Incremental Learning Versus Batch Learning for Classification of User's Behaviour in Medical Imaging

Carlos Viana-Ferreira, Sérgio Matos and Carlos Costa
Instituto de Engenharia Electrónica e Telemática de Aveiro, Universidade de Aveiro, Campus Universitário de Santiago, Aveiro, Portugal

Keywords: Batch Learning, DICOM, Incremental Learning, PACS, Pattern Recognition.

Abstract: Communication latency still hinders the adoption of Cloud computing paradigms in medical imaging environments where it could serve as a reliable technology to support repository outsourcing solutions or inter-institutional workflows, for instance. One way to overcome this is by implementing cache repositories and prefetching mechanisms. Nevertheless, such solutions are usually based on static rules that may inefficiently manage the cache storage capacity. For that reason, this paper compares a pattern recognition system using incremental learning versus batch learning, in order to assess which one could be more appropriately used in a medical imaging cache mechanism.

1 INTRODUCTION

Medical imaging is an important tool in medical practice, giving physicians valuable information for better diagnosis and treatment (Sylva, 2010, Rengier et al., 2010). Many medical imaging processes are supported by Picture Archiving and Communication Systems (PACS) (Huang, 2011), an umbrella term that embraces a set of technologies for acquisition, visualization, storage and distribution of medical imaging data (Valente et al., 2012). In order to do so, these systems rely on large IT infrastructures, comprising application servers, archives acquisition equipment (i.e. modality equipment) and network equipment, communicating through the Digital Imaging and Communications in Medicine (DICOM) standard (ACR-NEMA, 2011b).

Traditional PACS solutions are hosted in the healthcare institution and all equipment is connected in the physical layer by a Local Area Network (LAN). Nevertheless, with the proliferation of high-speed Internet connections, the PACS concept has broadened its horizons, embracing:

- **Infrastructure outsourcing** (Philbin et al., 2011, Chen and Sion, 2011), i.e. the moving of IT infrastructure from indoors to outdoors, reducing maintenance costs.
- **Institutional collaboration** (Marques Godinho et al., 2014, Sutton, 2011, Silva et al., 2013b), facilitating the remote access to examinations and reports (Costa et al., 2009) in response to the dispersion of patient's data that arises from their mobility between different institutions (Viana-Ferreira and Costa, 2014a).

In both cases, communication latency is a critical issue, because it is typically higher than in intra-institutional processes (Viana-Ferreira and Costa, 2014a).

This is emphasized by the nature of the data, since medical imaging examinations may reach volumes of hundreds of megabytes for some modalities (Yakami et al., 2011). To minimize this problem, there are two possible solutions: (1) cache, i.e. a small but fast repository hosted near the final consumer that stores a portion of the main repository data; and (2) prefetching, which consists in requesting images before users request them. However, the effectiveness of these solutions is highly dependent on their capability of predicting which data will be needed next. Most current solutions are based on static rules over specific parameters (Huang, 2011, Bui et al., 2001), considering the specific workflow of each institution. This tailoring constitutes a drawback of these solutions, as they may not be suitable for more dynamic scenarios, leading to a degradation of service quality or even denial of service in particular sets of conditions.

For these reasons, a pattern recognition solution

Viana-Ferreira C., Matos S. and Costa C..
Incremental Learning Versus Batch Learning for Classification of User's Behaviour in Medical Imaging.
DOI: 10.5220/0005219704310438
In *Proceedings of the International Conference on Health Informatics* (HEALTHINF-2015), pages 431-438
ISBN: 978-989-758-068-0

that could automatically adapt itself to user's behaviours and institutional workflows, while also giving special attention to situations that may become critical for the overall performance of the system, would be desirable. In this setting, this paper presents a comparison between incremental learning and batch training for such pattern recognition approach.

