



Slicer: Feature Learning for Class Separability with Least-Squares Support Vector Machine Loss and COVID-19 Chest X-Ray Case Study

David Charte¹(✉) , Iván Sevillano-García¹ , María Jesús Lucena-González³, José Luis Martín-Rodríguez³, Francisco Charte² , and Francisco Herrera¹ 

¹ Department of Computer Science and Artificial Intelligence, Andalusian Research Institute in Data Science and Computational Intelligence (DaSCI), University of Granada, Granada, Spain

{fdavidcl,isevillano}@ugr.es,herrera@decsai.ugr.es

² Department of Computer Science, Andalusian Research Institute in Data Science and Computational Intelligence (DaSCI), University of Jaén, Jaén, Spain
fcharte@ujaen.es

³ Hospital Universitario Clínico San Cecilio de Granada, Granada, Spain

Abstract. Datasets from real-world applications usually deal with many variables and present difficulties when modeling them with traditional classifiers. There is a variety of feature selection and extraction tools that may help with the dimensionality problem, but most of them do not focus on the complexity of the classes. In this paper, a new autoencoder-based model for addressing class complexity in data is introduced, aiming to extract features that present classes in a more separable fashion, thus simplifying the classification task. This is possible thanks to a combination of the standard reconstruction error with a least-squares support vector machine loss function. This model is then applied to a practical use case: classification of chest X-rays according to the presence of COVID-19, showing that learning features that increase linear class separability can boost classification performance. For this purpose, a specific convolutional autoencoder architecture has been designed and trained using the recently published COVIDGR dataset. The proposed model is evaluated by means of several traditional classifiers and metrics, in order to establish the improvements caused by the extracted features. The advantages of using a feature learner and traditional classifiers are also discussed.

Keywords: Feature learning · Class separability · Autoencoders

1 Introduction

Data classification [3] is one of the most studied problems in machine learning, and is applicable in many real-world contexts such as medicine, banking,

© Springer Nature Switzerland AG 2021

H. Sanjurjo González et al. (Eds.): HAIS 2021, LNAI 12886, pp. 305–315, 2021.

https://doi.org/10.1007/978-3-030-86271-8_26

robotics, natural sciences and other industries. Class complexity [5, 17] is a term which encompasses all the intrinsic traits in data that can hinder the performance of a classifier. There are several categories of metrics that can be used to gauge the complexity of a dataset: feature overlap, linearity, neighborhoods, dimensionality, class balance and network properties. Datasets that present high levels of complexity in some of these categories have been shown to cause poor classifier performance [14].

When it comes to addressing data complexity during a preprocessing [9] phase, several specific methods can be found on the literature, but they usually tackle high dimensionality (feature selection and extraction methods) and class imbalance (resampling methods). Little research has been published on how to address other complexity types during a preprocessing phase and most of it is centered around feature selection [20, 23].

In this work, we present Slicer (supervised linear classifier error reduction), an automatic feature extractor designed with linear class separability in mind. It is based on an autoencoder (AE) model [8], using a special loss function inspired by least-squares support vector machines (LSSVM) [18]. The objective of this model is to learn an alternative representation for each instance where classes are more easily distinguishable. Once trained, the model is able to project any new instance onto the learned feature space without knowing its class. This allows to work with compact representations of the samples instead of the original, high-dimensional ones, in a way that facilitates the work of traditional classifiers, which are usually hindered by high dimensionality [4, 6] unless they specifically select features internally. Extracting features can also help when it is necessary to combine variables from different sources (e.g. images and clinical data), and working with traditional classifiers makes it easier to understand the decision-making process.

The proposed model is applied in a specific use case, aiming both to analyze the level of performance that could be gained and to open new possibilities for combination of imagery and other data, as well as interpretability of classifiers. The chosen application is recognition of COVID-19 in chest X-ray images, using the COVIDGR dataset [19] for this purpose. The fitness of the set of features learned by Slicer is evaluated by means of classification metrics using several standard classifiers and is compared against using a basic AE and learning from the original, unmodified features. The results show a noticeable advantage of the Slicer-generated variables except when using support vector machines as classifier, even though specific traits of this dataset such as “apparently negative” positive samples might affect the learned representation.

The rest of this document is organized as follows. Section 2 describes the new feature learner named Slicer. Afterwards, Sect. 3 outlines the main aspects of the experimentation, including the dataset and the evaluation strategy. Section 4 discusses the results obtained in the experimentation above and, lastly, Sect. 5 draws some conclusions.

