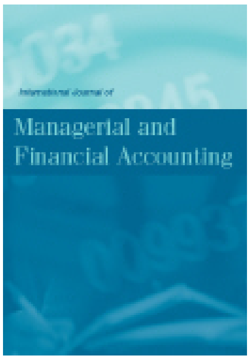

International Journal of Managerial and Financial Accounting

ISSN online: 1753-6723 - ISSN print: 1753-6715
https://www.inderscience.com/ijmfa

Competitive advantage in healthcare based on augmentation of clinical images with artificial intelligence: case study of the 'Sambias' project

Alessandra D'Amico, Michele Di Capua, Emanuel Di Nardo, Joanna Rosak-Szyrocka, Giuseppe Tortorella, Giuseppe Festa

Article History:

Received:	17 May 2023
Last revised:	25 July 2023
Accepted:	31 July 2023
Published online:	06 December 2023

Competitive advantage in healthcare based on augmentation of clinical images with artificial intelligence: case study of the 'Sambias' project

Alessandra D'Amico

Department of Radiology,
Casa di Cura Tortorella S.p.a.,
Salerno, Italy
Email: alessandra.damico@casadicuratortorella.it

Michele Di Capua and Emanuel Di Nardo

Research and Development Department,
Unlimited Technology S.r.l.,
Naples, Italy
Email: micheledicapua@gmail.com
Email: emanuel.dinardo@gmail.com

Joanna Rosak-Szyrocka

Department of Production Engineering and Safety,
Faculty of Management,
University of Technology,
Czestochowa, Poland
Email: joanna.rosak-szyrocka@wz.pcz.pl

Giuseppe Tortorella

General Management Department,
Casa di Cura Tortorella S.p.a., Salerno, Italy
Email: giuseppe.tortorella@casadicuratortorella.it

Giuseppe Festa*

Department of Economics and Statistics,
University of Salerno, Italy
Email: gfesta@unisa.it
*Corresponding author

Abstract: In the era of artificial intelligence, and particularly machine learning and deep learning models, the availability of large datasets is crucial to develop innovative and effective services, especially in the healthcare field. In this context, one essential requirement is access to verified information for contextualising/enriching the data. The SAMBIAS project analysed in this study involves the implementation of a software platform for data sharing in clinical scenarios, with the main objective of providing specific medical datasets to improve the competitiveness of the healthcare organisation from a general point of view. The platform, which is accessible via the web, provides on-demand, augmented sets of clinical situations, based on the enormous amounts of data that are collected by the health information systems of healthcare organisations. The case under investigation here is the Casa di Cura Tortorella s.p.a., Salerno, Italy. The implications of this platform are discussed in terms of more efficient performance.

Keywords: healthcare; artificial intelligence; machine learning; deep learning; data augmentation; business process management; business process improvement.

Reference to this paper should be made as follows: D'Amico, A., Di Capua, M., Di Nardo, E., Rosak-Szyrocka, J., Tortorella, G. and Festa, G. (2024) 'Competitive advantage in healthcare based on augmentation of clinical images with artificial intelligence: case study of the 'Sambias' project', *Int. J. Managerial and Financial Accounting*, Vol. 16, No. 1, pp.1–16.

Biographical notes: Alessandra D'Amico is the Scientific Director of the Department of Radiology at the Casa di Cura Tortorella S.p.a., Salerno, Italy. She is an Adjunct Professor at the School of Specialisation in Radio-Diagnostics of the University of Naples 'Federico II', Italy (since 2010). She is the author and co-author of more than 100 scientific publications in the field of radio-diagnostics.

Michele Di Capua is an Experienced Research Director with a demonstrated history of working in the computer software industry. She is skilled in OOP, mobile computing and machine learning (deep learning). She is a strong research professional graduate in Computer Science at the University of Salerno, Italy and with a PhD (in Computer Science) from the University of Milan 'Statale'. Since 2016, he is a teacher at the Apple Foundation Program (University of Naples 'Parthenope') in collaboration with Apple, Inc.

Emanuel Di Nardo is a PhD student in Computer Science at the University of Milan 'Statale'. His primary research interests are about machine learning and artificial intelligence, with focus on computer vision and visual object tracking using transformers networks. He conducted several studies on deep fuzzy neural networks, visual imagery, weather forecasting, naval engineering related problems, cancer prediction in bioinformatics, emotion recognition from faces and speech, and natural language processing for sentiment analysis.

Joanna Rosak-Szyrocka is an Erasmus+ coordinator at the Faculty of Management, Czestochowa University of Technology, Poland. She specialised in the fields of digitalisation, Industry 4.0, society 5.0, IoT, AI, quality management, and quality 4.0. She is a review editor, associate editor, and guest editor for several international journals. She is the author and co-author of more than 220 articles and 20 books at international level.

Giuseppe Tortorella is the Managing Director at the Casa di Cura Tortorella S.p.a., Salerno, Italy. He combines healthcare service practice with studies in the field of healthcare economics and management, with specific interests on risk management, quality management, and information systems management.

Giuseppe Festa is an Associate Professor of Management at the Department of Economics and Statistics of the University of Salerno (Italy). He holds a PhD in Economics and Management of Public Organisations from the University of

Salerno, where he is the Scientific Director of the Postgraduate Course in Wine Business and the Scientific Director of the Research Laboratory on 'Food and Wine Communication and Reputation Management'. He is the Chairman of the Euromed Research Interest Committee on Wine Business. His research interests mainly focus on wine business, information systems, corporate venture capital, and healthcare management.

