Role of unsupervised learning in distinguishing between robot's normal and faulty behaviours

Anton Vitko, Ladislav Jurišica , Michal Šavel

Abstract
In its first part the paper summarises main approaches to the unsupervised learning. The topic is then narrowed to the problem of finding a structure that is hidden in a set of sensor data with the aim to distinguish between robot's normal and faulty behaviours. The second part is focused on the application of Adaptive Resonance Theory – (ART) based clustering of "suspected" and possibly faulty behaviours of a walking robot.

Keywords: clustering, classification, robot's faults

Introduction

The unsupervised learning, that is learning without a teacher has long been considered to be one of the most fundamental models of learning and understanding and thus a striking feature of intelligent systems. The unsupervised learning paradigm can be found both at the sensory level of animals and at higher, cognitive level of humans. It still remains a hot topic of the neural networks research (see e.g. [1], [2], [3]).

In a view of the learning theory the clustering is a kind of unsupervised data driven learning. During learning the system tries to find a "structure" inside a (large) set of unlabeled information chunks. The "structure" is established by finding a certain kind of similarity between the members of the same cluster, or in other words, by finding intra-cluster similarity and inter-cluster dissimilarity.

The unsupervised clustering is even possible when a training set and class definition are, regardless of reasons, unavailable. However, one generally needs to know in advance, whether natural clusters actually exist in a given data set. The thing is that if such natural clusters do not exist, the clustering process may lead to finding artificial and meaningless structure. The process of determination whether the structure exists within data is a poorly studied problem, known as "determination of the clustering tendency" of a data set. [4]

The neural networks that are based on unsupervised learning may perform natural clustering on the basis of the similarity (with respect to a chosen similarity measure) between input patterns. As to the clustering enables discovering hidden similarities between items of a large data set it also enables grouping data into a small number of clusters. Due to this the clustering is an important tool in data mining (an umbrella term for variety of analytical techniques of knowledge discovery) and creates the basis for solving many advanced scientific and engineering problems where large amounts of apparently disordered data are to be succinctly comprehended.

As indicated above, the similarity is commonly evaluated on the basis of an adopted similarity criterion. The criterion may be either distance-based, or conceptual–based. Depending on the dimensionality of the data space or existence of the correlations between data items, the distance-based similarity is expressed through an appropriate distance measure, like Euclidean distance, Minkowski distance, Mahalanobis distance and others [5]. The typical representatives of the distance based algorithms are the K-means algorithm, allowing a crisp membership or C-means algorithm allowing partial memberships to two or more clusters. [5]

In the realm of neural networks, the typical representatives of the conceptual clustering are Kohonen's self organizing maps (SOM) [2] and Grossberg's adaptive resonance theory (ART) networks. [1]. The SOMs are low-dimensional grids of nods, each nod representing a model of particular observation. During learning the SOMs compress and convert nonlinear statistical relationships between high-dimensional data into low-dimensional displays while preserving the most important topological features of the primary data. After learning the neighboring nods represents similar observations. But the fact of compressing data while preserving topological features is in essence a kind of abstraction, or a kind of the creation of a concept, or a kind of perception.

The control community is familiar with the term of "intelligent control", connoting the abilities that the conventional control system cannot attain, like making complex decisions, adapt to new conditions, self-organizing, planning future activities, and the more. An autonomous robot is a particular example of the intelligent system. Its functionality relies on numerous disparate sensors through which the robot grasps consistent knowledge of what is going on around it. Therefore the robot is required to respond to instantaneous incentives coming form the surrounding environment. To this end it needs to handle wide range of unexpected events and distinguish between common (normal) and unusual contexts.

Soft computing techniques should be considered as mere building blocks or even "bricks" used for building up a "large house" of any intelligent system. What makes robots intelligent is just a synergic use of these techniques, which in time and space invoke, optimize and fuse elementary behaviours in order to produce an appropriate overall behaviour. System intelligence comes from the system architecture i.e.

from an inner modular structure of the both system elements and functionalities. One striking example of the robot intelligence is the subsumption architecture developed by Brooks [6]. Another example is the robot's ability to maintain its own "health" by early detection, identification and classification of the both imminent and existing faults and in such a way to prevent its functionality from fatal failures.

