DETECTION OF TRAFFIC SIGNS AND SEGMENTATION OF BRAIN TUMORS IN MAGNETIC RESONANCE IMAGES - COMPUTER VISION APPLICATIONS

Filipe Silva Carvalho[1]
a2021148720@isec.pt
Pedro Caridade[2, 3]
pedrojcaridade@uc.pt
Emília Bigotte de Almeida[1]
ebigotte@isec.pt
Verónica Vasconcelos[1,4]
veronica@isec.pt

[1]Polytechnic of Coimbra, Coimbra Institute of Engineering
[2] University of Coimbra
[3] Primelayer
[4] INESC-TEC: Institute for Systems and Computer Engineering, Technology and Science

Abstract

In this paper, two projects combining computer vision and artificial intelligence are presented. The first one focuses on traffic sign detection in static images and in real-time, which can enhance vehicle automation and increase road safety. The second one corresponds to the segmentation of magnetic resonance images for brain tumors. This procedure is essential for the early detection, diagnosis, therapy, and follow-up of brain tumor patients. Both projects yielded encouraging outcomes.

1 Introduction

Computer Vision (CV) is a field of artificial intelligence that develops techniques and algorithms enabling machines to interpret images and videos in a manner like humans. It extracts relevant information such as object recognition, pattern detection, object segmentation, and motion tracking, using techniques and algorithms including image processing and machine Learning (ML), among other methods. Its applications are extensive, covering medicine, industry, robotics, security, autonomous vehicles, and many others. With advancements in computational capacity and algorithms, CV continues to evolve as a research and development area, driven by large datasets.

Machine Learning is a subfield of Artificial Intelligence that focuses on the development of algorithms and techniques that enable computers to learn from data and make decisions without the need for explicit programming. Models are at the core of Machine Learning, representing mathematical structures that capture information and patterns from the data they are trained on. As they are fed more data and information, these models learn and generalize, becoming increasingly capable of making predictions or decisions based on new inputs. There are four types of Machine Learning: Supervised, Unsupervised, Semi-supervised, and Reinforcement [1].

Artificial Neural Networks (ANNs) are computer models inspired by the human nervous system, developed to solve complex problems, considering several variables. There are different types of ANNs, with the Convolutional Neural Network (CNN) being widely used in CV to identify, locate, and segment objects in images.

ANNs work in interconnected layers, including input and intermediate layers. Each processing unit in an ANN receives input signals, multiplies them by the corresponding weights (which represent the importance of the features), sums these weighted values and applies a non-linear activation function. This enables the network to learn and represent complex relationships and make decisions based on the results. TensorFlow, an open-source library developed by Google, is widely used for Machine Learning, optimizing calculations in ANNs, and taking advantage of specialized hardware, such as GPUs and TPUs, to speed up the process. The Python programming language is often used to implement these functionalities.

The described work applies CV to two different problems: the detection and segmentation of traffic signs in images or videos and the segmentation of tumors in brain images in Magnetic Resonance Imaging (MRI).

The detection and segmentation of traffic signs through CV are essential, especially in cities that aim to be sustainable. We can highlight some applications:

- Navigation and Traffic Alerts: In vehicle navigation systems can use traffic sign detection to provide real-time information on speed limits, overtaking restrictions, and other traffic information.
- Driver assistance: Driver assistance systems use CV to detect and recognize traffic signs, such as speed limit signs, stop signs, and traffic lights. This helps drivers keep up to date with road conditions and comply with traffic rules.
- Adaptive cruise control: Adaptive cruise control systems use traffic sign detection to automatically adjust the vehicle's speed according to the detected speed limits. This helps keep drivers within legal speed limits and reduces the likelihood of accidents.
- Detection of Vehicles at Intersections and Traffic Lights: Traffic sign detection can also be used along with vehicle detection to improve safety at intersections and traffic lights. These systems can identify traffic lights and their states to help coordinate traffic efficiently and to avoid collisions.
- Pedestrian safety: Traffic signal detection can be combined with pedestrian detection to ensure that vehicles give way to pedestrians at crosswalks in accordance with traffic signals.

These applications demonstrate how CV plays a key role in improving road safety and vehicle automation, making roads safer and more efficient. As technology continues to evolve, these systems are expected to become increasingly sophisticated and widely adopted.

