

# Adaptation of Autoencoder for Sparsity Reduction From Clinical Notes Representation Learning

Thanh-Dung Le, *Member, IEEE*, Rita Noumeir Ph.D., *Member, IEEE*,  
Jerome Rambaud M.D., Ph.D., Guillaume Sans M.D., and Philippe Jovet M.D., Ph.D.

**Abstract**—When dealing with clinical text classification on a small dataset recent studies have confirmed that a well-tuned multilayer perceptron outperforms other generative classifiers, including deep learning ones. To increase the performance of the neural network classifier, feature selection for the learning representation can effectively be used. However, most feature selection methods only estimate the degree of linear dependency between variables and select the best features based on univariate statistical tests. Furthermore, the sparsity of the feature space involved in the learning representation is ignored. *Goal*: Our aim is therefore to access an alternative approach to tackle the sparsity by compressing the clinical representation feature space, where limited French clinical notes can also be dealt with effectively. *Methods*: This study proposed an autoencoder learning algorithm to take advantage of sparsity reduction in clinical note representation. The motivation was to determine how to compress sparse, high-dimensional data by reducing the dimension of the clinical note representation feature space. The classification performance of the classifiers was then evaluated in the trained and compressed feature space. *Results*: The proposed approach provided overall performance gains of up to 3% for each evaluation. Finally, the classifier achieved a 92% accuracy, 91% recall, 91% precision, and 91% f1-score in detecting the patient's condition. Furthermore, the compression working mechanism and the autoencoder prediction process were demonstrated by applying the theoretic information bottleneck framework.

**Impact Statement**— An autoencoder learning algorithm effectively tackles the problem of sparsity in the representation feature space from a small clinical narrative dataset. Significantly, it can learn the best representation of the training data because of its lossless compression capacity as compared to other approaches. Consequently, its downstream classification ability can also be significantly improved, which cannot be done using deep learning models.

**Index Terms**—Clinical natural language processing, cardiac failure, autoencoder, sparsity.

## I. INTRODUCTION

CLINICAL decision support systems (CDSS) are continuously being developed and play a crucial role in promoting a personalized healthcare system, as more and more data are collected and stored continuously [1]. These data represent decisive points in advancing and enhancing the efficiency and effectiveness of CDSS operations. Predictive models have been developed based on the latter for preventive treatment and patient diagnosis, culminating in intelligent, precise, and timely healthcare improvement [2]. In one notable example, a recent study [3] analyzed the effect of CDSS on cardiovascular risk in 18,578 patients in 70 community health centers. In that case, CDSS significantly reduced the risk of cardiovascular disease among vulnerable high-risk patients.

Following the above successes, a CDSS was developed at CHU Sainte-Justine Research Center (CHUSJ). The system monitors the management of pediatric intensive care for all patients ranging in age from 0 to 18. Fig. 1 illustrates two fundamental processes in the CDSS workflow at CHUSJ,

which involve collecting and processing critical care data. First, clinical data are collected and stored in a clinical data warehouse. The data processing unit is then systematically aggregated and processed to convert raw data to a machine-readable form in the data processing unit. This process helps analyzing the unknown data interpretation and presentation. The CDSS can thus integrate the advanced analytic result of the data processing unit and learning algorithms; then clinicians can adequately use the CDSS to guide early intervention and prevention for healthcare management.

One of the goals of the CDSS system in CHUSJ is automatically screening the data from electronic medical records, chest X-rays and other data sources, which has the potential to increase the diagnosis rate and then improve the management of acute respiratory distress syndromes (ARDS) in real time. Usually, the diagnosis of ARDS was delayed or missed in two-thirds of patients, and the diagnosis was missed completely in 40% of patients [4]. To make the diagnosis of ARDS, three main conditions need to be detected: hypoxemia (low blood oxygenation), presence of infiltrates on chest X Ray and absence of cardiac failure [5]. Our research team has developed algorithms for hypoxemia [6], chest X-ray analysis [7], and identification of the absence of cardiac failure [8], [9]. Technically, we successfully performed extensive analyzes of machine learning algorithms (ML) aimed at detecting cardiac failure from clinical narratives using natural language processing based on such algorithms [8]. The study included the clinical notes of 1386 patients classified by two independent physicians using a standardized approach. It confirmed that the framework proposed herein yields an overall classification performance with 89% accuracy, 88% recall, and 89%

This work was supported in part by the Natural Sciences and Engineering Research Council (NSERC), in part by the Institut de Valorisation des données de l'Université de Montréal (IVADO), in part by the Fonds de la recherche en santé du Québec (FRQS), and in part by the Fonds de recherche du Québec-Nature et technologies (FRQNT).

Thanh-Dung Le is with the Biomedical Information Processing Lab, Ecole de Technologie Supérieure, University of Québec, Canada, and also with the Research Center at CHU Sainte-Justine, University of Montreal, Canada (Email: thanh-dung.le.1@ens.etsmtl.ca).

Rita Noumeir is with the Biomedical Information Processing Lab, Ecole de Technologie Supérieure, University of Québec, Canada.

Jerome Rambaud, Guillaume Sans, and Philippe Jovet are with the Research Center at CHU Sainte-Justine, University of Montreal, Canada.

precision by applying a multilayer perceptron neural network (MLP-NN) classifier in combination with a term frequency  $\times$  inverse document frequency (TF-IDF) learning representation.

