Incremental Learning for Classification of Objects of Interest

Gilles LE CHENADEC, Quentin OLIVAU, Benoit CLEMENT
Panagiotis PAPADAKIS, Karl SAMMUT

April 2019

Chapter 1

Context

1.1 Abstract

The proposed project resulting from a collaboration between Lab-STICC, Naval Group and Flinders University concerns incremental learning for the automatic classification of objects or particular environments represented by 2D or 3D images as well as sonar data. Three scenarios are considered in this project proposal:

- An aerial drone and the classification of ships from RGB images.
- The problem of an autonomous underwater seabed mapping and object detection form sonar data.
- A mobile, indoor robot that exploits RGB-D (colour and depth) sensory data.

In this context, incremental learning will focus on learning objects or events belonging to new categories and/or new instances of objects or events already included in the prediction model. This project aims to develop new robust incremental learning methods for embedded systems in a terrestrial, marine or underwater context.

1.2 Naval Group, Lab-STICC and Flinders University

This project is the result of a cooperation between Naval Group Research, Lab-STICC UMR CNRS 6285 in Brest (involving IMT Atlantique and ENSTA Bretagne) and Flinders University (Adelaide, Australia), under the program CIFRE (`http://www.anrt.asso.fr/fr/cifre-7843`), involving a shared stay between France and Australia.

1.3 Supervisory team

- Gilles Le Chenadec is an Associate Professor in ENSTA Bretagne, Brest. His research interests include seafloor classification, underwater acoustics and robotics and machine learning.
- Quentin Oliveau is an engineer in Naval Group.
- Panagiotis Papadakis is an Associate Professor of informatics and robotics in IMT Atlantique Bretagne/Pays de la Loire, Brest, with a background on computer vision for rigid and articulated 3D objects, cognitive robotics and robotic systems engineering. The main application areas are assistive and service robotics, notably, urban search and rescue (personal web-site `https://sites.google.com/site/pgpapadakis/`).

- Benoit Clement is a Professor of Control in ENSTA Bretagne, Brest. His current research focus are on Global Optimization used for structured robust control applied to robotics and also Machine Learning for adaptive control. Since 2017, he has served as deputy head of the Lab-STICC (CNRS UMR 6285).
- Karl Sammut is a Professor in Flinders University, Adelaide, Australia. He is currently a Professor with the Engineering Discipline in the College of Science and Engineering. His primary research activities are in the area of autonomous marine vehicles, including mission planning, vehicle navigation, guidance and control, and situation awareness. Prof Sammut is the Director of the Centre for Maritime Engineering at Flinders University.

1.4 Application

1.4.1 Candidate Profile

Holder of (or near graduation) of a postgraduate diploma, Master of research or engineer diploma in the domains of Computer Science, Robotics or equivalent. The candidate is expected to have a federating role between the collaborating teams and to be strongly motivated along with excellent communication skills.

Theoretical skills Machine learning, computer vision, image processing, cognitive robotics.
Technical skills C++/Python programming, numpy, Scikit, deep learning frameworks, OpenCV, PCL, etc.
Fluency in English is strongly required.

1.4.2 How to apply

Interested applicants should contact Quentin Oliveau (`quentin.oliveau@naval-group.com`) and Benoit Clement (`benoit.clement@ensta-bretagne.fr`) by email with the reference [ILCOI].

Applications must include:

- a detailed resume,
- a short one-page CV,
- academic transcripts,
- a list of (at least) two academic references.

Starting date: Sept./Oct 2019, for 3 years.
Application deadline: until filled.

Chapter 2

Scientific Project

2.1 Summary

While deep neural network models have been consistently improving prediction performance in many areas, their performance in marine and underwater application areas remains currently limited. This is partially due to restricted access to annotated data-sets, the fact that the marine environment is extremely fluctuating, poorly known globally and parsimoniously observed. Finally, the systematic use of marine or underwater drones involves limited computational resources and embedded memory making it difficult to develop effective adaptive models.

The proposed project resulting from a collaboration between Lab-STICC and Naval Group concerns incremental learning for the automatic classification of objects or particular environments represented by 2D or 3D images. Aiming for a solution embedded on a (aerial or underwater) drone, we seek to develop algorithms with a real-time update capability, using new potentially labeled images belonging to already learned or novel categories. This update must be performed without storing the examples seen previously (or by limiting this storage) and without the need for relearning from scratch. Incremental learning refers to learning a model of category prediction, for example from data provided to the drone as a mission, where the primary goal amounts to improving or preserving the performance of an initial prediction model potentially learned beforehand.

