Neural Networks 24 (2011) 906–916

Contents lists available at SciVerse ScienceDirect

Neural Networks

journal homepage: www.elsevier.com/locate/neunet

2011 Special Issue

A hierarchical ART network for the stable incremental learning of topological structures and associations from noisy data

Marko Tscherepanow [a,*], Marco Kortkamp [a], Marc Kammer [a,b]

[a] Bielefeld University, Applied Informatics, Universitätsstraße 25, 33615 Bielefeld, Germany
[b] Bielefeld University, Cognitive Interaction Technology – Center of Excellence, Universitätsstraße 21–23, 33615 Bielefeld, Germany

ARTICLE INFO

Keywords:
Topology learning
Associative memory
Incremental learning
Hierarchical representations
Adaptive Resonance Theory

ABSTRACT

In this article, a novel unsupervised neural network combining elements from Adaptive Resonance Theory and topology-learning neural networks is presented. It enables stable on-line clustering of stationary and non-stationary input data by learning their inherent topology. Here, two network components representing two different levels of detail are trained simultaneously. By virtue of several filtering mechanisms, the sensitivity to noise is diminished, which renders the proposed network suitable for the application to real-world problems. Furthermore, we demonstrate that this network constitutes an excellent basis to learn and recall associations between real-world associative keys. Its incremental nature ensures that the capacity of the corresponding associative memory fits the amount of knowledge to be learnt. Moreover, the formed clusters efficiently represent the relations between the keys, even if noisy data is used for training. In addition, we present an iterative recall mechanism to retrieve stored information based on one of the associative keys used for training. As different levels of detail are learnt, the recall can be performed with different degrees of accuracy.

1. Introduction

For numerous tasks, the traditional off-line learning approach with separate training, validation, and test phases is not sufficient. The diagnosis of genetic abnormalities (Vigdor & Lerner, 2006), interactive teaching of a humanoid robot (Goerick et al., 2009), and the subcellular localisation of proteins (Tscherepanow, Jensen, & Kummert, 2008) constitute several examples for such problems. As a consequence, incremental on-line learning has become more popular in recent years, since such machine learning techniques are required to gradually complete knowledge or adapt to non-stationary input distributions.

In this article, the TopoART network (Tscherepanow, 2010) is presented. It combines incremental and fast on-line clustering with topology learning. As TopoART originates from Adaptive Resonance Theory (ART) networks, in particular Fuzzy ART (Carpenter, Grossberg, & Rosen, 1991), TopoART creates stable representations while retaining its ability to learn new data. In order to render TopoART more suitable for real-world applications, it was designed in such a way that it becomes insensitive to noise. Furthermore, it creates a hierarchical representation of the input distribution reflecting different levels of detail.

TopoART can be extended to a hierarchical hetero-associative memory called TopoART–AM. Here, an iterative recall mechanism provides missing keys in decreasing order of confidence. Due to the properties inherited from TopoART, namely insensitivity to noise as well as the ability of incremental and fast on-line learning, this associative memory is particularly well-suited to real-world applications.

Related approaches are discussed in Section 2. Afterwards, details of TopoART and its extension TopoART–AM are introduced in Section 3. In Section 4, the results of TopoART and TopoART–AM applied to different types of datasets are compared to several state-of-the-art methods. Here, their ability to cope with noise and to incrementally learn new input data from non-stationary distributions will be shown. In addition, the iterative recall mechanism of TopoART–AM will be demonstrated. Finally, Section 5 summarises the most important points of this article.

2. Related work

As we intend to solve two different types of problems using TopoART, namely clustering and the learning of associations, we discuss related work from both research fields.

* Tel.: +49 521 106 12222; fax: +49 521 106 2992.
E-mail addresses: marko@techfak.uni-bielefeld.de, marko@tscherepanow.de (M. Tscherepanow), mkortkam@techfak.uni-bielefeld.de (M. Kortkamp), mkammer@cit-ec.uni-bielefeld.de (M. Kammer).

0893-6080/$ – see front matter © 2011 Elsevier Ltd. All rights reserved.
doi:10.1016/j.neunet.2011.05.009

2.1. Clustering techniques

The k-means algorithm (MacQueen, 1967), which constitutes a very well-known unsupervised learning technique, determines a partitioning of an input distribution into *k* regions or rather clusters. Each cluster is represented by a reference vector. The choice of the number of required clusters constitutes a crucial problem. For this reason, the Linde–Buzo–Gray (LBG) algorithm (Linde, Buzo, & Gray, 1980) was developed. Based on a fixed training set, it successively computes sets of reference vectors of increasing size until a stopping criterion is fulfilled. The topological structure of the input data is not considered by this algorithm.

In 1982, the Self-Organising Feature Maps (SOFMs), which map input data to a lattice of neurons, were introduced by Kohonen. Here, the reference vectors are encoded by the weights of the neurons. The lattice possesses a predefined topological structure, the dimension of which is usually lower or equal to the dimension of the input space. If the input distribution is not completely known in advance, an appropriate lattice structure is difficult to choose. This problem was solved by the Growing Neural Gas (GNG) algorithm (Fritzke, 1994). It allows for the incremental incorporation of new neurons and the learning of the input distribution's topology by adding and deleting edges between different neurons.