These results were made possible by the contributions of the feature selection process, also known as SelectKBest. The advantage of the process was proven for supervised models as the classifier performance brought overall improvements of up to 3-4% over the case without the feature selection. It is obvious to understand because there are fewer misleading features; after selecting the best  $K$  features, the classifier accuracy effectively improves. Unfortunately, the SelectKBest feature selection continues to have certain limitations in the proposed framework. One reason is that the feature selection method is based on a statistical test that estimates the degree of linear dependency between random variables. Then, it removes irrelevant features and ignores the correlation between data elements. As a result, more samples are required for an accurate estimation and avoidance of overfitting, which is not possible in our case [10]. Furthermore, SelectKBest does not deal mainly with the sparsity of the feature space in the note representation matrix [11]. Consequently, the sparsity that characterizes the learning representation space is ignored.

In the health care field, the autoencoder algorithm (AE) has lived up to its promises and has shown its effectiveness when it comes to improving outcomes for efficient clinical decision making. First, AE can find informative transformed feature vectors through the compressed latent representation. For example, a study [12] demonstrates an efficient framework for automatically learning compact representations from heterogeneous raw data sources from patient health data. In addition, AE can improve the predictability of the six different learning models to detect Parkinson's classification [13]. Another study [14] shows that AE improved the performance of a novel outlier detection mechanism by retrofitting word vectors for the biomedical ontology matching task. Second, having rich and accurate clinical data is very challenging [15] because the acquisition and sharing of medical data faces a significant obstacle in the form of privacy issues and the sensitive nature of the data. Fortunately, AE can be applied for sparsity reduction in clinical representation feature space to allow to tackle problems related to limited data availability. It could effectively discover the low dimensional embeddings, and reveal the underlying effective manifold structure from a sparse high dimensional document-term matrix [16].

Therefore, the present study examines alternatives to feature selection and focuses mainly on compressing data without loss of information by employing an AE algorithm. First, we aim to achieve a better feature space without any sparsity. We are interested in compressing the sparse TF-IDF matrix and reducing its dimensions to improve the efficiency of the feature space representation. Notably, we incorporate a neural network to learn efficient codings of unlabeled data to address the issues caused by sparse vectors generated from the TF-IDF representation feature space for clinical notes. Then, the compressed vector space from the TF-IDF matrix is fed into the classifiers as a refined input. Finally, ML classifiers conduct the learning process to draw comparative results, which are then used to evaluate the classification performance.

Our study confirms that AE effectively compresses the vector space of the TF-IDF representation for clinical narratives into a lower dimension. The proposed approach can retain the critical feature by capturing the correlation between attributes during the training process, hence; the downstream classification task can generally be increased to 2-3% for each evaluation criterion. Furthermore, we also highlight the value of AE behaviors in a limited data set. We analyze the working mechanism of the AE, which explains how the AE works to compress data through the encoder and decoder. Based on the information-theoretic framework, the working mechanism of the AE is to optimize the information bottleneck during the compression and prediction process, respectively. As a result, the behavior of AE in limited data is exactly in harmony with such cases where there are much larger data availability.

Section II will discuss the materials and methods. The experimental results and discussion then will be discussed in section III, IV. Finally, section V provides concluding remarks.

## II. MATERIALS AND METHODS

### A. Data Sparsity Challenges

In numerical analysis, a sparse matrix or sparse array is a matrix in which most elements are zero [17]. The number of zero-valued elements divided by the total number of elements (e.g.,  $m \times n$  for a  $m \times n$  matrix) is called the matrix sparsity (equal to 1 minus the density of the matrix). Using these definitions, a matrix will be sparse when its sparsity is more significant than 0.5. In our case, after the research protocol was approved by the research ethics board from Research Center of the Sainte-Justine University Hospital, we have more than 580000 (unigrams) word count from 5444 single lines of notes with 1941 positive cases (36% of total) and 3503 negative cases. We applied the SelectKBest to select top best 'k=20000' of the vectorized features for the TF-IDF presentation learning feature space. Finally, we have a matrix of features of (5444  $\times$  20000). It is calculated by the Eq. 1, and the sparsity of the TF-IDF matrix is greater than 0.9.

It confirms that the representation matrix from the TF-IDF is sparse because every word is treated separately. Hence, the semantic relationship between separated entities is ignored, which would cause information loss. Although the combination of TF-IDF and MLP-NN consistently outperformed other combinations with overall performance and was the most stable under all circumstances [8], the sparsity remains. Therefore, the motivation is how we can compress the sparse, high-dimensional data by reducing the dimension from the TF-IDF feature space of clinical notes representation.

$$\text{sparsity} = 1 - \frac{\text{count\_nonzero}(\text{TF-IDF})}{\text{total\_elements\_of\_}(\text{TF-IDF})} \quad (1)$$

### B. Autoencoder Learning Algorithm

An AE was originated by [18] to solve a nonlinear dimensional reduction; later AE was famously promoted by training a MLP-NN with a small central layer to reconstruct high-dimensional input vectors [19], [20]. Technically, AE takes an input  $X \in \mathcal{R}^{N \times D}$  and, maps it to a latent representation