2 MEDICAL IMAGING LABORATORIES

Most medical imaging services, from acquisition and storage, to transmission and visualization of medical imaging data are manage by Picture Archiving and Communication Systems (PACS) (Huang, 2011). Most systems of this kind are intrinsically complex, as they are responsible for handling all medical imaging data of a healthcare institution. Figure 1 shows an example of a PACS instance, composed by several modalities (i.e. image acquisition devices), the repository, a PACS server, workstations, printers and a Radiology Information System (RIS), all linked by a Local Area Network (LAN).

2.1 Digital Imaging and Communications in Medicine (DICOM)

Currently in version 3, the DICOM standard (ACR-NEMA, 2011b) is composed of twenty parts, defining a wide set of processes related to medical imaging, such as: network communication layers, service commands, encoding and data structures and visualization processes (Pianykh, 2011). The wide range of processes and its versatility made DICOM a well-accepted standard, being currently followed by virtually all medical imaging equipment. For this work, the most important aspect of the standard is related to the DICOM services (ACR-NEMA, 2011a), including:

- C-Store service is used to push DICOM objects into the repository.
- C-Get is for requesting objects by their identifiers from an archive.
- C-Move service is for copying an object from one repository into another.
- C-Find is used to query an archive about objects that match a query.

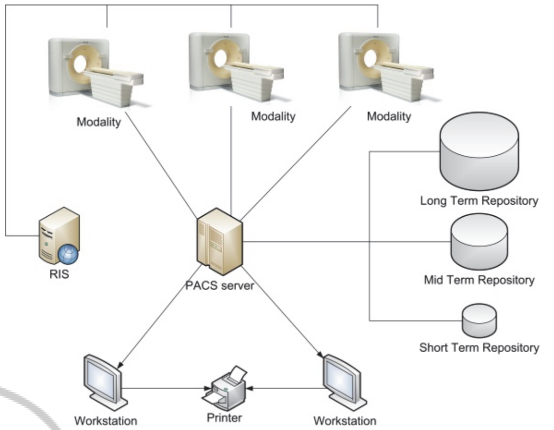

Figure 1: Typical PACS instance architecture (Viana-Ferreira and Costa, 2014a).

2.2 Federation of Healthcare Institutions

With the spread of fast Internet connections, the PACS concept has reached different settings. One example of this is the PACS described in (Silva et al., 2013a).

This PACS serves two healthcare institutions at the same time: Institution A and Institution B. While both institutions have image acquisition devices, only Institution A has a repository that stores all produced images. This strategy allowed the institutions to reduce costs with IT infrastructure, while promoting medical cooperation. However, the main drawback of this architecture is related to the quality of the service provided to Institution B. While Institution A accesses data via a fast LAN, Institution B must rely on an Internet connection, which is significantly slower than the LAN of Institution A, to access the repository. This means that any object exchange between the institutions will be limited by the upstream bandwidth, which in this case is 12 megabits per second. Although this is already able to provide a satisfactory quality of service, it cannot compete with the quality of service provided by traditional indoor solutions (Philbin et al., 2011), hindering the adoption of these federated approaches.

One way to reduce this problem is to endow Institution B with a cache and a prefetching mechanism that populated the cache with examinations that will be needed in a close future.

3 CACHE AND PREFETCHING

Cache is a small and fast repository that is used to hide the impact of communication latency, by temporarily storing objects that more likely will be needed soon. The population of this mechanism can be carried out in two ways: (1) in a passive mode, in which the cache is populated with the last used objects; (2) in a more active way, recurring to prefetching, i.e. predicting the objects that will be needed and requesting them before the users do. Either way, caches have a limited capacity, which leads to the need of discarding some stored objects when they are full. For this reason, one characteristic of cache repositories is their replacement policy, which discards the objects that are less probable of being needed (Smith, 1982). There are numerous cache replacement policies, the most traditional ones being: Least Recently Used (LRU) (Ali et al., 2011), by size (Williams et al.), First In First Out, by predicting when they will be needed and discarding the last ones (Jaleel et al., 2010), Least Frequently Used (LFU) (Podlipnig and Boszormenyi, 2003), by a decision function (Cao and Irani, 1997) and randomized (Psounis and Prabhakar, 2001).