2 Class-Informed Autoencoder for Complexity Reduction

This section is dedicated to introducing a new model designed to learn features with improved class separability. The proposed model, Slicer, is based on the minimization of the error of a linear classifier, at the same time that it attempts to maximize its reconstruction abilities. As a result, the learned features are influenced both by the overall information within the data as well as their relation to the class.

2.1 Autoencoder Fundamentals

An AE is an artificial neural network that is trained to reconstruct the inputs at its output [8]. It includes a certain bottleneck where the representation of the data is somehow restricted: e.g. it is lower dimensional, more sparse or robust against noise. This prevents the AE from simply copying the input instance throughout the network. Instead, an encoder f learns a new representation for the data while a decoder g must be able to recover the original features. This transformation is learned by optimizing the reconstruction error $\mathcal{J}_{\text{RE}}(x, (g \circ f)(x))$ which typically measures a distance between the original samples and their reconstructions, but different penalties and regularizations can allow to influence other aspects of the learned representation.

AEs are usually unsupervised tools, in the sense that they do not receive any information about the labels of the data nor the desired encodings for each instance. Their learning mechanism is, as a result, self-supervised [13]. This means that they can be applied in many contexts where label information is not necessarily available or just partially so: anomaly detection, semantic hashing, data compression, among others [7]. Nonetheless, some AEs do use label information [16], even though the objective is other than learning more separable features. Our objective is to regularize an AE so that it learns from the class labels during training, but does not need them during the prediction phase, and thus facilitates classification tasks.

2.2 Slicer Model: The Loss Function

Slicer is an AE model regularized by a special penalty function which takes class separability into account. In order to do this, an additional component is introduced to simultaneously fit a LSSVM to the encoded samples as well as evaluate the encoding based on the fitness of said classifier.

Figure 1 shows a diagram with the main components of Slicer. f represents the encoder, which is a neural network that transforms the inputs onto encodings. g refers to the decoder, a similar network whose objective is to reconstruct the original data inputs out of the encodings. These encodings are also fed to a single fully connected layer with no activation function, which acts as the support vector machine. Each of the decoder and this last layer are evaluated with their loss functions: the reconstruction error and the LSSVM minimization objective, respectively.




Fig. 1. Schematic illustration of the Slicer model

The reconstruction error chosen for the purposes of modeling input samples is cross entropy: a measure of disagreement between two probability distributions, for the case of Bernoulli distributions. It is usually the better option when all values in each instance are in the $[0, 1]$ interval. Its formulation for n instances with k variables is shown in Eq. 1.

$$\mathcal{J}_{\text{RE}}(x, \theta) = -\frac{1}{nk} \sum_{i=1}^n \sum_{j=1}^k x_j^{(i)} \log \left[(g \circ f) \left(x^{(i)} \right)_j \right] + \left(1 - x_j^{(i)} \right) \log \left[1 - (g \circ f) \left(x^{(i)} \right)_j \right]. \tag{1}$$

If the variables are not scaled to the $[0, 1]$, other reconstruction errors like the mean squared error could be used.

For its part, the LSSVM objective assumes that labels are in $\{-1, 1\}$ and is simply a sum of quadratic errors e_i^2 subject to the equality constraint $1 - e_i = y^{(i)} - w^T f(x^{(i)}) + b$, resulting as formulated in Eq. 2.

$$\mathcal{J}_{\text{LSSVM}}(x, y, \theta) = \frac{\mu}{2} w^T w + \frac{\zeta}{2} \sum_{i=1}^n \left(y^{(i)} - w^T f \left(x^{(i)} \right) + b \right)^2. \tag{2}$$

In the previous equation, f usually refers to the kernel used in the model but, in this case, it represents the encoder of the neural network. w holds the weights of the SVM, which are associated to a penalty with coefficient μ . The term that compares classes (y) to the LSSVM output is weighted by ζ .

The resulting loss function for the Slicer model is the sum of both the reconstruction error and the LSSVM loss:

$$\mathcal{J}(x, y, \theta) = \mathcal{J}_{\text{RE}}(x, \theta) + \mathcal{J}_{\text{LSSVM}}(x, y, \theta) \tag{3}$$

3 Experimental Framework

The objective of this experimentation is to apply the proposed complexity reduction method in a real world practical case. In particular, we aim to improve the

performance of simple classifiers when dealing with a chest X-ray image dataset for COVID-19 classification. This would open several promising research lines, such as the combination of these extracted features with other clinical and laboratory variables or the possibility of using easily interpretable classifiers with the generated features.

