Novelty-Organizing Classifiers Applied to Classification and Reinforcement Learning: Towards Flexible Algorithms

Danilo Vasconcellos Vargas
Kyushu University
Fukuoka, Japan
vargas@cig.ees.kyushu-u.ac.jp

Hirotaka Takano
Kyushu University
Fukuoka, Japan
takano@cig.ees.kyushu-u.ac.jp

Junichi Murata
Kyushu University
Fukuoka, Japan
murata@cig.ees.kyushu-u.ac.jp

ABSTRACT

It is widely known that reinforcement learning is a more general problem than supervised learning. In fact, supervised learning can be seen as a class of reinforcement learning problems. However, only a couple of papers tested reinforcement learning algorithms in supervised learning problems. Here we propose a new and simpler way to abstract supervised learning for any reinforcement learning algorithm. Moreover, a new algorithm called Novelty-Organizing Classifiers is developed based on a Novelty Map population that focuses more on the novelty of the inputs than their frequency. A comparison of the proposed method with Self-Organizing Classifiers and BioHel on some datasets is presented. Even though BioHel is specialized in solving supervised learning problems, the results showed only a trade-off between the algorithms. Lastly, results on a maze problem validate the flexibility of the proposed algorithm beyond supervised learning problems. Thus, Novelty-Organizing Classifiers is capable of solving many supervised learning problems as well as a maze problem without changing any parameter at all. Considering the fact that no adaptation of parameters was executed, the proposed algorithm's basis seems interestingly flexible.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*intelligent agents, languages and structures, multi-agent systems*

Keywords

Novelty, Novelty Map, Reinforcement Learning, Structured Evolutionary Algorithms, Supervised Learning, Self Organizing Classifiers, Novelty Organizing Classifiers

GECCO'14, July 12–16, 2014, Vancouver, BC, Canada.
ACM 978-1-4503-1964-5/14/07.
http://dx.doi.org/10.1145/2598394.2598429 .

1. INTRODUCTION

Supervised learning is a relatively close problem to reinforcement learning (RL). Although the similarities exist, there are few RL methods that were applied without any changes to supervised learning. The hurdle is that the discount factor used for multi-step RL problems is not suited for single-step problems.

To the knowledge of the authors, the first complete reformulation of the supervised learning problem for RL arrived only recently in [4]. However, the employed supervised learning abstraction is too complex. This paper presents a more straightforward abstraction featuring only a resample of the data inside a loop and we show that this is enough to enable good results. In fact, the abstraction is similar to how students memorize a couple of words to a language class.

This article also proposes an algorithm called Novelty-Organizing Classifiers. It is based on a novelty map population which evolves a feedforward neural model. The new novelty map population has a dynamic that focuses more on the novelty of the inputs than their frequency, resulting in less changes in the cells' positions and, consequently, better specialization of the cells. To evaluate the algorithm, it is compared with both Self-Organizing Classifiers and BioHel in many famous classification datasets. Moreover, without any modification, the Novelty-Organizing Classifiers algorithm is run over a continuous input-continuous output maze problem to verify its flexibility.

2. RELATED METHODS

Self-Organizing Classifiers (SOC) uses a population structure which dynamically self-organizes itself with the input. The dynamics used by the population are the same as the self-organizing map (SOM) and therefore it is called self-organizing map population. Inside every cell of the population structure there is a subpopulation.

When an input is presented to the algorithm the cells compete for it. The winner cell (cell which is closest to the input) chooses randomly an individual from its subpopulation to act on the environment. A Q-learning reinforcement style is used to reward classifiers. Moreover, the evolutionary algorithm is local therefore it runs on every cell independently.

3. PROPOSED METHOD

This work employs the novel Novelty Organizing Classifiers, an algorithm somewhat similar to self-organizing classifiers (SOC) [3, 2]. The main difference lies in the population structure (Novelty Map population is used instead of SOM population) and the evolutionary algorithm (EA).

The EA is even closer to a structured version of differential evolution algorithm where each parent inside a subpopulation have one child and the selection procedure happens only between them (i.e., each parent can only be eliminated if surpassed by their own child). This type of selection procedure aids in the preservation of diversity. The Novelty Map Population is explained in the following subsections.

