

Active Learning using Adaptive Curiosity

Alexis Bondu*

Vincent Lemaire*

*France Telecom R&D TECH/EASY/TSI

2 avenue Pierre Marzin 22300 Lannion, France

(alexis.bondu)(vincent.lemaire)@orange-ftgroup.com

Abstract

Exploratory activities seems to be crucial for our cognitive development. According to psychologists, exploration is an intrinsically rewarding behaviour. That explains the autonomous and active development of children. The developmental robotics aim to design computational systems that are endowed with such an intrinsic motivation mechanism. There are possible links between developmental robotics and classical machine learning. Active learning strategies aim to the most informative examples and adaptive curiosity allows a robot to explore its environment in an intelligent way. In this article, the adaptive curiosity framework is reformulated in terms of active learning terminology, and compared directly to existing algorithms in this field. The main contribution of this article is a new criterion evaluating the potential interestingness of zones of the sensorimotor space.

1. Introduction and notation

Human beings develop in an autonomous way, carrying out exploratory activities. This phenomenon is an intrinsically motivated behaviors. Psychologists (White, 1959) have propose theory which explain exploratory behaviors as a source of self rewarding. Building such a robot is a great challenge of developmental robotics. The ambition of this field is to build a computational system that try to capture curious situations. Adaptive curiosity (Oudeyer and Kaplan, 2004) is one possility to aim this objective. This approach push a robot towards situations in which it maximizes its learning progress. The robot first spends time in situations that are easy to learn, then shifts progressively its attention to more difficult situations, avoiding situations in which nothing can be learnt.

This article does a bridge between developmental robotic and classical machine learning. Active learning strategies allow a predictive model to construct its training set in interaction with an expert. The learning starts with few labelled examples. Then the model selects examples (with no label) which considers the most informative and asks their associated output to the expert. The model learns faster thanks to active learning strategies, reaching the best performances using less data. These

approaches minimize the labeling cost inducted by the training of a model.

On the one hand, active learning brings into play a predictive model that explore the space of unlabelled examples, in order to find the most informative ones. On the other hand, adaptive curiosity allow a robot to explore its environment in an intelligente way, and tries to deal with the dilemma exploration / exploitation. This paper proposes to fit adaptive curiosity to supervised active learning. The organization of this paper is as follow : in section 2., adaptive curiosity is presented in a generic way and original implementation choices are described. The next section shows a possible implementation of adaptive curiosity for classification. The behavior of this strategy is examined on a toy example. Considering the obtained results, a new strategy of adaptive curiosity is defined in section 4.. This new strategy is then compared with two other active learning strategies. Finally, possible improvements of adaptive curiosity are discussed.

Notations : $\mathcal{M} \in \mathbb{M}$ is the predictive model that is trained with an algorithm \mathcal{L} . $\mathbb{X} \subseteq \mathbb{R}^n$ represents all possible input examples of the model and $x \in \mathbb{X}$ is a particular examples. \mathbb{Y} is the set of possible outputs (answers) of the model; $y \in \mathbb{Y}$ refers to a class label which is associated to $x \in \mathbb{X}$.

The point of view of selective sampling¹ is adopted (Castro et al., 2005) in this paper. The model observes only one restricted part of the universe $\Phi \subseteq \mathbb{X}$ which is materialized by training examples with no label. The image of a “bag” containing examples for which the model can ask for associated labels is usually used to describe this approach. The set of examples for which the labels are known (at one step of the training algorithm) is called L and the set of examples for which the labels are unknown is called U with $\Phi = U \cup L$ and $U \cap L = \emptyset$.

The concept which is learnt can be seen as a function, $f : \mathbb{X} \rightarrow \mathbb{Y}$, with $f(x_1)$ the desired answer of the model for the example x_1 and $\hat{f} : \mathbb{X} \rightarrow \mathbb{Y}$ the obtained answer of the model; an estimation of the concept. The elements of L and the associated labels constitute a training set T . The training examples are pairs of input vectors and desired labels such that $(x, f(x))$.

¹In practice, the choice of selective (Roy and McCallum, 2001) or adaptive (Singh et al., 2006) sampling depends primarily on the applicability where the model is authorized, or not, “to generate” new examples.