3.1 COVIDGR Dataset

The COVIDGR dataset of X-ray chest images was introduced in [19]. These images were collected under a collaboration with expert radiologists of the Hospital Universitario San Cecilio in Granada, Spain. In total, 852 images were annotated under a strict protocol: positive images correspond to patients who have been tested positive for COVID-19 using RT-PCR within a time span of at most 24h between the X-ray image and the test. Every image was taken using the same type of equipment and always with the posterior-anterior view. It is important that all images are consistent since, otherwise, classifiers could find cues to distinguish COVID-positive samples from negative ones different from the intended aspects of the X-ray that characterize the pneumonia associated to the disease [15]. Figure 2 includes one positive example and a negative one.




Fig. 2. A negative sample (left) and a positive one (right)

More information about class distribution is provided in Table 1. In the following experiments, the exact same partitions used in [19] are employed, in order to ease comparisons with previous results.

3.2 Evaluation Strategy

Since the final objective of the proposed model is to improve classification performance, the evaluation framework will consist in a variety of simple classifiers that will be trained with either the original features or the encoded ones. Standard classification metrics will be computed using the predictions over test subsets. A 5-fold cross validation scheme will be applied 5 times for a total of 25 train-test

Table 1. Class distribution in COVIDGR dataset. Normal-PCR+ refers to X-rays where COVID-19 was not detected by the experts but the patients tested positive.

| Class | #images | Severities | |
|----------|---------|-------------|-----|
| Negative | 426 | | |
| Positive | 426 | Normal-PCR+ | 76 |
| | | Mild | 100 |
| | | Moderate | 171 |
| | | Severe | 79 |

runs, so as to prevent errors from statistical chance. Table 2 lists the available feature sets and every classifier and evaluation metric included in the experiment. Each column is independent, in the sense that every feature set has been tested with each one of the classifiers and the performance has always been assessed with all four metrics.

Table 2. Evaluation framework: available feature sets, tested classifiers and evaluation metrics. TP, TN, FP and FN denote true positives, true negatives, false positives and false negatives, respectively.

| Feature sets | Classifiers | Evaluation metrics | |
|--------------|------------------------------|--------------------|-----------------------------------------------------------------------------------------|
| Original | Decision tree (DT) | Accuracy | $\frac{TP + TN}{TP + TN + FP + FN}$ |
| Basic AE | k nearest neighbors (kNN) | Precision | $\frac{TP}{TP + FP}$ |
| Slicer | Support vector machine (SVM) | Recall | $\frac{TP}{TP + FN}$ |
| | Gaussian process (GP) | F1-score | $\frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$ |

3.3 Architecture of the Slicer Model Used with the COVIDGR Dataset

The proposed Slicer model has been implemented in the Python language on top of the Tensorflow [1] library. Since the model needs to deal with image data, the specific architecture makes use of convolutional layers for the encoder and deconvolutional (or transposed convolutional) layers for the decoder.

More specifically, most of the AE is composed of residual blocks such as the ones in the ResNet-V2 architectures [11], as can be seen in Fig. 3. The left side of this diagram shows the detailed architecture of the AE-based model, with the encoder ranging from the input to the dense layer with 128 units, and the decoder from there to the last deconvolutional layer. The classification component consists in the fully connected (dense) layer that maps the encoding

to one variable and is then connected to the LSSVM loss. The overall loss is simply the sum of both error measures, as explained above.