Healthcare institutions may store such huge amounts of medical imaging data that it becomes financially unfeasible to make all data accessible at the best quality of service. For that reason, as depicted in Figure 1, they store all data in long-term repositories and only data that is more likely to be needed is replicated in faster repositories, i.e. mid-term and short-term repositories (Huang, 2011).

In this environment, prefetching is traditionally carried out through static rules over predefined parameters (Huang, 2011). Nevertheless, such solutions are usually especially designed for each situation or are too generic, causing the prefetching of too many objects and overloading the network with useless traffic. As an example, in (Bui et al., 2001), a prefetching mechanism with static rules based on multiple information sources is describes. The tests carried out by the authors indicated a recall of 100%, but only 50% of precision. This means that, although all needed data was prefetched, only half of the prefetch data were relevant.

The authors believe that machine learning and pattern recognition can lead to more effective cache and prefetching mechanisms. Nevertheless, it is a relatively unexplored field, with only residual references found in the literature. An example is the work described in (Liu Sheng et al., 2000), in which neural networks and decision trees were tested to predict which patient's images would be needed. However, we did not find any solution that took into account distinct usage patterns.

4 PATTERN RECOGNITION

Pattern recognition has been an active research field for the last decades (Pal and Pal, 2001) and it consists on the development of algorithms for automatic decision-making processes (Maji and Pal, 2011), using data to infer patterns (Yegnanarayana, 2009, Duda et al., 2012). This has been applied in a wide range of scenarios, such as: rivers bio-assessment (Feio et al., 2013), computer-aided diagnosis (Ramírez et al., 2013), content based image retrieval (Valente et al., 2013), stock market index prediction (Guresen et al., 2011) and computer vision (Chen et al., 2010).

Pattern recognition embraces a set of tasks, such as: pattern association, pattern classification, pattern mapping, pattern grouping, and feature mapping, among others (Yegnanarayana, 2009). In this article, we are focused in the pattern classification problem that consists on the use of a set of patterns and their labels and finding the distinctions between patterns of distinct labels (Duda et al., 2012).

4.1 Artificial Neural Networks

One of the most well-known machine learning methods for pattern classification is the artificial neural networks (ANN) (Yegnanarayana, 2009). Due to its versatility, its ability to detect nonlinear relationships between variables and because of being able to update with new samples, the algorithm used in this work is based on ANN.

Basically, it consists on a group of processing units (or neurons) that are linked in a determined way. One of the possible topologies is the multilayer perceptron (MLP), where the processing units are organized in layers, and usually each one of them receives the output of all processing units of the previous layer. For example, Figure 2 shows a representation of a MLP that receives 4 inputs and returns 2 outputs, having one hidden layer with 5 neurons.

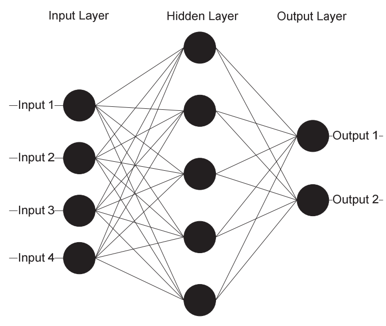

Figure 2: Representation of a multilayer perceptron.

4.2 Models, Classes and Features

For this work, the authors applied a previously developed algorithm that predicts which class of objects will be needed, considering the C-Finds sent to the PACS server. This algorithm considers four distinct usage pattern classes:

- **Pattern revising (class 1):** user is revising multiple studies of a single patient, for instance in a patient's appointment or when a clinician is evaluating the history of a patient.
- **Modality revising (class 2):** user is revising studies of a specific modality in a time window.
- **Inconsequent query (class 3):** this pattern class is for some queries which do not result in C-Move requests. One example of this is when a user erroneously introduces the search parameter. Another example is produced by some DICOM viewers that repeatedly send C-Find requests to refresh the interface.
- **"Other usage" (class 4):** this usage pattern is for all usage patterns that do not relate directly to the healthcare provision service itself in healthcare institutions, for instance, a data auditing.

