How Captain Amerika Uses Neural Networks to Fight Crime

Steven K. Rogers, Matthew Kabrisky, Dennis W. Ruck and Mark E. Oxley

Air Force Institute of Technology
Department of Electrical and Computer Engineering
2950 P Street, Wright-Patterson AFB, OH 45433-7765
15 February 1994

Abstract

Artificial neural networks models can make amazing computations (some of which are applicable to fighting crime: recognition of faces; speaker identification; fingerprint recognition). Those models will be explained along with the application of those models into problems associated with fighting crime. Specific problems addressed are identification of people using face recognition, speaker identification as well as fingerprint and handwriting analysis (biometric authentication).

I   Introduction

Before getting started it is common to explain the Captain Amerika connection. Captain America comic books describe the superhero as: "born in the U.S.A," that obviously applies to the authors; "endowed with a superhuman physique," once you see the authors at the conference you will make the obvious connection with this point; and finally "fights an ongoing battle for liberty, justice, and the American dream!", who needs Ross Perot? Oh, by the way, you might also notice in the comic book that Captain America's secret identity is "Steve Rogers". The "k" in Captain Amerika is just a copyright infringement worry of that author.

This lecture covers the application of artificial neural network techniques for fighting crime. For example the image of a suspect might be provided to some law enforcement agency for processing, possibly to recognize the person in the image. Image processing usually consist of three stages. The first is the location of regions of interest within the image (segmentation-find the face). The second step is the extraction of a set of numbers which characterize the regions that are extracted (feature extraction-describe the face). The last step is the processing of the features for decision making (classification-decide who it is).

II   Crime Fighting Problems

An enormous part of crime fighting is recognition of faces. We will use this problem to demonstrate the application of artificial neural networks to real world problems. During the lecture other problems like fingerprint identification, speaker identification and handwriting analysis will also be addressed. From automatic mugshot matching to border crossing monitoring, law enforcement agencies need an autonomous

231

face recognition capability. Such a system could also be used to verify users of automatic teller machine cards, or control of login into sensitive computer systems. This capability has also been used to interface handicapped people to computers. To be honest this last application is the one that our group is the most excited about. In this case a young Chicago lady (13 years old) who has cerebral palsy was interfaced to her personal computer by recognizing her facial expressions.

III   Segmentation

The finding of regions of interest in an image is called segmentation-find the face in the image. Any errors in this step are preferred to be false acceptance, (passing pixels that may not contain parts of the face), but not false negatives (miss regions that might contain parts of the face). The same concept applies to processing sound. For example, when trying to identify a speaker's voice, sound is recorded. The parts of the recording that need to be identified must be segmented from the rest of the recording. To be of any benefit, this step must significantly reduce the number of pixels or periods of the recording that the next steps of feature extraction and classification must deal with. The processing of the raw pixels to find the regions that might contain the face may be the toughest of the image processing stages. To reduce the amount of computation necessary for the subsequent processing the system should only look in those regions of space, time, frequency, intensity or texture where the face is likely to be located. A one-pass segmentation algorithm filters the raw data to eliminate obvious nonface regions (a function of neighborhood calculations).

Before feature extraction, image preprocessing is usually necessary. The most common preprocessing is some form of energy normalization. The preprocessing is necessary because images have characteristically low contrast and lots of irrelevant structure. To be effective for real world images, the energy normalization is usually based on local neighborhood information. Most segmentation techniques are based on morphological operations, texture analysis and local intensity comparisons or spatial frequency information processing that allow discrimination of regions of interest from the rest of the pixels.

Single neurons can be probed by electrodes and stimulus response measurements made. The results of such measurements show that the system cares about local orientation information and motion direction. Similar more recent measurements have expanded this idea to localized texture information as being the critical first step. To get information from multiple locations, radioactive dyes have been used and clearly show the mapping of the real world onto the visual cortex. One problem with these experiments is that the animal has to volunteer to have its metabolism reduced to zero for the measurements. Only volunteer animals are used of course. Using VLSI technology, multiplexed array cortical electrodes have recently been made and implanted directly onto cortex.