1 Introduction

When reasoning about the adoption of artificial intelligence (AI) in healthcare, and particularly in radiology, one of the most important medical specialties in terms of health impact and related costs, a significant problem arises: the use of medical images in the related healthcare process involves patient privacy issues, making the dissemination of clinical data very complex. This is because the procedures for anonymisation, storage, and dissemination are not always optimal, and are sometimes even impossible, due to strict internal policies or emergency circumstances (Houfani et al., 2022). Obtaining permission from the patient to disseminate their clinical data, even if anonymously, is often problematic; for instance, in the European Union (EU), it is still often mandatory to comply with the General Data Protection Regulation (GDPR No. 2016/679), although in practice it is not requested any more in most cases.

In addition to these considerations, it is worth mentioning the difficulty of making certain types of diagnoses, regardless of their rate of occurrence, due to the possibility of imbalance between positive and negative samples in the dataset. In this case, further obstacles to the efficiency of healthcare processes may arise that are related to the categorisation of data; even when diagnostic facilities are able to perform the entire dissemination procedure, the data may not be labelled for a specific type of pathology or technical intervention, for example in the segmentation of pathological areas (Khan et al., 2020).

To solve these problems, several methodologies based on the most advanced information and communication technologies (ICTs) can be used, for example to increase the available number of samples in the dataset. The most common techniques in this area are related to data augmentation (DA); these methods aim to artificially increase the available number of samples, as mentioned above, and can be applied to one-dimensional data represented by signals of single or multiple types, such as measurements made using an electrocardiogram (Nita et al., 2022).

Other techniques are specific to two-dimensional data such as images, and often aim to manipulate chromatic components or to perform geometric operations through affine transformations. These solutions generally operate on the data under investigation while leaving the original data still valid, and generate further data with proper care, meaning that the new data are generally not required to be re-certified by an expert.

In recent years, thanks to the widespread use of AI, numerous solutions have emerged to create artificial data from real data with very high fidelity. Naturally, concerns over 'deep fakes' and the generation of synthetic images more generally are very different situations from the context of medical data; in the former case, if a generation error occurs, meaning that the fidelity and the quality of the generated image are not satisfactory, it is not necessary to solve that problem as a matter of maximum priority. In addition, the generation of appropriately labelled medical data may not eliminate the need for an expert who can validate those data.

Based on these considerations, a project focusing on AI in healthcare that aimed to enable DA in the field of medical images, entitled 'SAMBIAS', was carried out in the Casa di Cura Tortorella s.p.a., Salerno, Italy. The underlying idea of this project was to verify whether and how AI in general (and DA more specifically) could help in improving the global effectiveness of the healthcare processes. The managerial and financial points of view were also taken into consideration in the current study (Cannavale et al., 2022; Jia et al., 2022), with particular reference to healthcare activities connected with the treatment of medical images in the area of radiology.

2 Scientific background

There are various types of diagnostic technologies in the medical field, of which the most well-known are computed tomography (CT) and magnetic resonance imaging (MRI). These techniques construct a view of the organ to be analysed. CT scans allow for three-dimensional scans of the affected area so that the body part can be viewed in sections, to give a localised view of the pathology, whereas MRI cannot provide a three-dimensional image but allows for differentiation based on tissue types and images of body sections in three different planes (axial, coronal, and sagittal).

These data are considered sensitive, because despite being difficult to interpret, they provide important information about an individual's health status, and this means that dissemination of data is difficult. To address this problem, scholars and professionals have started generating data that are similar to those of real patients.

The generation of artificial data from real data has become widespread in recent years. Given the excellent performance of systems of this type, they have also begun to be used on sensitive images such as medical images, in an attempt to increase the number of samples available without raising the privacy issues discussed above that may be associated with individual patients.

Some systems are specific to a single area of the body, for example those that target liver lesions (Frid-Adar et al., 2018) or malignant and benign nodules in the lung (Chuquicusma et al., 2018). The results are reviewed by a group of experts to verify the level of reliability of the processed image with respect to real or fake data. Other techniques attempt to generate entire portions of CT scans by using *a priori* knowledge about the data, extrapolated via an autoencoder, with generative knowledge. In this way, artificial neural networks (ANNs) can be forced to generate content that is similar to real observations, but with random differences that cannot be predicted.

In a study by Frid-Adar et al. (2018), only lesion images from CT scans were considered for analysis; neither the three-dimensional image representing the CT nor the whole section of the liver image were used, but only the point at which the lesion appeared. More specifically, the lesions examined in their study were cysts, metastases, and haemangiomas (Figure 1 and Table 1).

As explained in further detail below, the classical structure of a deep convolutional generative adversarial network (DCGAN) was adopted in this study to generate the augmented images. A DCGAN is a direct extension of a generative adversarial network (GAN) and is similar except that it explicitly uses convolutional and convolutional-transpose layers in the discriminator and generator, respectively.