The algorithm taken as the basis of this work uses five MLPs for each workstation in the PACS:

- Four MLPs, one for each class, with 26 input perceptrons, 250 perceptrons in the hidden layer and 1 output perceptron. As features, these MLPs uses three kinds of features: (1) time features that describe the pattern according to its temporal location; (2) history features that describe the pattern according to the user history until the moment of the C-Find of this pattern; and (3) the type of query.
- One MLP with 4 input perceptrons, 20 perceptrons in the hidden layer and 4 output perceptrons. Since the other four MLPs are trained independently, this MLP is used to take into account the outputs of the others and reach a conclusion about the actual class.

5 EXPERIMENTAL PROCEDURE

In this work, we compared the use of incremental learning to the use of a previously trained model (i.e. batch learning) when classifying the usage pattern when a C-Find request is detected. The objective of this pattern recognition step is to help infer which set of objects will more likely be requested afterwards, which in turn would allow developing and improving cache replacement and prefetching mechanisms.

5.1 Oracle

The oracle is a module that provides the classification of previous usage patterns, based on information about C-Move requests produced after the C-Find requests. This is a key component of the system, since it gives the actual classification of the patterns, to be used in training and updating the models that will then be applied online.

The labelling of the patterns is carried out in the following way:

- If there is only one C-Move between two C-Finds, it uses also the previous and the next patterns.
- If no studies were requested between two C-Finds, this patter is classified as "Inconsequent query" (class 3).
- If (almost) all requested studies are from the same patient, then the pattern is assigned as "Patient revising" (class 1).
- If it does not pass the previous test and (almost) all requested studies are of the same modality, then the pattern is assigned as "Modality revising" (class 2).
- In case a pattern seems ambiguous, i.e. if considering only the first C-Move requests the pattern would be assigned as one class, but if considering only the last C-Move requests the pattern would be assigned as other class, then the oracle will only consider the first C-Move requests.
- If the pattern failed all previous evaluations then the pattern is assigned as class 3.

5.2 Real-world Dataset

The real-world datasets is divided in two parts: a XML file and an index. The XML file contains anonymized information about 5186 DICOM messages that were sent to and from the PACS server in a period of roughly 3 months, while the index has data about the studies stored in the clinics' database. In both parts of the dataset, data was anonymized using hash functions to guarantee patient's privacy, while enabling the reproduction of the queries and their respective results.

After processing the messages with the oracle, we concluded that the real-world dataset consists of 17% patterns of class 1; 4% patterns of class 2; 29% patterns of class 3; and 50% patterns of class 4.

5.3 Synthesized Dataset

Ideally, we would use only real-world data for the tests, nevertheless, due to bureaucratic and ethical issues real-world datasets are not easy to obtain. Moreover, even when they are obtained, the range of distinct situations is usually limited.

In order to complement the results obtained with the real-world dataset, we have used a synthesizer of DICOM traffic based on behaviour profiles (Viana-Ferreira and Costa, 2014b). With this tool we have simulated the behaviour of three workstations in a one-year period:

- **Workstation A** has a regular behaviour along the experiment.
- **Workstation B** is most exclusively used to review studies of a given modality in the first six months, but in the following six months is used also for patient appointments.
- **Workstation C** behaves without a notion of timetable, being used indistinctly along time. This represents a workstation in a volatile scenario.

Table 1 shows the distribution in percentage of the samples among the 4 classes in the three simulated workstations. It also includes the distribution of the whole dataset with the three workstations combined.