IV   Feature Extraction

The processing of the data to extract a set of measurements (describe the face) that represent the gestalt of the information required to decide who is in the image is called feature extraction. There can be no information gained by this step; its purpose is to increase the ratio of pertinent information to irrelevant data. If a perfect classification stage could be accomplished on the raw data, it would achieve the lowest error possible. But, in the problems of interest here, image processing for face recognition, the processing of the raw data (the original images) is not always feasible. The dimensionality alone of such a task make it not an option for some applications. For each region of interest segmented, a set of features must be found to represent the region for classification.

There are several popular methods for obtaining the features to be used. The first is to ask experts in the field of interest. For example in the problem of target recognition some common features include: length-to-width ratio; hot spot intensity; or complexity. Similarly, relevant expert extracted features are used in face recognition, such as the distance between anthropometrically significant features. The distance between the eyes or from the bridge of the nose to the chin. No one believes that computer aides for recognition are useful if human extracted features have to be keyed in. Finding the important parts of the face by using artificial neural networks is a key first step.

The second alternative is to have the segmented regions processed directly by the neural feature extractor. One common neural feature extraction technique uses a layer of artificial neurons with receptive fields in the input raw data. This is similar to the processing discovered in visual striate cortex, V1. The Nobel Prize winning results of Hubel and Wiesel clearly demonstrated that orientation selectivity and motion direction selectivity within the receptive field of a striate neuron exists. The weights for these artificial neurons are either found using a gradient search based learning algorithm, hardwired based on some a priori knowledge (such as a Hubel and Wiesel or the later work of Jones and Palmer) of types of feature extraction that might be useful.

Quite often after classification, questions are asked about which features caused a particular decision to be made. That is, the question of why a particular region of a photograph was called President Clinton and another called Ross Perot. It's not the shoes. It's got to be the ears! A related question is: of the many features that may have been suggested as useful for a given problem which ones are the most important ones for the task of interest? The answer to this question is often used to reduce the set of feature measurements (vector) to a smaller dimension. This is critical in applications where there are only a limited amount of training data available. To reduce the feature vector, the most common statistical and trial-and-error techniques have been augmented with neural feature saliency techniques. Conventional statistical correlation ideas are the most common technique to find how features are related. The discovery of nonobvious relationships between features may be one of the great contributions of neural networks. One of the early applications of neural networks was in loan analysis. The data on the application for the loan were fed into a neural network and the network that had been

trained on historical data on loan defaults would predict whether you would default. For litigation reasons the users of such networks had to be able to determine the application information that the network considered to be the indicator of you eventually defaulting. There also currently exists artificial neural network systems that monitor credit card transactions to detect fraud. They are trained on historical transaction data and analyze current transactions to detect fraudulent transactions.

As a side note, using the biological insight a good set of candidate features can often be found. In the application of speaker identification, measurements of the processing of the pinna and frequency extraction as a function of distance along the cochlea have resulted in models that have been demonstrated useful in sound localization and speaker identification.

V   Classification

Once the features that are to be used to decide whether a particular region of interest requires further attention are extracted, they are submitted to the classification stage. This is the area where neural techniques have proven to be most useful. The most common neural techniques require an enormous amount of labeled data. Labeled data has to be hand labeled by experts. It is the experience of these experts that the classification step must learn to encode in the interconnection weights. In the application of face recognition, some expert must feed the network with images and tell the network the identity of the face. Similarly, someone must identify the voice from a training recording before the system can identify the person from a later recording.

It has been proven many times in the literature that the common neural techniques perform as approximators of the Bayes optimal decision elements (minimum probability of error). This allows the user to know that if correctly engineered there are no first order statistical techniques which will outperform the neural algorithms with respect to accuracy. Even with this knowledge the comparison of the neural classification algorithms with statistical techniques such as regression or quadratic discriminant function analysis is useful to ensure that the neural technique is correctly engineered.