Different prediction models will be explored in this project. The use of deep neural networks will be investigated for their performance and structural complexity even if they are prone to the phenomenon of "catastrophic forgetting". Other models for representing the data in a compact way will also be investigated, such those based on wide margins or decision trees.

When a drone operates with limited resources and sensory data arriving at constant flux, a strategy of adaptation of the model is required. The question of "when and how to adapt the model to evolve and/or create a new category" will also be a central issue of this project. The creation of a new category begins with the definition of decision criteria for the appearance of this new category and a modification of the training architecture. This suggests to regard the problem of automatic classification firstly as a novelty detection problem allowing to optimize the classification not only between the known categories but also with respect to the unknown categories.

Three scenarios are considered in this project proposal. The first scenario concerns an aerial drone and the classification of ships; a communication link with a human allows him to refine the knowledge of the drone and the respective model. The second scenario considers the problem of an autonomous underwater seabed mapping and object detection drone that has to adapt to the appearance of new events. The third one concerns a terrestrial robot that exploits RGB-D (colour and depth) sensory data. In these scenarios, incremental learning will focus on two possible levels: learning objects or events belonging to new categories and/or new instances of objects or events already included in the prediction model. As such, this thesis project aims to develop new robust incremental learning methods for autonomous or semi-autonomous embedded systems in a terrestrial, marine or underwater context.

These algorithms will be implemented and evaluated in real embedded systems with limited resources (Raspberry pi, NVIDIA JETSON TX2, etc).

2.2 Assumptions, posed questions, identification of blocking points

The main challenges that will have to be jointly addressed during the development of an incremental learning algorithm [11] concern:

- The choice of the classification algorithm (deep neural network, SVM, random forest, etc.) and the representation of the information that the algorithm implies (and thus the memory allocated in the drone).
- The evolution of the data statistics during the mission, resulting in an adaptation or improvement of the model for the existing categories or a possible creation of new categories (anomaly detection or novelty detection).
- The adaptation of the current model of prediction (when and how the model is adapted), especially the choice of the complexity (the architecture) of the model (adaptive or fixed) with the aim of maintaining the performance level of the previous models (i.e. alleviating catastrophic forgetting).
- The definition of the meta-parameters necessary to employ the prediction models (e.g. the learning rate in the case of a neural network).
- These issues need to be calibrated according to the system and its resources.

Different families of algorithms will be considered to address the above challenges, as posed by the envisioned scenarios and on two or three architectures (with different resources) of embedded systems. Three scenarios (use cases) will be considered, as identified by corresponding data-sets that will allow the evaluation of the developed methods, namely:

- An aerial drone embeds a classification algorithm of ships (sailboat, cargo, etc) and is equipped with a camera. While operating on a mission and upon spotting of a ship, the drone will assign a category and transmit this result as well as its confidence to a land-based operator that inspects the items returned by the drone. The operator, by validating (or invalidating) the classification made by the drone, will automatically re-inject these examples of newly labeled images and update in real time the algorithm implemented in the drone.
- An underwater drone incorporates a sonar sensor and an algorithm for classifying seabed and submarine objects. As the mission progresses, the drone is meant to autonomously improve its detection and classification of the encountered events (refine and create categories) and thus update its model in real time.
- A mobile robot identifies and classifies encountered 3D objects in order to offer services such as object grasping and inspection or semantic mapping of the environment. Incremental learning here can amount to the learning of objects belonging to new classes or new instances of already known object categories. To resolve ambiguities due to increased levels of uncertainty, the robot should be able to update its prediction model by soliciting the end-user, for example via relevance feedback.

Finally, the targeted embedded architectures will be defined to evaluate the potential and the performance of the developed methods in an operational framework. The Raspberry Pi, NVIDIA Jetson and AGX Xavier cards are considered as representative candidates. The systems' inputs will first be analyzed to automate the decision of when and how to relearn a model. Incremental learning methods based on Convolutional Neural Network (CNN) algorithms and their adaptation to the problem will be investigated. In view of the previously described operational use-cases, algorithms other than CNNs are likely to be more suitable in an incremental learning context (see [10] for the robotics domain) and they will minimally be the object of a comparative study.

2.3 Workplan

This work will be based and motivated via a comprehensive bibliographic synthesis on incremental learning methods adapted to the problem. It will be orientated on novelty detection and incremental learning, with a particular interest torward Convolutional Neural Network (CNN) algorithms. Novelty detection will be required for scenarios where the drone / robot is not supervised by an operator for the designation of a new class, thus requiring an automatic strategy. The classification model embedded by the robot will condition the way of detecting the novelty, for example, through one-class SVM, kPCA, Neural Networks or others.