Figure 1 Comparison of generated and real images: (a) classification accuracy, (b) operating augmentation only with the addition of synthesised images (see online version for colours)

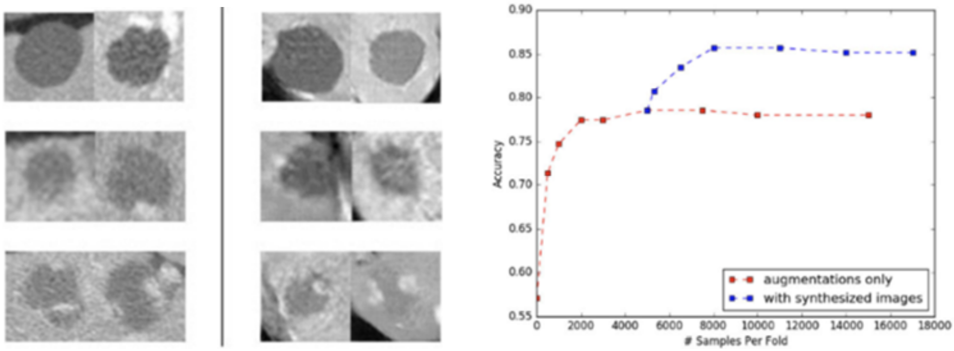

(a) (b)

Source: Frid-Adar et al. (2018)

Table 1 Confusion matrices for classifiers using (a) only real data, (b) real and synthetic data

True/auto	*Cyst*	*Met*	*Hem*	*Sensitivity*
Cyst	52	1	0	98.1%
Met	2	44	18	68.7%
Hem	0	18	47	72.3%
Specificity	98.4%	83.9%	84.6%	

(a)

True/auto	*Cyst*	*Met*	*Hem*	*Sensitivity*
Cyst	53	0	0	100.00%
Met	2	52	10	81.2%
Hem	1	13	51	78.5%
Specificity	97.7%	89.0%	91.4%	

(b)

Source: Frid-Adar et al. (2018)

This type of software was first described in a study by Radford et al. (2015), in which the input to the generator was a random vector of 100 elements and the output was a 64x64 image. The data were collected from the databases of Sheba Medical Centre, and were annotated by experienced radiologists, with the dataset containing 53 cysts, 64 metastases, and 65 haemangiomas.

In Figure 1(b), the variation in the accuracy depends on the number of samples and on the increase in their variance. In the final analysis, both the real and synthetic images were submitted to two expert radiologists, in order to verify the human analysis capability and the extent to which the synthetic images differed from real ones.

The first expert reported results of 78% accuracy for the real data and 77.5% accuracy for the generated data, while the second expert assessed the percentage of correctly classified items as 69.2% for both types of data. A different DCGAN was then trained for each class.

Tables 1(a) and 1(b) show the classification results from the use of simple augmentation and after enriching the dataset with synthetic data, respectively. It is evident that this treatment increased the performance of the classification methodology by providing more data, which were highly faithful, for use in the training and learning phases.

In a study by Chuquicusma et al. (2018), images of benign/malignant nodules were generated using a DCGAN and submitted to two expert radiologists. A visual Turing test

(VTT) was then carried out in which the experts, before proceeding to the relative classification, were asked to determine whether the submitted image was real or synthetic. This methodology allowed the authors to investigate whether the quality of the images was reliable enough to deceive an expert, and if so, whether the quality of the generated image approximated that of the real data that were then classified by the expert.

Figure 2 Nodule generation results: (a) low-grade nodules, (b) high-grade nodules, (c) malignant nodules, (d) benign nodules (see online version for colours)

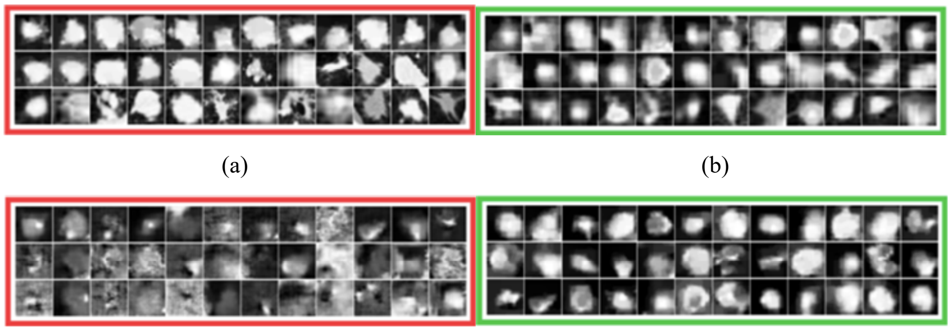

(a) (b) (c) (d)

Source: Authors' elaboration from Chuquicusma et al. (2018)

In that study, the authors used as input a random vector of 100 elements, while the output image had a size of 56 × 56. A public dataset entitled Lung Image Database Consortium – Image Database Resource Initiative (LIDC-IDRI) was used, which consisted of 1,018 CT thoracic screening images, with annotations and with staging information from one (least aggressive) to five (most aggressive). As shown in Figure 2, nodules with a diameter of 3 mm or more were annotated by at most four radiologists; the chosen nodules were those with an annotation from at least three radiologists, while nodules with an aggressiveness value of three were excluded.

In that study, a total of 1,145 nodules were obtained after pre-processing, of which 635 were benign and 510 were malignant. Tests were conducted using a class-specific network to be generated: a DCGAN for benign samples (with 114,000 iterations), a DCGAN for malignant samples (with 110,000 iterations), and a mixed test in which only one DCGAN was used for both classes (with 99,000 iterations).

Figures 2(a) and 2(b) show a comparison of the generated nodules, with the low-quality ones shown in red and the high-quality ones in green. From this comparison, it can be seen that there is no difference between the classes. In addition, it can be observed that the high-quality nodules have well-defined contours and no pixelation effects.