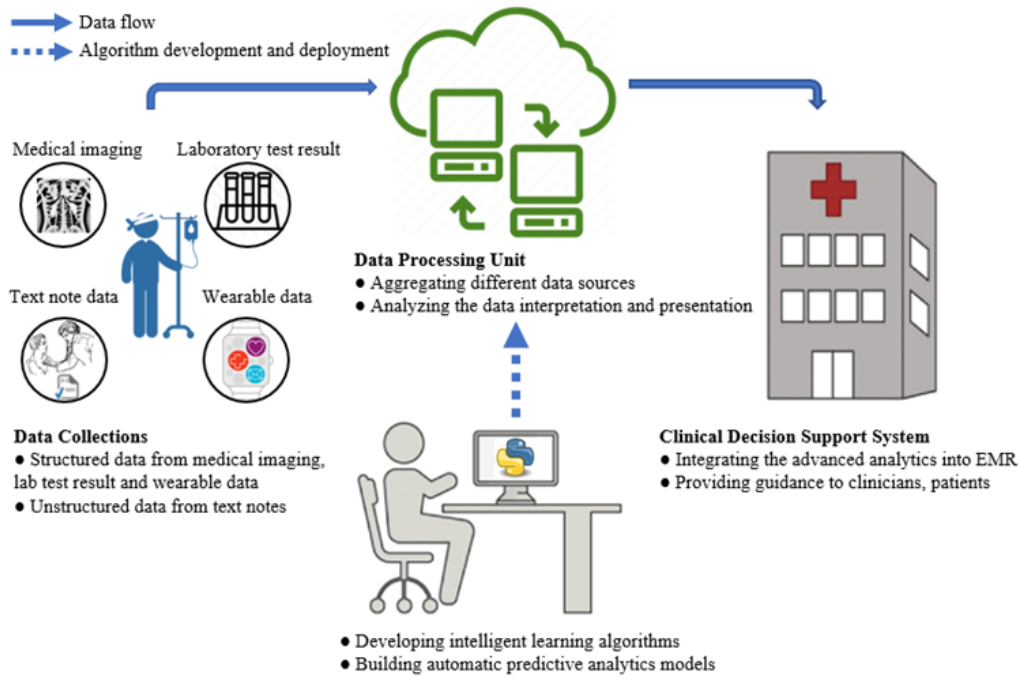


Fig. 1. Workflow demonstration of a clinical decision-support system at CHU Sainte-Justine hospital.

$Z \in \mathcal{R}^{N \times M}$  via a nonlinear mapping. Let us call  $x \in X$ , and  $z \in Z$ , we will have:

$$z = g(Wx + b) \quad (2)$$

$W$  is a weight matrix during training,  $b$  is a bias vector, and  $g(\cdot)$  stands for a nonlinear function, such as the logistic sigmoid function or a hyperbolic tangent function. The encoded feature representation  $x$  is then used to reconstruct the input  $x$  by reverse mapping, leading to the reconstructed input  $x'$ :

$$x' = f(W'z + b') \quad (3)$$

where  $W'$  is usually limited to the form of  $W' = W^T$ , i.e. the same weight is used to encode the input and decode the latent representation.  $f(\cdot)$  is also a non-linear function. The AE tries to learn a function  $f_{W',b'}(x) \approx x'$ . In other words, it is trying to learn an approximation of the identity function for the output  $x'$  that is similar to  $x$ . Still, by placing constraints on the network, such as limiting the number of hidden units, we can discover interesting data structures. Then, the reconstruction error is defined as the Euclidean distance between  $x$  and  $x'$  that is constrained to approximate the input data  $x$  (that is, minimizing  $\|x - x'\|^2$ ).

$$\begin{aligned} \mathcal{L}(x, x') &= \|x - x'\|^2 \\ &= \|x - f(W'(g(Wx + b)) + b')\|^2 \end{aligned} \quad (4)$$

For the reconstruction evaluation between the original data  $x$ , and the reconstructed output  $x'$ , we will apply the statistical measure  $R_i^2$  for the  $i^{th}$  variable of  $x_i$ . It can be computed as:

$$R_i^2 = 1 - \frac{\sum_{j=1}^m (x_{j,i} - x'_{j,i})^2}{\sum_{j=1}^m x_{j,i}^2} \quad (5)$$

Since  $R^2 = 1$  will be a perfect reconstruction; we will evaluate the reconstruction by how much the value of  $R^2$  is close to 1.

Ideally, an effective AE can be designed and trained based on the minimization of reconstruction error from Eq. 4 and maximization of the reconstructed effectiveness from Eq. 5; however, it is substantially based on its width (number of neuron units or latent representation dimension  $M$ ) and its depth (number of hidden layers). First, conventional AE relies on the dimension of the latent representation  $z$  being smaller than that of the input  $x$  ( $M < D$ ), which means that it tends to learn a low-dimensional compressed representation. The study [21] presents methods to learn the decoder function  $f(\cdot)$  as a learnable function through the reconstruction error in Eq. 4 in several representation learning approaches. It is concluded that the compression depends on dimension  $M$  but less on dimension  $D$ . Second, it has been shown that training a neural network-based by increasing the number of hidden layers (in combination with an increase in the number of neuron units per layer) achieves less consistent results [22]. Therefore, in our case, we use a small and simple AE. We employ an AE with three layers (one input layer, one hidden layer, and one output layer). Mainly, to reduce the parameters from the latent space of the AE, we apply the regularization technique from study [23] to remove redundant parameters.