In the case where the selected incremental learning approach is based on the selection of examples representative of their respective categories, one will try to determine appropriate examples selection strategies and estimate the impact of the number of selected examples on the final performance. Indeed, we can consider that the most representative examples of a category are examples close to the average of all the examples (see [3] for example). These "average" examples are also often the easiest to classify. It can be assumed that integrating examples close to the boundaries of the classifier, and therefore more difficult to classify, could make it possible to improve the robustness of the algorithm. To verify this hypothesis, we can rely on the state-of-the-art for sample selection methods implemented in applications like "active learning" and integrate these methods into an existing incremental learning strategy such as that proposed by Castro et al. [3]. Instead of storing in memory a limited number of examples of interest corresponding to the old categories [3, 5, 6] other methods rely on the generation of synthetic characteristic examples of these old categories [4, 9].

Given the embedded nature of the application cases, the work will focus on the development of algorithms that either do not change the structure of the initial neural network or do so minimally (as opposed to [7] for example). This will safeguard against an explosion in the amount of memory available for the network, or even allow to benefit from an effective implementation of its initial architecture (FPGA for example). This strategy, on the other hand, has the constraint of possibly imposing an upper limit on the number of total categories that the system can learn. For methods that are based on the extraction of pre-learned features [8], we can explore a learning method integrating discriminant feature learning and classification. To limit the consequences of the phenomenon of catastrophic forgetting, the implementation of incremental learning algorithms will go through the integration of a knowledge distillation step in the process of learning new classes.

Upon specification of the evaluation protocols (namely, addition of new examples of known categories or addition of examples of new categories) and provision of associated data, it will be possible to perform rapid testing of existing methods. In this respect, popular data-sets in the image processing community are expected to be used as benchmarks (imageNet, CIFAR for 2D images, modelNet, SUN RGBD for 3D objects, etc). The scenarios of interest can be structured with an increasing level of complexity, namely, by firstly considering a semi-autonomous situation where a human operator certifies the label of the images before relearning and finally, pushing ahead the level of autonomy by a strategy for autonomously deciding for the creation of a new category.

Bibliography

[1] Ian J. Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville and Yoshua Bengio, “An Empirical Investigation of Catastrophic Forgetting in Gradient-Based Neural Networks”, ICLR 2014

[2] James Kirkpatrick et al. “Overcoming catastrophic forgetting in neural networks”, PNAS 2017

[3] Francisco M. Castro, Manuel J. Marin-Jimenez, Nicolas Guil, Cordelia Schmid and Karteek Alahari, “End-to-End Incremental Learning”, ECCV 2018

[4] Ronald Kemker and Christopher Kanan, “FearNet: Brain-Inspired Model for Incremental Learning”, ICLR 2018

[5] Rebuffi, S.-A., Kolesnikov, A., Sperl, G. and Lampert, C. H., “iCaRL: Incremental Classifier and Representation Learning”, CVPR 2017

[6] Shmelkov, K., Schmid, C. and Alahari, K., “Incremental Learning of Object Detectors without Catastrophic Forgetting” ICCV 2017

[7] Sarwar, S. S., Ankit, A. and Roy, K. “Incremental Learning in Deep Convolutional Neural Networks Using Partial Network Sharing” 2017

[8] Hacene, G. B., Gripon, V., Farrugia, N., Arzel, M., Jezequel, M. “Transfer Incremental Learning using Data Augmentation”, 2018

[9] Atkinson, C., McCane, B., Szymanski, L. and Robins, A. “Pseudo-Recursal: Solving the Catastrophic Forgetting Problem in Deep Neural Networks”, 2018

[10] Ulrich Nehmzow, Yiannis Gatsoulis, Emmett Kerr, Joan Condell, Nazmul Siddique, T. Martin McGuinnity, « Novelty Detection as an Intrinsic Motivation for Cumulative Learning Robots », in Book Ch., Intrinsically Motivated Learning in Natural and Artificial Systems, Springer, pp. 185-207, 2013.

[11] Gepperth, A. and Hammer, B. , “Incremental learning algorithms and applications “, Computational Intelligence, p. 12, 2016.

[12] Pfulb et al., A Comprehensive Application-oriented Study of Catastrophic Forgetting in DNNs, International Conference on Learning Representations 2019, https://openreview.net/pdf?id=BkloRs0qK7