Figures 2(c) and 2(d) show the malignant nodules and the benign ones, respectively. Of these, the nodules numbered 1 to 18 in both classes were generated by the DCGAN, while the remainder are the real ones. Eighteen VTTs were conducted by one radiologist with 14 years' experience and another with four years' experience. Each test contained 36 high-quality samples, and the real and generated samples were pre-processed using an anisotropic diffusion filter (Perona-Malik). The experiments that were conducted are summarised as follows (see Figure 3 for further information about the false recognition rate, FRR):

- Experiments 1, 2, 3, 16, 17, and 18 (benign/malignant): the experts were asked to identify the generated nodules and the real ones.
- Experiments 4, 5, 6, 13, 14, 15 (benign): the experts were asked to identify the generated nodules and the real ones, and to identify which of the nodules were benign, if any.
- Experiments 7, 8, 9, 10, 11, 12 (malignant): the experts were asked to identify the generated nodules and the real ones, and to identify which of the nodules were malignant, if any.

The results for each experiment in terms of the FRR are shown in Figure 3, with the exception of experiments 2, 5, 8, 11, 14, and 17, which used real data. The most experienced radiologist assessed synthetic nodes as real in 69% of the cases, while the least experienced radiologist was wrong in more than 99% of the cases. These tests indicate that the algorithm was able to generate data that highly resembled the originals, or at least resembled them enough to deceive two experts in the field.

In a study by Wang et al. (2022), a priori knowledge of the CT was adopted and the spatial, shape, and size information was learned via an autoencoder. This methodology allowed the authors to increase the representativeness of the data, which were then used to train a conditional GAN (CGAN). More specifically, this study exploited the ability of GANs to perform domain-adaptation (which differs from the concept of transfer learning). The network in question was composed of multiple modules (autoencoder + GAN) and a random element; in this case, the generator was aided by knowledge about the real image when generating a dummy image, while the random vector helped to avoid mode collapse of the network on the same images (i.e., equal image deriving from different generations).

In this basic structure, the first module consisted of a convolutional auto-encoder (CAE) that performed feature space encoding, both spatially and shape-wise. This methodology, as mentioned above, was adopted to ensure that the neural network could find a compact representation of the input data and relied on the ability to reconstruct the original data.

Figure 3 FRR for the experiments on the generated nodules, where the percentage of nodules recognised as synthetic is indicated (see online version for colours)

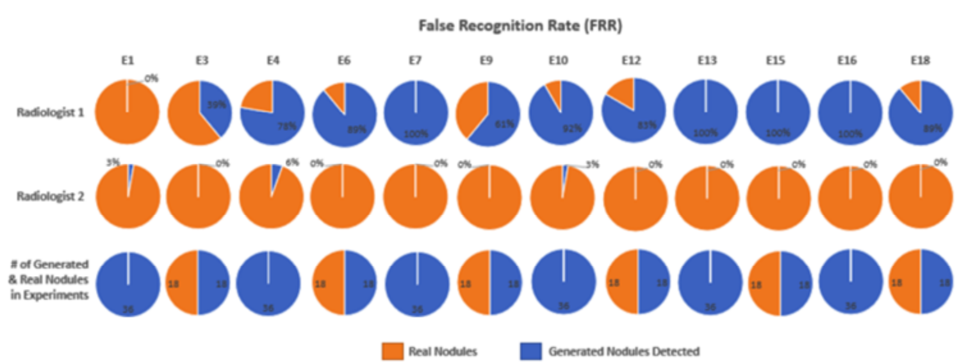

Source: Chuquicusma et al. (2018)

The encoding was then used to extract meaningful features for the remaining modules of the ANN. These features were refined before being used in the second module, which formed the heart of the generator. Since the goal was the generation of new images, a random vector was also given as input; this was changed from time to time, to provide the possibility of generating ever-changing data. This approach allowed the random distribution of the vector to be mapped onto a latent subspace that was the result of the learning process that the generator can express in a GAN.

The two inputs, i.e., the encoding of the real data provided by the CAE and the expansion of the random vector, needed to be aggregated, so that they could form a single input to the data generation process. To further strengthen the ability of the generator to output ever-changing data, two additional random vectors were inserted into the reconstruction process, while the discriminator was a simple classifier composed of convolutional layers, which was used to recognise real data from fake data. The GAN was trained using Wasserstein's loss, which provides a natural notion of dissimilarity for probability measures, together with the gradient penalty technique described by Zhang et al. (2019).

3 Research methodology

The design of the processing pipeline included an initial stage in which filtering and cleaning were applied to the image data coming from heterogeneous databases, with the intention of normalising the information contained in the digital archives. In the next phase of augmentation, other data (metadata) that were used in the subsequent investigation were added to the retrieved clinical information (e.g., geo-referencing the patient's territory of origin, adding data on correlations with other family pathologies, standard classifications of pathologies) (Khan et al., 2020; Shaikh and Ali, 2020; El Samad et al., 2022). A generative ANN (GANN) is a generative artificial neural network that attempts to learn the data distribution of real samples. In this case, two ANNs are adopted: one called the generator (which generates the dummy data) and another one called the discriminator (which discriminates true data from false data). The two networks are in competition with each other according to the MiniMax game policy, as described in game theory.

In essence, the generator attempts to trick the discriminator. Whenever the discriminator guesses correctly whether the data are real or fictitious, this information is used by the generator to change the type of images it produces, thus bringing the distribution of the data closer and closer to the real one and generating further data that are more and more similar to real data. A graphical schema of a GANN is shown in Figure 4. This type of network aims to achieve a Nash equilibrium, which is a fair trade-off between the ability of the generator to create bogus data and the ability of the discriminator to determine whether the submitted data are real or not.