3.1 Novelty Map

Novelty Map is a table of cells. Each cell has one of the most novel individuals, according to a novelty metric, as a weight array. When a new input is presented to the map, a competition takes place where the cell with the closest weight array wins. The table is updated by substituting the cell with the smallest novelty with the input array if and only if the input array has higher novelty. This way the table is always kept up to date.

3.2 Novelty Map Population

The Novelty Map Population has the same behavior as Novelty Map. Additionally, inside every cell there is a subpopulation of individuals evolved by Novelty Organizing Classifiers.

4. PROPOSED ABSTRACTION

This work also proposes a supervised learning abstraction for reinforcement learning. The abstraction consists of resampling the training dataset many times and providing these samples as input to the reinforcement algorithm. As usual, the input of the algorithm is the sample itself and the output is the classes' weights. Therefore, the output size is equal to the number of classes and the output index related to the highest output value inside the array indicates the class which the sample pertains. Rewards were set to zero or one when the algorithm makes respectively incorrect or correct predictions.

5. EXPERIMENTS

Table 1 shows the parameters for Novelty-Organizing Classifiers (NOC). BioHel uses the same settings from [1]. To apply the supervised learning tests a new abstraction of supervised learning to reinforcement learning was created (See Section 4). The results of the comparison are presented in Table 2. Moreover, NOC was observed to solve a maze problem without changing its settings.

The supervised learning results show only a trade-off between both algorithms even when BioHel is a specialized algorithm made to solve this class of problems. In fact, the same NOC that was applied to the single-step problems (supervised learning) is used without any changes in a multi-step problem (maze problem), demonstrating the flexibility of the algorithm. Furthermore, this demonstrates that the abstraction proposed is sufficient to convert supervised learning problems to reinforcement learning problems.

6. ACKNOWLEDGMENT

This work was supported in part by JSPS KAKENHI Grant Number 24560499.

Table 1: Parameters for Novelty-Organizing Classifiers

NN	Hidden Nodes	10
DE	CR	0.1
	F	0.5
Novelty Map	Number of Cells	100
	Novelty Metric	Uniqueness
General	Widrow-hoff coefficient	0.01
	β	10
	ν	10
	ι	1000
	Discount factor	0.99
	Novel initial fitness	-10
	Best initial fitness	0

Table 2: Comparison over UCI and other famous datasets. Each number is the mean accuracy measured over three stratified 10-fold cross-validation results, resulting in a total of 30 runs per dataset. Missing values were treated naively as zero for all tests. A result is marked in bold when the Mann-Whitney test give a p-value equal or bigger than 0.01. This means that the results are significantly different.

	BioHel [%]	NOC [%]
Diabetes	67.31(±3.65)	67.92(±3.54)
Glass	54.32(±11.02)	**69.66(±7.72)**
Iris	92.71(±6.88)	92.74(±6.86)
Sonar	63.86(±10.50)	**78.93(±11.01)**
Vehicle	**68.06(±3.27)**	61.17(±5.31)
Vowel	74.15(±3.29)	73.03(±5.37)
Zoo	91.18(±8.21)	91.09(±10.10)

7. REFERENCES

[1] M. A. Franco, N. Krasnogor, and J. Bacardit. Gassist vs. biohel: critical assessment of two paradigms of genetics-based machine learning. *Soft Computing*, pages 1–29, 2013.

[2] D. V. Vargas, H. Takano, and J. Murata. Self organizing classifiers and niched fitness. In *Proceeding of the fifteenth annual conference on Genetic and evolutionary computation conference*, pages 1109–1116. ACM, 2013.

[3] D. V. Vargas, H. Takano, and J. Murata. Self organizing classifiers: first steps in structured evolutionary machine learning. *Evolutionary Intelligence*, 6(2):57–72, 2013.

[4] M. A. Wiering, H. van Hasselt, A.-D. Pietersma, and L. Schomaker. Reinforcement learning algorithms for solving classification problems. In *Adaptive Dynamic Programming And Reinforcement Learning (ADPRL), 2011 IEEE Symposium on*, pages 91–96. IEEE, 2011.