After training, we use the weight matrix from the hidden layer as a pre-trained tool. A classifier subsequently use this pre-train latent space representation to perform the binary classification, as shown in Fig. 2. For the classifiers, it is essential to have consistency in evaluating the proposed approach's performance. Then, we employ six different ML classifiers, including Random Forest (RF), Multinomial Naive Bayes (MultinomialNB), Logistic Regression (LR), Support Vector Machine (SVC), Gaussian Naive Bayes (GaussianNB)

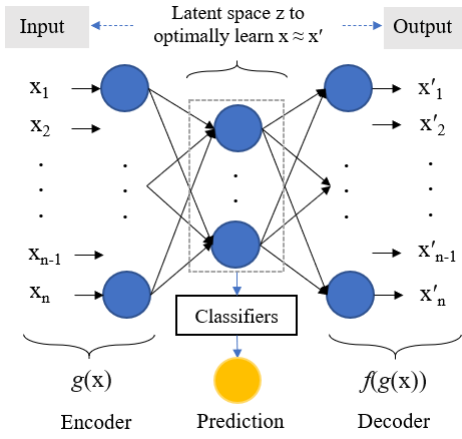


Fig. 2. Schematic structure of an AE-based for compression and prediction.

and Multilayer Perceptron Neural Network (MLP-NN).

Furthermore, to understand the dynamics of learning and the behavior of AE; especially, in our case with limited data, we also analyze the behavior of AE during the training process from the encoder and decoder. Technically, we capture to understand how the AE can retain the information during the compression process. To do that, we apply the information-theoretic quantities and their estimators based on information-theoretic learning, which compute and optimize information-theoretic descriptors named mutual information. The information-theoretic framework [24], [25], [26] has been utilized for a detailed theoretical explanation of an AE. These studies rely on the “information bottleneck” [27], [28] to understand and estimate how the AE works by quantifying its information plane coordinates. The information bottleneck can be used as an optimal bound that maximally compresses the input  $x$ , for a given mutual information on the desired output  $x'$ . There are comprehensive overviews of recent studies [29], [30], [31]. Technically, we first bin the output activation as stated in [28], and we treat each hidden layer  $i$  ( $1 \leq i \leq K$ ) as a single variable  $T_i$ . Then we will be able to estimate the mutual information between all the hidden layers and the input/output layers by estimating the joint distribution  $P(X, T_i)$  and  $P(T_i, X')$ , and use them to calculate the mutual information of the encoder (between the input  $X$  and the hidden layer  $T_i$ ), and the mutual information of the decoder (between the hidden layer  $T_i$  and the desired output  $X'$ ) using the following equations Eq. 6, 7. Finally, we can learn the good representation  $T(X)$ , which is characterized by its encoder and decoder distribution  $P(T|X)$ , and  $P(X'|T)$ , respectively, to effectively map the input patterns  $X$  to a good prediction of the desired output  $X'$ .

$$I(X; T_i) = \sum_{x \in X, t \in T_i} P(x, t) \log \left( \frac{P(x, t)}{P(x)P(t)} \right) \quad (6)$$

$$I(T_i; X') = \sum_{t \in T_i, x' \in X'} P(t, x') \log \left( \frac{P(t, x')}{P(t)P(x')} \right). \quad (7)$$

TABLE I  
HYPERPARAMETERS SUMMARY FOR AE TRAINING

Hyperparameter	Ranges
Hidden layers	1-3
Neurons	100-500
Activation	Sigmoid
Kernel initializer	GlorotNormal
Optimizers	SGD, ADAM
Learning rate	0.001 - 0.01
$\beta_1$	0.9
$\beta_2$	0.999
$\epsilon$	$e^{-8} - e^{-7}$

### III. RESULTS

To assess the performance of our method, metrics including accuracy, precision, recall (or sensitivity), and F1 score were used [32]. These metrics are defined as follows.

$$\text{Accuracy (acc)} = \frac{TP + TN}{TP + TN + FP + FN}$$

$$\text{Precision (pre)} = \frac{TP}{TP + FP}$$

$$\text{Recall/Sensitivity (rec)} = \frac{TP}{TP + FN}$$

$$\text{F1-Score (f1)} = \frac{2 * \text{Precision} * \text{Recall}}{\text{Precision} + \text{Recall}}$$

where TN and TP stand for true negative and true positive, respectively, and are the number of negative and positive patients correctly classified. FP and FN represent false positives and false negatives, respectively, and represent the number of positive and negative patients incorrectly predicted.

For implementation, we used the same hyperparameters from our previous study [8] for all classifiers so that we can have a consistent evaluation for the performance. The data was also divided into 60% training, 20% validation, and 20% testing. The implementation was done using Python Scikit learn [33] and Keras [34]. We performed a grid search for up to three hidden layers and 500 neurons per layer, and other hyperparameters are summarized in Table I for AE training. For the optimizers, we used the Stochastic Gradient Descent (SGD) and Adaptive Moment Estimation (ADAM) with small scalar  $\epsilon$ , and the forgetting factors for gradients and second moments of gradients,  $\beta_1$  and  $\beta_2$ . Then, a combination with the highest estimations was considered the best performance.

### IV. DISCUSSION

To deal with the sparsity, many researchers simply focus on dimension reduction. There are two most popular techniques, namely Linear Discriminant Analysis (LDA) and Principal Component Analysis (PCA), for their simplicity among other dimension reduction techniques [35], even with a large dataset [36]. Especially when the training data set is small, the PCA-supervised discriminative approach can outperform; it is also less sensitive to the variability of the training sets [37]. The study [38] shows that PCA can increase the performance of different ML classifiers for the prediction of cardiac failure.

