (d) classification, when the input to each algorithm is the CICs obtained by DNN. The bar height and error bar length show the mean and standard deviation of 5-fold cross-validation results, respectively. (Right) Four different dimensionality reduction algorithms are compared with the CICs obtained by the DNN. For each algorithm, the dimension of mRNA EP is reduced to 8 (the same dimension of CIC). Given the 8-dimensional vector of each algorithm as the input, an Ensemble learning classifier is trained to predict the tissue (e) or disease (f) type. Bar heights and error bar lengths represent mean and standard deviation of accuracy in 5-fold cross validation. PCA = Principal Components Analysis; Kernel PCA = Non-linear version of PCA using kernels; Sparse PCA = implementation of PCA that finds a sparse set of components that can optimally reconstruct the data; Incremental PCA = PCA using Singular Value Decomposition (SVD) of the data, keeping only the most significant singular vectors, KNN = k-nearest neighbors, SGD = Stochastic Gradient Descent, SVM = Support Vector Machine. See scikit-learn manual for more details about these algorithms.





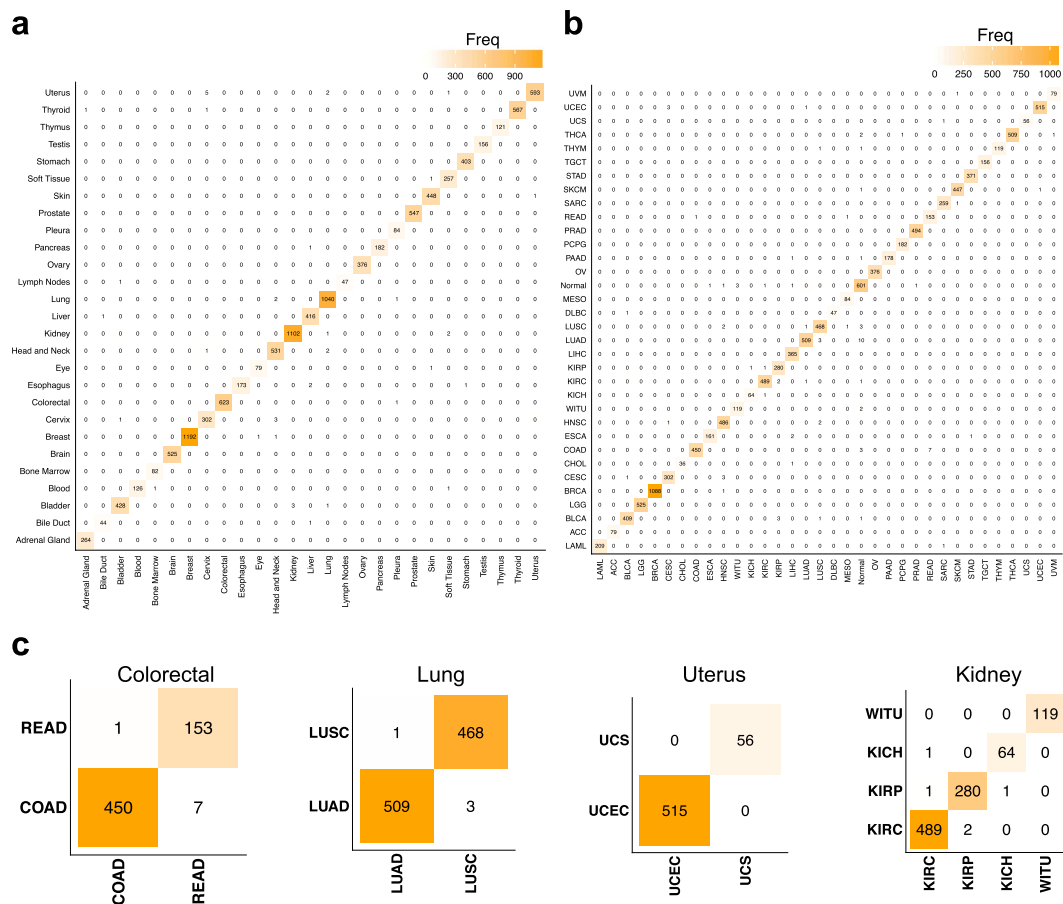
**Figure 3.** Discrimination of disease state of samples using the original mRNA expression profiles vs. Cell Identity Codes (CIC). **(a)** PCA plot of the original mRNA expression profiles of all samples. Each dot and its color show a sample and its disease state, respectively. **(b)** PCA plot of the 8-dimensional CIC space. **(c)** t-SNE plot of the original mRNA expression profiles. **(d)** t-SNE plot of the 8-dimensional CIC space.