Table 1: Distribution of the samples among the distinct classes (1 – Patient revising, 2 – Modality revising, 3 – Inconsequent query and 4 – Other usages) in the 3 synthetized workstations (A, B and C).and in the whole synthesized dataset (Combined).

Class	Workstation A	Workstation B	Workstation C	Combined
1	20.7 %	12.3 %	63.5 %	37.5 %
2	4.5 %	72.7 %	24.3 %	42.3 %
3	7.2 %	3.8 %	4.1 %	4.3 %
4	67.6 %	11.2 %	8.1 %	15.9 %

5.4 Experimental Tests and Discussion

The experimental tests were done with the two datasets: real-world and synthesized ones. Each one was tested under 4 distinct scenarios:

- **Train 25:** batch learning with the 25% oldest samples of the dataset, while the other 75% of the dataset is used to test them.
- **Train 50:** the 50% oldest samples were used to train the models, while the others were used to test them.
- **Train 75**: 75% oldest sample were used to train the models while the others 25% were used to test the models.
- **Incremental Learning:** only the first week was used to train the models. From then on, the samples of each week were used, firstly, to test the models and, secondly, to update the model.

Each test condition was executed ten times and the results averaged to mitigate the noise caused by random initialization of MLPs. In order to compare the performance of each learning method we chose two measures: (1) the accuracy which is a ratio between the number of times the prediction was right and the total number of samples; and (2) the F-Measure of each class which is calculated as shown in equation 1.

$$F - Measure(C) = \frac{2\,TC}{2 \times TC + FC + F\bar{C}} \qquad (1)$$

In equation 1, F-Measure(C) is the F-Measure of class C, TC is the number of times the method predicted the class C correctly, FC is the number of times the method wrongly predicted the sample belonged to class C, and $F\bar{C}$ is the number of times the method wrongly labeled the sample as not belonging to class C.

6 RESULTS

In this section, the results of the experiments are presented, divided in real-world and synthesized datasets.

6.1 Real-world Dataset

Figure 3 is a graph with the accuracy and the f-measures for each class, in each testing condition with the real-world dataset.

From the analysis of the graph, we can conclude that the algorithm behaved worse for the least representative class, i.e. class 2, in every condition. Nevertheless, it must be highlighted that this is a consequence of very few data about that usage

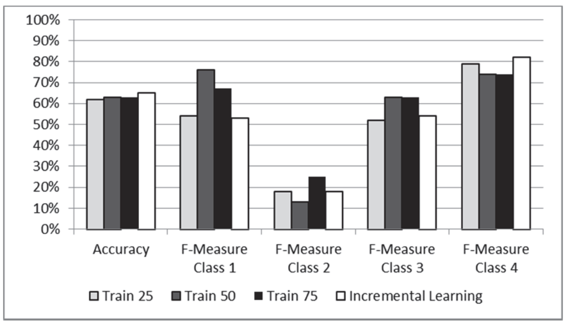

Figure 3: Graph with accuracy and F-measures for the real dataset.

pattern class, only 4% of class representation in the dataset. Besides that, there are no evident differences between the different testing conditions, i.e. Train 25, Train 50, Train 75 and Incremental learning. This can be explained by the limited time window represented by the dataset that did not include situations like changes of medical staff in the institution.

6.2 Synthesized Dataset

Figure 4 contains graphs of the f-measures for each class and the overall accuracy. The first three graphs, i.e. (a), (b) and (c), show the evaluation measures for each workstation of the synthesized dataset, while the last one, i.e. (d), shows the evaluation measures for the whole synthesized dataset.

From the analysis of the graphs, emphasis to the f-measure for class 1 in workstation B, where all training/testing scenarios had less than 15%, while the incremental learning achieved more than 50%. This can be explained by the nature of the synthesized data for this workstation, with a change in behavior during the experiment. The results clearly show that the incremental learning was the only training method capable of adapting the classifiers for this situation. Moreover, we can conclude that incremental learning was only worse than the batch learning conditions in workstation C which represents a very volatile scenario.