VI   Future Work

The most important future area of research is in field test and demonstration. Large scale tests will determine whether anything useful will come out of the preliminary exciting results. It will only be by statistically significant improvement in real world applications such as crime fighting that this technology will be proven.

Fundamental work on generalization predictions is also necessary. The question is how much datawill be required in a given application to allow the system to be fielded with some confidence on how well it will perform. How much shrinkage should be expected from the accuracy rate seen in training to the rate that is expected in the real world.
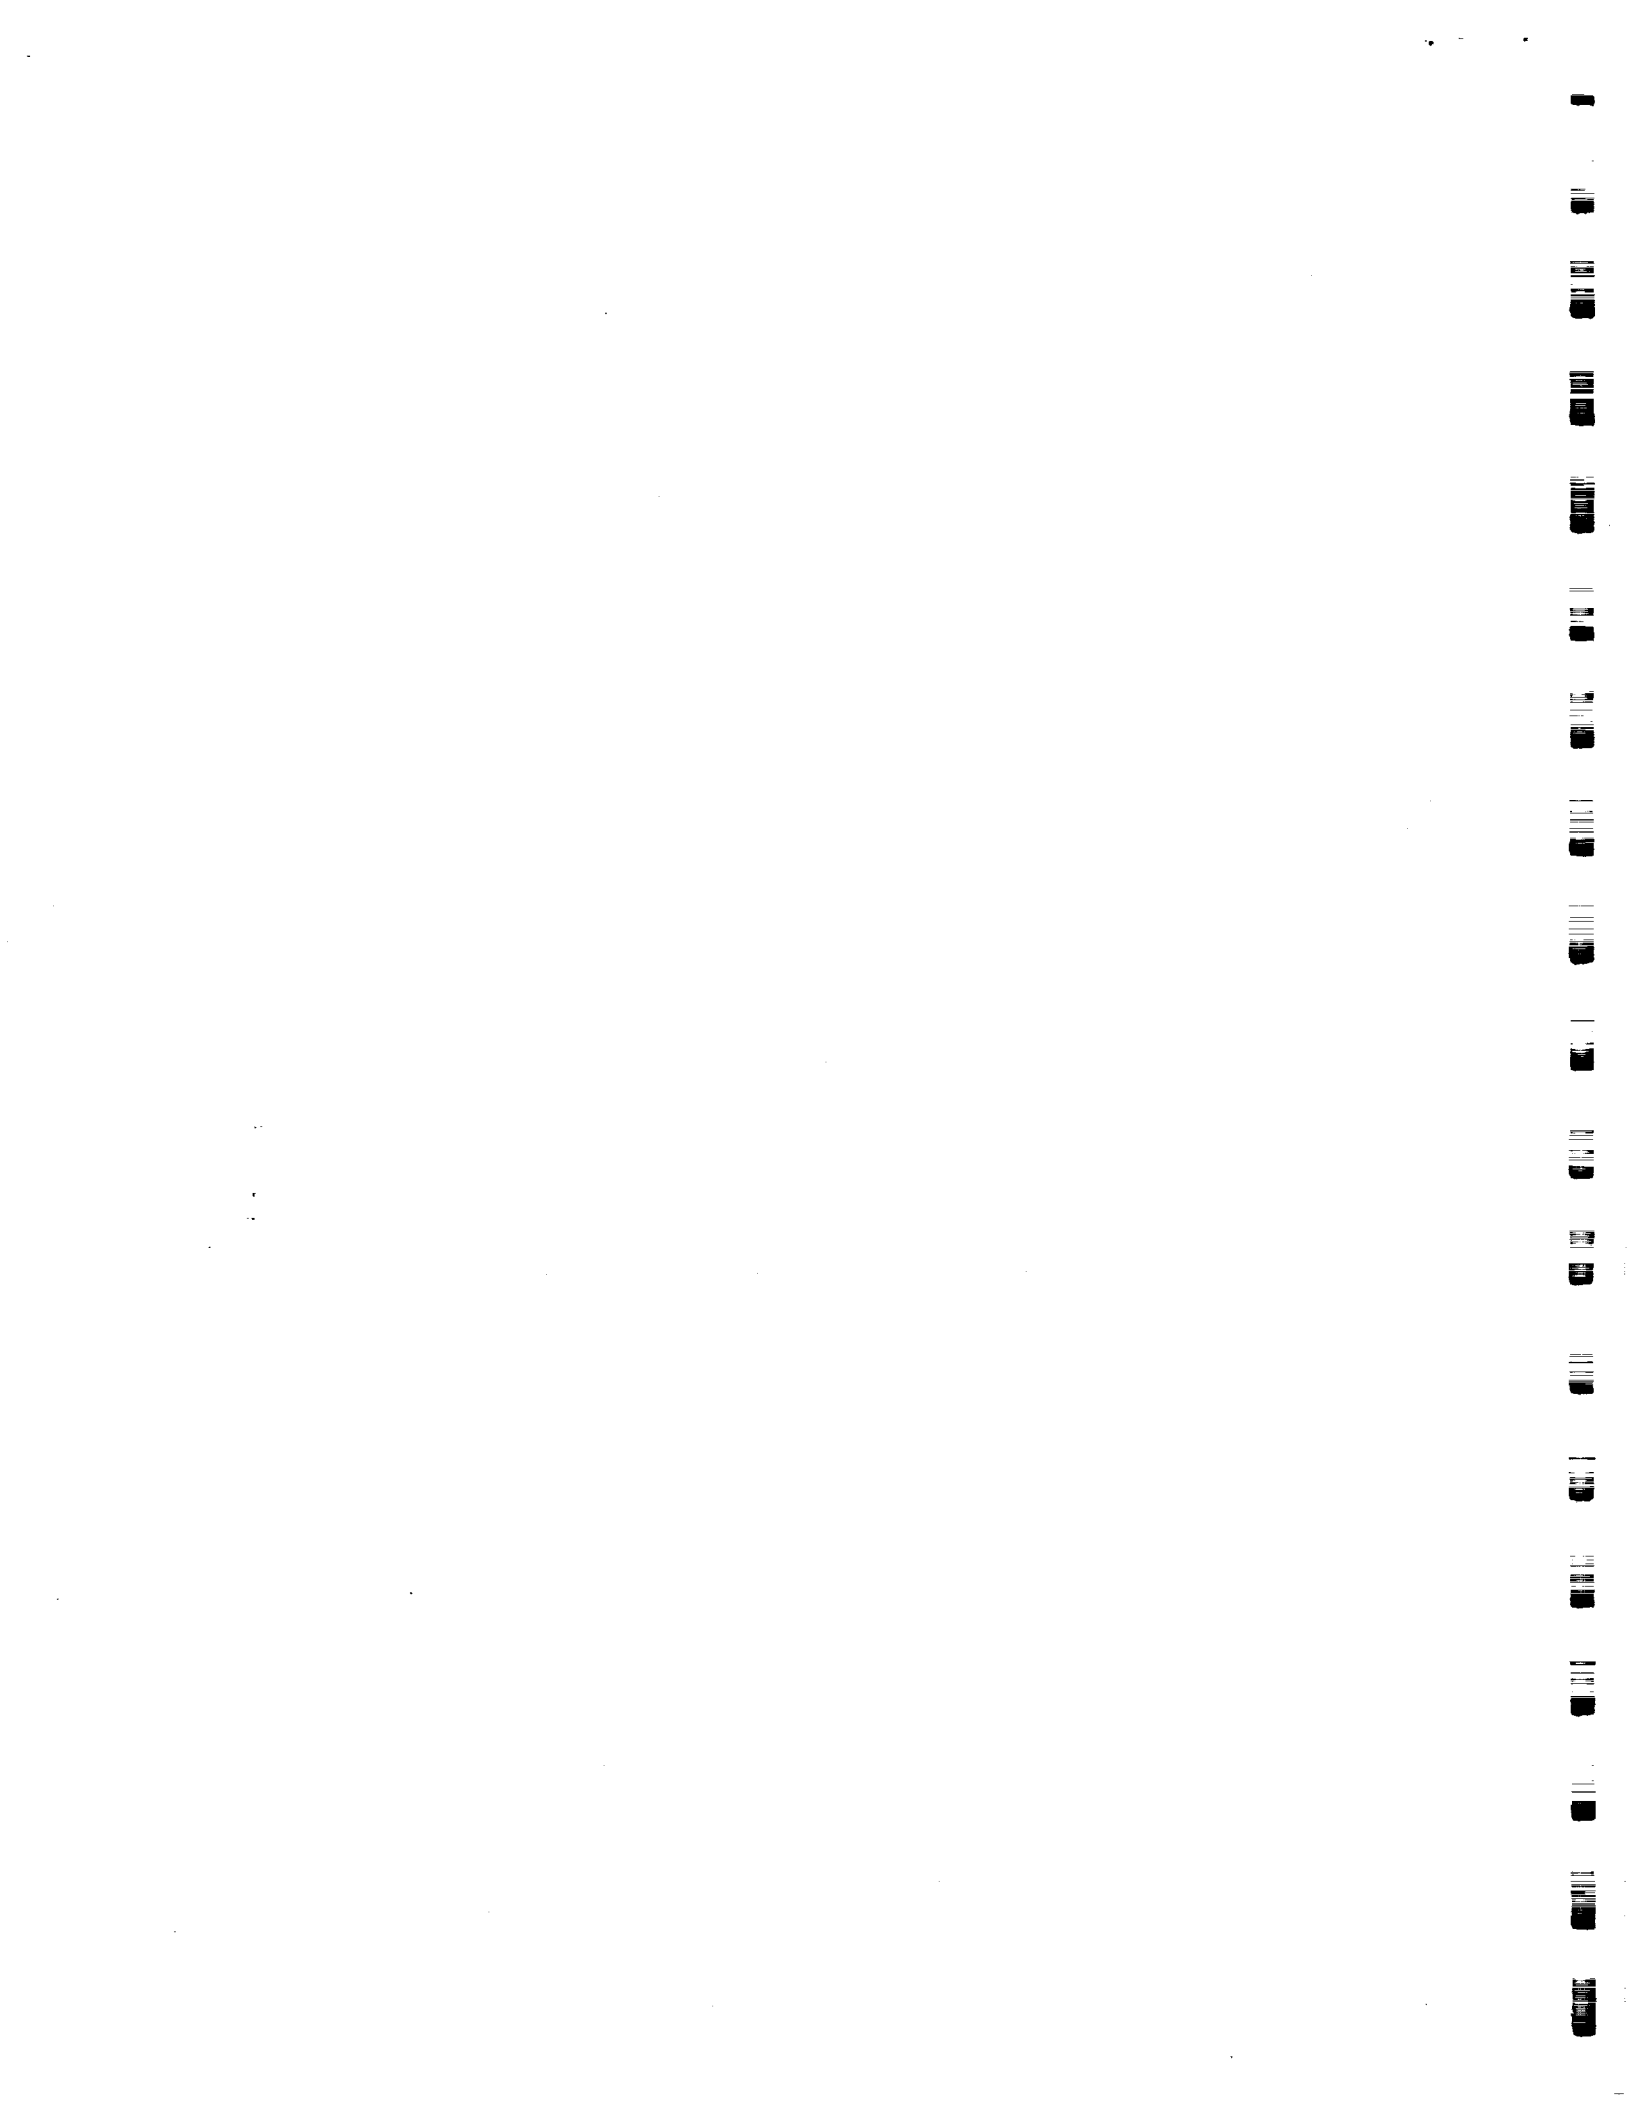
The combination of neural with fuzzy and expert system techniques will also play a key role in driving these solutions to useful applications. Joint conferences, such as the IEEE World Congress on Computational Intelligence, may allow a quick improvement in this area.