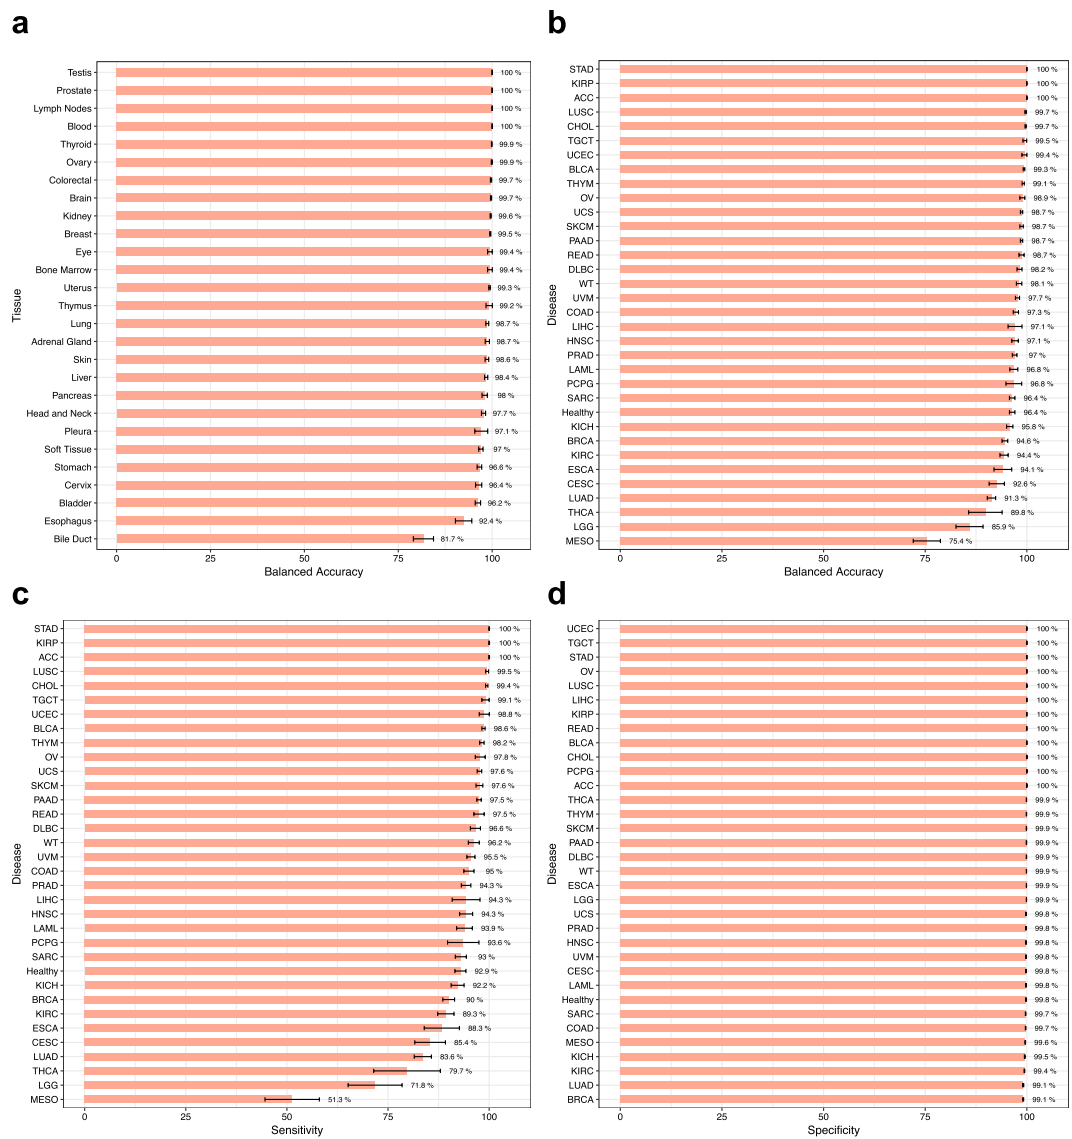
of the disease state of the samples than the original gene expression profiles. Doing the same analyses for studying tissue-type discrimination provided similar results (Fig. S5). Taken together, these results have two important conclusions: (I) There is no significant batch effect in the original mRNA expression profiles that can discriminate different studies. (II) The 8-dimensional CIC space is learning features from the transcription profiles that can very well discriminate different tissues and diseases.