6.3 Overall Discussion

For what concerns accuracies, all testing conditions demonstrated to achieve roughly the same accuracy. Nevertheless, the Train 75 scenario was slightly better for the synthesized dataset, while the Incremental Learning was the best for the real-world dataset.

Concerning f-measures, incremental learning has proven to be more effective in classes with less representation in the dataset, while only slightly worse for classes with more representation in the dataset.

Nevertheless, it must be highlighted that the results of the incremental learning includes all predictions of the model, starting from the second week of data, when the models were in a very immature state. This means that with only a slight degradation of performance, we could launch the solution with only one week of data, instead of 3 months, which is represented by the Train 25 condition in the synthesized dataset.

7 CONCLUSIONS

In this paper, we tested a pattern recognition system that is based on machine learning for classification

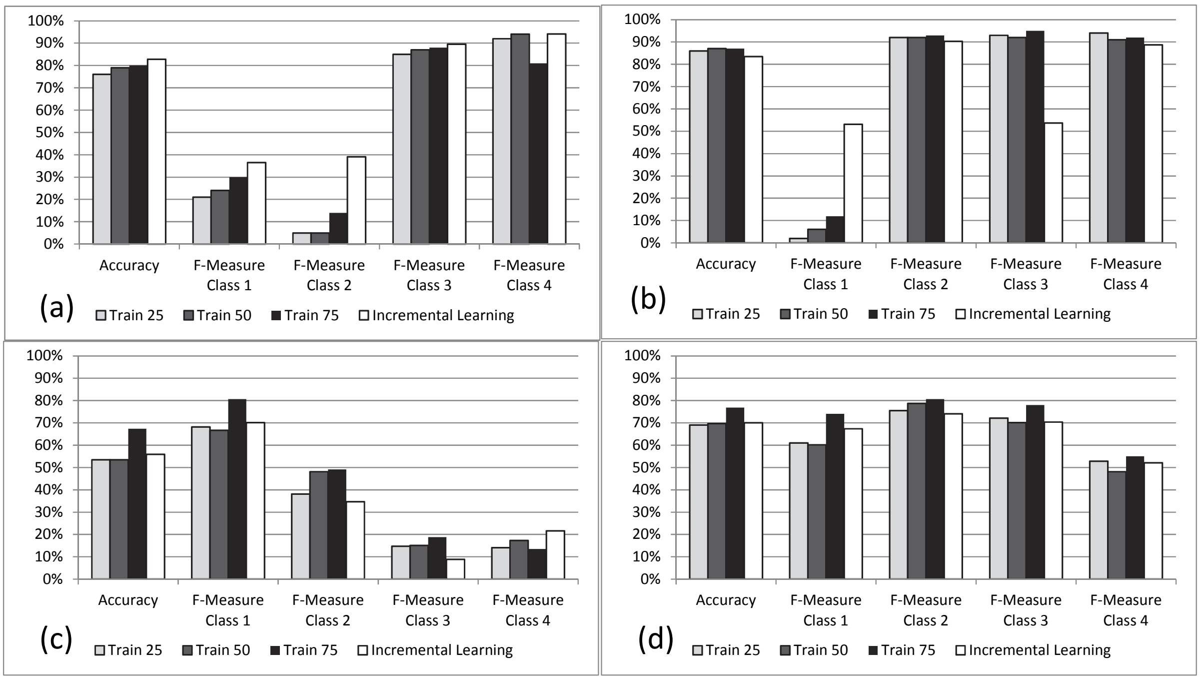

Figure 4: Graphs with the accuracy and F-measures for each synthesized workstation and for the combined dataset: (a) workstation A; (b) workstation B; (c) workstation C; (d) combined.

of users' behaviours.