One of the most interesting areas of research is in consciousness. Real brains, of course, think about being real brains. The idea of self-awareness as a computation going on within your brain is controversial but true. How does a piece of meat think about being a piece of meat? Could meat ever understand how it does it? Why does human meat seem to be different from that of other animals even though all mammalian brains are constructed to the same basic plan using the same basic parts? There are fundamental limits to the computational capability of the human brain. One way to see the limitations is by the concept of Miller's magical number seven plus or minus two. The human brain is limited to keeping track of about seven things. If keeping track of more than seven things is required to build a stable world society then we have a problem. In the context of this lecture if more "chunks" (more than seven) are required to understand self-awareness then we will never understand how we do it. A puppy dog has fewer chunks than the seven. How many does a chimp have? How can we measure the number of "chunks" for nonverbal animals or if they also can compute their own existence? Series of delay-non-matching-to-sample tests may work here.

The illusion of self awareness is aided and abetted by a series of tricks and lies perpetrated by the human sensory systems; the world is not quite the way it looks, not at all the way it sounds, and the sense of the flow of time is a total confabulation which runs about 200 milliseconds behind real time. The purpose of the brain is to construct as accurate a model of the world as it can given the inevitable limitations of being made out of meat. The results, though, are really amazing; we live inside our own private bags of life which are equipped with a seemingly high fidelity stereo sound system, a 3-dimensional movie display and complete cognizance of touch and smell. We have an enormous content-addressable memory and can keep track of about seven things simultaneously. We can manipulate arbitrary symbols and create the illusion that we are aware of our own existence (and thus compute that it will someday end). Some of the neural hardware forming the sensory systems was described in this lecture but a complete description of how it all works does not exist nor is there any reason to imagine that a human brain could understand it if it did.

VII   Conclusions

It has been shown in several areas that artificial neural networks can make a significant impact in fighting crime. The biometric authentication systems are being fielded. The application of neural technology to other crime-related problems is necessary. This will require a joint effort between experts in the law enforcement area with signal processing people. Participation at the professional meetings of each group by the other is critical.

235

1. Report No. 94-10	2. Government Accession No.	3. Recipient's Catalog No.

4. Title and Subtitle	5. Report Date
A Decade of Neural Networks: Practical Applications and Prospects	May 11, 1994
	6. Performing Organization Code

7. Author(s)	8. Performing Organization Report No.
Sabrina Kemeny	

9. Performing Organization Name and Address	10. Work Unit No.
JET PROPULSION LABORATORY	
California Institute of Technology	11. Contract or Grant No.
4800 Oak Grove Drive	NAS7-918
Pasadena, California 91109	13. Type of Report and Period Covered

12. Sponsoring Agency Name and Address	
NATIONAL AERONAUTICS AND SPACE ADMINISTRATION	
Washington, D.C. 20546	14. Sponsoring Agency Code
	RF 238 0001 PX-644-11-00-00-0

15. Supplementary Notes

16. Abstract

The Jet Propulsion Laboratory Neural Network Workshop, sponsored by NASA and DoD, brings together sponsoring agencies, active researchers, and the user community to formulate a vision for the next decade of neural network research and application prospects. While the speed and computing power of microprocessors continue to grow at an ever-increasing pace, the demand to intelligently and adaptively deal with the complex, fuzzy, and often ill-defined world around us remains to a large extent unaddressed. Powerful, highly parallel computing paradigms such as neural networks promise to have a major impact in addressing these needs. Papers in the workshop proceedings highlight benefits of neural networks in real-world applications compared to conventional computing techniques. Topics include fault diagnosis, pattern recognition, and multiparameter optimization.

17. Key Words (Selected by Author(s))	18. Distribution Statement
Electronics and Electrical Engineering, Mathematical and Computer Sciences, Computer Systems, Information Theory	

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	235	