1. Detection, clustering and classification of robot's faults

Due to the extensive use of complex mechanical components like arms, legs, actuators, gears, clutches, grippers etc., the robot's mechanical parts suffer from significantly higher fault rates than pure electric and electronic circuitry. These faults must be detected, identified and classified in accordance with their criticality, and appropriate measures to compensate them must be taken in order to prevent the system from failure. The system should be able to anticipate possible faults on the basis of certain "pathologic" behaviours, that is those, which are novel and suspected in comparison with normal functioning. In this view the incorporation of an appropriate mechanism of detection of novelties becomes necessary. Imminent failures are often manifested through the declined values of system parameters and variables or their fused complexes. The idea is to identify any deviation from normal behaviour. The component degradation, like wear, increased friction, stiction due to contamination, corrosion etc., is related to an observable effect on the system performance (higher vibrations, increased friction, decreased positioning precision etc). These relationships may change as the process of performance degradation progresses.

Over the past three decades numerous approaches to fault management have been developed. They range from checking the limit values or trends to those based on the state and/or parameter estimation, fault trees, and the artificial intelligence based systems, like expert systems, case-based reasoning systems, and fuzzy and neural learning approaches. It is beyond the scope of this paper to even summarize all these solutions. Therefore, in the sequel will be described a neural, unsupervised-based learning mechanism of novelty detection clustering and classification

It is known that one serious problem with neural classification is that, in real situations, the problem domain does not always behave well in a sense that if some unexpected and strongly different input patterns appear the neural system cannot manage it appropriately, because it has no built-in mechanism to recognize and store them without deteriorating performance. Said it in a different way, the system should preserve previously learned patterns (so called problem of stability) while keeping its ability to learn new patterns (so called problem plasticity). This phenomenon is known as a stability-plasticity dilemma. An elegant solution to this problem provides a family of the neural networks based on the "adaptive resonance theory" (ART), developed by Grossberg and Carpenter [1].

The ART family of neural networks perform competitive learning. Its architecture is able to cluster input patterns on the basis of a given measure of similarity. In particular, the ART1 network used in the authors' experiment, allows the incremental learning of prototypes, rather then instantaneous input exemplars. The clusters of similar inputs are updated by using information from the currently presented input pattern. In this way every cluster preserves main features of the accepted similar input patterns.

The next paragraph briefly describes the way of using the ART neural network for clustering "suspected" torques that appear in the robot leg joints when the robot walks. Besides the competitive learning used as a basis of a learning philosophy the ART network is further improved to be able to deal effectively with the stability–plasticity dilemma, i.e. it should be sufficiently stable in retaining already learned inputs, while sufficiently susceptible (plastic) to acquire new inputs. Too much stability means that the network is "stubborn" when learning new inputs, while too much plasticity could cause that newly learned inputs may deteriorate those learned previously. The ART neural networks resolve this problem by introducing the phenomenon of resonance, meaning that the current input is compared with all already stored prototypes. If it does not match sufficiently any of stored prototypes (resonance does not take place) and the ART creates new prototype. In this way the previously learned inputs are not deteriorated by the newly accepted ones. Significant advantage of the ART is an ability to discover concepts of various levels of abstractions, which are hidden in the input data. This is achieved with the *vigilance parameter*, which determines whether the currently presented input pattern should be recognized as an already known concept or as a novel pattern.

The ART1 shown in Fig.1 consists of two fully connected layers. The pattern-representation layer F_1 accepts an input pattern and through the button-up weighted connections b_{ij} (initially set to one) sends it to the cluster-representation layer F_2. The (binary coded) input pattern "I" is created as a concatenation of particular sensor outputs and/or flags of communication errors. Due to a competition-based "winner-takes-all" paradigm that is evoked by lateral negative feedback ($-\varepsilon$), a neuron in the F_2 layer, which receives the highest button-up activity is declared a winner. Its output is set to the unit value and projected back to F_1 through the top-down weights t_{ij}.