**DNN accurately discriminates different cancer subtypes.** Different cancer types or subtypes might arise from the same tissue (e.g. lung squamous cell carcinoma vs. lung adenocarcinoma). Each cancer type/subtype might have a specific therapeutic strategy. Hence, correct discrimination of cancer types/subtypes that arise from the same tissue is a critically important and sometimes challenging task of cancer pathology. To scrutinize this, we first visualized the confusion matrix of the 5-fold cross-validation classification results for all tissues (Fig. 4a) and disease states (Fig. 4b). Extended results are provided in Supplementing Tables S3 and S4. It is also crucial to restate that all the datasets utilized for statistical analysis of Tables S3 and S4 have been resulted in combining all corresponding samples generated by each fold of cross-validation.

Then, we focused on all 4 tissue types that had more than one type of cancer in this dataset: Colorectal, Lung, Uterus and Kidney. As shown in Fig. 4c, the confusion matrices for these cancer types confirm there are only 17 misclassifications out of 3120 total samples, which provides us with 99.4% accuracy for cancer subtype classification. This accuracy ranges from 98.7% (Colorectal cancer subtypes) to 100% (Uterus). Discriminating cancer subtypes is clinically very important for deciding a correct therapeutic strategy, and is often challenging due to histological similarity of cancer subtypes and their heterogeneity. This finding paves the road of using high-throughput molecular data to address challenging pathology problems.

**Stable classification of different tissues and cancer types.** To ensure the results are not overfitted to a particular subset of the data, we performed a 5-fold cross-validation and measured different accuracy criteria for tissue and disease classification. The data was shuffled and each sample was randomly assigned to one of 5