The tested aimed to compare incremental learning with batch learning conditions, to assess if incremental learning is advantageous or not for this scenario.

We have concluded that despite of a minor degradation of the results in some cases, incremental learning is advantageous for pattern recognition since it has a smaller length of time for deployment. Besides, even the slight degradation of performance may be explained with the premature start of result extraction from the incremental learning testing condition.

Based on these results, as future work the authors will use incremental learning for the pattern recognition algorithm that aims at giving information to prefetching and cache replacement agents about which subset of images will be probably needed in a close future.

ACKNOWLEDGEMENTS

This work has received support from the EU/EFPIA Innovative Medicines Initiative Joint Undertaking (EMIF grant n° 115372). Carlos Viana-Ferreira is funded by the FCT grant SFRH/BD/68280/2010. Sérgio Matos is funded under the FCT Investigator programme.

REFERENCES

ACR-NEMA 2011a. Digital Imaging and Communications in Medicine. *Part 4: Service Class Specifications.* Rosslyn, VA, USA: NEMA.

ACR-NEMA 2011b. Digital Imaging and Communications in Medicine (DICOM). Rosslyn, VA: National Electrical Manufacturers Association.

Ali, W., Shamsuddin, S. M. & Ismail, A. S. 2011. A survey of Web caching and prefetching. *Int. J. Advance. Soft Comput. Appl,* 3**,** 18-44.

Bui, A. A., Mcnitt-Gray, M. F., Goldin, J. G., Cardenas, A. F. & Aberle, D. R. 2001. Problem-oriented prefetching for an integrated clinical imaging workstation. *Journal of the American Medical Informatics Association,* 8**,** 242-253.

Cao, P. & Irani, S. Cost-aware WWW proxy caching algorithms. Proceedings of the 1997 USENIX Symposium on Internet Technology and Systems, 1997.

Chen, C.-H., Pau, L.-F. & Wang, P. S.-P. 2010. *Handbook of pattern recognition and computer vision*, World Scientific.

Chen, Y. & Sion, R. 2011. To cloud or not to cloud?: musings on costs and viability. *Proceedings of the 2nd ACM Symposium on Cloud Computing.* Cascais, Portugal: ACM.

Costa, C., Freitas, F., Pereira, M., Silva, A. & Oliveira, J. L. 2009. Indexing and retrieving DICOM data in disperse and unstructured archives. *International Journal of Computer Assisted Radiology and Surgery,* 4**,** 71-77.

Duda, R. O., Hart, P. E. & Stork, D. G. 2012. *Pattern classification*, John Wiley & Sons.

Feio, M. J., Viana-Ferreira, C. & Costa, C. 2013. Combining Multiple MAChine Learning Algorithms to Predict Taxa Under Reference Conditions For Streams Bioassessment. *River Research and Applications***,** n/a-n/a.

Guresen, E., Kayakutlu, G. & Daim, T. U. 2011. Using artificial neural network models in stock market index prediction. *Expert Systems with Applications,* 38**,** 10389-10397.

Huang, H. 2011. *PACS and imaging informatics: basic principles and applications*, Wiley-Blackwell.

Jaleel, A., Theobald, K. B., Simon C. Steely, J. & Emer, J. 2010. High performance cache replacement using re-reference interval prediction (RRIP). *Proceedings of the 37th annual international symposium on Computer architecture.* Saint-Malo, France: ACM.

Liu Sheng, O. R., Wei, C.-P., Hu, P. J.-H. & Chang, N. 2000. Automated learning of patient image retrieval knowledge: neural networks versus inductive decision trees. *Decision Support Systems,* 30**,** 105-124.

Maji, P. & Pal, S. K. 2011. *Rough-Fuzzy Pattern Recognition: Applications in Bioinformatics and Medical Imaging*, John Wiley & Sons.