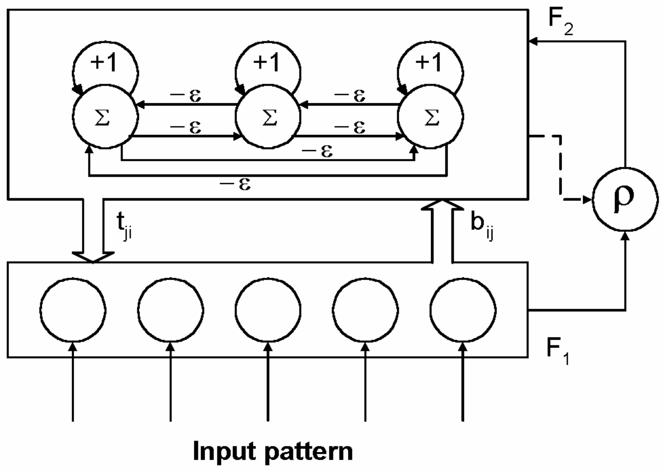

Fig.1 The ART-1 network

If the similarity between the projected winner and the input pattern proves to be greater then a given value of the *vigilance ρ*, a "*resonant state"* occurs and the weights t_{ij}, b_{ij} are (by different ways) moved closer to the input pattern (learning). If the resonant state does not occur, the winner is disqualified and the process searches for the second best matching neuron, which is then submitted to the vigilance test. Searching repeats until either vigilance test passes or no more neurons are available for testing. It is just the set value of the vigilance that calibrates how much novelty the neural network can tolerate before it clusters and classifies the input pattern into particular class. The experiments have shown that the ART-1 is fully justified for using as a means for the novelty detection and classification of patterns of the behaviours.

2. Experimental results

Efficiency of the developed neural classifier was verified by both simulation and by experimentation with the specially developed legged robot. The simplest and the most evident faults like those related to control sequences that control the movement of joints or the faults appearing during switching between robot gaits or incorrect coordination of legs due to improper timing (fall out of phase or fall out of step and the like), were easily detected and classified by a deterministic final-state machine. Contrary to the erroneous control sequences, much more complex faults may be caused by the increased friction in bearings, slipping or dragging clutches, lack of a lubrication or partial losses of the energy delivery to the leg's joints. Malfunctions of this kind may remain hidden for longer time and may gradually lead to fatal failures, like the total destruction of bearings or drives, lagging legs movement, which could jeopardize the walking stability or even cause instability of the robot. Such faults are commonly manifested through abnormal trajectories of the joint torques or forces.

The joints of the experimental robot are equipped with torque sensors, which sense current time dependences of the joint torques. Every leg can be either in a stance state, when it supports the robot body or in a swing state, when it moves in air to the position where it can begin a new stance. A time-course of the normal (faultless) torque exerted in a femur joint is shown in Fig. 2. One complete step cycle is performed in three phases, each lasting one second. As seen from the figure, these three phases can be easily observed from the torque-time dependence. Particular phases are supplemented with imbedded sub-figures depicting the leg configuration that corresponds to the current phase. During the first phase the leg remains in a flexed configuration in the stance. The femur joint exerts the torque value about 30 Nm, which maintains an attitude of the robot body. The second phase starts at one second. The leg is uncoupled from the ground and starts its movement in a direction of walking. While the torque exerted in the femur joint causes raising the leg, the coxa joint is rotating the leg about the vertical axis and the tibia joint is extending the leg. When reaching the highest position the femur joint exerts maximum torque. Just after the third second the femur torque slightly decreases so as to make the foot go down until it reaches the ground. At this moment (at about the fourth second) the leg is entering into its stance state again, and supports the robot body.

During learning, the neural network ART1 is first taught to learn the normal torque time course. As a result, the neural network appoints the normal torque course as the centre of a receptive field of the cluster of all "approximately normal" torque patterns. This is done by adaptation of the bottom-up weights leading to the most left neuron in the layer F_2. From this time on the unit output of this neuron will indicate that the current input belongs to the cluster of "approximately normal" torque courses and this cluster will represent a class of normal torque courses. Then a training list, i.e. a series of faulty torque patterns was repeatedly presented. The experimental results have shown that the learning task may be considered as accomplished after presentation of about 5 or 6 epochs. After learning the neural network becomes able to successfully classify any other set of faulty courses.

