

# Learning audio and image representations with bio-inspired trainable feature extractors

Nicola Strisciuglio\*

*\*Johann Bernoulli Institute for Mathematics and Computer Science, University of Groningen, Netherlands*

*This work was carried out at the University of Groningen (Netherlands) and at the University of Salerno (Italy)*

Received 24th Jul 2017; accepted 24th Nov 2017

## Abstract

Recent advancements in pattern recognition and signal processing concern the automatic learning of data representations from labeled training samples. Typical approaches are based on deep learning and convolutional neural networks, which require large amount of labeled training samples. In this work, we propose novel feature extractors that can be used to learn the representation of single prototype samples in an automatic configuration process. We employ the proposed feature extractors in applications of audio and image processing, and show their effectiveness on benchmark data sets.

---

## 1 Introduction

Since when very young, we can quickly learn new concepts, and distinguish between different kinds of object or sound. If we see a single object or hear a particular sound, we are then able to recognize such sample or even different versions of it in other scenarios. For example, if one sees an iron chair and associates the object to the general concept of “chairs”, he will be able to detect and recognize also wooden or wicker chairs. Similarly, when we hear the sound of a particular event, such as a scream, we are then able to recognize other kinds of scream that occur in different environments. We continuously learn representations of the real world, which we then use in order to understand new and changing environments.

In the field of pattern recognition, traditional methods are typically based on representations of the real world that require a careful design of a suitable feature set (i.e. a data representation), which involves considerable domain knowledge and effort by experts. Recently, approaches for automated learning of representations from training data were introduced. Representation learning aims at avoiding engineering of hand-crafted features and providing automatically learned features suitable for the recognition tasks. Nowadays, widely popular approaches for representation learning are based on deep learning techniques and convolutional neural networks (CNN). These techniques are very powerful, but are computationally expensive and require large amount of labeled training data to learn effective models for the applications at hand.

In this paper we report the main achievements included in the doctoral thesis titled ‘Bio-inspired algorithms for pattern recognition in audio and image processing’, in which we proposed novel approaches for *representation learning* for audio and image signals [5].

---

Correspondence to: Nicola Strisciuglio <n.strisciuglio@rug.nl>

Recommended for acceptance by Anjan Dutta and Carles Sánchez

<http://dx.doi.org/10.5565/rev/elcvia.1128>

ELCVIA ISSN:1577-5097

Published by Computer Vision Center / Universitat Autònoma de Barcelona, Barcelona, Spain

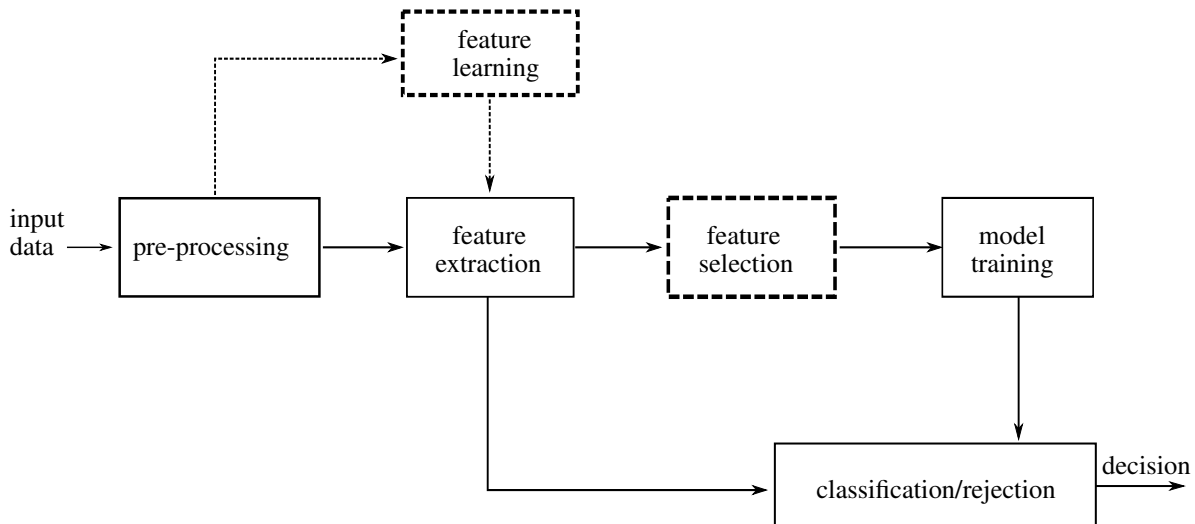


Figure 1: Overview of a pattern recognition system. The input data are pre-processed and then features are computed to extract important properties from such data. The features to be computed can be determined by an engineering process or can be learned from the data (*representation learning*). Feature selection procedures can be employed to determine a sub set of discriminant features that are then used to train a classifier, which determines a model of the training data. Such model is then used in the operating phase of the system, while a classifier takes decisions on the input data.