2. Adaptive Curiosity

2.1 General remarks

Adaptive curiosity (Oudeyer and Kaplan, 2004) is the ability for a robot to choose appropriate situations² according to its learning³. Indeed, the robot can be in a trivial state (or on the contrary, in a too difficult state) in which it can not learn anything. The objective of the robot is to maximize its progress carrying out the good actions in its environment.

Y. Nagai (Nagai et al., 2002) shows that a robot can learn faster considering situations where the difficulty progressively increases. The aim of adaptive curiosity is to make the robot autonomous in the choice of learnt situations. In the best case, the robot is interested by more and more difficult situations, and leaves situations for which there is nothing to learn.

The first intuition for robot's progress assessment is to compare successive performances. If the robot carries out a task in a better way than previously, one considers it makes progress. With such training rules, the robot can adopt aberrant behaviors.

To illustrate that point, Y. Nagai (Nagai et al., 2002) uses the example of a robot that learns to estimate its own position after a move. The robot believes to make big progress alternating a collision with an obstacle and immobility. Indeed, "immobility" is the action that allows the robot to predict its next position with the more important precision. Comparing this performance with the previous state (the collision), the progress is maximum.

Adaptive curiosity compares similar situations (and not successive situations) (Oudeyer and Kaplan, 2004) to measure robot's progress. Several sub-models which are specialized in certain types of situations are trained at the same time. The aim of adaptive curiosity is to make the robot autonomous in the discovery of the environment.

2.2 Generic Algorithm

Adaptive curiosity (Oudeyer and Kaplan, 2004) involves a double strategy. The first strategy makes a recursive partitioning of \mathbb{X} , the input space of the model. The second strategy selects zones to be fed with labelled examples (and to be split by recursive partitioning). It is an active learning as far as the selection of a zone defines the subset of examples which can be labelled (those which belong to the zone). Adaptive curiosity is described below in a generic way and illustrated by an algorithm.

The input space \mathbb{X} is recursively partitioned in zones (some of them are included in others). Each zone corresponds to a type of situations the robot must learn. Adaptive curiosity uses a criterion to select zones and preferentially splits area of input space \mathbb{X} in which learning improves. The main idea is to schedule learnt situations in order to accelerate the robot's training.

²A situation is defined as the state of the whole of sensors.

³The robot is learning to carry out a task in its environment.

Each zone is associated with a sub-model which is trained with examples belonging to the zone only. Sub-models are trained at the same time on disjointed examples sets. The partitioning of the input space is progressively realized, at the same time new examples are labelled. Just before the partitioning of a zone, the sub-model of the "parent" zone is duplicated in "children" zones. Duplicated sub-models continue independently their learning thanks to the examples which appear in their own zones.

Algorithm (1) shows the general steps of adaptive curiosity. It is an iterative process during which examples are selected and labelled by an expert. A first criterion chooses a zone to be fed with examples (stage A). The following stage consists in drawing an example in the selected zone (stage B). The expert gives the associated label (stage C) and the sub-model is trained with an additional example (stage D). A second criterion determines if the current zone must (or must not) be partitioned. In this case, one seeks (in the "parent" zone) adequate separations to create "children" zones (stage i). Lastly, the sub-model is duplicated into "children" zones (stage ii).

```
Given :
• a learning algorithm  $\mathcal{L}$ 
• a set  $\mathcal{M} = \{m_1, m_2, \dots, m_n\}$  of  $n$  predictive sub-models
•  $U = \{u_1, u_2, \dots, u_n\}$ ,  $n$  subsets of unlabelled examples
•  $L = \{l_1, l_2, \dots, l_n\}$ ,  $n$  subsets of labelled examples
•  $T = \{t_1, t_2, \dots, t_n\}$  the training subsets corresponding to sub-models, with  $t_i = \{(x, f(x))\} \forall x \in l_i$ 

 $n \leftarrow 1$ 
Repeat
  (A) Choose a sub-model  $m_i$  to be fed with examples
  (B) Draw a new example  $x^*$  in  $u_i$ 
  (C) Label the instance  $x^*$ ,  $t_i \leftarrow t_i \cup (x^*, f(x^*))$ 
  (D) Train the sub-model  $m_i$  thanks to  $\mathcal{L}$ ,  $U$  and  $t_i$ 
  If the split criterion is satisfied then
    (i) Separate  $l_i$  in two sub-sets  $l_j$  and  $l_k$  the most homogeneous as possible
    (ii) Duplicate  $m_i$  into two sub-models  $m_j$  and  $m_k$ 
    (iii)  $n \leftarrow n + 1$ 
  end If
until  $U \neq \emptyset$ 
```

Algorithm 1: Adaptive Curiosity

The main purpose of this algorithm is to seek interesting zones in the input space, at the same time the machine discovers data to learn. The algorithm chooses (as soon as possible) the examples belonging to the zones where there is possible progress. Five questions appear :

- How to decide if a zone must be partitioned?
- How to carry out the partitioning?
- How many "Children" zones?
- How to choose zones to be fed in examples?
- What kind of sub-models must be used?