The application of brain tumor segmentation in magnetic resonance images (MRI) is extremely important in the medical field and has several significant implications:

- Early and accurate diagnosis: Brain tumor segmentation helps radiologists and oncologists identify the presence and location of tumors in the brain. This enables faster and more accurate diagnosis, allowing patients to receive appropriate treatment earlier, which is crucial for increasing the chances of successful treatment.
- Surgical planning: Segmentation of brain tumors is fundamental for surgical planning. It provides detailed information on the location and extent of the tumor, allowing surgeons to plan more precise and less invasive surgical procedures, minimizing damage to surrounding healthy tissue.
- Radiotherapy, targeted, and monitoring treatment: For treatments such as radiotherapy and targeted therapies, accurate segmentation of tumors is essential. This helps to concentrate the radiation or therapy on the tumors, minimizing the exposure of normal tissues and reducing side effects. Throughout treatment, the segmentation of tumors in MRI images is used to monitor the tumor's response to therapies. This allows doctors to assess whether the treatment is effective or whether adjustments are needed.
- Medical research: The segmentation of brain tumors is also fundamental for clinical and scientific research. It helps to better understand the progression of tumors, the effectiveness of different treatments, and the specific characteristics of each type of tumor, contributing to the development of more effective therapies in the future.
- Personalization of treatment: With precise segmentation, treatments can be customized according to the individual characteristics of each tumor, considering its location, size, and biological characteristics. This is key to optimizing treatment results and minimizing side effects.
- Improved quality of life: By identifying and treating brain tumors effectively, MRI image segmentation helps to improve patients' quality of life, reducing the risk of neurological complications and ensuring proper treatment.

In summary, the application of brain tumor segmentation in MRI images plays a crucial role in the early detection, diagnosis, treatment, and monitoring of patients with brain tumors. It contributes significantly to the success of treatment, the quality of life of patients, and the advancement of medical research in the field of oncology.

2 Detection of Traffic Signs

A system was developed capable of detecting traffic signs in still images and videos in real-time. To do that, it was used the Darknet version of the YOLO v4 (You Only Look Once) algorithm [2]. YOLO is an object detection approach that efficiently identifies and delimits objects in images in real-time. It creates a bounding box in turn of each detected signal with the class name and a value between 0.0 and 1.0. These values represent the model's confidence level in detecting each object, where 1.0 is the maximum confidence value. In addition to tests being carried out on videos and images acquired by the authors, a dataset is also used with 4928 images of size 416 x 416 pixels of Traffic Signs, divided into training and test sets. The training set contains 90% of the images and the test set contains 10% of the data to test the model, provided by the Roboflow platform [3].

The MAP (Mean Average Precision) metric was used to evaluate the model, which is common in the detection and localization of objects in tasks of image processing and CV. Furthermore, MAP considers average accuracy for different confidence levels. The YOLO algorithm was trained to detect traffic signs for 2000 iterations. The results are shown in the graph in Figure 1. After training, the model achieved a MAP of 94.4% and an average loss of 0.3174. A lower loss is desirable, indicating closer proximity to correct predictions.

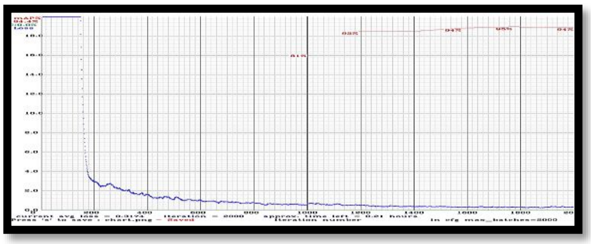

Figure 1: Training progress of the YOLO algorithm with 2000 iterations.

The promising results reflect the power of the YOLO algorithm and the success of the training approach, which can be put into practice on images such as those shown in Figure 2 or on videos acquired in real time.

Figure 2: Results of the application of the traffic sign detection algorithm. (image from Coimbra).

3 Segmentation of Brain Tumors in MRI

The U-Net architecture [4] was used to segment tumors in MRI. This architecture is a CNN widely used in image segmentation and medical image processing tasks and is designed to handle situations in which is important to extract detailed and contextualized information from the images. It used a dataset of 2828 MRI-Brain images with ground-true, available in the Kaggle platform [5]. Training the model took around 8.5 hours on a 6GB NVIDIA-1660 Ti GPU. A script was used to divide the dataset into 90% for training and 10% for testing the model.