## 2 Motivation and contribution

Motivated by the fact that we can learn effective representations of a new category of objects or sounds from a single example and successively generalize to a wide range of real-world samples, we studied the possibility of learning data representations from small amounts of training samples. We investigated the design of feature extractors that can be automatically trained by showing single prototype samples, and employed them into pattern recognition systems to solve practical problems.

We proposed representation learning techniques for audio and image processing based on novel trainable feature extractors. The design and implementation of the proposed feature extractors are inspired by some functions of the human auditory and visual systems. The structure of the proposed feature extractors is learned from training samples in an automatic configuration step, rather than fixed a-priori in the implementation [5]. We employed the newly designed methodologies into systems for audio event detection and classification in noisy environments and for delineation of blood vessels in retinal fundus images. The contributions of this work can be listed as: *a*) novel bio-inspired trainable feature extractors for representation learning of audio and image signals, respectively called COPE and *B*-COSFIRE; *b*) a system for audio event detection based on COPE feature extractors; *c*) the release of two data sets of audio events of interest mixed to various background sounds and with different signal to noise ratios (SNR); *d*) a method for delineation of elongated and curvilinear patterns in images based on *B*-COSFIRE filters; *e*) feature selection mechanisms based on information theory and machine learning approaches.

## 3 Methods

We introduced a novel approach for representation learning, based on trainable feature extractors. We extended the traditional scheme of pattern recognition systems with feature learning algorithms (dashed box at the top of Figure 1), which construct a suitable representation of training data by automatic configuring a set of feature extractors.

We proposed trainable COPE (Combination of Peaks of Energy) feature extractors for sound analysis, that can be trained to detect any sound pattern of interest. In an automatic configuration process performed on a single prototype sound pattern, the structure of a COPE feature extractor is learned by modeling the constellation of peak points in a time-frequency representation of the input sound [10]. In the application phase, a COPE feature has high value when computed on the same sound used for configuration, but also to similar or corrupted versions of it due to noise or distortion. This accounts for generalization capabilities and robustness of detection of the patterns of interest. The response of a COPE feature extractor is computed as the combination of the weighted score of its constituent constellation of energy peaks. For further detail we refer the reader to [10]. For the design of COPE feature extractors, we were inspired by some functions of the cochlea membrane and the inner hair cells in the auditory system, which convert the sound pressure waves into neural stimuli on the auditory nerve. We employed COPE feature extractors together with a multi-class Support Vector Machine (SVM) classifier to perform audio event detection and classification, also in cases where sounds have null or negative SNR.

We proposed *B*-COSFIRE (that stands for Bar-selective Combination of Shifted Filter Responses) filters for detection of elongated and curvilinear patterns in images and apply them to the delineation of blood vessels in retinal images [1, 8]. The *B*-COSFIRE filters are trainable, that is their structure is automatically configured from prototype elongated patterns. The design of the *B*-COSFIRE filters is inspired by the functions of some neurons, called simple cells, in area V1 of the visual system, which fire when presented with line or contour stimuli. A *B*-COSFIRE filter achieves orientation selectivity by computing the weighted geometric mean of the output of a pool of Difference-of-Gaussians (DoG) filters, whose supports are aligned in a collinear manner. Rotation invariance is efficiently obtained by appropriate shiftings of the DoG filter responses. For further detail we refer the reader to [1].

After configuring a large bank of *B*-COSFIRE filters selective for vessels (i.e. lines) and vessel-endings (i.e. line-endings) of various thickness (i.e. scale), we proposed to use several approaches based on information theory and machine learning to select an optimal subset of *B*-COSFIRE filters for the vessel delineation task [7, 9]. We indicate this procedure with the dashed box named ‘feature learning’ in Figure 1. We consider the selected filters as feature extractors to construct a pixel-wise feature vector which we used in combination with a SVM classifier to classify the pixels in the testing image as *vessel* and *non-vessel* pixels.

## 4 Experiments and Results

We released two data sets for benchmark of audio event detection and classification methods, namely the MIVIA audio events [3] and the MIVIA road events [2] data sets. We reported baseline results (recognition rate of 86.7% and 82% on the two data sets) by using a real-time method for event detection that is based on an adaptation of the *bag of features* classification scheme to noisy audio streams [3, 2]. The results that we achieved by using COPE feature extractors show a considerable improvement with respect to the ones of the bag of features approach. We obtained a recognition rate of 95.38% on the MIVIA audio event and of 94% (with standard deviation on cross-validation experiments equal to 4.32) on the MIVIA road event data sets. We performed *t*-Student tests and observed a statistically significant improvement of the recognition rate with respect to baseline performance on both data sets.

