Proceedings of the Twenty-Ninth International
Florida Artificial Intelligence Research Society Conference

Creation of a Student Research Environment in Artificial Intelligence with Industrial Involvement

**Sambit Bhattacharya,[1] Bogdan Czejdo,[1] John Martellaro,[2]
Aisha Bowe,[2] Kajal Pancholi,[3] and Sue Hall [3]**

[1]Fayetteville State University
Fayetteville, NC, USA
{sbhattac, bczejdo}@uncfsu.edu

[2]STEMBoard
Crystal City, VA, USA
{john, aisha}@stemboard.com

[3]Millennium Integration and Enginering Company
Arlington, VA, USA
{kpancholi, shall}@meicompany.com

Abstract

In this paper we describe ongoing efforts in creating and using an environment for undergraduate research projects in Artificial Intelligence (AI). Our goal is to provide hands-on and real-world experience to students. Industry involvement in this environment provides partnership and supervision, based on deliverables. Students learn real-world skills related to AI in industry and this makes them employable in the job market which is experiencing growing demand for AI applications. The initial direction of the undergraduate research supported by this environment is related to ground based surveillance from stationary cameras where various aspects of human detection, human movement detection, and human movement classification are key components. The central theme of the research is how to use machine learning and computer vision to accomplish these goals and what hardware and software libraries should be selected. More specifically, we describe the research environment that consists of a laboratory with different stationary and remotely controlled cameras, including workstations with high end Graphical Processing Units (GPUs). Open source software libraries for machine learning and computer vision are used for advanced algorithm development in this environment. As a result our research environment has the following characteristics – it is geared towards attractive areas of machine learning and computer vision; and has preconfigured hardware and software components for experiments with data.

Introduction

Machine Learning (ML) and Computer Vision (CV), within Artificial Intelligence (AI) are undergoing dramatic changes with the discovery of new approaches almost every day. This rapid change is being brought about by big data and better understanding of existing approaches like artificial neural networks and how they can take advantage of fast and parallel hardware and get trained on extremely high amounts of data (Schmidhuber, J., 2015). Industry is also having a strong impact in these developments by intensifying their own AI research due the potentially high number of commercial applications. As an example Google Research reported improvements in pedestrian detection (Angelova et. al, 2015), which is related to the research focus of the environment described in this paper. Due to these rapid changes and advanced nature of the work, it is challenging to include recent developments in undergraduate coursework in a meaningful way that teaches the real world impact of ML and CV. In general, many practical solutions have been proposed for including more active forms of student learning. Undergraduate research has been identified as one of the highest impact practices (Lopatto, 2003). Student research is one of the best forms of active learning that is associated with significant improvement of students' comprehension and retention (Erkan, Newmark and Ommen, 2009).

The practical implementation of the undergraduate research programs in AI, however, has encountered many challenges. In our observations we identified three main requirements affecting success of our student research: (a) need to connect with real-world applications to build strong motivation for students (b) need to scaffold the learning of AI, especially during the initial research activities (c) need to produce meaningful results of student's efforts. In order to alleviate challenges and satisfy the above requirements we developed a research environment for AI.

Students are strongly motivated to learn when they are faced with real-world and challenging applications that need solutions. Some areas, which are rich sources of such applications, are real-time surveying and mapping, persistent surveillance, remote inspection and monitoring, high altitude imaging, standoff adversary identification, and advanced video analytics. These areas are full of surprises for beginners in AI since the differences between computer analysis of visual data and that natural process of human vision is not obvious. On one hand tracking and understanding the behavior of a few persons in a limited area is effortless for humans whereas CV is still faced with challenging and unsolved problems which start with detecting the presence of humans, counting how many are present based on detections, their locations, their tracks built out of the sequence of locations, and their identities which are derived from tracks (Teixeira et. al, 2010). On the other hand natural human vision is overwhelmed when tasked with detecting and tracking extremely high numbers of moving persons in a wide area with multiple video feeds. However in this situation CV tends to do better due to the high data processing capabilities of modern computers.

We have therefore engaged in university-industry cooperation between Fayetteville State University (FSU) and companies Millennium Integration, and Engineering Company (MEI) and STEMBoard (SB), to address requirements (a) and (c) stated before. We find that the industry relationship is essential for students to get the opportunity to work on real-world problems in AI (Fincher, Petre, and Clark. 2001). Due to similar reasons this connection is necessary for students to have first-hand experience of the impact of their research. Our students receive feedback from industry experts in bi-weekly meetings. This feedback is an extremely important enhancement to academic evaluation, which they normally receive. Students are also exposed to nuances of intellectual property rights in software and the proper methods to ensure strict conformation to such requirements.