After training, significant progress was seen in the following performance metrics:

- Binary Accuracy: Measures the proportion of correct predictions made by the model in relation to the total number of predictions (Figure 3). The value obtained was 99.79%.
- Validation Binary Accuracy: This metric counts the percentage of accurate predictions the model makes on a dataset that was not used during training. With a score of 99.78%, the model is successfully predicting unknown data.
- Dice Coefficient: Calculates how similar two regions are. It is particularly helpful for determining how much a prediction region (the segmentation) overlaps with the reference region. The total percentage of 87.57% indicates that the locations that the model identified are fairly like the ground truth.
- Jaccard Index: It is also a common metric in image segmentation tasks, measuring the similarity between the area predicted and the area segmented by the model. A value of 80.18% was obtained.

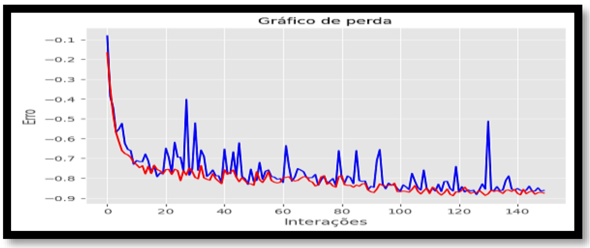

Figure 3: Training progress based on incorrect predictions.

Based on the results, it can be concluded that the U-Net model trained with TensorFlow showed good segmentation of tumors in MRI images, as can be seen in Figure 4 and from the results presented above.

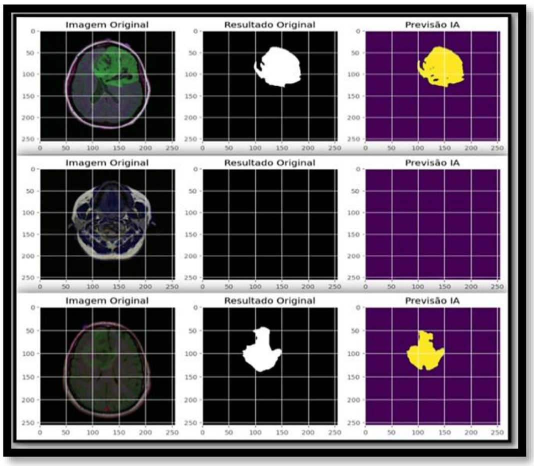

Figure 4: MRI-brain segmentation results. The MRI images are in the first column, the regions segmented by an expert are in the second column, and the regions predicted by the model are in the third column.

Conclusion

In this paper, two projects were presented that combine AI with CV. The first one focuses on the recognition of traffic signs in static images as well as in real-life scenarios. Road safety could be improved, and vehicle automation could advance thanks to this technology. The segmentation of brain tumors in magnetic resonance images is the focus of the second project. The early detection, diagnosis, therapy, and monitoring of patients with brain tumors can be improved by this procedure. The positive outcomes of both projects highlight the important contributions that CV and AI have made in solving these critical issues.

References

[1] Z. Akhtar. 4 Types of Machine Learning (Supervised, Unsupervised, Semi-supervised & Reinforcement) - DatabaseTown. 2020. URL https://databasetown.com/types-of-machine-learning/. Access date: 2023-08-07.

[2] J. Redmon. *Darknet: Open-source neural networks in C.* *URL*: https://pjreddie.com/darknet/. Access date: 2023-07-07.

[3] Roboflow: Give your software the power to see objects in images and video. (n.d.). URL: https://roboflow.com/. Access date: 2023-08-07.

[4] U-NET: Convolutional Networks for Biomedical Image Segmentation. (n.d.). URL https://lmb.informatik.uni-freiburg.de/people/ronneber/u-net/. Access date: 2023-08-07.

[5] *Kaggle: Your Machine Learning and Data Science Community*. (n.d.). URL https://www.kaggle.com/. Access date: 2023-08-06.

[6] TensorFlow. *URL https://www.tensorflow.org/overview*. Access date: 2023-08-07.