Marques Godinho, T., Viana-Ferreira, C., Bastiao Silva, L. & Costa, C. 2014. A Routing Mechanism for Cloud Outsourcing of Medical Imaging Repositories. *Biomedical and Health Informatics, IEEE Journal of,* PP**,** 1-1.

Pal, S. K. & Pal, A. 2001. *Pattern recognition: from classical to modern approaches*, World Scientific.

Philbin, J., Prior, F. & Nagy, P. 2011. Will the Next Generation of PACS Be Sitting on a Cloud? *Journal of Digital Imaging,* 24**,** 179-183.

Pianykh, O. S. 2011. *Digital imaging and communications in medicine (DICOM)*, Springer.

Podlipnig, S. & Boszormenyi, L. 2003. A survey of Web cache replacement strategies. *ACM Comput. Surv.,* 35**,** 374-398.

Psounis, K. & Prabhakar, B. A randomized Web-cache replacement scheme. INFOCOM 2001. Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE, 2001 2001. 1407-1415 vol.3.

Ramírez, J., Górriz, J. M., Salas-Gonzalez, D., Romero, A., López, M., Álvarez, I. & Gómez-Río, M. 2013. Computer-aided diagnosis of Alzheimer's type dementia combining support vector machines and discriminant set of features. *Information Sciences,* 237**,** 59-72.

Rengier, F., Mehndiratta, A., Tengg-Kobligk, H., Zechmann, C. M., Unterhinninghofen, R., Kauczor, H.

U. & Giesel, F. L. 2010. 3D printing based on imaging data: review of medical applications. *International Journal of Computer Assisted Radiology and Surgery,* 5**,** 335-341.

Silva, L. A. B., Costa, C. & Oliveira, J. L. 2013a. An agile framework to support distributed medical imaging scenarios. *IEEE International Conference on Healthcare Informatics 2013 (ICHI 2013).* Philadelphia, USA.

Silva, L. B., Costa, C. & Oliveira, J. 2013b. DICOM relay over the cloud. *International Journal of Computer Assisted Radiology and Surgery,* 8**,** 323-333.

Smith, A. J. 1982. Cache memories. *ACM Computing Surveys (CSUR),* 14**,** 473-530.

Sutton, L. N. 2011. PACS and diagnostic imaging service delivery—A UK perspective. *European Journal of Radiology,* 78**,** 243-249.

Sylva, P. 2010. A Situation Analysis on PACS prospects for a Developing Nation. *Sri Lanka Journal of Bio-Medical Informatics,* 1**,** 112-117.

Valente, F., Costa, C. & Silva, A. 2013. Dicoogle, a Pacs Featuring Profiled Content Based Image Retrieval. *PLoS ONE,* 8**,** e61888.

Valente, F., Viana-Ferreira, C., Costa, C. & Oliveira, J. L. 2012. A RESTful Image Gateway for Multiple Medical Image Repositories. *Information Technology in Biomedicine, IEEE Transactions on,* 16**,** 356-364.

Viana-Ferreira, C. & Costa, C. 2014a. Challenges of Using Cloud Computing in Medical Imaging. *In:* Ramachandran, M. (ed.) *Advances in cloud Computing Research.*

Viana-Ferreira, C. & Costa, C. 2014b. DICOM Traffic Generator based on behavior profiles. *In:* IEEE (ed.) *IEEE-EMBS International Conferences on Biomedical and Health Informatics.* Valencia, Spain.

Williams, S., Abrams, M., Standridge, C. R., Abdulla, G. & Fox, E. A. Removal Policies in Network Caches for World-Wide Web Documents.

Yakami, M., Ishizu, K., Kubo, T., Okada, T. & Togashi, K. 2011. Development and Evaluation of a Low-Cost and High-Capacity DICOM Image Data Storage System for Research. *Journal of Digital Imaging,* 24**,** 190-195.

Yegnanarayana, B. 2009. *Artificial neural networks*, PHI Learning Pvt. Ltd.