2.3 Original choices (Oudeyer and al, 2004)

2.3.1 Partitioning

A zone must be partitioned when the number of labelled examples exceeds a certain threshold. Partitioned zones are those which were preferentially chosen during previous iterations. These zones are interesting to be partitioned when more populated. Associated sub-models have done important progress.

To cut a "parent" zone into two "children" zones, all dimensions of the input space \mathbb{X} are considered. For each dimension, all possible cut values are tested using the sub-model to calculate the variance of example's predictions (on both sides of the separation). During this stage, observable data Φ is used. This criterion⁴ consists in finding a dimension to cut and a cut value minimizing the variance. This criterion elaborates preferentially pure zones to facilitate the learning of associated sub-models. Another constraint is added by the authors, the cut has to separate labelled examples into two subsets whose cardinalities are about balanced.

2.3.2 Zones selection

At every iteration, the sub-model which most improves is considered as having the strongest potential of improvement. Consequently, adaptive curiosity needs an estimation of sub-model's progress. Firstly, performances of sub-models are measured on labelled data. The choice of a performance measure is required. Secondly, sub-model's performances are evaluated on a temporal window. The sub-model which realizes the most important progress is chosen to be fed with new examples uniformly drawn.

3. Implementation for classification

In this section the relevance of the adaptive curiosity approach is evaluated. A toy example is used to examine the behavior of this approach within the active learning framework.

3.1 Transposition of original choices

3.1.1 Used model

A logistic regression implemented by a neural network is used (Sarle, 1994), its architecture is represented in figure 1. This perceptron has two output neurons (O_1 and O_2) which are dedicated to both classes. This model consists in a single hidden neuron (H). The weights vector $[w_1..w_5]$ gathers parameters which are adjusted during the training stage. The first two network's input (x_1 and x_2) correspond to co-ordinates of the instance $x \in l$. Network's skew is an additional input whose value is 1.

⁴This recursive partitioning playing a discretization method. For a state of the art on discretization methods, interested readers, can refer to (Bouille, 2006).

It makes possible to vary ordinate at the origin of the linear separating which is learnt by the model. The outputs of this model are normalized by a soft max function in the interval $[0, 1]$. Outputs correspond to probabilities of observing classes, conditionally to the instance which is placed as input of the model. Neural network's training is stopped when the training error does not decrease more than 10^{-8} , and the training step is fixed to 10^{-2} .




Figure 1: Neural network for logistic regression (when input vector has 2 dimensions).

Logistic regression is used as a global model (m_* on figure 2) which is trained independently of the input space partitioning, using examples which are selected by sub-models ($m_1, m_2...m_5$. on figure 2 represent sub-models which are associated with each zone). Sub-models play a role in the selection of zones and in the selection of instances to be labelled only. m_* is trained after using these examples. m_* allows to make a coherent comparison between adaptive curiosity and stochastic strategies. Performances of the global model report only the quality of selected examples.




Figure 2: Local and global models

3.1.2 Partitioning

Zones containing at least 30 labelled examples are split. A cut separates labelled examples into two $\pm 25\%$ balanced subsets (according to the criterion of section 2.3.1). These arbitrary choices are preserved for all experiments in this paper.

3.1.3 Zones selection

The original criterion (section 2.3.2) which selects interesting zones in \mathbb{X} is modified to transpose the adaptive curiosity to classification problems. The objective is to estimate sub-model's progresses in each zone using a measure of performance. The area under ROC curves (Fawcett, 2003) (AUC) is used to evaluate performances of sub-models on labelled examples which belong to the zone (l).

