Deep Learning Techniques for Spacecraft Telemetries Anomaly Detection

Giorgio De Magistris[1], Carlo Ciancarelli[2], Salvatore Cognetta[2]
Luca Iocchi[1], Christian Napoli[1], Daniele Nardi[1]

[1] Department of Computer, Automation and Management Engineering, Sapienza University of Rome
[2] Thales Alenia Space Italia

{cnapoli, demagistris, iocchi, nardi}@diag.uniroma1.it; carlo.ciancarelli@thalesaleniaspace.com

Abstract

The purpose of this project is to devise new deep learning techniques for the analysis of spacecraft telemetries aimed at identifying anomalies. In the context of space missions, the task of anomaly detection is of paramount importance. If anomalies are not handled properly the life of the spacecraft could be shortened, causing mission failures and huge losses to the company and the society. Existing anomaly detection systems are based on simple algorithms, and the introduction of deep learning techniques is a hot research topic in the field.

1 Objectives of the Project

Anomaly detection is a well studied problem in many fields of computer science, as well as it requires solutions that make use of machine learning, deep learning and data mining. The detection of anomalous behaviors can become a crucial asset for many applications such as network intrusion detection, credit card fraud detection, business processes, medical diagnosis, etc. The difficulty of this task arises from the unavailability of training examples, because anomalies are by definition rare events. While much research has been devoted to anomaly detection on record data and symbolic sequences, less attention has been dedicated to univariate and multivariate time series [Chandola, 2009]. This project aims to move a step forward to bridge this gap with the introduction of new anomaly detection techniques for multivariate time series applied to spacecraft telemetries anomaly detection. In particular, a new trend in anomaly detection based on the Generative Adversarial Networks (GANs) [Goodfellow *et al.*, 2014] has emerged in the last few years, but its employment on time series is still limited. In this direction we propose the application of this framework on alternative representations of time series, in order to devise anomalous temporal correlations jointly with advanced classification systems.

The GANs Framework, firstly introduced in [Goodfellow *et al.*, 2014], is composed by two networks: a generator and a discriminator, often referred as G and D. The generator learns a mapping from the latent space Θ_z, usually the set of k-dimensional standard normal vectors, to the data space Θ_{data} and the discriminator learns a categorical probability distribution over the generated and real samples to discriminate between real and fake samples. G and D optimize the same criterion in opposite directions, following a two player minimax game. At the optimum the two distributions $x \sim G(z)|_{z \sim \Theta_z}$ and Θ_{data} are the same and the expected value returned by the discriminator over real and generated samples is always one half.

The general idea behind the application of GANs to anomaly detection consists in using the output of the discriminator and the reconstruction error to assign an anomaly score to a data sample x. The output of the discriminator is the probability that a sample is “real", hence its inverse can be directly interpreted as an anomaly score. The second part is less obvious and its justification relies on the structure of the latent space. In [Radford *et al.*, 2015] it is shown that the space learned by the generator has smooth transitions, because walking on the learned manifold results in semantic changes to the generated image. This encouraged the usage of the GANs framework as an unsupervised features extractor through an inverse mapping from data space to the latent space. For example, in [Schlegl *et al.*, 2017] the inverse mapping is defined as an iterative search in the latent space, where the optimization criterion is the maximization of a similarity measure between x and $G(z)$ w.r.t. z. The authors used this similarity, called "reconstruction error", to assign an anomaly score to medical images such that samples with higher reconstruction errors got a higher anomaly score. This was based on the assumption that the generator G, trained only on normal images, learns the manifold of normal samples, hence it is not able to reproduce the anomalies. The main drawback of this approach is that it requires each time an iterative search to assign an anomaly score, hence it is not adequate for online or real-time applications. However different solutions have been devised to come up with this problem by introducing directly in the GANs framework an inverse mapping from data to latent space [Zenati *et al.*, 2018][Geiger *et al.*, 2020].

While these GAN-based approaches have been originally designed to work for images, on the other hand, the original GANs framework does not impose any restriction on the architectures of G and D, hence they can be adapted for sequential data. In particular [Mogren, 2016] and [Esteban *et al.*, 2017] are successful examples of sequential data generation with GANs. They both used Recurrent Neural Networks (RNN) for the generator and the discriminator. These works encouraged the extension of the GAN-based anomaly detection techniques for time series. For example [Li *et al.*, 2018] and [Geiger *et al.*, 2020] proposed two unsupervised anomaly detection systems for multivariate time series where the generators and discriminators are recurrent networks and the anomaly scores are obtained from the output of the discriminator and from the reconstruction error.

The research on spatial representations of time series is justified by the success of the *Convolutional Neural Networks* (CNN) in computer vision and by their reduced training time with respect to *Recurrent Neural Networks* (RNN). For these reasons this project aims at combining existing and custom spatial representations of time series with this GANs framework for anomaly detection.

For example the *spectrogram*, that is one of the outmost basilar 2D representation of a time series, has been applied in literature for time series classification [Huang *et al.*, 2019][Costa *et al.*, 2017].

Another example is the novel spatial representation introduced in [Wang e Oates, 2015], where the authors obtained state of the art results for supervised classification using CNNs.