It can be said that the classifiers performed better after applying LDA to the linear data set. So, in the case of linear data, LDA can reduce the dimensionality and be used

in different classification tasks [39]. However, the TF-IDF enhanced with the LDA approach did not allow the classifier to score high accuracy compared to the other two methods when smaller datasets were fed [40]. One of the reasons was explained in [36]; the results showed that ML algorithms with PCA produce better results when the dimensionality of the data sets is high. When the dimensionality of datasets is low, the ML algorithms without dimensionality reduction yield better results. Another possible way is using an unsupervised generative Latent Dirichlet allocation to estimate the topic distribution (topics) by using observed variables (words). Latent Dirichlet allocation shows the effectiveness of overcoming the sparsity from the feature space matrix of TF-IDF [41]. It can also help to make texts more semantically focused and reduce sparseness [42]. However, its selection of characteristics does not improve performance with small data [43].

We explored the possibility of PCA for sparsity reduction because of their advantages mentioned above. The training was tuned and performed, and the best performance was achieved with a decrease to 2 principal dimensions. The completed test has an accuracy of 88%; this is less than the performance of SelectKBest with 89%. Furthermore, following the recommendation of [44], we also tested a statistical method, Neighborhood Component Analysis (NCA) [45], to reduce the dimensions of the data set. NCA has shown that it works well on a small dataset for the medical domain. However, the result is slightly better than PCA; NCA only achieves an accuracy of 89% (the same as SelectKBest). Consequently, neither PCA nor NCA can improve the classification accuracy. It confirms the limitation of these approaches by linearly approximating a feature subspace to maximize class separability.

Furthermore, non-linear activation function AE (AE) shows its best performance on compression the sparse TF-IDF representation space. We compare the effectiveness of reconstruction based on the reconstruction evaluation from Eq. 5 between PCA, linear activation function AE (LAE), AE, and stacked AE (SAE) [46]. The results confirm that the PCA and LAE have the same performance, achieving about 80% of the reconstruction. When the activation of AE is linear, then PCA and LAE are identical. There is no improvement if the SAE is used to extract the features in cases of limited data. Besides, the effectiveness of non-linear activation in AE is proved, when it can maximally reconstruct up to 86% compared to the original spare data. It is one of the advantages of nonlinear transformation from AE, which is trained by a neural network, superior to linear transformation from other approaches.

Overall, the downstream classification performances are effectively improved by feeding the compressed feature space output from the AE to ML classifiers. Fig. 3 shows the loss during the training and validation process by optimizing the loss function from Eq. 4 for training the AE; the loss converges perfectly. After successfully training the AE, we have a pre-trained compressed, low-dimension feature space. We continue to the machine learning classifiers to perform the classification and evaluate the performance. Instead of performing on MLP-NN, LR, and GaussianNB, we also tested with other classifiers such as Random Forest (RF), Multinomial Naive Bayes, and Support Vector Machine. Fig. 5 shows the comparison, using a




Fig. 3. Loss for training and validation for the AE algorithm.




Fig. 4. Confusion matrix of the MLP-NN classifier, showing the classification of positive (1) and negative (0) between predicted and actual labels.

box plot, of the 5-fold cross-validation. Again, MLP-NN gives the best performances; LR follows right after; GaussianNB is comparatively similar to LR. And all other classifiers are less effective. The best performance from MLP-NN is achieved at 92%, 91%, 91% and 91%, respectively, for accuracy, precision, recall, and f1 score. And, the detailed confusion matrix showing the classification of positive cases (1) and negative cases (0) between predicted and actual labels is shown in Fig. 4. The experimental results are improved to 2-3 % for each evaluation criterion from [8], which had a general classification performance in a sparse TF-IDF feature space at 89% accuracy, 88% recall, and 89% precision. It confirms that the AE method can deal with the sparsity by compressing the TF-IDF feature space. Consequently, it finally improves the downstream task performance of the MLP-NN classifier, and is also more robust than other methods. Recent work [47] also confirmed a similar effect, but it was applied to a different dataset type and larger data availability. These results confirm the effectiveness of compressing the feature representation learning space into a low-dimensional representation using the AE algorithm. Notably, the robust transformation can outplay the deep learning models with limited data resources.

Furthermore, an important aspect of performance analysis is that the proposed approach still shows its advantageous capacity to increase data availability. We investigated the effectiveness of AE for compressing feature space, and studied how algorithm performance varies with the increasing of training examples from the compressed feature space. We assessed the performance of two classifiers GaussianNB and MLP-




Fig. 5. A comparison evaluation of the box plot 5-fold cross-validation results for classifiers performance.