Measure of performances: ROC curves plot the rate of good predictions against the rate of bad predictions on a two dimensional space. These curves are build sorting instances of test set according to the output of the model. ROC curves are usually built considering a single class. Consequently, $|\mathbb{Y}|$ ROC curves are considered. AUC is computed for each ROC curve, and the global performance of the model is estimated by the mathematical expected value of AUC, over all classes : $AUC_{global} = \sum_{i=1}^{|\mathbb{Y}|} P(y_i) \cdot AUC(y_i)$

Measure of progress: Progresses of sub-models are estimated on a temporal window which is constituted by two successive iterations. Progresses are defined as follow, with $l \in L$ the subset of labelled examples : $Progress(l) = AUC_{global}^t(l) - AUC_{global}^{t-1}(l)$

3.2 Experimental conditions

3.2.1 Stochastic strategy

The "stochastic" strategy handles a global model and uniformly selects examples according to their probability distribution. This strategy plays a role of reference and is used to measure the contribution of adaptive curiosity.

3.2.2 Toy example

The toy example is a binary classification problem in a two dimensional space $\mathbb{X} = x \times y$. We consider two classes that are separated by the boundary $y = \sin(x^3)$, on intervals $x \in [-2, 2]$ and $y \in [-2, 2]$ (see figure 4). In the following experiments, we use 2000 training examples (Φ) and 30000 test examples that are uniformly generated over the space \mathbb{X} .

3.2.3 Protocol

Beforehand, data is normalized using mean and variance. At the beginning of experiments, the training set contains only two labelled examples which are randomly chosen among available data. At every iteration, a single example is drawn in the current zone to be labelled and added to the training set. Active learning stops when 250 examples are labelled⁵.

The used model is a logistic regression implemented by a neural network (section 3.1.1). Two criteria (section 3.1.3) evaluating zones, which are respectively based on the mean square error and the empirical risk, are tested during two series of experiments. Adaptive curiosity is compared to stochastic strategy (section 3.2.1) in a third serie of experiments.

These experiments evaluate the average performance of the system, according to the number of labelled examples. Each experiment has been done ten times in order to obtain an average provided with its variance, for every point of results curves.

⁵After 250 labelled examples, results does not vary significantly.

3.3 Results and discussion

3.3.1 Performances

The criterion which is used below to evaluate strategies on the test sets is the AUC (see section 3.1.3).

In this part, performances of the global model (m_* on figure 2) are presented for adaptive curiosity approche. Figure 3 draws AUC of global model, against the number of labelled examples. Natches on curves represent variance of the 10 experiments ($\pm 2\sigma$). Performances of "stochastic" strategy also appears on figure 3. We notice that adaptive curiosity gives better performances than the stochastic strategy, nevertheless both strategies are very close. Results on figure 3 show this first implementation of adaptive curiosity does not improve significantly the quality of selected examples.

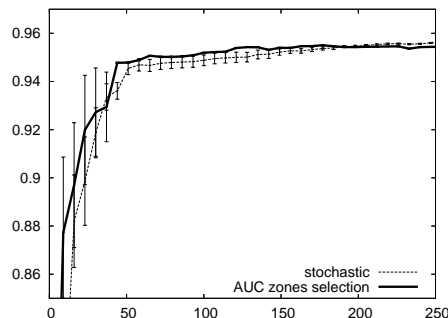


Figure 3: AUC versus number of examples

3.3.2 Selected examples

Figure 4 shows examples which have been selected during an experiment evaluating zones using AUC. The partitioning of input space and the choice of examples are relatively uniform; even if a little more populated area can be noticed for each classes (at the top right and at the middle bottom of figure 4). This strategy is unsatisfactory because areas which contain most labelled examples are not organized around the hidden pattern.




Figure 4: AUC zones selection in \mathbb{X} , with "o" points of first classe, and "•" points of second classe

4. A new criterion of zones selection

Adaptive curiosity tries to deal with the dilemma exploration / exploitation drawing new examples in zones where progress is possible. To take in consideration this dilemma in a better way, a new criterion of zones selection is proposed in this section. The rest of the adaptive curiosity method is not modified. The new criterion is composed by two terms which respectively correspond to the exploitation and the exploration. A compromise between both terms is provided by the new criterion.

4.1 Exploitation : Mixture rate

Among existing splitting criteria (Breiman, 1996), we use the entropy as a mixture rate. The function $MixRate(l)$ (equation 1) use labels of examples $l \subseteq L$ (which belong to the zone) to calculate the entropy over classes.