2 Impact and Innovation

In recent times data have become very important resources. In the world there are billions of IoT devices that continuously produce streams of data monitoring industrial processes, personal health, smart cities, smart homes, etc. These data become valuable only if a solution can be developed in order to extract meaningful information. When the amount of data is huge and it needs to be processed in very short time, such a solution must be based on fully automatized processes. When the data are collected for monitoring purposes, then the anomalous patterns are the information that we need to extract. The anomaly detection in temporally correlated data is very common, some examples are intrusion or cyber attack detection from a sequence of instruction, fraud detection in bank transactions, anomaly detection in climatic data, just to cite a few. Sometimes the data assumes the form of a multivariate time series, this is a common situation in industrial systems where the data are acquired by multiple sensors.

In the context of space missions, the task of anomaly detection is of paramount importance. In particular, spacecrafts are complex machines, with thousands of components and are designed to remain in continuous operation for a long period of time. If anomalies are not handled properly the life of these devices could be shortened, causing mission failures and huge losses to the company and the society. In fact, repairs are so expensive that it is often preferred to abort the mission. Given the amount of monitoring data produced by such devices, a manual analysis is not feasible. Since a few years ago, the most widely used approaches were based on simple Out of Limit (OOL) checks [Martinez, 2012], meaning that when the signal exceeded some predefined upper and lower bounds an alarm was triggered. Later on, more advanced solutions introduced clustering techniques on multidimensional vectors obtained by manually extracted features, like the *Inductive System Health Monitoring* developed at NASA [Iverson, 2004] and the *Automated Telemetry Health Monitoring System* (ATHMoS) developed at the German Space Operation Center (GSOC) [OMeara *et al.*, 2016]. However the application of deep learning techniques in this field is still limited [OMeara *et al.*, 2018][Martinez e Donati, 2018][Hundman *et al.*, 2018].

Riferimenti bibliografici

[Chandola, 2009] Varun Chandola. Anomaly detection for symbolic sequences and time series data. 2009.

[Costa *et al.*, 2017] Yandre MG Costa, Luiz S Oliveira, e Carlos N Silla Jr. An evaluation of convolutional neural networks for music classification using spectrograms. *Applied soft computing*, 52:28–38, 2017.

[Esteban *et al.*, 2017] Cristóbal Esteban, Stephanie L Hyland, e Gunnar Rätsch. Real-valued (medical) time series generation with recurrent conditional gans. *arXiv preprint arXiv:1706.02633*, 2017.

[Geiger *et al.*, 2020] Alexander Geiger, Dongyu Liu, Sarah Alnegheimish, Alfredo Cuesta-Infante, e Kalyan Veeramachaneni. Tadgan: Time series anomaly detection using generative adversarial networks. *arXiv preprint arXiv:2009.07769*, 2020.

[Goodfellow *et al.*, 2014] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, e Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.

[Huang *et al.*, 2019] Jingshan Huang, Binqiang Chen, Bin Yao, e Wangpeng He. Ecg arrhythmia classification using stft-based spectrogram and convolutional neural network. *IEEE Access*, 7:92871–92880, 2019.

[Hundman *et al.*, 2018] Kyle Hundman, Valentino Constantinou, Christopher Laporte, Ian Colwell, e Tom Soderstrom. Detecting spacecraft anomalies using lstms and nonparametric dynamic thresholding. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 387–395, 2018.

[Iverson, 2004] David L Iverson. Inductive system health monitoring. In *IC-AI*, pages 605–611, 2004.

[Li *et al.*, 2018] Dan Li, Dacheng Chen, Jonathan Goh, e See-kiong Ng. Anomaly detection with generative adversarial networks for multivariate time series. *arXiv preprint arXiv:1809.04758*, 2018.

[Martinez e Donati, 2018] Jose Martinez e Alessandro Donati. Novelty detection with deep learning. In *2018 SpaceOps Conference*, page 2560, 2018.

[Martinez, 2012] Jose Martinez. New telemetry monitoring paradigm with novelty detection. In *SpaceOps 2012*, page 1275123. 2012.

[Mogren, 2016] Olof Mogren. C-rnn-gan: Continuous recurrent neural networks with adversarial training. *arXiv preprint arXiv:1611.09904*, 2016.

[OMeara *et al.*, 2016] Corey OMeara, Leonard Schlag, Luisa Faltenbacher, e Martin Wickler. Athmos: Automated telemetry health monitoring system at gsoc using outlier detection and supervised machine learning. In *14th International Conference on Space Operations*, page 2347, 2016.

[OMeara *et al.*, 2018] Corey OMeara, Leonard Schlag, e Martin Wickler. Applications of deep learning neural networks to satellite telemetry monitoring. In *2018 SpaceOps Conference*, page 2558, 2018.

[Radford *et al.*, 2015] Alec Radford, Luke Metz, e Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.

[Schlegl *et al.*, 2017] Thomas Schlegl, Philipp Seeböck, Sebastian M Waldstein, Ursula Schmidt-Erfurth, e Georg Langs. Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In *International conference on information processing in medical imaging*, pages 146–157. Springer, 2017.

[Wang e Oates, 2015] Zhiguang Wang e Tim Oates. Imaging time-series to improve classification and imputation. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.

[Zenati *et al.*, 2018] Houssam Zenati, Chuan Sheng Foo, Bruno Lecouat, Gaurav Manek, e Vijay Ramaseshan Chandrasekhar. Efficient gan-based anomaly detection. *arXiv preprint arXiv:1802.06222*, 2018.