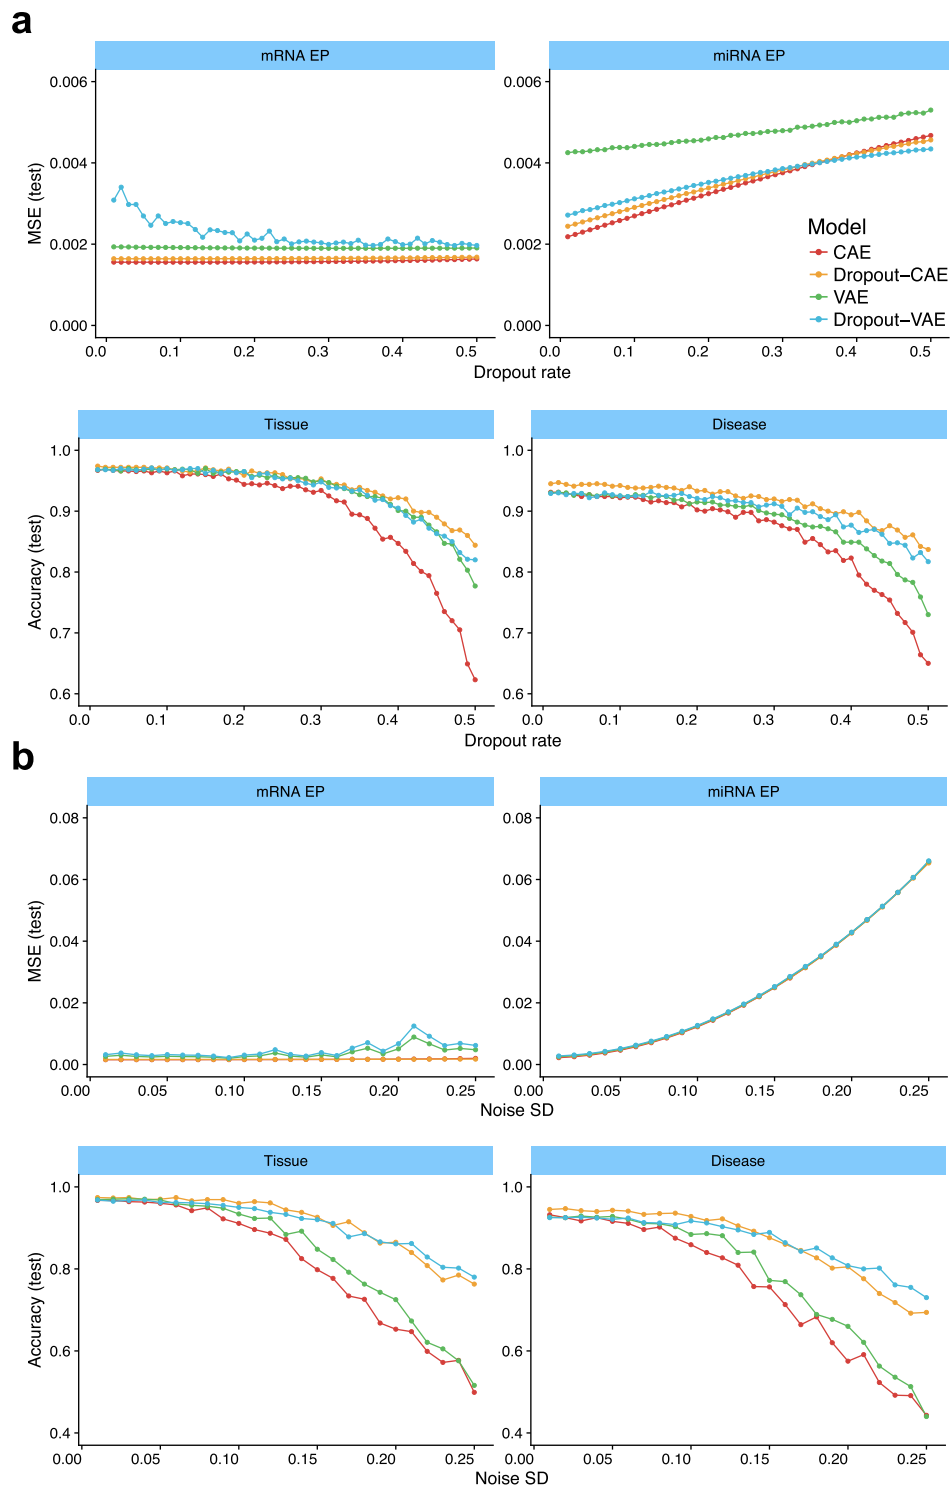


**Figure 5.** Performance of DNN in the classification of different tissues and diseases. **(a)** Balanced accuracy of tissue classification. **(b)** Balanced accuracy of disease classification. Healthy samples are depicted as “Healthy”. **(c,d)** Sensitivity and specificity of disease classification. All analyses are performed using 5-fold cross-validation. Error bars indicate standard error (SE).

50% of the genes has a negligible effect on reproducing mRNA expression profiles, with almost no MSE change for all architectures excepting Dropout-VAE. Interestingly, the increased dropout rate caused an improved MSE for Dropout-VAE. Also, increasing the dropout rate elevated the MSE of predicting miRNA EP for all DNN architectures. But even at 50% dropout, the MSE of all architectures was around 0.004, which is quite small.

The classification accuracy of all DNNs, particularly the Dropout-CAE and Dropout-VAE had small changes by increasing the missing values from 0 to 20%. Even at 30% dropout rate, Dropout-CAE was 95% accurate. Disease classification accuracy of both Dropout-CAE and Dropout-VAE had also small changes by a dropout rate up to 20%. These experiments showed that the DNN architectures that contained Dropout layer had the highest resistance to missing values.