Part "A" of equation 1 corresponds to the entropy of classes that appear in a zone. Probabilities of classes $P(y_i)$ are empirically estimated by a counting of examples which are labelled with the considered class.

The entropy belongs to the interval $[0, \log |\mathbb{Y}|]$ (with \mathbb{Y} the number of classes). Part "B" of equation 1 normalizes mixture rate in the interval $[0, 1]$.

$$MixRate(l) = - \underbrace{\sum_{y_i \in \mathbb{Y}} P(y_i) \log P(y_i)}_A \times \underbrace{\frac{1}{\log |\mathbb{Y}|}}_B \quad (1)$$

$$\text{with } P(y_i) = \frac{|x \in l, f(x) = y_i|}{|l|}$$

Mixture rate is the "exploitation" term of the proposed criterion. By choosing zones which have strongest entropy, the hidden pattern is locally clarified thanks to new labelled examples which are drawn in these zones. The model becomes very precise, on certain area of the space. Figure 5 shows an experiment which is realized on the toy example, using entropy to select interesting zones. Selected examples are grouped around the boundary, but there is a large part of the space which is not explored.




Figure 5: Selected examples using Mixture Rate only in \mathbb{X} , with "o" points of first classe, and "•" points of second classe

4.2 Exploration : Relative density

Relative density is the proportion of labelled examples among available examples in the considered zone. Equation 2 expresses relative density, with $\phi \subseteq \Phi$ the subset of observable examples which belong to the zone. As mixture rate, relative density varies in the interval $[0, 1]$.

$$RelativeDensity(l, \phi) = \frac{|l|}{|\phi|} \quad (2)$$

Relative density is the "exploration" term of the criterion. The homogeneity of drawn examples over the input space is ensured by choosing zones which have lowest relative density. This strategy is different than a random sampling because homogeneity of drawn examples is forced. Figure 6 shows an experiment which is realized on the toy example, using relative density to select interesting zones. Input space partitioning and examples drawing are homogeneous.




Figure 6: Selected examples using Relative Density only in \mathbb{X} , with "o" points of first classe, and "•" points of second classe

4.3 Compromise Exploitation vs. Exploration

The criterion evaluates the interest of zones, taking into account both terms; mixture rate and relative density. Equation 3 shows how each term is used. The parameter $\alpha \in [0, 1]$ corresponds to a compromise between exploitation of already known mixture zones and exploration of new zones.

$$Interest(l, \phi) = (1 - \alpha) MixRate(l) + \alpha (1 - RelativeDensity(l, \phi)) \quad (3)$$




Figure 7: Selected examples with $\alpha = 0.5$ in \mathbb{X} , with "o" points of first classe, and "•" points of second classe

The notion of progress is included in the criterion : the relative density (which increases at the same time new examples are labelled) forces the algorithm to leave zones in which mixture rate does not increase quickly. If there is no thing else to discover in a zone, the criterion naturally avoids it. In certain cases, the criterion prefers none mixed zones which are not enough explored. This criterion does not need a temporal window to evaluate the progress of sub-models (see paragraph 2.3.2). So its implementation is easier than original adaptive curiosity approach. Figure 7 shows an experiment which is realized on the toy example, using the criterion with $\alpha = \frac{1}{2}$. Input space partitioning and examples drawing are organized around the boundary without leaving any region of space.

4.4 Results and discussion

In this section, the toy example (section 3.2.2) as well as the experimental protocol (section 3.2.3) is re-used. Several series of experiments are realized for $\alpha = [0, \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, 1]$. The purpose of this part is to estimate the influence of this parameter on performances, and to compare the obtained results to "stochastic" strategy.

Figure 8 shows performances of the proposed strategy for various values of α . The first curve represents the "stochastic" strategy. When $\alpha = 0$ only mixture rate is considered by the criterion. In this case, the observed performances are significantly lower than the "stochastic" strategy considering less than 100 examples. This phenomenon can be intuitively interpreted by a strong exploitation of detected mixture zones, to the detriment of the remaining space. When $\alpha = 1$ only relative density is considered. In this case, adaptive curiosity gives lower performances than the "stochastic" strategy considering less than 70 examples. The best performances are observed for $\alpha = 0.25$. In this case, the maximum AUC is reached very early (with 60 labelled examples). Observed performances are superior to "stochastic" strategy for all number of learnt examples. This value obviously offers a good compromise between exploration and the exploitation.