The second important technique used in this study is known as an autoencoder'; it is a structure composed of an ANN that attempts to reconstruct input data. Differently from a GANN, the vanilla autoencoder cannot be used to generate new data; in fact, as it can be observed in Figure 5, given as input a real data, the reconstructed output must be as equal to as possible (more specifically, an encoding of the input into a subspace representing a simplified, but more meaningful representation, takes place).

Figure 4 Structure of a GANN (see online version for colours)

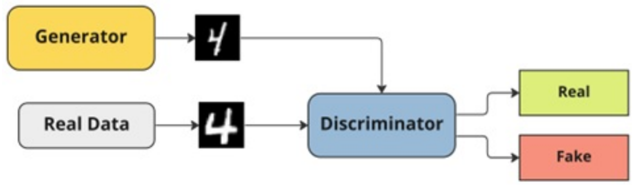

Source: Authors' elaboration

Figure 5 Architecture of the autoencoder (see online version for colours)

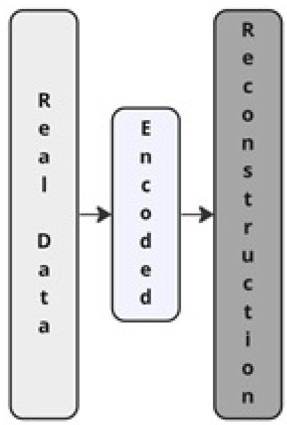

Source: Authors' elaboration

This encoded element is then decoded to reconstruct the input. The encoding process may involve many encoding elements, the last of which represents the most compact and meaningful representation of the initial feature space. This type of representation is called 'latent representation' and is typically used for inclusion in other methodologies to have a strong characterisation of the problem and the input to be processed.

As a starting point for this study, the approach proposed by Wang et al. (2022) was adopted. In this methodology, there are two stages in the learning phase: the first is needed to create a deep and compact representation of the input images, and relies on an autoencoder, while in the second stage, the GANN is trained with the generator, which combines a random vector network with the learned representation to force the generator network to create images that are of the same type as the input, but with diverse graphical representations.

To strengthen the learning of the latent representation, several modifications are made to the autoencoder and the discriminator; of these, the most reliable is the use of skip connections, which enhance the representation of the data while preserving the semantic meaning and providing an improved understanding of the spatial features in the reconstruction phase. Furthermore, the GANN was modified to form a CGAN, to which class constraints were added. The conditioning of a GAN is useful to steer the generation of images for a specific image, and in combination with the autoencoder latent representation exerts a strong influence over the quality and the ability of the neural network to generate a valid and specific image.

4 Data analysis

The first stage of the analysis concerned the processing of the medical image data. CT scans were used for this purpose, as this approach allows for three-dimensional scans of the affected area, meaning that the body part can be divided into different sections for investigation and the pathology can be located efficiently.

Image data were made available by the Casa di Cura Tortorella, a private hospital located in Salerno, Italy. This organisation was elected for this study since it uses state-of-the-art technological tools and procedures in its clinical processes.

All data were provided in Digital Imaging and COmmunication in Medicine (DICOM) format, an international standard that is used to manage medical images such as CT and MRI scans. Each DICOM image was categorised in a series. It generally equates to a specific type (modality) of data, or the position of a patient on the acquisition device. For the specific purposes of the experimental application in this study, it was of interest to select, among the available series, the number 202, since it is the one capable of reporting information regarding vasculopathy, in which there is strong research interest. All scans taken prior to 2020 were removed, as there was a change in the machinery and the acquisition technique used in the hospital.

Table 2 Quantity of samples by class

Class	*Quantity*
Vasculopathy	6,492
No vasculopathy	4,700

Source: Authors' elaboration

The dataset under analysis contained data on a total of 41 patients. From these data, images that could not be used (because they represented summary diagnostic information, or involved orientations different from the axial one or a region of the brain that did not provide information about the disease under study) were filtered out, leaving a total of 11,192 images. All files were divided into two categories based on the disease status of the patient.

The dataset imbalance about the 'Vasculopathy' class was manageable, with a difference in size between the two classes of less than 30%. Table 2 shows the total number of samples in each class.

The data were randomly split, with 85% of the samples forming the training set and 15% forming the testing set. To further validate the methodology, a validation process was applied at each step. The data used for validation were drawn from the training set, creating a new split of 10% of the data. The final training set therefore contained about 75% of the whole dataset of images.

The images were initially pre-processed. The first type of processing applied to this type of image involved transforming the intensity value of the component elements; originally, these images had values in terms of Hounsfield units (the so-called Hounsfield scale, deriving its name from Sir Godfrey Hounsfield, is a quantitative scale for describing radiodensity, and is frequently used in CT scans). These values can range from –1,024 to +3,071, where each range of values represents a specific physiological feature, for example bone, soft tissue, or brain tissue. For our case study of vasculopathy, a central window of values was chosen based on the internal acquisition statistics reported in the DICOM files. It is possible to find a prevalence of a window of values centred on +40, with a range of ±40 units; this range specifically represents the so-called 'grey matter'.