We also measured the resistance of DNNs to a noisy input (Fig. 6b). A layer just after the input added a zero-mean Gaussian noise to the input GEPs. The magnitude of the noise was controlled by increasing its standard deviation (SD) from 0 to 0.25 with 0.01 steps. Both CAE and Dropout-CAE architectures were quite resistant to noise in reproducing the mRNA EP, and their MSEs were almost unchanged at SD = 0.25. All networks could predict miRNA EP with MSE  $\leq 0.01$  when the SD of Gaussian noise was at most 0.08. But their errors were increased by increasing the noise magnitude. The tissue and disease classification accuracy of Dropout-CAE were almost unchanged for a noise with SD  $\leq 0.05$ . Dropout-CAE and Dropout-VAE outperformed the other architectures in resisting against Gaussian noise.



**Figure 6.** The resistance of different DNN architectures against missing values and noise. **(a)** In each plot, the x-axis shows the fragment of randomly-selected input values that are set to zero (dropout), and y-axis shows either regression MSE or classification accuracy for the test dataset. **(b)** The x-axis shows the standard error (SD) of a zero-centered Gaussian noise which was added to the input values, and y-axis shows MSE or accuracy for the test dataset. Colors indicate different DNN architectures (see the figure legend).

Collectively, our results indicate the power of DNNs in obtaining biologically and clinically important information from transcriptome profiles. Our networks compress the transcriptome profile into a thumbnail CIC, and obtain tissue and disease type and miRNA expression profiles out of it. This process is greatly robust against noisy and missing data and outperforms baseline algorithms in accuracy. We suggest employing DNNs in inferring the outcome of molecular cancer pathology.

**Future works.** Due to the black-box nature of deep neural networks, it is difficult to extract easy to understand patterns from the model and discover causality relationship between the original dataset and outputs<sup>46,47</sup>. The authors, however, have been exerting efforts to utilize statistical methods to extract such information from the trained models. These methods are usually based on gradient variation and are categorized as sample-based and model-based approaches, including Shap<sup>48,49</sup>, DeepLift<sup>50</sup>, LIME<sup>51</sup>, and Interpret<sup>52</sup>. This can pave the way for finding novel biomarkers and regulatory interactions of different tissues and disease states.

## Data availability

All data can be obtained from Genomic Data Commons (<https://gdc.cancer.gov>). All data pre-processing and analysis source codes are available at <https://github.com/SharifBioinf/DeePathology>.

Received: 13 December 2018; Accepted: 23 October 2019;