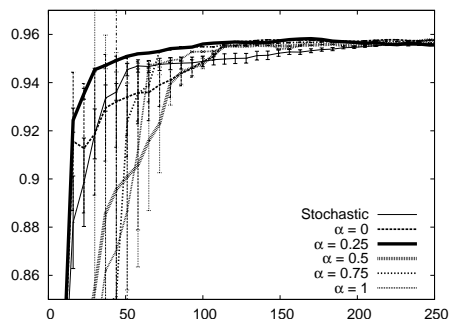


Figure 8: AUC vs number of examples

These results show that adaptive curiosity can be beneficially used in active learning framework, with the pro-

viso of using an adapted zones selection strategy. Moreover, the new strategy of zones selection is only based on data typology. Sub-models are only used to carry out the partitioning and not to choose interesting zones.

5. Comparison with two active strategies

The objective of this section is to compare the previously obtained results with active learning approaches which come from the literature. Two active strategies are considered in this paper : "uncertainty sampling" and "error reduction sampling".

5.1 Uncertainty sampling

Uncertainty sampling is an active learning strategy (Thrun and Möller, 1992) which is based on a confidence measure associated by the model to its prediction. The used model must be able to produce an output and to estimate the relevance of its answers. The logistic regression estimates the probability of observing each class, given an instance $x \in \mathbb{X}$. The model selects the one that maximizes $\hat{P}(y_j|x)$ (with $y_j \in \mathbb{Y}$) among all possible classes. A prediction is considered as uncertain when the probability to observe predicted class is weak. This strategy of active learning selects unlabelled examples which maximize the uncertainty of the model. The uncertainty can be expressed as follows :

$$Incertain(x) = \frac{1}{\operatorname{argmax}_{y_j \in \mathbb{Y}} \hat{P}(y_j|x)} \quad x \in \mathbb{X}$$

5.2 Sampling by risk reduction

The purpose of this approach is to reduce the generalization error, $E(\mathcal{M})$, of the model (Roy and McCallum, 2001). It chooses examples to be labeled so as to minimize this error. In practice this error cannot be computed because the distribution of instances in \mathbb{X} is unknown. Nicholas Roy (Roy and McCallum, 2001) shows how to bring this strategy into play since all the elements of \mathbb{X} are unknown. He uses an uniform prior for $P(x)$:

$$\hat{E}(\mathcal{M}^t) = \frac{1}{|L|} \sum_{i=1}^{|L|} \mathcal{L}oss(\mathcal{M}^t, x_i)$$

In this paper, one estimates the generalization error ($E(\mathcal{M})$) using the empirical risk (Zhu et al., 2003) :

$$\hat{E}(\mathcal{M}) = R(\mathcal{M}) = \sum_{i=1}^{|L|} \sum_{y_j \in \mathbb{Y}} \mathbb{1}_{\{f(x_i) \neq y_j\}} P(y_j|x_i) P(x_i)$$

where f is the model which estimates the probability that an example belong to a class, $P(y_i|x_i)$ the real probability to observe the class y_i for the example $x_i \in L$, $\mathbb{1}$ the indicating function equal to 1 if $f(x_i) \neq y_j$ and equal to 0 else. Therefore $R(\mathcal{M})$ is the sum of the probabilities that

the model makes a bad decision on the training set (L). Using a uniform prior to estimate $P(x_i)$, one can write :

$$\hat{R}(\mathcal{M}) = \frac{1}{|L|} \sum_{i=1}^{|L|} \sum_{y_j \in \mathbb{Y}} \mathbb{1}_{\{f(x_i) \neq y_j\}} \hat{P}(y_j|x_i)$$

In order to select examples, the model is re-trained several times considering one more “fictive” example. Each instance $x \in U$ and each label $y_j \in \mathbb{Y}$ can be associated to constitute this supplementary example. The expected cost for any single example $x \in U$ which is added to the training set is then:

$$\hat{R}(\mathcal{M}^{+x}) = \sum_{y_j \in \mathbb{Y}} \hat{P}(y_j|x) \hat{R}(\mathcal{M}^{+(x,y_j)}) \quad \text{with } x \in U$$

