Image Classification and Its Applications in Law Enforcement

Emily Morehouse

Florida State University

Email: EEM09C@my.fsu.edu

Abstract:

Significant research has been performed on the topic of image classification. Techniques used to solve this problem include neural networks, decision trees, multimodal space modeling and extreme learning machines. There are significant possible applications of this and a system that could be more effectively implemented using one of the above techniques could produce an extremely useful tool for law enforcement use. Theoretically, an image classifier could one day be able to even determine whether a person appears in an image even over significant lengths of time.

I. Introduction

One of the most widely appealing things about the world of computer science is how incredibly big it is. There is an endless supply of problems that can be solved, solutions to be found, and new ways to apply knowledge found from these solutions. One topic with a particularly wide application set is image classification. Contained within image classification is the act of determining what is depicted in an image (i.e. whether or not a picture contains a cat) or even facial detection and recognition of specific persons. Additionally, the ways in which image classification is implemented is perhaps even as broad as the topic itself. Namely, we will look at image classification as it is implemented using neural networks, decision trees, multimodal space modeling and extreme learning machines as well as the ways in which this can be applied and additional issues that stem from image classification. The areas in which this research can be applied are very wide reaching, from digital cameras to law enforcement software and tools.

II. Background / Prior Work

The most recent buzz with image classification has come from ImageNet's use of convolutional neural networks in order to perform classification. A neural network, in its most simple form with ImageNet's usage, is essentially a circuit that consists of "neurons" used to imitate the structure of a human brain in a sense. This technique is far more sophisticated than simply trying to use straightforward logic to find solutions to problems. Each neuron within the network can, at a minimum, implement logic gates but they can also compute significantly more sophisticated functions. With such a large circuit of neurons, it is much more precise to be able to change the weights between said neurons in order to learn a particular data set. Since data sets for images are particularly large due to both the number of images required as well as the large resolution of the included images, neural networks become huge in order to handle these constraints (Krizhevsky, Sutskever and Hinton).

Additional projects have used other various techniques in order to achieve more efficient image classifications. One project used both pixel and object-based image analysis along with machine learning algorithms to classify agricultural landscapes (Duro, Franklin and Dube). However, their results deemed that their techniques were not statistically significant. This is a common problem with image classifications that use more simplistic techniques and algorithms (namely decision trees and a support vector machine) to attempt to solve this problem. Another project used multimodal subspace learning in order to achieve a significantly higher accuracy. This approach is able to take multiple *different* aspects of the image into consideration when analyzing its contents (Yu, Lin and Seah). Following along the lines of striving to achieve higher and higher accuracies, another project utilized effective extreme learning machine to reduce dimensionality while using extreme k-means for improved clustering.

One of the most interesting projects that influenced this research is one that realized an additional problem with image classification: unlabeled data sets are abundant, however labeled ones are not. On top of that, the learning process is a time consuming and expensive operation. They worked on a system that would actually be able to analyze an image and deem which parts of the image were more likely to be relevant in the learning stage of image classification (Shen, Ju and Jiang).

III. My Work

As many possibilities for how image classification can be implemented, there are just as many (and probably even more) applications for the use of image classification. One of the most practical and useful ways of utilizing image classification would be to increase its use in law enforcement and government agencies. In an ideal system, one would be able to identify a person in not only pictures (including mug shots, crime scene photos, etc) but also in videos or even based on physical descriptions without the photos needing to be directly tagged. Additionally, image classification could take into consideration the possible differences in ages between pictures, haircuts, facial hair changes, etc. This is a very large scope that would require significant experience and time in order to complete. What I attempted to do was simply learn off a small set of pictures in order for a user to determine if a particular object existed within a specified photo.

The way that this problem was approached was using vector analysis to define lines within the images that could be used to form objects. OpenCV has built in modules that handle this, however successfully implementing them is extremely difficult.

IV. Method of Evaluation / Experiment

As with most projects of this nature, a lot of trial and error occurred. If one attempt did not work, another approach was taken until it did work (or maybe until it didn't and another issue would be moved to the forefront). Certainly there are a vast amount of resources for OpenCV and image classification, but due to variances in system setups and available resources, there is a lot of code that is published as examples that could not run, did not produce the desired results, or simply would not compile entirely. The way in which the project was developed was simply to create the system piece by piece from the bottom up and add pieces as they worked together.

V. My Results

Unfortunately, my results were not very positive. With very limited experience with OpenCV in general and absolutely no experience with programming image verification or classification, it was a very large task to even begin to attempt a successful implementation of this (especially one to be run from the command line without a GUI to actually display the images and results). Thus, perhaps this context was not the most appropriate to implement and demonstrate a working prototype of this system as it would be far better suited with more graphical support.

VI. Conclusion & Future Work

Obviously, a full implementation of this topic is out of the scope of the implementation required for this class and additionally, a completed primary run of the code would need to be better implemented and tested before expanding its scope. However, if a system were able to be completed that could encompass all of the above specifications in their most accurate forms possible it certainly would be very useful, particularly in law enforcement and government agencies. The project has a great deal of potential future development to use these tactics to create a very succinct, efficient tool for many applications in identifying persons in photos and videos.

VII. References

Cao, Feilong, Bo Liu and Dong Sun Park. "Image classification based on extreme learning machine." Elsevier (2012).

de la Calleja, Jorge and Olac Fuentes. "Machine learning and image analysis for morphological galaxy classification." (2004).

Duro, Dennis C, Steven E Franklin and Monique G Dube. "A comparison of pixel-based and object-based image analysis with selected machine learning algorithms for the classification of agricultural landscapes using SPOT-5 HRG imagery." Elsevier (2011).

Krizhevsky, Alex, Ilya Sutskever and Geoffrey Hinton. "ImageNet Classification with Deep Convolutional Neural Networks." (n.d.).

Shen, Jianfeng, et al. "Column subset selection for active learning in image classification." (2011).

Yu, Jun, et al. "Image classification by multimodal subspace learning." Elsevier (2012).