

Online Learning Strategies for Classification of Static Data Streams

M. MILLÁN-GIRALDO J. S. SÁNCHEZ

Universitat Jaume I

Dept. Llenguatges i Sistemes Informàtics

Av. Sos Baynat s/n, 12071 Castelló de la Plana

SPAIN

{mmillan,sanchez}@uji.es

Abstract: This paper addresses the problem of online learning for static streaming data. The ultimate objective is to compare a number of very simple learning methods, mainly taken from the literature. We include a straightforward time-weighted strategy for forgetting obsolete objects from the reference set. Experiments are conducted on ten real data sets and using five different classifiers in order to identify which online learning model is the most suitable in terms of classifier performance.

Key-Words: Data mining, Online learning, Static streaming data, Forgetting

1 Introduction

Classification is one of the key tasks in many pattern recognition and data mining applications. The essence of classification is to use previously observed data to construct a model that is able to predict the categorical or nominal value of a dependent variable (the class) given the values of the independent variables (the features or attributes). Obtaining a high accuracy in classification is usually the primary goal. Another important objective is comprehensibility, which refers to the ability of a human expert to understand the classification model. The third aim is compactness, which relates to the size of the model. Classifiers usually try to trade-off these three objectives.

Most traditional learning algorithms assume the availability of a set of labelled (training) examples $T = \{e_1, e_2, \dots, e_n\}$. Each training example e_i is a pair formed by a feature vector \vec{x}_i and a discrete value y_i (class label) taken from a finite set Y . In many domains, however, collecting labelled training objects may be costly, time-consuming or even dangerous, while it is relatively easy to obtain unlabelled objects. This has provoked significant interest in semi-supervised learning [3] and other closely related areas (incremental learning, online learning, streaming data).

In a growing number of real applications, the data are not available as a batch but comes one object at a time (called *streaming data*). From the processing point of view, a data stream is an infinite flow of highly rapid generated objects that challenge our computing systems to store, process and transmit [7]. Examples of applications with data streams include

Internet peer-to-peer downloads, radar derived meteorological data, banking and credit transactions, classification of stock data, and intrusion detection in computer networks, for which it is not possible to collect all relevant input data before using the classification algorithm. Under these circumstances, the learning systems have to operate continuously (*online systems*) and process each data item in near-real time.

It has been argued that a good online classifier should have the following characteristics [4, 5, 12]: (i) Single pass through the data (the classifier must be able to learn from each data point without revisiting it); (ii) Limited memory and processing time (each data point should be processed in a constant time regardless of the number of points processed in the past); and (iii) Any-time-learning (if stopped at time t , the current classifier should be equivalent to a classifier trained on the batch data up to time t). These are also valid for online learning of streaming data, but yet another *desiderata* is generally accepted for this: At any time t in the data stream, we would like the per-item processing time, storage as well as the computing time to be simultaneously $o(N; t)$, being N the number of data items processed.

The focus of this paper is on evaluating several simple strategies for online versions of traditional learning algorithms applied to static streaming data. We empirically compare three existing strategies taken from the literature when applied to the 1-Nearest Neighbor classifier, a linear discriminant, a quadratic classifier, a decision tree, and the Naïve Bayes classifier. The aim of this experimental study is to conclude the most appropriate online learning model for streaming data.

2 Strategies for Online Learning

One of the most important decisions in designing an online classifier is how to maintain the reference set. The three trivial possibilities with respect to the memory size are [6]:

1. *Full memory*, in which the learner retains all training objects.
2. *Partial memory*, in which it retains only some of the training objects.
3. *No memory*, in which it retains none.

Some online classifiers where stored examples are used directly to form predictions should be in the partial memory group. Ideally, we need to keep a (small) reference set which, when deemed necessary, is expanded or shrunk within given limits. This means that there must be some mechanism to forget (remove) objects. How to forget is a difficult problem that can be tackled through several strategies.

- *Passive forgetting* [9] (also called time-weighted forgetting [11] and implicit forgetting [5, 11]) is based only on the time elapsed since the object was added to the reference set. It assumes that the importance of data decreases over time. Passive forgetting acts as a moving window where the reference set is the last data batch. Its size is a parameter of the algorithm.
- *Active forgetting* [9] (also referred as to explicit forgetting [5, 11]) implies that more information from data is used to decide which objects should be dropped. Two alternatives are possible for active forgetting:
 - (i) *Density-based forgetting* follows the intuition of the “life” game. If a region is too crowded, we sieve out some objects (locally-weighted forgetting [11]). On the other hand, if a region is too distant, and not providing nearest neighbors, it is removed altogether [11]. In the jargon of data editing, the former strategy corresponds to condensing, while the latter corresponds to editing.
 - (ii) *Error-based forgetting* is perceived as the most successful of the forgetting heuristics [1, 9, 11]. In this case, each object in the reference set has a classification record. The more streaming data it labels correctly, the stronger its record becomes. The objects with weak records are cleared at regular intervals.