We evaluated the performance of the proposed *B*-COSFIRE filters on four data sets of retinal fundus images for benchmarking of blood vessel segmentation algorithms, namely the DRIVE, STARE, CHASE\_DB1 and HRF data sets. The results that we achieved (DRIVE: Se = 0.7655, Sp = 0.9704; STARE: Se = 0.7716, Sp = 0.9701; CHASE\_DB1: Se = 0.7585, Sp = 0.9587; HRF: Se = 0.7511, Sp = 0.9745) are higher than the ones reported by many state-of-the-art methods based on filtering approaches. The filter selection procedure based on supervised learning that we proposed in [9] contributes to a statistically significant increase of performance results, which are higher than or comparable to the ones of other methods based on machine learning techniques.

We extended the application range of the *B*-COSFIRE filters to aerial images for the delineation of roads and rivers, natural and textured images [4], and to pavement and road surface images for the detection of cracks and

damages [6]. The results that we achieved are better than or comparable to the ones achieved by existing methods, which are usually designed to solve specific problems. The proposed *B*-COSFIRE filters demonstrated to be effective in various applications and with different types of images (retinal fundus photography, aerial photography, laser scans) for delineation of elongated and curvilinear patterns.

We studied the computational requirements of the proposed algorithms in order to evaluate their applicability in real-world applications and the fulfillment of real-time constraints given by the considered problems. The MATLAB implementation of the proposed algorithms are publicly released for research purposes\*.

## 5 Conclusions

In this work, we proposed novel trainable feature extractors and employed them in applications of sound and image processing. The trainable character of the proposed feature extractors is in that their structure is learned directly from training data in an automatic configuration process, rather than fixed in the implementation. This provides flexibility and adaptability of the proposed methods to different applications. The experimental results that we achieved, compared to the ones of other existing approaches, demonstrated the effectiveness of the proposed methods in various applications.

This work contributes to the development of techniques for *representation learning* in audio and image processing, suitable for domains where there is no availability of large amount of labeled training data.

## References

- [1] Azzopardi, G., Strisciuglio, N., Vento, M., Petkov, N.: Trainable COSFIRE filters for vessel delineation with application to retinal images. *Medical Image Analysis* 19(1), 46 – 57 (2015)
- [2] Foggia, P., Petkov, N., Saggese, A., Strisciuglio, N., Vento, M.: Audio surveillance of roads: A system for detecting anomalous sounds. *IEEE Trans. Intell. Transp. Syst.* 17(1), 279–288 (2016)
- [3] Foggia, P., Petkov, N., Saggese, A., Strisciuglio, N., Vento, M.: Reliable detection of audio events in highly noisy environments. *Pattern Recogn. Lett.* 65, 22 – 28 (2015)
- [4] Strisciuglio, N., Petkov, N.: Delineation of line patterns in images using b-cosfire filters. In: 2017 International Conference and Workshop on Bioinspired Intelligence (IWOB). pp. 1–6 (July 2017)
- [5] Strisciuglio, N.: Bio-inspired algorithms for pattern recognition in audio and image processing. University of Groningen (2016), [http://www.cs.rug.nl/~nick/strisciuglio\\_phd.pdf](http://www.cs.rug.nl/~nick/strisciuglio_phd.pdf)
- [6] Strisciuglio, N., Azzopardi, G., Petkov, N.: Detection of curved lines with b-cosfire filters: A case study on crack delineation. In: CAIP 2017, pp. 108–120 (2017)
- [7] Strisciuglio, N., Azzopardi, G., Vento, M., Petkov, N.: Multiscale blood vessel delineation using B-COSFIRE filters. In: CAIP, LNCS, vol. 9257, pp. 300–312 (2015)
- [8] Strisciuglio, N., Azzopardi, G., Vento, M., Petkov, N.: Unsupervised delineation of the vessel tree in retinal fundus images. In: VIPIMAGE, pp. 149–155 (2015)
- [9] Strisciuglio, N., Azzopardi, G., Vento, M., Petkov, N.: Supervised vessel delineation in retinal fundus images with the automatic selection of B-COSFIRE filters. *Mach. Vis. Appl.* pp. 1–13 (2016)
- [10] Strisciuglio, N., Vento, M., Petkov, N.: Bio-inspired filters for audio analysis. In: BrainComp 2015, Revised Selected Papers. pp. 101–115 (2016)

---

\*The code is available from the GitLab repositories at <http://gitlab.com/nicstrisc>