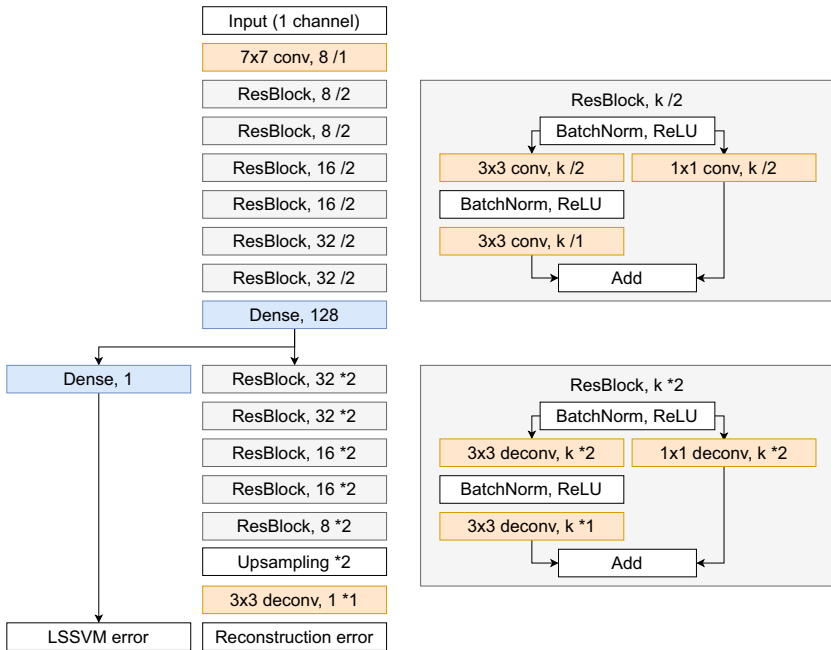


Fig. 3. Structure of the Slicer model used with the COVIDGR dataset. Each convolutional and deconvolutional layer is followed by the number of filters as well as an indication of the stride: /2 indicates a stride of 2 for a convolutional layer (the side of the image is halved), and *2 indicates a stride of 2 for a deconvolutional layer (the side of the image is doubled). Each of the residual blocks unrolls like the corresponding diagram on the right.

Using the previous architecture, a model was trained for each training partition of the total of 25 runs, using the parameters detailed in Table 3. Images were resized to a common resolution of 512×512 pixels (in total, 262144 variables in the range $[0, 1]$), and the encoding size was of 128 variables, which gives a reduction ratio of 1:2048, or equivalently, it represents the images using just 0.049% of the original variables. An equivalent basic AE was also trained using the same architecture (except for the LSSVM layer and loss) and partitions.

4 Results and Discussion

In this section, the main results of the experimentation are analyzed. Table 4 contains average metrics for the total of 25 runs that were performed. We can observe drastically different behaviors according to the classifier that was used:

Table 3. Relation of hyperparameters. DA refers to data augmentation techniques (only the techniques shown were applied).

| Parameter | Value |
|---------------------|------------------|
| μ | 0.01 |
| ζ | 0.1 |
| Encoding dimension | 128 |
| Epochs | 50 |
| Batch size | 8 |
| Optimizer | Adam |
| DA: random rotation | $\leq 5^\circ$ |
| DA: horizontal flip | Yes |
| kNN: k | 5 |
| DT: maximum depth | 10 |
| Image size | 512×512 |

Table 4. Average classification metrics over 25 runs (5 times 5-fold cross validation). Results from convolutional neural networks are reported in [19].

| Classifier | Features | Accuracy | Precision | Recall | F1-score |
|------------|-------------|---------------|---------------|---------------|---------------|
| DT | Original | 58.098 | 58.259 | 57.470 | 57.665 |
| DT | Autoencoder | 58.377 | 58.889 | 55.919 | 57.196 |
| DT | Slicer | 58.593 | 58.837 | 57.422 | 57.988 |
| GP | Original | 50.024 | 45.333 | 1.266 | 2.452 |
| GP | Autoencoder | 50.024 | 4.000 | 0.047 | 0.093 |
| GP | Slicer | 62.656 | 64.339 | 56.673 | 60.032 |
| kNN | Original | 62.585 | 65.804 | 52.780 | 58.358 |
| kNN | Autoencoder | 61.837 | 65.091 | 51.469 | 57.194 |
| kNN | Slicer | 62.326 | 63.624 | 57.284 | 60.092 |
| SVM | Original | 67.329 | 66.611 | 69.920 | 67.931 |
| SVM | Autoencoder | 67.072 | 67.006 | 67.622 | 67.057 |
| SVM | Slicer | 65.987 | 66.235 | 65.025 | 65.393 |

the decision tree had similar performance independently of the set of features that were provided, while the Gaussian process only was competitive when using the features learned by Slicer.

Several deductions can be made out of the results in Table 4. First, the classifiers that take the most advantage from the Slicer-generated features are kNN and GP. In fact, the latter struggles to find an acceptable model of the data using either the original features or the autoencoded ones. For its part, the DT shows little variance with respect to the set of variables it uses, although the

Slicer-generated has a slight lead in F1 score and accuracy. The SVM, however, does not benefit from the more separable features and loses performance with respect to the original and autoencoded ones.