Then, as is common for this type of task, several pre-processing techniques were applied to process the images and to allow the machine learning (ML) model to work to its full potential. After the data preparation phase, training of the feature extractor and the generative neural network were performed; the training phase of the feature extractor containing the so-called a priori knowledge was performed using a common ML metric known as binary cross entropy (Vincent et al., 2008).

The network was then modified, as described above, by the use of skip connections. This kind of neural network learns a latent space in which the most relevant information about the input is represented in a compact form; following this, the learned representation is passed as part of the input to the generator, together with a random vector, which results in a secondary input branch.

The GANN was trained in the usual way, with a combination of a generator and discriminator, using a Wasserstein distance with a gradient penalty (Gulrajani et al., 2017). The output took the form of augmented data that preserved some features of the input data with random differences in the details.

5 Results and discussion

All tests were performed on the newly obtained dataset with the specific data enrichment deriving from the Casa di Cura Tortorella clinical databases, which were contextually necessary. Results are reported here for the learning of the feature extractor and the ability of the GANN to find an equilibrium between its inner components.

The first stage involved evaluating the quality of the data reconstruction by the autoencoder. Figure 6 shows the behaviour of the loss value, which represents the variation in the error (where a lower value is better), and it can be seen from the curves whether the training phase has achieved a significant result (or not yet).

Figure 6 Learning curves for the autoencoder with skip connections: the training loss is shown in orange, and validation loss in blue (see online version for colours)

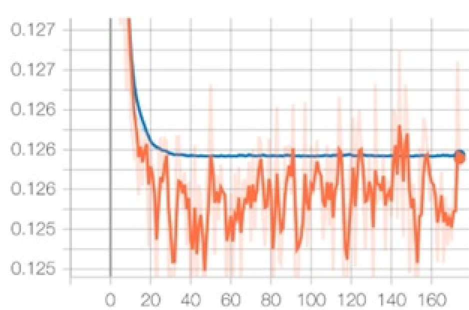

Source: Authors' elaboration

It can be observed that the two curves tend to converge towards a small number in the range [0.125–0.126]. Furthermore, their close proximity during learning means it is likely that there are no overfitting or underfitting problems.

The oscillations in the orange curve are due to the high variance in the data from shape to pixel values. These parameters are the most important elements in understanding the quality of the proposed neural network.

Figure 7 Learning curves for the GANN architecture: (a) generator, (b) discriminator (see online version for colours)

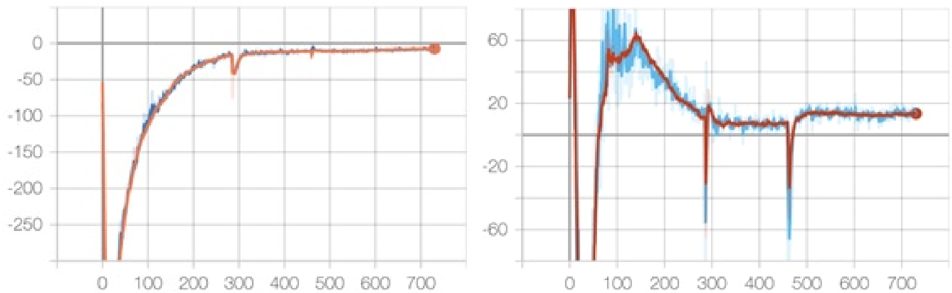

(a) (b)

Source: Authors' elaboration

The knowledge obtained in the first stage was used in the second stage as input to the GANN, and the learning curves for the generator and for the discriminator can be seen in Figures 7(a) and 7(b), respectively. The aim is to find an equilibrium where both plots are near to zero, and lower oscillations in the curves can be observed for both networks.

The number of total iterations to reach a good equilibrium is higher than in the previous model; for the generator, the distance from zero tends to reduce at each step, with few fluctuations, while the discriminator generally approaches zero with variations on similar steps. The most significant difference is that as training proceeds, the distance from the optimal equilibrium does not seem to increase for the discriminator. These observations can be explained by the fact that the generator is becoming expert in generating artificial images that are very similar to the original ones, and at a certain point, the discriminator is deceived by these data.

Again, the proximity of the training and the validation curves is important. Although these plots are rich in information, the most important test is the visualisation of the images generated by the GAN, as shown in Figure 8.

Figure 8 CT images generated by CGAN

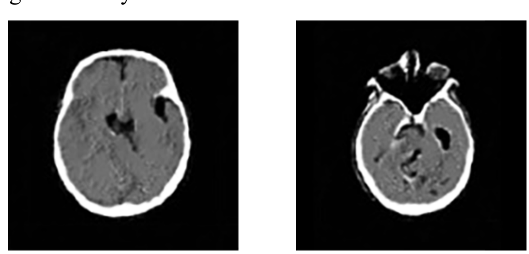

(a) (b)

Source: Authors' elaboration

Figure 8 CT images generated by CGAN (continued)

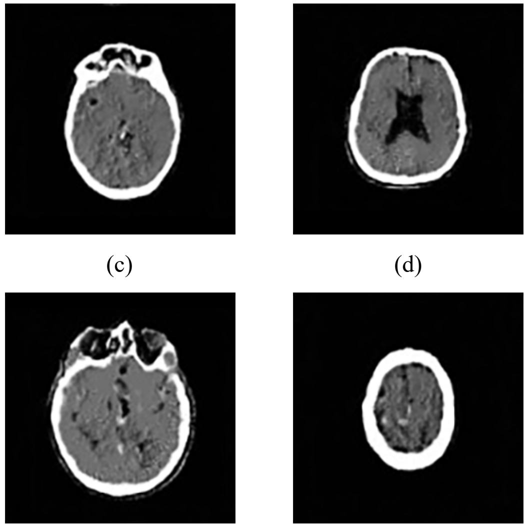

(c) (d)

(e) (f)

Source: Authors' elaboration

6 Scientific and managerial implications

The evidence resulting from the experiments conducted here may have several important implications for the adoption of AI models in radiology, in terms of increased competitiveness of the related processes. Of the potential range of implications, some seem to be particularly significant, from both a scientific and a managerial perspective.