5.3 Results on the toy example

Once again, the same toy example (section 3.2.2) and the same experimental protocol (section 3.2.3) are used. Experiments bring into play active strategies which were presented in sections 5.1 and 5.2, using a global model. As shown on figure 9, our adaptive curiosity strategy (with $\alpha = 0.25$) is the best active learning strategy. The uncertainty sampling gives a very high variance (for a question of legibility, natches on curve represent $\pm \frac{\sigma}{5}$ only for uncertainty sampling). Moreover, the average performance of this approach is very low in comparison to stochastic sampling. So uncertainty sampling is a very bad strategy for the considered toy example. Sampling by error reduction gives better results than the other active strategy, but the observed performances are always lower than stochastic sampling and our adaptive curiosity strategy.

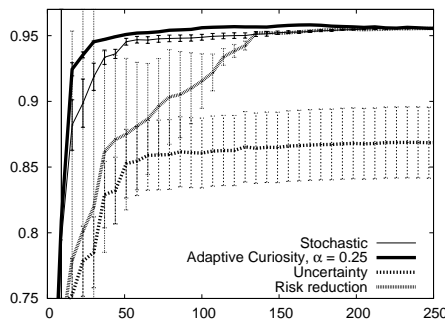


Figure 9: AUC of active learning methods

5.4 Results on real data

Experiments are conducted using also two public data files coming from the “*UCI repository*” (D.J. Newman and Merz, 1998). The datasets used are the following :

Diabetes tracking: “Pima” data file deals with detection of diabetes problems for patients who are older than

21 years. The 786 subjects (Training : 354, Test : 354) of this dataset are characterized by 9 medical indicators such as blood pressure or body mass index. The considered problem is a binary classification between individuals who have (or not) diabetes problems. Figure

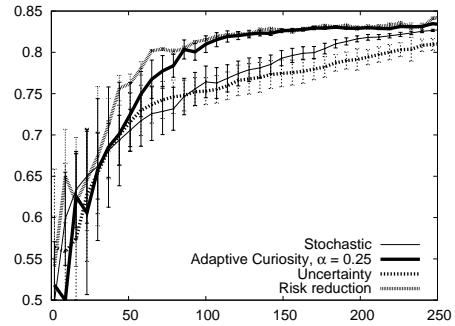


Figure 10: AUC on “Pima”

10 shows performances of different strategies on “Pima”, according to the number of labelled examples. On this dataset, sampling by risk reduction gives the best results. The AUC values are highest, for all considered number of examples. Only one curve of adaptive curiosity is shown (this curve corresponds to the best value of α). In this case, adaptive curiosity gives good performances very close to sampling by risk reduction. Moreover, adaptive curiosity gives very low variance. Finally, uncertainty sampling is the worse strategy, with AUC values which are largely lower than stochastic strategy.

Credit approval: “Australian” dataset concerns credit approvals. The 690 instances (Training : 345, Test : 345) of this dataset are defined by 14 attributes. All attribute names and values have been changed to meaningless symbols to protect confidentiality of the data. The considered problem is a binary classification on the acceptance of credits. Figure 11 shows performances

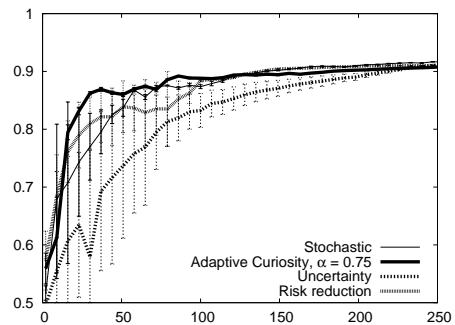


Figure 11: AUC on “Australian”

of different strategies on “Australian”. On this dataset, adaptive curiosity gives the best performances. The maximum AUC value (0.9) is reached with few labelled examples (about 80). When the number of labelled examples is greater than to 120, performances of “stochastic” strategy, sampling by error reduction and adaptive curiosity are very close. Once again, uncertainty sampling is the worse strategy.

Remarks : These results show that adaptive curiosity behaves similarly on the toy example and on real data. In both cases, the trend is the same : uncertainty sampling gives bad performances (worse than stochastic strategy); sampling by risk reduction and adaptive curiosity give close performances. However sampling by risk reduction generate a computing time 7 times higher than adaptive curiosity. Adaptive curiosity seems to be an efficient active learning strategy, with the proviso of properly adjusting the parameter α using a probabilistic estimation.