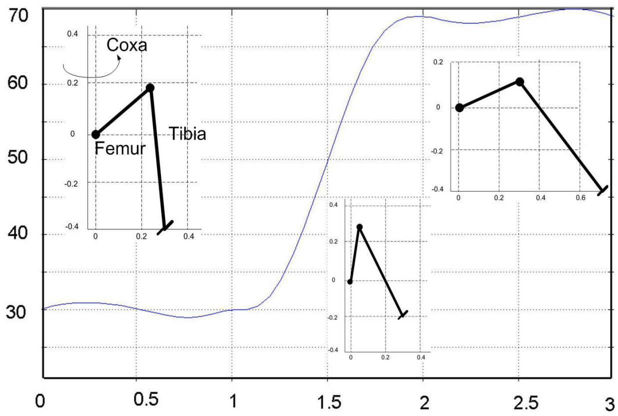

Fig.2 Normal torque in the femur joint

The results obtained are briefly summarized in Fig.3. Faulty torque patterns were subsequently presented to the previously learned ART1 network. Each of them corresponded to a particular fault as indicated by the text under graphs. For the vigilance "ρ" set to the value 0.8, the input patterns were classified into five distinctive classes. Raising the vigilance to 0.9 meant that the system was no longer willing to tolerate so much novelty (dissimilarity to the normal pattern) as before. A direct consequence of the increased vigilance was creation of as many as eight classes that, for the sake of brevity, are not shown.

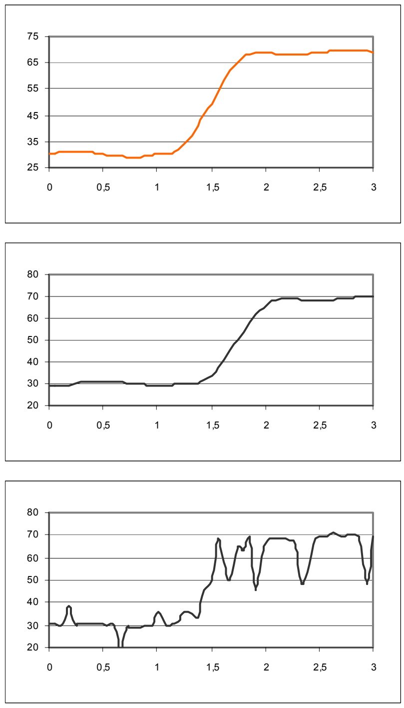

Class 1: Normal torque for slightly dragging bearing
Normal torque delayed by 0.24 s
Normal torque with imposed strong noise caused by dragging and slipping clutch

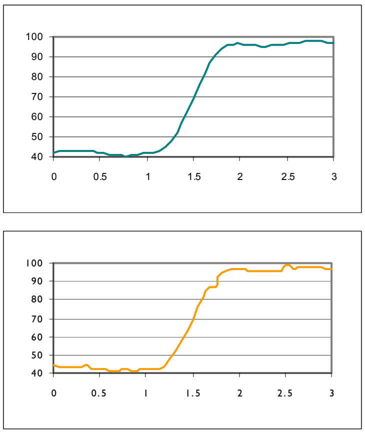

Class 2: Increased torque due to increased both friction in bearing and load

Increased and noisy torque due to increased friction in bearing ; clutch is slightly dragging

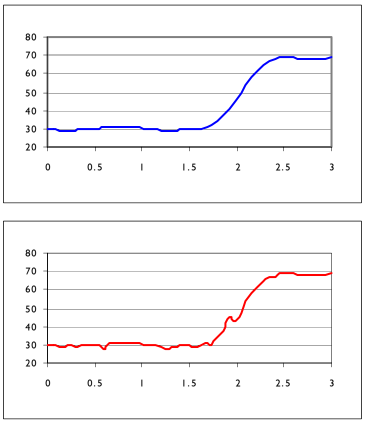

Class 3: Normal torque delayed by 0.54 s

Normal torque for slightly dragging bearing and delayed by 0.54 s

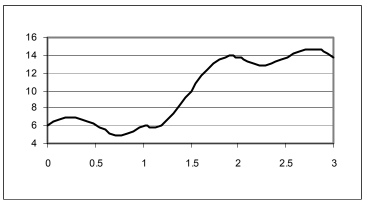

Class 4: Torque for the foot sinking the into a soft ground (almost zero load)