Overall, it is not very surprising that the classifiers that are typically more affected by the quality of features are those which benefit more from the encodings provided by Slicer, whereas classifiers that internally perform their own feature selection or transformations either see small improvements or even decrease their performance.

It is important to note that there exist specific deep learning architectures designed for COVID-19 classification in chest X-rays, such as COVIDNet [21], COVID-CAPS [2] and COVID-SDNet [19]. Comparing the classification performance with these is out of scope for the present work, since we only aim to assess how useful the features extracted by the Slicer model are for traditional classifiers, not to find the best COVID-19 classifier.

As a summary, Table 5 displays the average metrics across all 4 classifiers, for each feature set. The average ranking that each one achieved for each metric is also shown. From this, we can conclude that the Slicer model produces feature sets that are consistently superior to a basic AE. Furthermore, it is able to preserve or even improve the quality of the original features, while simultaneously reducing drastically the dimensionality.

Table 5. Average classification metrics per feature set provided to the classifiers. The average ranking achieved by each feature set in each of the tests is shown in parentheses (lower is better).

| Features | Accuracy | Precision | Recall | F1-score |
|-------------|---------------|---------------|---------------|---------------|
| Original | 59.509 (2.06) | 59.002 (1.97) | 45.359 (1.92) | 46.602 (1.89) |
| Autoencoder | 59.327 (2.09) | 48.747 (2.16) | 43.764 (2.38) | 45.385 (2.29) |
| Slicer | 62.391 (1.85) | 63.259 (1.87) | 59.101 (1.71) | 60.876 (1.82) |

5 Conclusions and Future Work

This work has presented a novel framework for class separability enhancement using an AE-based model with a linear classification component that contributes to the loss function. The model has been implemented as a convolutional AE for the transformation of chest X-ray images onto a more manageable number of variables in order to employ simple classifiers in COVID-19 classification. An exhaustive experimentation has shown that the proposed model improves classification performance over a basic AE with no regularizations and maintains or even improves the performance compared to using the unprocessed data, even though the number of variables is dramatically reduced.

The promising results lead to consider several ways of continuing the work for practical real-world uses:

- Analyze the impact of severity levels on the adequacy of the learned representation. For example, Normal-PCR+ samples appear to look just like negative ones, although they are positive, which could affect the behavior of the model. Removing these images could provide better separation abilities as a consequence.
- Learned features can be combined with other variables that do not come from the chest X-ray, that is, clinical and laboratory data about each patient such as age, gender, comorbidities, etc. Several scores for this kind of data have been proposed but they do not take advantage of the full chest X-ray information [10, 12].
- Several ways to provide meaning to the extracted variables, such as feature disentanglement [22], in combination with transparent classifiers like decision trees, would enable more interpretable pipelines for COVID-19 classification, where users could trace predictions back to the original features.
- Combining the proposed loss function with more advanced AE models such as variational or adversarial AEs could add more potential of improving subsequent classification tasks.

Acknowledgments. D. Charte is supported by the Spanish Ministry of Science under the FPU National Program (Ref. FPU17/04069). F. Charte is supported by the Spanish Ministry of Science project PID2019-107793GB-I00/AEI/10.13039/501100011033. F. Herrera is supported by the Spanish Ministry of Science project PID2020-119478GB-I00 and the Andalusian Excellence project P18-FR-4961. This work is supported by the project COVID19RX-Ayudas Fundación BBVA a Equipos de Investigación Científica SARS-CoV-2 y COVID-19 2020.

References

1. Abadi, M., et al.: TensorFlow: large-scale machine learning on heterogeneous systems (2015). <https://www.tensorflow.org/>
2. Afshar, P., Heidarian, S., Naderkhani, F., Oikonomou, A., Plataniotis, K.N., Mohammadi, A.: COVID-CAPS: a capsule network-based framework for identification of COVID-19 cases from x-ray images. *Pattern Recogn. Lett.* **138**, 638–643 (2020)
3. Aggarwal, C.C.: *Data Classification*, pp. 285–344. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-14142-8_10
4. Aggarwal, C.C., Hinneburg, A., Keim, D.A.: On the surprising behavior of distance metrics in high dimensional space. In: Van den Bussche, J., Vianu, V. (eds.) *ICDT 2001*. LNCS, vol. 1973, pp. 420–434. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-44503-X_27
5. Basu, M., Ho, T.K.: *Data Complexity in Pattern Recognition*. Springer, Heidelberg (2006). <https://doi.org/10.1007/978-1-84628-172-3>
6. Beyer, K., Goldstein, J., Ramakrishnan, R., Shaft, U.: When is “nearest neighbor” meaningful? In: Beeri, C., Buneman, P. (eds.) *ICDT 1999*. LNCS, vol. 1540, pp. 217–235. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-49257-7_15
7. Charte, D., Charte, F., del Jesus, M.J., Herrera, F.: An analysis on the use of autoencoders for representation learning: fundamentals, learning task case studies, explainability and challenges. *Neurocomputing* **404**, 93–107 (2020). <https://doi.org/10.1016/j.neucom.2020.04.057>