In this paper we describe ongoing efforts in creating and using an environment for undergraduate research projects that are not only directed towards attractive real-world problems but also to scaffolding the learning of AI by providing tools and structures in CV projects. Our environment supports student activities by making available to them the preconfigured hardware and software platform for CV experiments. More specifically our platform contains different stationary cameras, remotely controlled cameras and robots carrying cameras. The environment also includes workstations with GPUs that have open source software libraries in CV and ML installed. Hardware and software components of our research environment are provided to students to scaffold their learning so they can proceed with research experiments and thus avoid unnecessary impediments like having to configure hardware and software. The advanced software libraries provide the building blocks for students to develop their own algorithms. In addition these libraries enable them to quickly obtain meaningful results.

We have observed the many benefits for our undergraduate researchers: learning in-depth some topics in AI relatively early; developing advanced computational models; systematic data collection and programming on multiple platforms; learning to use appropriate research methodology, specifically, formulating hypotheses and specifying methods for its verification; applying knowledge to a real situation; learning to analyze data and transform it to various representations; appreciation of scientific experiments and learning what scientific research actually entails.

Student Research Environment

A student research environment needs to satisfy the requirements for efficient student research. Our research environment supports such requirements: 1) is geared towards an attractive area of ML and CV 2) has structure for industry involvement 3) has preconfigured hardware and software components for experiments along with advanced software library modules that can be used individually or integrated to create advanced computational models.

Application Area

This section describes the general application area around which our research environment is created. This area is rich in important problems that attract research and development efforts from the industry. In general remote sensing camera systems and object-tracking algorithms can provide valuable information in desired regions of interest, while assisting humans and possibly protecting from an unnecessary direct engagement. Multiple cameras can be controlled together to provide low-cost persistent surveillance and continuous monitoring of a region with a wide area of coverage. In many cases, it is desirable to detect and track individuals and other stationary and dynamic objects across these areas of interest from live video to understand social dynamics taking place. The accurate detection and tracking of humans from a camera system poses an interesting problem, given the processing needed to detect unique human characteristics and behaviors. The significance of this problem and opportunity is that once the camera can detect and track individuals, it would then be able to provide 'tracks" data in real-time, keeping a continuous record of a person of interest's motion path. One would then have a full picture of all movements that have taken place in a given area.

After these movements were identified, advanced analytics could be employed to understand dynamics in crowds, and determine certain patterns as they were taking place. If these algorithms were coupled with thermal sensors, then one could begin to detect abnormal behavior before an event takes place. For example, a person's body temperature maybe very high indicating some illness and thus can signal to an operator to look at the previous behavior. With the tracks data that has been produced, one could follow their path back in time to see where they were and who they have met with. If the person of interest had actually an infectious disease, it could be found more easily, potentially avoiding or limiting the spread of the disease.

Industry Involvement

We have engaged in university-industry cooperation between FSU and companies MEI and SB to provide the environment with industry partnership and supervision, based on deliverables; relevant exposure to industry practice and concerns; direct input from industry and to find a good balance between industry and academic perspectives. Our objectives are to provide real-world skills relevant to AI in industry and to make our graduates more employable in the growing AI application-based industry. Based on guidelines from (Fincher, Petre, and Clark. 2001), the student learning objectives are 1) how to tackle problems with industrial relevance and currency 2) the methods and standards of industry and how to conform to them 3) completing project with practical time constraints 4) developing professional skills 5) developing communication and presentation skills and 6) how to effectively work in a team. The deliverables required by industry in our collaboration are 1) design documents 2) test reports and statistics 3) software listing and documentation 4) user manual and 5) demonstrations.

The companies MEI and SB are participants in the Department of Defense's prestigious Mentor-Protege Program. The multi-year, multi-million dollar partnership provides mutually beneficial developmental assistance to SB while exposing MEI to new and diverse opportunities and clients. The agreement also provides funding to Historically Black Colleges and Minority Institutions such as FSU to not only enhance SB's technical capabilities but also to expose traditionally underserved students to opportunities within AI and related areas. Since the beginning of the project, MEI and SB's staff have met the students and faculty mentor regularly (about twice a month) to discuss research and to collectively brainstorm on solving current problems and initiating new directions.