2.1 The models for the experiment

The first strategy we will here experiment with is a full memory model, in which all new objects have to be incorporated into the reference set. The other two lie in the group of partial memory, that is, only some new objects will be added to the reference set. In the case of partial memory, one model employs a passive forgetting strategy in order to remove “obsolete” examples from the reference set.

1. *All*, which corresponds to the full memory option. It assumes that all new examples are valid and equally important. It is clear that the use of this model may result in a huge reference set, making impossible its practical application in most real problems.
2. *Every n objects*, in which every n 'th object will be added to the reference set to retrain the classifier. It considers that the ‘All’ strategy will retrain too often. On the other hand, this model overcomes the storage problem of the ‘All’ approach.
3. *Window of fixed size*. The reference set will be of fixed size with the last data, assuming that the most relevant information is at the last processed objects. Besides, the a priori class distributions are kept for each reference set in order to avoid a class to be emptied.

Although very simple strategies for online learning, they will allow to compare the partial memory and full memory approaches in a static streaming data scenario.

3 Experiment

In this section we present the experiments carried out in order to compare the simple online learning strategies previously described. The aim of this empirical study is to analyze the behavior of each pair (online learning strategy, classifier).

3.1 The data sets

Twelve real data sets were employed in the present experiment (a summary can be seen in Table 1). Data were normalized in the range $[0, +1]$. All features (attributes) were numerical and there were no missing values.

Although there is no strict guideline about what a sufficient data size is, the common wisdom [8] is that the size of the training data should be around $10 \times d \times c$, where d is the number of features and

Table 1: Characteristics of the real data sets used in the experiment

Data set	Features	Classes	Objects	Source
iris	4	3	150	UCI ¹
wine	13	3	178	UCI
crabs	6	2	200	Ripley ²
sonar	60	2	208	UCI
thyroid	5	3	215	UCI
voice3	10	3	238	Library ³
breast	9	2	277	UCI
intubation	17	2	302	Library
liver	6	2	345	UCI
spect	44	2	349	Library
laryngeal3	16	3	353	Library
australian	42	2	690	UCI

¹UCI [2]

²Ripley [10]

³Library http://www.informatics.bangor.ac.uk/~kuncheva/activities/real_data_full_set.htm

c is the number of classes in a problem. Our small initial reference set was of size $1 \times d \times c$.

The experimental set-up was as follows:

- 100 runs were carried out with 90% of the data used for training and 10% used for testing. The splits were done using stratified sampling.
- From each training part of the data, a random stratified sample of $N_l = 1 \times n \times c$ was taken as the initial labelled references set.
- The remaining part of the training data was used as the new coming online data. To simulate an i.i.d. sequence, the data was shuffled before each of the 100 runs.
- One point from the online data was fed to the system at a time. The point was processed according to the respective strategy to handle online learning. The classification error was evaluated on the testing set. In this way we created a “progression curve” which is the classification error as a function of the number of online objects seen by the classifier.
- The results were averaged across the 100 runs, giving a single progression curve for each data set.

The strategies we compare through this experiment are those introduced in Section 2, that is, ‘All’, ‘Every n objects’ ($n = 5$), and ‘Window of fixed

size’ (equal to the original size). All these have been applied to five different classifiers: the 1-Nearest Neighbor classifier (1-NN), a linear discriminant, a quadratic classifier, a decision tree, and the Naïve Bayes classifier.

For each of the 100 runs of the experiment, all methods received the same partitions of the data into initial, online and testing sets. The online data was presented to all methods in the same order.

3.2 The results

The results for the iris, wine, liver, spect, laryngeal3, and wbc are displayed in Figure 1, whereas those for breast, intubation, crabs, sonar, thyroid, and australian are in Figure 2. The x-axis corresponds to the number of processed unlabelled samples and the y-axis is the progression of the classification error, evaluated on the testing sets and averaged across 100 runs.