6. Conclusion

This paper shows that adaptive curiosity can be used as an active learning strategy in machine learning framework. Adaptive curiosity is a strategy which is not dependent of the predictive model. This strategy can be applied on numerous real problems and is easy to use with existing systems.

We have defined a new zones selection criterion which gives good results on the considered toy example and on real data. However, this criterion balances exploitation and exploration using a parameter. Future works will be done to make the algorithm autonomous to adjust this parameter (Osugi et al., 2005).

Adaptive curiosity was initially developed to deal with high dimensionality input spaces, where large parts are unlearnable or quasi-random. Future works will be realized to estimate the interest of our new criterion in such conditions. The influence of the complexity of the problem to be learnt (that is say, the number of examples necessary to solve it) will be studied to.

The partitioning step of adaptive curiosity has a $O(n^3)$ complexity and is prohibitive to treat high dimensionality datasets. Moreover, the cut criterion involves two parameters : the maximum number of labelled examples belonging to a zone, and the maximum balance rate of labelled examples subsets of a zone split. The use of non parametric discretization method (Bouille, 2006) could be an efficient way to decide "when" and "where" a zone has to be split. This aspect will be considered in future works.

References

- Bouille, M. (2006). MODL: A bayes optimal discretization method for continuous attributes. *Machine Learning*, 65(1):131–165.
- Breiman, L. (1996). Technical note: Some properties of splitting criteria. *Machine Learning*, 24(1):41–47.
- Castro, R., Willett, R., and Nowak, R. (2005). Faster rate in regression via active learning. In *NIPS (Neural Information Processing Systems)*, Vancouver.
- D.J. Newman, S. Hettich, C. B. and Merz, C. (1998). UCI repository of machine learning databases. <http://www.ics.uci.edu/~mllearn/MLRepository.html>, University of California, Irvine, Dept. of Information and Computer Sciences.
- Fawcett, T. (2003). Roc graphs: Notes and practical considerations for data mining researchers. T. Fawcett. ROC Graphs: Notes and Practical Considerations for Data Mining Researchers. Technical Report HPL-2003-4, HP Labs, 2003.
- Nagai, Y., Asada, M., and Hosoda, K. (2002). Developmental learning model for joint attention. In *Proceedings of the 15th International Conference on Intelligent Robots and Systems (IROS)*, pages 932–937.
- Osugi, T., Kun, D., and Scott, S. (2005). Balancing exploration and exploitation: A new algorithm for active machine learning. In *Proceedings of the Fifth IEEE International Conference on Data Mining (ICDM'05)*.
- Oudeyer, P.-Y. and Kaplan, F. (2004). Intelligent adaptive curiosity: a source of self-development. In Berthouze, L., Kozima, H., Prince, C. G., Sandini, G., Stojanov, G., Metta, G., and Balkenius, C., (Eds.), *Proceedings of the 4th International Workshop on Epigenetic Robotics*, volume 117, pages 127–130. Lund University Cognitive Studies.
- Roy, N. and McCallum, A. (2001). Toward optimal active learning through sampling estimation of error reduction. In *Proc. 18th International Conf. on Machine Learning*, pages 441–448. Morgan Kaufmann, San Francisco, CA.
- Sarle, W. S. (1994). Neural networks and statistical models. In *Proceedings of the Nineteenth Annual SAS Users Group International Conference, April, 1994*, pages 1538–1550, Cary, NC. SAS Institute.
- Singh, A., Nowak, R., and Ramanathan, P. (2006). Active learning for adaptive mobile sensing networks. In *IPSN '06: Proceedings of the fifth international conference on Information processing in sensor networks*, pages 60–68, New York, NY, USA. ACM Press.
- Thrun, S. B. and Möller, K. (1992). Active exploration in dynamic environments. In Moody, J. E., Hanson, S. J., and Lippmann, R. P., (Eds.), *Advances in Neural Information Processing Systems*, volume 4, pages 531–538. Morgan Kaufmann Publishers, Inc.
- White, R. (1959). Motivation reconsidered: The concept of competence. *Psychological Review*, 66:297–333.
- Zhu, X., Lafferty, J., and Ghahramani, Z. (2003). Combining active learning and semi-supervised learning using gaussian fields and harmonic functions. In *ICML (International Conference on Machine Learning)*, Washington.