Published online: 11 November 2019

## References

1. Raab, S. S. & Grzybicki, D. M. Quality in Cancer Diagnosis. *CA: A Cancer Journal for Clinicians* **60**, 139–165 (2010).
2. Bruner, J. M., Inouye, L., Fuller, G. N. & Langford, L. A. Diagnostic discrepancies and their clinical impact in a neuropathology referral practice. *Cancer* **79**, 796–803 (1997).
3. Staradub, V. L., Messenger, K. A., Hao, N., Wiley, E. L. & Morrow, M. Changes in breast cancer therapy because of pathology second opinions. *Annals of surgical oncology* **9**, 982–987 (2002).
4. Hamady, Z. Z. R., Mather, N., Lansdown, M. R., Davidson, L. & MacLennan, K. A. Surgical pathological second opinion in thyroid malignancy: impact on patients' management and prognosis. - PubMed - NCBI. *European Journal of Surgical Oncology (EJSO)* **31**, 74–77 (2005).
5. Elmore, J. G. *et al.* Pathologists' diagnosis of invasive melanoma and melanocytic proliferations: observer accuracy and reproducibility study. - PubMed - NCBI. *BMJ* **357**, j2813 (2017).
6. Maor, D., Ondhia, C., Yu, L. L. & Chan, J. J. Lichenoid keratosis is frequently misdiagnosed as basal cell carcinoma. *Clinical and Experimental Dermatology* **42**, 663–666 (2017).
7. Xie, Y. *et al.* Hepatoid adenocarcinoma of the stomach is a special and easily misdiagnosed or missed diagnosed subtype of gastric cancer with poor prognosis but curative for patients of pN0/1: the experience of a single center. *International Journal of Clinical and Experimental Medicine* **8**, 6762–6772 (2015).
8. Emori, M. *et al.* A typical presentation of primary pulmonary epithelioid sarcoma misdiagnosed as non-small cell lung cancer. *Pathology International* **67**, 222–224 (2017).
9. Lapuk, A. V. *et al.* From sequence to molecular pathology, and a mechanism driving the neuroendocrine phenotype in prostate cancer. *The Journal of pathology* **227**, 286–297 (2012).
10. Li, T., Zhang, C. & Ogihara, M. A comparative study of feature selection and multiclass classification methods for tissue classification based on gene expression. *Bioinformatics* **20**, 2429–2437 (2004).
11. Bhat, R. R., Viswanath, V. & Li, X. DeepCancer: Detecting Cancer through Gene Expressions via Deep Generative Learning. *arXiv.org* 1612.03211v2 (2016).
12. Khalili, M., Majid, H. A. & Khodakaram, S. Prediction of the thromboembolic syndrome: an application of artificial neural networks in gene expression data analysis. *Journal of Paramedical Sciences* **7**, 15–22 (2016).
13. Yang, L., Liu, Z., Yuan, X., Wei, J. & Zhang, J. Random Subspace Aggregation for Cancer Prediction with Gene Expression Profiles. *BioMed Research International* **2016**, 1–10 (2016).
14. Patil, S., Naik, G., Pai, R. & Gad, R. Stacked autoencoder for classification of glioma grade iii and grade iv. *Biomedical Signal Processing and Control* **46**, 67–75 (2018).
15. Lê Cao, K.-A., Bonnet, A. & Gadat, S. Multiclass classification and gene selection with a stochastic algorithm. *Computational Statistics & Data Analysis* **53**, 3601–3615 (2009).
16. Chai, H. & Domeniconi, C. An evaluation of gene selection methods for multi-class microarray data classification. In *Proceedings of the Second European Workshop on Data Mining and Text Mining for Bioinformatics (in conjunction with ECML/PKDD)* (2004).
17. Yang, S. & Naiman, D. Q. Multiclass cancer classification based on gene expression comparison. - PubMed - NCBI. *Statistical applications in genetics and molecular biology* **0**, 477–496 (2014).
18. Grossman, R. L. *et al.* Toward a Shared Vision for Cancer Genomic Data. *dx.doi.org* **375**, 1109–1112 (2016).
19. Ding, L. *et al.* Perspective on Oncogenic Processes at the End of the Beginning of Cancer Genomics. *Cell* **173**, 305–320.e10 (2018).
20. Sanchez-Vega, F. *et al.* Oncogenic Signaling Pathways in The Cancer Genome Atlas. *Cell* **173**, 321–337.e10 (2018).
21. Hoadley, K. A. *et al.* Cell-of-Origin Patterns Dominate the Molecular Classification of 10,000 Tumors from 33 Types of Cancer. *Cell* **173**, 291–304.e6 (2018).
22. Malta, T. M. *et al.* Machine Learning Identifies Stemness Features Associated with Oncogenic Dedifferentiation. *Cell* **173**, 338–354.e15 (2018).
23. Xiao, Y., Wu, J., Lin, Z. & Zhao, X. A semi-supervised deep learning method based on stacked sparse auto-encoder for cancer prediction using rna-seq data. *Computer methods and programs in biomedicine* **166**, 99–105 (2018).
24. Yuan, Y. *et al.* DeepGene: an advanced cancer type classifier based on deep learning and somatic point mutations. *BMC Bioinformatics* **17**, 476 (2016).
25. Li, Y. *et al.* A comprehensive genomic pan-cancer classification using The Cancer Genome Atlas gene expression data. *BMC Genomics* **18**, 508 (2017).
26. Colaprico, A. *et al.* TCGAAbiolinks: an R/Bioconductor package for integrative analysis of TCGA data. *Nucleic acids research* **44**, e71 (2016).
27. Charte, D., Charte, F., Garcia, S., del Jesus, M. I. A. J. & Herrera, F. A practical tutorial on autoencoders for nonlinear feature fusion: Taxonomy, models, software and guidelines. *Information Fusion* **44**, 78–96 (2018).
28. Rifai, S., Glorot, X., Bengio, Y. & Vincent, P. Adding noise to the input of a model trained with a regularized objective. *arXiv.org* 1–13 (2011).
29. Vincent, P., Larochelle, H., Bengio, Y. & Manzagol, P.-A. Extracting and composing robust features with denoising autoencoders. *Proceedings of the 25th international conference on Machine learning - ICML '08* 1096–1103 (2008).
30. Goodfellow, I., Bengio, Y. & Courville, A. Regularization for Deep Learning. *Deep Learning* 216–261 (2016).
31. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I. & Salakhutdinov, R. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research* **15**, 1929–1958 (2014).
32. Pedregosa, F. *et al.* Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* **12**, 2825–2830 (2012).
33. Keras deep-learning python package., <https://github.com/keras-team/keras>.
34. Abadi, M. I. N. *et al.* TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems. *arXiv.org* (2016).
35. The Theano Development Team *et al.* Theano: A Python framework for fast computation of mathematical expressions. *arXiv.org* 1–19 (2016).