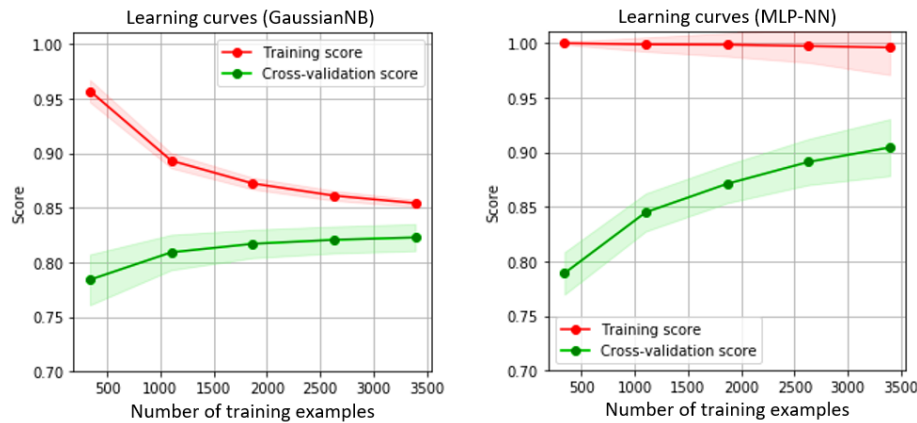


Fig. 6. Performance of classifiers in case of increasing the training size: GaussianNB (left) and MLP-NN (right).

NN to evaluate the effectiveness. When it possibly increases data availability in the future, whether the classifier improves performance or not. In this case, study [48] confirms that when the number of training examples increases, the generative model based on Naive Bayes would expect to perform better. However, our results are in contrast to that confirmation. Fig. 6 shows the GaussianNB training and validation score when increasing the number of training examples. Technically, the GaussianNB reaches a plateau of performance after around 2000 training examples with the same dataset size, and the cross-validation score could not improve. We should expect that this is one of the limitations of GaussianNB, namely the linear discrimination characteristic for a real-world dataset, discussed in [49]. In contrast, the MLP-NN shows improvement with the increasing size of the dataset. Its cross-validation score gradually increases and shows no signs of reaching a maximum point. In short, we can confirm that when data is possibly increased, our approach still perfectly works and continually improves its performance for the classification.

Moreover, the behavior of AE in limited data is in harmony with more significant data cases based on the information-theoretic framework. We analyzed the behavior of AE based on an information-theoretic framework, as mentioned in Eq.

6, and 7. We want to understand how the AE behaves during the compression process by analyzing the mutual information of each hidden layer from the encoder and decoder. Generally, this type of analysis has been performed for a larger data set and has mainly focused on other data sources compared to our case; such as computer vision [50], medical imaging [51], and genetics [52]. We performed the analysis for two AE models with respecting to varying number of hidden layers (three hidden layers, and five hidden layers). As shown in Fig. 7, there are two phases of the information plane in each hidden layer of the three-layer and five-layer cases. It is noted that from left to right it illustrates the behavior for each hidden layer. And in each hidden layer, from top to bottom, it captures the mutual information for each training epoch. Finally, all trajectories seem to follow a similar path during the learning process, then eventually converge and get closer to the optimal points in the theoretic information bottleneck bound.

Specifically, it can be divided into two phases for the working mechanism of AE in Fig. 7. The first phase is called the drift phase, where the AE attempts to learn the latent representation  $T(X)$  with a smaller dimension compared to the original data  $X$ . During the compression, there will be information lost, that is why we can see the trend of decreasing

the mutual information of encoder  $I(X;T)$ . At the end of this step, we will have a compressed latent representation  $T(X)$ , and an optimal mutual information  $I(X;T)$ . Then, the second phase is named the diffusion phase. Within this step, the AE tries to find the reconstructed data  $X'$ , which are optimally close to the original data  $X$ . The AE maps the latent representation  $T(X)$  to the reconstructed data  $X'$  by maximizing the mutual information of the decoder  $I(T;X')$ . By doing that, we can see the increasing trend of  $I(T;X')$ ; until  $I(T;X')$  reaches its optimal bound for each layer. One more interesting point is that the optimal mutual information will get smaller, when AE has more hidden layers. In case of three hidden layers, the optimal mutual information of the encoder  $I(X,T)$  is larger by 6.0, but is maximum at 5.5 for five hidden layers. It is the same for the optimal mutual information of the decoder  $I(T,X')$  at nearly 7.0 and 6.5 for three hidden layers and five hidden layers, respectively. These results illustrate the mechanism of an AE is to optimize the information bottleneck trade-off  $T(X)$  during compression and prediction, respectively, for each layer. Remarkably, it is trained on a small and sparse dataset; still, it proves its effectiveness by compressing and maximizing the mutual information from the TF-IDF feature space.

### V. CONCLUSION

First, this study has shown that the participation of an AE in training can effectively compress the feature space of TF-IDF. The AE with a nonlinear activation function can achieve the reconstruction capacity at 86% compared to the original data. It outperforms other approaches such as PCA, NCA, LAE (AE with linear activation function), and stacked AE. It concludes that AE can learn the best representation of the training data due to its lossless compression capacity.

Additionally, the AE also works well with a small clinical dataset; especially, in harmony with the information-theoretic mechanism of an AE for a larger dataset, and from different data sources. It has two phases of learning: the drift phase of the encoder by trying to compress the data. The second phase is related to the diffusion phase by maximizing the mutual information process in the decoder. Consequently, it shows the effectiveness of losses information in compressing the data. By doing so, we also capture the interpretability, comprehensibility, and transparency of the proposed model for decision making in our CDSS system recommended by [53].

The second step involves the use of an MLP-NN to predict the health status based on the compressed feature space. It has been shown that the sparsity reduction for the feature space strongly affects the classifier performance in downstream task. AE learning algorithm effectively leverages the sparsity reduction. As a result, it helps the MLP-NN classifier achieve 92% accuracy, 91% recall, 91% precision, and 91% f1-score. This efficient ensemble model can outperform all alternative approaches: GaussianNB, LR, RF, MultimonialNB, and SVC.