The graphs are especially meant to visualize the direction of the error curves rather than the details. From these, several typical patterns can be observed. The error rates have different trends when comparing the ‘All’ and ‘Every n objects’ strategies with the ‘Window of fixed size’ approach. While using the ‘All’ and ‘Every n objects’ methods gives a decrease in error as more unlabelled samples are processed, the third model produces the opposite effect.

Probably, the degradation of the window-based approach is mainly due to the fact that the reference set is gradually updated with new objects labelled by the own classifier; if this errs, then misclassified objects will be added to the reference set. As stated in Section 2, passive forgetting is based only on the time elapsed since the object was added to the reference set. This means that no test will be used to evaluate the goodness of each new object labelled and consequently, noise and errors may be incorporated into the reference set. Thus the quality of the reference set gradually deteriorates with more unlabelled objects being processed.

The quadratic classifier clearly shows the behavior just described. Contrary to the case of the window-based model, the ‘All’ and ‘Every n objects’ strategies start from a high error rate, but this rapidly decreases with processing of new objects. This pattern, however, is not matched on all data sets. There are two databases (liver and crabs) where the errors of the quadratic classifier increase with the processing of new samples by means of the ‘All’ and ‘Every n objects’ approaches.

Although the general behavior of the rest of classifiers is similar to that of the quadratic, it worths pointing out that the error rates of 1-NN, the linear discriminant, the decision tree, and the Naïve Bayes

classifier keep quite steady along time in the case of the 'All' and 'Every n objects' models. It seems that in general, small changes in the reference set do not affect the classifier performance. Nevertheless, when the whole set is updated as a result of the forgetting mechanism, it produces a significant degradation in performance.

As a further confirmation of the comments on these results, we should compare the error rates of each model at the initial time $t = 0$ and at the end, when all the streaming data have already been seen. However, the lack of space does not allow to do this analysis in the present paper.

4 Conclusions and Further Extensions

In this paper, we have compared a number of simple strategies for online learning of streaming data. Two models belong to the group of partial memory (only some new objects are retained in the reference set), whereas the third is a full memory method (all new unlabelled samples are kept). Besides, one of the partial memory models includes a time-weighted forgetting strategy in order to remove "obsolete" objects from the reference set.

The empirical study has employed five classifiers with very different properties so that one can conclude which strategy is more suitable for each learning model. The experiments have shown that the error rates of the 'Window of fixed size' (partial memory with passive forgetting) approach increases with the processing of new samples, suggesting that a drastic update of the reference set may significantly deteriorate the classifier performance. From the results, it can be concluded that the inclusion of some forgetting mechanism will be especially useful in the case of non-stationary data streams.

Future work will focus on a more exhaustive study of a larger number of strategies for online learning, especially addressed to devise more elaborated forgetting methods. Another topic for further study will be to determine the optimal number of training examples to keep in the reference set. Also, extending the present study to dynamic streaming data constitutes one of the most important lines for future research.

Acknowledgements: This work has been supported in part by the Spanish Ministry of Education and Science under grants DPI2006-15542 and CSD2007-00018 (Consolider-Ingénio 2010).

References:

- [1] D. Aha, D. Kibler, and M.K. Albert, Instance-based learning algorithms, *Machine Learning* 6, 1991, pp. 37-66.
- [2] A. Asuncion and D. J. Newman, UCI Machine Learning Repository, School of Information and Computer Science, University of California, Irvine, CA, 2007. <http://www.ics.uci.edu/~mllearn/MLRepository.html>.
- [3] O. Chapelle, B. Schölkopf, and A. Zien, *Semi-Supervised Learning*, MIT Press, Cambridge, MA, 2006.
- [4] P. Domingos and G. Hulten, A general framework for mining massive data streams, *Journal of Computational and Graphical Statistics* 12, 2003, pp. 945-949.
- [5] F. J. Ferrer-Troyano, J. S. Aguilar-Ruiz, and J. C. Riquelme, Incremental rule learning and border examples selection from numerical data streams, *Journal of Universal Computer Science* 11, 2005, pp. 1426-1439.
- [6] M. A. Maloof and R. S. Michalski, Selecting examples for partial memory learning, *Machine Learning* 41, 2000, pp. 27-52.
- [7] S. Muthukrishnan, Data streams: algorithms and applications, *Foundations and Trends in Theoretical Computer Science* 1, 2005, pp. 117-236.
- [8] G. Nagy, Classifiers that improve with use, In *Proc. Conf. on Pattern Recognition and Multimedia*, Tokyo, Japan, 2004, pp. 79-86.
- [9] H. Nakayama and K. Yoshii, Active forgetting in machine learning and its application to financial problems, In *Proc. Intl. Joint Conf. on Neural Networks*, Como, Italy, 2000, pp. 123-128.
- [10] B. D. Ripley, *Pattern Recognition and Neural Networks*, Cambridge University Press, 1996.
- [11] M. Salganicoff, Density-adaptive learning and forgetting, In *Proc. 10th Intl. Conf. on Machine Learning*, Amherst, MA, 1993, pp. 276-283.
- [12] W. N. Street and Y. S. Kim, A streaming ensemble algorithm (SEA) for large-scale classification, In *Proc. 7th ACM SIGKDD Intl. Conf. on Knowledge Discovery and Data Mining*, Santa Barbara, CA, 2001, pp. 377-382.