The GNG algorithm is contained as a special case in a recently proposed extension, which is called the limited branching tree Growing Neural Gas (lbTreeGNG) (Kortkamp & Wachsmuth, 2010). It creates hierarchical codebooks that locally preserve the topology of the input space, while allowing a very efficient mapping from input samples to codewords and avoiding overfitting during training.

However, the above-mentioned methods do not directly employ mechanisms that deal with the *stability–plasticity dilemma* (Grossberg, 1987). A continuing presentation of input data results in a continuing adaptation of the neurons' weights, i.e. the reference vectors, and the network topology. Thus already-learnt structures may get altered or even lost. This can occur, for instance, if the input distribution is complex or due to small changes of the input probabilities. The sequencing of the input data may cause a similar effect.

Adaptive Resonance Theory (ART) networks have been proposed as a solution to the stability–plasticity dilemma (Grossberg, 1987). These networks learn top–down expectations which are matched with bottom–up input. The expectations, which are called categories, summarise sets of input data into clusters. Depending on the type of ART network, the categories exhibit different shapes such as a hyperspherical shape (Anagnostopoulos & Georgiopoulos, 2000), a hyperelliptical shape (Anagnostopoulos & Georgiopoulos, 2001), or a hyperrectangular shape (Carpenter et al., 1991). Besides enabling ART networks to create stable and plastic representations, the categories allow for an easy novelty detection. But in contrast to SOFMs and GNG, ART networks do not capture the topology of the input data. Furthermore, their ability of stable learning leads to an increased sensitivity to noise.

In 2006, the Self-Organising Incremental Neural Network (SOINN) was introduced by Furao and Hasegawa. Similar to GNG, SOINN clusters input data by incrementally adding neurons, the weights of which represent reference vectors, and the topology is reflected by edges between the nodes. But it has several additional features. First, SOINN has a two-layered structure representing the input distribution at different levels of detail. Additionally, this structure reduces the sensitivity to noise. The second layer is trained after the training of the first layer has been finished. Second, novelty detection can be performed based on an adaptive threshold. Third, each neuron has an individual learning rate which decays if the amount of input samples that it represents increases. In this way, a more stable representation is achieved. But the weights of the neurons do not stabilise completely. Furthermore, a high number of relevant parameters (eight parameters per layer) has to be set in order to apply SOINN.

The Enhanced Self-Organising Incremental Neural Network (ESOINN) (Furao, Ogura, & Hasegawa, 2007) solves some of the above-mentioned problems: by removing the second layer and one condition for the insertion of new neurons, the number of required parameters is considerably reduced (4 in total). Furthermore, the whole network can be trained on-line. But similar to SOINN, the weights do not stabilise completely. Moreover, ESOINN loses the ability to create hierarchical representations.

TopoART combines the advantages of ART and topology-learning networks (see Section 3.1). From its ART ancestors, it inherits the ability of fast and stable on-line learning using expectations (categories). These categories are extended by edges reflecting the topology of the input distribution. Therefore, they enable the formation of arbitrarily shaped clusters. In addition, TopoART adopts the ability to represent input data at different levels of detail from SOINN; but unlike SOINN, it learns both levels simultaneously.

2.2. Associative memories

There exist several approaches to associative memories, which are based on clustering methods. Some examples are the bidirectional hetero-associative memories of Chartier, Giguère, and Langlois (2009) and of Ichiki, Hagiwara, and Nakagawa (1993), which incorporate SOFMs, as well as SOIAM (Sudo, Sato, & Hasegawa, 2009), an associative memory based on a simplified version of SOINN. In contrast to traditional approaches such as Hopfield networks (Hopfield, 1982) and bidirectional associative memories (BAMs) (Kosko, 1988), they do not have to be trained with noise-free input patterns and perform information compression: the underlying clusterer summarises similar input samples to clusters, which may be considered as a simple type of categorisation. As a consequence, these approaches reduce the amount of data to be stored which is a major aspect of the principle of *cognitive economy* (Goldstone & Kersten, 2003). This is particularly beneficial for artificial agents such as robots operating in real-world environments, as they have to process large amounts of noisy and corrupted data.

The capacity of Hopfield networks and BAMs depends on the size of the associative keys (Hopfield, 1982; Kosko, 1988). After the maximum capacity has been reached, further training results in forgetting the previously learnt data. SOFM-based associative memories suffer from a similar problem, although they are capable of generalisation, which increases the capacity. Since the application of SOFMs requires the topology and network size to be chosen in advance (e.g., Chartier et al., 2009; Ichiki et al., 1993), the capacity of these methods is limited as well. Furthermore, SOFMs do not create stable representations. Hence, catastrophic forgetting might result from training with non-stationary data. In contrast, the capacity of SOIAM is not limited, as it is an incremental network. Its capacity rather fits the learnt knowledge. But similar to SOINN, the knowledge is not completely stable. Furthermore, since SOIAM is based on a one-layered version of SOINN, no hierarchical clustering is performed. This hierarchical clustering might have been beneficial for real-world tasks, as it enables the representation of further abstraction levels.

Another important aspect, which needs to be considered, is the type of information which can be processed. While Hopfield networks require binary input, BAMs allow for the storage of real-valued data. Associative memory models incorporating clustering techniques can be applied to real-valued data as well. But one data type, which typically occurs in real-world scenarios, is