The proposed approach is still proving successful in cases where data availability is increased. The MLP-NN is effective in achieve a better performance after the GaussianNB reaches its maximum capacity. In future work, we will choose the




Fig. 7. The evolution of the layers with epochs in the information plane for three hidden layers (top) and five hidden layers (bottom).

optimal parameters and validate our method on more datasets. We will explore the weak supervision approach that recently proved its effectiveness in 4,000 cardiac magnetic resonance sequences with imperfect labels [54]; because it can maximize the use of unlabeled data at scale, which is costly to annotate.

### ACKNOWLEDGMENT

Clinical data were provided by the Research Center of CHU Sainte-Justine hospital, University of Montreal. The authors thank Dr. Sally Al Omar, Dr. Michael Sauthier, Dr. Rambaud Jerome and Dr. Sans Guillaume for their data support of this research. This work was supported by a scholarship from the Fonds de recherche du Quebec-Nature et technologies (FRQNT) to Thanh-Dung Le, and the grants from the Natural Sciences and Engineering Research Council (NSERC), the Institut de valorisation des donnees (IVADO), and the Fonds de la recherche en sante du Quebec (FRQS).

REFERENCES

[1] M. A. Musen and et. al., "Clinical decision-support systems," in *Biomedical informatics*. Springer, 2021, pp. 795–840.

[2] R. T. Sutton and et. al., "An overview of clinical decision support systems: benefits, risks, and strategies for success," *NPJ digital medicine*, vol. 3, no. 1, pp. 1–10, 2020.

[3] R. Gold and et. al., "Effect of clinical decision support at community health centers on the risk of cardiovascular disease: A cluster randomized clinical trial," *JAMA Network Open*, vol. 5, 2022.

[4] G. Bellani and et. al., "Epidemiology, patterns of care, and mortality for patients with acute respiratory distress syndrome in intensive care units in 50 countries," *Jama*, vol. 315, no. 8, pp. 788–800, 2016.

[5] P. A. L. I. C. C. Group et al., "Pediatric acute respiratory distress syndrome: consensus recommendations from the pediatric acute lung injury consensus conference," *Pediatric critical care medicine: a journal of the Society of Critical Care Medicine and the World Federation of Pediatric Intensive and Critical Care Societies*, p. 428, 2015.

[6] M. Sauthier and et. al., "Estimated pao<sub>2</sub>: A continuous and noninvasive method to estimate pao<sub>2</sub> and oxygenation index," *Critical care explorations*, vol. 3, no. 10, 2021.

[7] N. Zaglam and et. al., "Computer-aided diagnosis system for the acute respiratory distress syndrome from chest radiographs," *Computers in biology and medicine*, vol. 52, pp. 41–48, 2014.

[8] T. D. Le and et. al., "Detecting of a patient's condition from clinical narratives using natural language representation," *IEEE Open Journal of Engineering in Medicine and Biology*, 2022.

[9] T.-D. Le and et. al., "Machine learning based on natural language processing to detect cardiac failure in clinical narratives," in *36e Congres de la recherche au CHU Sainte-Justine*. CHU Sainte-Justine, 2021.

[10] D. Jain and et. al., "Feature selection and classification systems for chronic disease prediction: A review," *Egyptian Informatics Journal*, vol. 19, no. 3, pp. 179–189, 2018.

[11] G. Forman and et. al., "An extensive empirical study of feature selection metrics for text classification," *J. Mach. Learn. Res.*, vol. 3, 2003.

[12] C. Zhou and et. al., "Optimizing autoencoders for learning deep representations from health data," *IEEE J. Biomed. Health Inform.*, 2018.

[13] Y. Xiong and et. al., "Deep feature extraction from the vocal vectors using sparse autoencoders for parkinson's classification," *IEEE Access*, vol. 8, pp. 27 821–27 830, 2020.

[14] P. Kolyvakis and et. al., "Biomedical ontology alignment: an approach based on representation learning," *J. Biomed. Semantics*, vol. 9, 2018.

[15] J. C. Quiroz and et. al., "Challenges of developing a digital scribe to reduce clinical documentation burden," *NPJ Digital Medicine*, 2019.

[16] M. L-Abadi and et. al., "Denoising autoencoder as an effective dimensionality reduction and clustering of text data," in *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, 2017.

[17] N. Hurley and et. al., "Comparing measures of sparsity," *IEEE Transactions on Information Theory*, vol. 55, no. 10, pp. 4723–4741, 2009.

[18] M. A. Kramer, "Nonlinear principal component analysis using autoassociative neural networks," *AIChE journal*, vol. 37, pp. 233–243, 1991.

[19] G. E. Hinton and et. al., "Reducing the dimensionality of data with neural networks," *science*, vol. 313, no. 5786, pp. 504–507, 2006.

[20] Y. Wang and et. al., "Auto-encoder based dimensionality reduction," *Neurocomputing*, vol. 184, pp. 232–242, 2016.

[21] S. Garg and et. al., "Functional regularization for representation learning: A unified theoretical perspective," *Advances in Neural Information Processing Systems*, vol. 33, 2020.