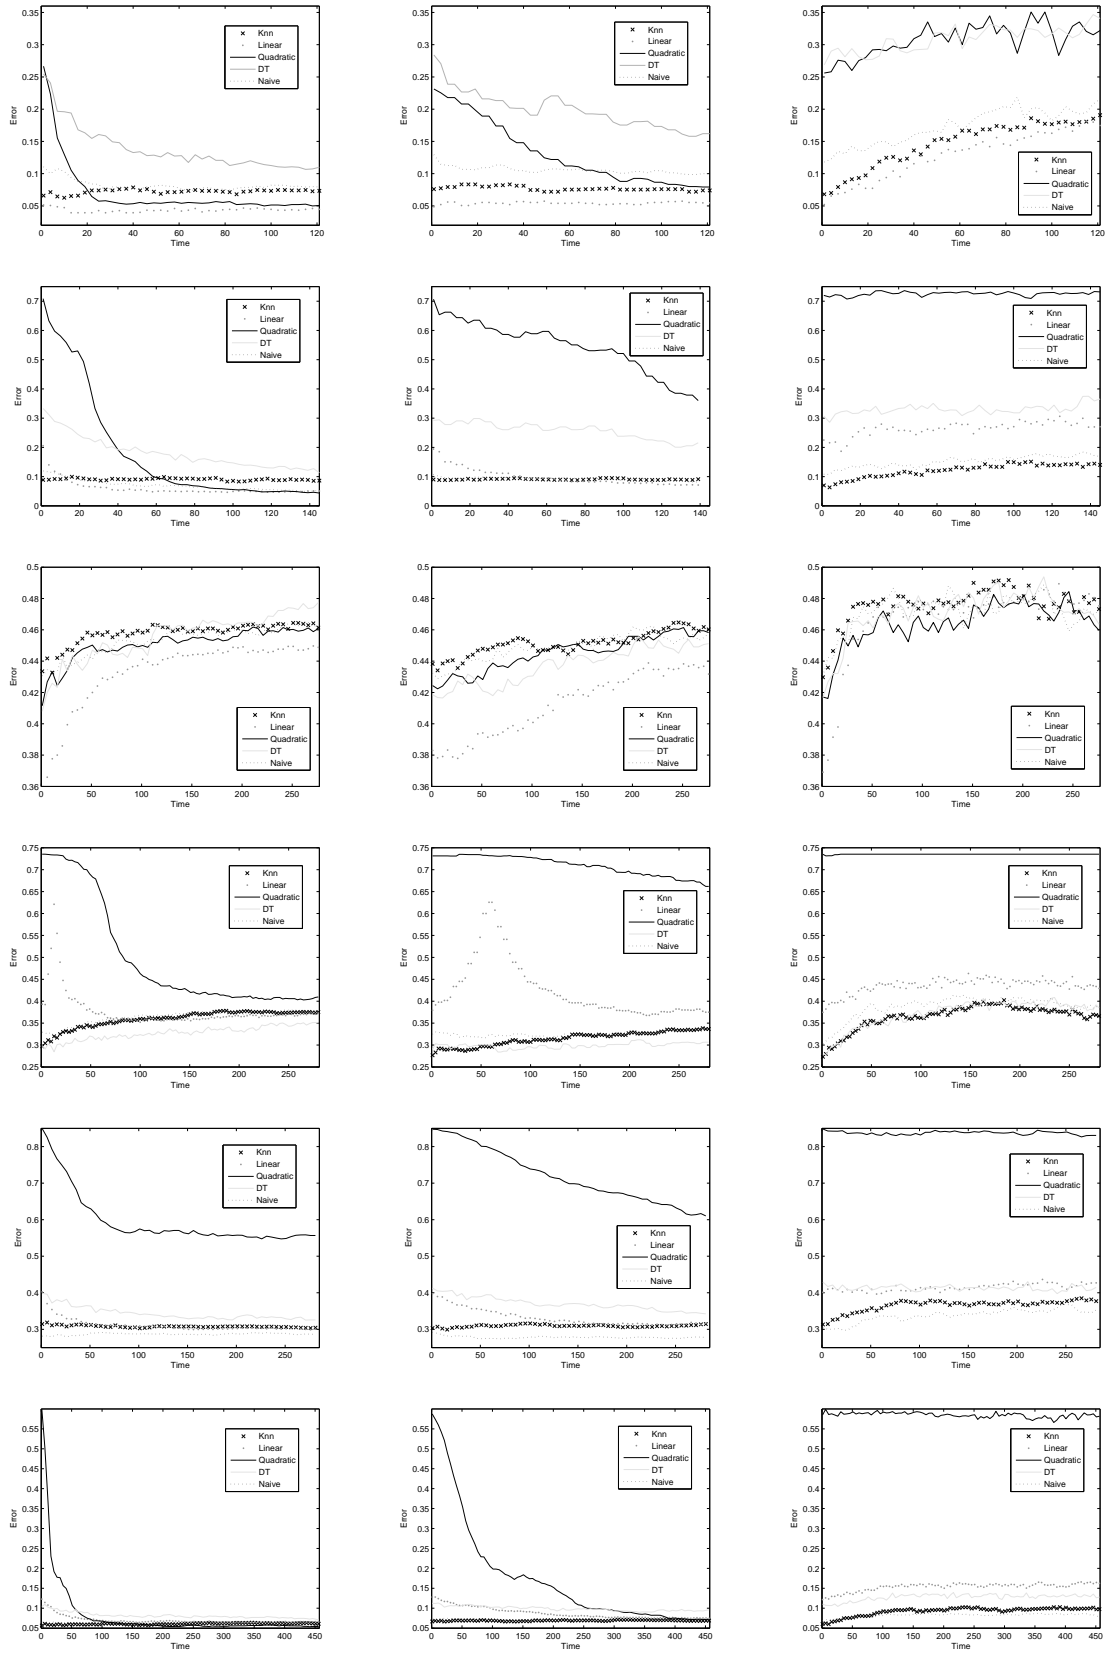


Figure 1: Error progression with sequential processing of new unlabelled data through the three online learning strategies: All (left), every $n = 5$ objects (middle), and window of fixed size (right). From top to bottom, the figures correspond to iris, wine, liver, spect, laryngeal3, and wbc databases.