36. Ioffe, S. & Szegedy, C. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift, <http://arxiv.org/abs/1502.03167>. 1502.03167, (2015).
37. Shahriari, B., Swersky, K., Wang, Z., Adams, R. P. & De Freitas, N. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE* **104**, 148–175 (2016).
38. Dewancker, L., Clark, S., Com, I. A. N. S., Com, M. S. & Com, S. S. Bayesian Optimization Primer 2–5 (2015).
39. Snoek, J., Larochelle, H. & Adams, R. Practical Bayesian Optimization of Machine Learning Algorithms. *Nips* 1–9 (2012).
40. Bergstra, J., Yamins, D. & Cox, D. D. Hyperopt: A python library for optimizing the hyperparameters of machine learning algorithms 13–20 (2013).
41. Nair, V. & Hinton, G. E. Rectified Linear Units Improve Restricted Boltzmann Machines. *Proceedings of the 27th International Conference on Machine Learning* 807–814 (2010).
42. Ramachandran, P., Zoph, B. & Le, Q. V. Searching for Activation Functions. *arXiv.org* 1–12 (2017).
43. Klambauer, G., Unterthiner, T., Mayr, A. & Hochreiter, S. Self-Normalizing Neural Networks. *arXiv.org* (2017).
44. Glorot, X., Bordes, A. & Bengio, Y. Deep sparse rectifier neural networks. *Journal of Machine Learning Research* **15**, 315–323 (2011).
45. Maaten, L. V. D. & Hinton, G. Visualizing Data using t-SNE. *Journal of Machine Learning Research* **9**, 2579–2605 (2008).
46. Ghorbani, A., Abid, A. & Zou, J. Y. Interpretation of neural networks is fragile. *CoRR* abs/1710.10547 (2018).
47. Zhang, Z. *et al.* Opening the black box of neural networks: methods for interpreting neural network models in clinical applications. *Annals of translational medicine* **6**(11), 216 (2018).
48. Lundberg, S. M. & Lee, S.-I. A unified approach to interpreting model predictions. In Guyon, I. *et al.* (eds) *Advances in Neural Information Processing Systems 30*, 4765–4774, <http://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions.pdf> (Curran Associates, Inc., 2017).
49. Lundberg, S. M. *et al.* Explainable machine-learning predictions for the prevention of hypoxaemia during surgery. *Nature Biomedical Engineering* **2**, 749 (2018).
50. DeepLIFT: Deep Learning Important Features. <https://github.com/kundajelab/deeplift>.
51. Ribeiro, M. T., 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, San Francisco, CA, USA, August 13–17, 2016*, 1135–1144 (2016).
52. InterpretML: python package for training interpretable models and explaining blackbox systems., <https://github.com/microsoft/interpret>.

## Acknowledgements

Authors would like to acknowledge creative comments and ideas by Dr. S. M. Ali Eslami and Dr. Mahmoud Ghandi. We are also grateful to Hamid Malki for his artworks in Figures 1 and S4.

## Author contributions

A.S.Z. and H.C. developed the idea of using DNNs for processing mRNA profiles. AS developed a pipeline for integration and preprocessing of the data. B.A. developed the deep learning framework and hyperparameter optimization, generated the results and created the figures. All authors contributed in discussing the results and writing the manuscript. H.C. provided the computational resources.

## Competing interests

The authors declare no competing interests.

## Additional information

**Supplementary information** is available for this paper at <https://doi.org/10.1038/s41598-019-52937-5>.

**Correspondence** and requests for materials should be addressed to A.S.-Z.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://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 license, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons license 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 license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2019