[22] C. Steinmeyer and et. al., "Sampling methods and feature selection for mortality prediction with neural networks," *Journal of Biomedical Informatics*, vol. 111, p. 103580, 2020.

[23] Y. Shi and et. al., "Learning robust auto-encoders with regularizer for linearity and sparsity," *IEEE Access*, vol. 7, pp. 17 195–17 206, 2019.

[24] S. Yu and et. al., "Understanding autoencoders with information theoretic concepts," *Neural Networks*, vol. 117, pp. 104–123, 2019.

[25] N. I. Tapia and et. al., "On the information plane of autoencoders," in *IEEE International Joint Conference on Neural Networks*, 2020.

[26] S. Lee and et. al., "Information flows of diverse autoencoders," *Entropy*, vol. 23, no. 7, p. 862, 2021.

[27] N. Tishby and et. al., "The information bottleneck method," *arXiv preprint physics/0004057*, 2000.

[28] R. Shwartz-Ziv and et. al., "Opening the black box of deep neural networks via information," *arXiv preprint arXiv:1703.00810*, 2017.

[29] B. C. Geiger, "On information plane analyses of neural network classifiers—a review," *IEEE Trans. Neural Netw. Learn. Syst.*, 2021.

[30] B. C. Geiger and et. al., "Information bottleneck: Theory and applications in deep learning," 2020.

[31] M. A. Alomrani, "A critical review of information bottleneck theory and its applications to deep learning," *arXiv:2105.04405*, 2021.

[32] C. Goutte and et. al., "A probabilistic interpretation of precision, recall and f-score, with implication for evaluation," in *European conference on information retrieval*. Springer, 2005, pp. 345–359.

[33] F. Pedregosa and et. al., "Scikit-learn: Machine learning in python," *The Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

[34] F. Chollet and et. al., "keras," 2015.

[35] F. Anowar and et. al., "Conceptual and empirical comparison of dimensionality reduction algorithms (pca, kpca, lda, mds, svd, lle, isomap, le, ica, t-sne)," *Computer Science Review*, vol. 40, p. 100378, 2021.

[36] G. T. Reddy and et. al., "Analysis of dimensionality reduction techniques on big data," *IEEE Access*, vol. 8, pp. 54 776–54 788, 2020.

[37] A. M. Martinez and et. al., "Pca versus lda," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 2, pp. 228–233, 2001.

[38] A. K. Gárate-Escamila and et. al., "Classification models for heart disease prediction using feature selection and pca," *Informatics in Medicine Unlocked*, vol. 19, p. 100330, 2020.

[39] J. Ghosh and et. al., "Improving classification model's performance using linear discriminant analysis on linear data," in *IEEE International Conference on Computing, Communication and Networking Technologies*, 2019.

[40] R. Dzisevič and et. al., "Text classification using different feature extraction approaches," in *IEEE Open Conference of Electrical, Electronic and Information Sciences*, 2019.

[41] S.-W. Kim and et. al., "Research paper classification systems based on tf-idf and lda schemes," *Human-centric Computing and Information Sciences*, vol. 9, no. 1, p. 30, 2019.

[42] Q. Chen and et. al., "Short text classification based on lda topic model," in *IEEE International Conference on Audio, Language and Image Processing*, 2016.

[43] S. J. Fodeh and et. al., "Classification of patients with coronary microvascular dysfunction," *IEEE/ACM Trans. Comput. Biol. Bioinform.*, vol. 17, 2019.

[44] S. Laghmati and et. al., "Classification of patients with breast cancer using neighbourhood component analysis and supervised machine learning techniques," in *IEEE International Conference on Advanced Communication Technologies and Networking*, 2020.

[45] J. Goldberger and e. a. Hinton, "Neighbourhood components analysis," in *NeurIPS*, 2005, pp. 513–520.

[46] J. Gehring and et. al., "Extracting deep bottleneck features using stacked auto-encoders," in *IEEE international conference on acoustics, speech and signal processing*, 2013.

[47] I. D. Mienye and et. al., "Improved sparse autoencoder based artificial neural network approach for prediction of heart disease," *Informatics in Medicine Unlocked*, vol. 18, p. 100307, 2020.

[48] A. Jordan and et. al., "On discriminative vs. generative classifiers: A comparison of logistic regression and naive bayes," *Advances in neural information processing systems*, vol. 14, no. 2002, p. 841, 2002.

[49] J.-H. Xue and et. al., "Comment on discriminative vs. generative classifiers: A comparison of logistic regression and naive bayes," *Neural processing letters*, vol. 28, no. 3, pp. 169–187, 2008.

[50] P. Viola and et. al., "Alignment by maximization of mutual information," *International journal of computer vision*, vol. 24, pp. 137–154, 1997.

[51] J. P. Pluim and et. al., "Mutual-information-based registration of medical images: a survey," *IEEE Trans. Med. Imaging*, vol. 22, no. 8, 2003.

[52] C. Olsen and et. al., "On the impact of entropy estimation on transcriptional regulatory network inference based on mutual information," *Journal on Bioinformatics and Systems Biology*, vol. 2009, 2008.

[53] C. Rudin, "Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead," *Nature Machine Intelligence*, vol. 1, no. 5, pp. 206–215, 2019.

[54] J. A. Fries and et. al., "Weakly supervised classification of aortic valve malformations using unlabeled cardiac mri sequences," *Nature Communications*, vol. 10, no. 1, pp. 1–10, 2019.