The most important benefits of the proposed DA could affect several aspects of the global healthcare process. These include more accurate and timely diagnoses (increased efficacy), reduction of misinterpretation (increased efficacy in terms of safety, and increased effectiveness in terms of judiciary problems), workflow optimisation (increased effectiveness), personalised treatments (increased efficacy, also in terms of related business intelligence, for example using data mining), and access to radiological care via distance evaluations (increased effectiveness) (Talha et al., 2010; Marcarelli, 2018; Pascarella et al., 2021; Kumar et al., 2023; Singh et al., 2023).

At the same time, the most relevant implication for AI in healthcare, and more specifically radiological activity, which is under investigation in the current study, is the constant interaction of the software platform with the competence of the radiologist. ML and deep learning (DL) models can offer considerable support in regard to increasing the effectiveness of the overall clinical process, but only if continuously adopted under the supervision, coordination, and responsibility of healthcare specialists (Yang, 2022).

7 Limitations and future work

This study is subject to several constraints. Potential solutions to the limitations described here could be reached by extending the investigation, suggesting possible avenues for improving the global impact of the current research, which in its present form undoubtedly represents an exploratory experiment.

Firstly, this work considered as a case study the Casa di Cura Tortorella in Salerno, Italy, which was chosen as an extreme case since it had a very advanced health information system. Consequently, it is not possible to generalise the results of the research, meaning that enlarging the number of observed healthcare structures would be necessary to gain a deeper understanding of whether AI techniques, and ML and DL models in particular, could be widely used to improve the effectiveness of radiology activities.

Secondly, the study was conducted in the form of an experiment, and a specific software platform was developed that accepted real medical images and created additional synthetic medical images that were as close as possible to real images with other specific characteristics, using additional clinical data provided by the Casa di Cura Tortorella health information system to contextualise and enrich the generated images. Another experiment (i.e., another software platform) may provide different results, and it therefore seems necessary, in future research, to extend the current investigation with other tests that could be performed by other software platforms.

Thirdly, the study was conducted in the form of simulation, without carrying out checks between the generation of the augmented medical images (i.e., the simulated images) and the real medical images corresponding to the actual health status of the patient under analysis in the clinical situation. In other words, a comparison between/among real images, artificial images, and real images, with longitudinal verification, seems necessary for the global evolution of the experiment.

8 Conclusions

The use of AI techniques in the healthcare sector is still in its initial stages (Apell and Eriksson, 2023). The medical field, and radiology in particular, has traditionally been one of the areas that is most sensitive to the introduction of innovative technologies, and this is especially true in view of the tremendous impacts that various solutions arising from the world of AI may have (Vishwakarma et al., 2023).

In the specific case of medical images, several methodologies based on ML and DL models may help in achieving valid and reliable DA (from a clinical point of view), which may translate into a more effective organisation (from a managerial and financial point of view). In the healthcare sector, the quality of medical treatment remains the fundamental pillar of the entire process, but the careful management of operations, processes, and activities is also essential, since we must balance the need for treatments that are ever more expensive with the need for affordable healthcare for the community.

AI can offer support in this regard, as demonstrated by the results of the experiment conducted in this study (Feng et al., 2022). It is probable that in the very near future, debate will no longer centre on the possibility of using AI in a valid and reliable way in the medical field, but rather on how to use it to continually increase the effectiveness of the global healthcare process. The healthcare organisations that adopt AI in their processes are likely to gain a competitive advantage with respect to their competitors, although the priority is naturally the impact on the health of the patient.

Acknowledgements

The SAMBIAS project, started in 2022 and completed in 2023, was funded by POR Campania FESR 2014-2020 (CUP: B57H22003260007).

References

Apell, P. and Eriksson, H. (2023) 'Artificial intelligence (AI) healthcare technology innovations: the current state and challenges from a life science industry perspective', *Technology Analysis & Strategic Management*, Vol. 35, No. 2, pp.179–193.

Cannavale, C., Esempio, A., Leone, D. and Schiavone, F. (2022) 'Innovation adoption in inter-organizational healthcare networks – the role of artificial intelligence', *European Journal of Innovation Management*, Vol. 25, No. 6, pp.758–774.

Chuquicusma, M.J.M., Hussein, S., Burt, J. and Bagci, U. (2018) 'How to fool radiologists with generative adversarial networks? A visual Turing test for lung cancer diagnosis', *Proceedings of the IEEE 15th International Symposium on Biomedical Imaging (ISBI)*, Washington, DC, USA, pp.240–244.

El Samad, M., El Nemar, S., Sakka, G. and El-Chaarani, H. (2022) 'An innovative big data framework for exploring the impact on decision-making in the European Mediterranean healthcare sector', *EuroMed Journal of Business*, Vol. 17, No. 3, pp.312–332.