8. Charte, D., Charte, F., García, S., del Jesus, M.J., Herrera, F.: A practical tutorial on autoencoders for nonlinear feature fusion: taxonomy, models, software and guidelines. *Inform. Fusion* **44**, 78–96 (2018). <https://doi.org/10.1016/j.inffus.2017.12.007>
9. García, S., Luengo, J., Herrera, F.: *Data Preprocessing in Data Mining*, vol. 72. Springer, Heidelberg (2015). <https://doi.org/10.1007/978-3-319-10247-4>
10. Gong, J., et al.: A tool for early prediction of severe coronavirus disease 2019 (COVID-19): a multicenter study using the risk nomogram in Wuhan and Guangdong, China. *Clin. Infect. Dis.* **71**(15), 833–840 (2020)
11. He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) *ECCV 2016*. LNCS, vol. 9908, pp. 630–645. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46493-0_38
12. Knight, S.R., et al.: Risk stratification of patients admitted to hospital with COVID-19 using the ISARIC WHO clinical characterisation protocol: development and validation of the 4C mortality score. *bmj* **370**, 1–13 (2020)
13. Liu, X., et al.: Self-supervised learning: generative or contrastive. *arXiv preprint arXiv:2006.08218* **1**(2) (2020)
14. Luengo, J., Fernández, A., García, S., Herrera, F.: Addressing data complexity for imbalanced data sets: analysis of smote-based oversampling and evolutionary undersampling. *Soft. Comput.* **15**(10), 1909–1936 (2011)
15. Maguolo, G., Nanni, L.: A critic evaluation of methods for COVID-19 automatic detection from x-ray images. *Inform. Fusion* **76**, 1–7 (2021). <https://doi.org/10.1016/j.inffus.2021.04.008>
16. Makhzani, A., Shlens, J., Jaitly, N., Goodfellow, I., Frey, B.: Adversarial autoencoders. *arXiv preprint arXiv:1511.05644* (2015)
17. Pascual-Triana, J.D., Charte, D., Arroyo, M.A., Fernández, A., Herrera, F.: Revisiting data complexity metrics based on morphology for overlap and imbalance: snapshot, new overlap number of balls metrics and singular problems prospect. *Knowl. Inf. Syst.* **63**, 1961–1989 (2021)
18. Suykens, J.A., Vandewalle, J.: Least squares support vector machine classifiers. *Neural Process. Lett.* **9**(3), 293–300 (1999)
19. Tabik, S., Gómez-Ríos, A., Martín-Rodríguez, J.L., Sevillano-García, I., Rey-Area, M., Charte, D., et al.: COVIDGR dataset and COVID-SDNet methodology for predicting COVID-19 based on chest x-ray images. *IEEE J. Biomed. Health Inform.* **24**(12), 3595–3605 (2020). <https://doi.org/10.1109/JBHI.2020.3037127>
20. Wang, L.: Feature selection with kernel class separability. *IEEE Trans. Pattern Anal. Mach. Intell.* **30**(9), 1534–1546 (2008)
21. Wang, L., Lin, Z.Q., Wong, A.: COVID-Net: a tailored deep convolutional neural network design for detection of COVID-19 cases from chest x-ray images. *Sci. Rep.* **10**(1), 1–12 (2020)
22. Yu, X., Chen, Y., Li, T., Liu, S., Li, G.: Multi-mapping image-to-image translation via learning disentanglement. *arXiv preprint arXiv:1909.07877* (2019)
23. Zhang, Y., Li, S., Wang, T., Zhang, Z.: Divergence-based feature selection for separate classes. *Neurocomputing* **101**, 32–42 (2013). <https://doi.org/10.1016/j.neucom.2012.06.036>