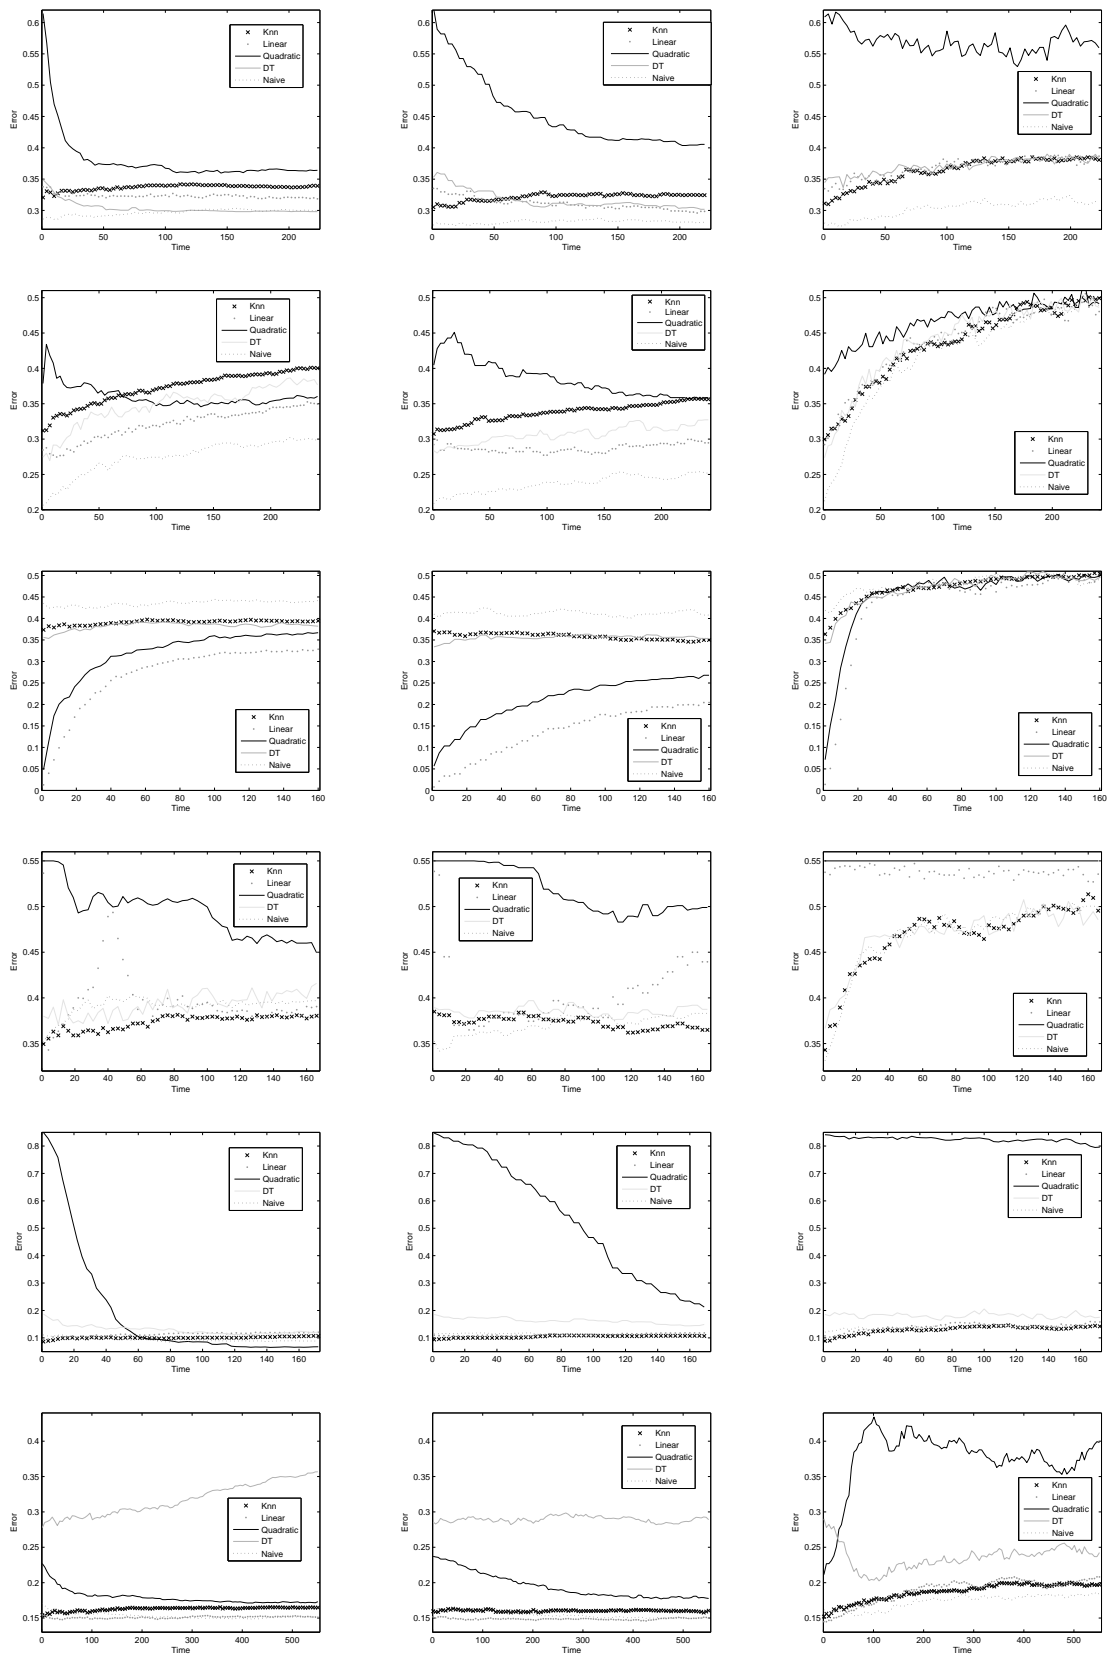


Figure 2: Error progression with sequential processing of new unlabelled data through the three online learning strategies: All (left), every $n = 5$ objects (middle), and window of fixed size (right). From top to bottom, the graphs are for breast, intubation, crabs, sonar, thyroid, and australian databases