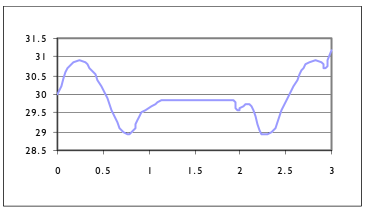

Class 5: The foot was raised and laid to the ground without rotation in the coxa joint what caused a plateau in the middle

Fig.3 Results of classification for the vigilance value 0.8

Conclusion

An intelligent robot operating autonomously in an unknown environment is a particular instantiation of an intelligent system. As such, it is expected to exhibit abilities beyond those attainable by traditional robots. To this end, sensor data should be fused into information-rich patterns, which are further clustered and classified into classes corresponding to various contexts. (Issues of the sensor fusion were not included here). Due to the classification of the contextual information the robot is able to distinguish between normal and erroneous patterns of behaviour and take appropriate measures. The methodology described above was used in the development of the learning neural-based fault detection and classification.

The results obtained provide sufficient evidence that the ART1 neural network is a very flexible and reliable means for detection and classification of the novelties that appear in the robot behaviour. Contrary to other neural or statistical approaches, there is no need to specify number of classes in advance. Based on value of the vigilance “ρ”, the network classifies input patterns into so many classes, how many is required for separating dissimilar input patterns. In more complex cases the system can detect and classify even contexts i.e. a composite states of the robot together with its environment.

Acknowledgment

The work has been supported by the grant VEGA 3120.

References

[1] Carpenter. G.A, Grossberg. S. A. (1987) Massively Parallel Architecture for a Self-organizing Neural Pattern Recognition Machine, Comp. Vision, Graphics and Image Proc. vol. 37, pp 54-115

[2] Kohonen, T. (1982) Self-organized Formation of Topologically Correct Feature Maps. Biological Cybernetics, 43, 59-69

[3] Oja, T. (1989) Neural Networks, Principal Components, and Subspaces. Int. J. Neural Systems

[4] Massey, L., Determination of Clustering Tendency with ART Neural Networks.
http://citeseer.ist.psu.edu/698457.html

[5] Jain, A.K., Murtty, M.N.and Flynn, P.J. (1999) Data clustering: A Review, ACM Computing Surveys, Vol.31, No.3, Sept. 1999

[6] Brooks R. A. (1989) A Robot that Walks: Emergent Behaviours from a Carefully Evolved Network. Proceedings of IEEE Conf. on Robotics and Automation, 1989, 692-696

Abstract

The paper in its first part the paper summarises main approaches to the unsupervised learning. The presented topic is then narrowed to the problem of finding a structure tahat is hidden in a set of sensor data with the aim to distinguish between robot’s normal and faulty behaviours. The second part is focused on a possible application of the Adaptive Resonance Theory (ART) - based clustering of faulty behaviours of a walking robot.

Doc. Ing. Anton Vitko, PhD.

Slovak University of Technology Bratislava
Faculty of Mechanical Engineering
Department of automation, information
and instrumentation technology
Námestie slobody 17
812 31 Bratislava
Tel.: 02-57294 558
E-mail: vitko@kam.vm.stuba.sk

Prof. Ing. Ladislav Jurišica, PhD.

Slovak University of Technology,
Faculty of Electrical Engineering
and Information Technology
Institute of Control and Industrial Informatics
Ilkovičova 3
812 19 Bratislava

Ing. Michal Šavel

Slovak University of Technology Bratislava
Faculty of Mechanical Engineering
Department of automation, information
and instrumentation technology
Námestie slobody 17
812 31 Bratislava