Feng, J., Phillips, R.V., Malenica, I., Bishara, A., Hubbard, A.E., Celi, L.A. and Pirracchio, R. (2022) 'Clinical artificial intelligence quality improvement: towards continual monitoring and updating of AI algorithms in healthcare', *NPJ Digital Medicine*, Article 66, Vol. 5, No. 2022, pp.1–9.

Frid-Adar, M., Klang, E., Amitai, M., Goldberger, J. and Greenspan, H. (2018) 'Synthetic data augmentation using GAN for improved liver lesion classification', *Proceedings of the IEEE 15th International Symposium on Biomedical Imaging (ISBI)*, Washington, DC, USA, pp.289–293.

Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V. and Courville, A.C. (2017) 'Improved training of wasserstein gans', *Proceedings of the Conference on Advances in neural information processing systems 30 (NIPS 2017)*, Long Beach, CA, USA, pp.5767–5777.

Houfani, D., Slatnia, S., Kazar, O., Saouli, H. and Merizig, A. (2022) 'Artificial intelligence in healthcare: a review on predicting clinical needs', *International Journal of Healthcare Management*, Vol. 15, No. 3, pp.267–275.

Jia, Q., Zhu, Y., Xu, R., Zhang, Y. and Zhao, Y. (2022) 'Making the hospital smart: using a deep long short-term memory model to predict hospital performance metrics', *Industrial Management & Data Systems*, Vol. 122, No. 10, pp.2151–2174.

Khan, H., Srivastav, A. and Mishra, A.K. (2020) 'Use of classification algorithms in health care', in Tanwar, P., Jain, V., Liu, C-M. and Goyal, V. (Eds.): *Big Data Analytics and Intelligence: A Perspective for Health Care*, pp.31–54, Emerald, Bingley, UK.

Kumar, P., Kumar Sharma, S. and Dutot, V. (2023) 'Artificial intelligence (AI)-enabled CRM capability in healthcare: the impact on service innovation', *International Journal of Information Management*, Article 102598, Vol. 69, No. 2023, pp.1–15.

Marcarelli, G. (2018) 'An integrated network model for performance management: a focus on healthcare organisations', *International Journal of Managerial and Financial Accounting*, Vol. 10, No. 2, pp.163–180.

Nita, S., Bitam, S., Heidet, M. and Mellouk, A. (2022) 'A new data augmentation convolutional neural network for human emotion recognition based on ECG signals', *Biomedical Signal Processing and Control*, Article 103580, Vol. 75, No. 2022, pp.1–18.

Pascarella, G., Rossi, M., Montella, E., Capasso, A., De Feo, G., Botti Sr., G., Nardone, A., Montuori, P., Triassi, M., D'Auria, S. and Morabito, A. (2021) 'Risk analysis in healthcare organizations: methodological framework and critical variables', *Risk Management and Healthcare Policy*, Vol. 14, No. 2021, pp.2897–2911.

Radford, A., Metz, L. and Chintala, S. (2015) *Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks*, Working Paper, pp.1–16, DOI: 10.48550/arXiv.1511.06434.

Shaikh, T.A. and Ali, R. (2020) 'Computer-aided big healthcare data (BHD) analytics', in Tanwar, P., Jain, V., Liu, C-M. and Goyal, V. (Eds.): *Big Data Analytics and Intelligence: A Perspective for Health Care*, pp.115–138, Emerald, Bingley, UK.

Singh, A., Madaan, G., Swapna, H.R. and Kumar, A. (2023) 'Impact of artificial intelligence on human capital in healthcare sector post-COVID-19', in Tyagi, P., Chilamkurti, N., Grima, S., Sood, K. and Balusamy, B. (Eds.): *The Adoption and Effect of Artificial Intelligence on Human Resources Management, Part A (Emerald Studies in Finance, Insurance, and Risk Management)*, pp.47–69, Emerald, Bingley, UK.

Talha, M., Christopher, S.B. and Kamalavalli, A.L. (2010) 'Sensitivity of profitability to working capital management: a study of Indian corporate hospitals', *International Journal of Managerial and Financial Accounting*, Vol. 2, No. 3, pp.213–227.

Vincent, P., Larochelle, H., Bengio, Y. and Manzagol, P.A. (2008) 'Extracting and composing robust features with denoising autoencoders', *Proceedings of the 25th International Conference on Machine Learning*, Helsinki, Finland, pp.1096–1103.

Vishwakarma, L.P., Singh, R.K., Mishra, R. and Kumari, A. (2023) 'Application of artificial intelligence for resilient and sustainable healthcare system: systematic literature review and future research directions', *International Journal of Production Research*, pp.1–23, DOI: 10.1080/00207543.2023.2188101.

Wang, X., Yu, Z., Wang, L. and Zheng, P. (2022) 'An enhanced priori knowledge GAN for CT images generation of early lung nodules with small-size labelled samples', *Oxidative Medicine and Cellular Longevity*, Article 2129303, Vol. 2022, No. 2022, pp.1–9.

Yang, C.C. (2022) 'Explainable artificial intelligence for predictive modeling in healthcare', *Journal of Healthcare Informatics Research*, Vol. 6, No. 2022, pp.228–239.

Zhang, T., Li, Z., Zhu, Q. and Zhang, D. (2019) 'Improved procedures for training primal Wasserstein GANs', *Proceedings of the IEEE Conference on SmartWorld, Ubiquitous Intelligence & Computing, Advanced & Trusted Computing, Scalable Computing & Communications, Cloud & Big Data Computing, Internet of People and Smart City Innovation*, Leicester, UK, pp.1601–1607.