Preconfigured Hardware and Software Components

In order for our research environment to provide proper support for student research a physical lab consisting of appropriate hardware and software was created. We equipped our lab with several stationary cameras and several cameras on our robots. We have analyzed the usefulness of other possible cameras including pan-tilt-zoom cameras. Generally we created a multi-camera system since we want our students to have a highly mobile and modular testing system that has the capabilities to adequately track humans in many environments.

Our environment also includes workstations with high end GPUs. Real-time video processing will require a cluster of GPUs to analyze the video with strict time constraints. In the future, depending on resources available, we will consider video analysis in an Amazon EC2 cluster, or we will configure a local cluster of GPUs to target real-time video algorithms. Currently we use open source software libraries in CV (openCV) and ML (Sci-kit Learn) to build our algorithms. The openCV library (ref) provides implementations of most advanced algorithms for computer vision such as Histogram of Oriented Gradients (HOG) (Qiang, Yeh, and Kwang-Ting and Avidan, 2006), different methods of foreground versus background separation (Z.Zivkovic, and Van der Heijden, 2006), connected components and shape and color based analysis which are necessary for the experimental research described earlier. The Sci-kit Learn library provides a comprehensive implementation of standard ML algorithms for classification, clustering, neural networks etc. In combination, OpenCV and Sci-kit Learn create a powerful development environment for implementation of complex tasks. E.g. HOG and Support Vector Machines (SVM) can be trained on home-grown datasets for detecting persons from different perspectives.

Student Activities within our Research Environment

Let us describe and analyze some of the activities of a group of students that has participated in our undergraduate research activities within this environment. The research objective was to develop and demonstrate algorithms capable of off-line human detection from standalone video processing systems, as well as extracting simple (linear time and movement) movement data.

Once student researchers became acquainted with the research platform, they began collecting video from the camera systems. This required a series of parameters to bind the problem, including videos from multiple heights, multiple modes of operation, and varying scenes with different groups of present humans. This was a critical step in creating a corpus of data to be analyzed.

The next task was to process video and develop human detection computational models. Initially, models based on individual software library modules served to detect which frames contained humans. The experiments determined some limitations to how well humans were detected. After analyzing limitations, the human detection models went through several iterations. A demonstrations have shown the effectiveness of the algorithms to detect humans in video, as well as characterize the different types of interactions that were taking place, such as the rough speed of the individual and if they are moving or stationary. As research progressed, a number of other attributes were identified based on the progress made. The intent of this research was to be able to detect and characterize multiple individuals separately and differentiate their movement. Table 1 shows the hierarchy of properties used in this research and specific algorithms used by the students, which formed the components of the overall software product created by them.

Hierarchy of Properties ↓	**Spatio-temporal properties**	**Specific algorithms**
	Presence	HOG + SVM
	Count	Non-maxima suppression
	Location	Geometric attributes of detection window
	Track	Sequence of locations and Kalman Filter for smoothing
	Identity	Consistent labeling of persons through matching windows

Table 1. Hierarchy of properties in object sensing (Teixeira et. al, 2010) and specific algorithms for human detection tracking used in our student projects (Fincher, Petre, and Clark. 2001). The hierarchy presents a clear approach which when adapted to the human detection and tracking problem, specific algorithms need to be selected.

Summary

In this paper we described an environment for undergraduate research in AI that has been created with involvement from industry. It contains research problems in attractive areas of machine learning and computer vision; has a formal structure for industry involvement; and provides preconfigured hardware and software components for experiments. We have observed that our students benefited from industry cooperation by learning how to tackle problems with industrial relevance and currency while conforming to their methods and standards, completing project with practical time limits. They had training in developing professional skills, communication and presentation skills and how to effectively work as a team. There were many additional benefits of using our research environment for our undergraduate students: developing proficiency in laboratory practice, specifically, systematic data collection and software on multiple platforms with easy first steps; learning a topic area in AI relatively early using preprogrammed advanced software library modules and meaningful applications; learning how to carry out scientific experiments including appreciation of the role of refinement in AI research.

Our environment and the area of application is a rich source of problems. For instance we plan to develop more algorithms to detect spatial and temporal relations between tracks / trajectories in video. This topic is interesting because certain relations show events occurring in the real world through group interactions (Dodge et. al, 2008).

Acknowledgements

We acknowledge generous financial support and expert advice of scientific and technical staff from Millennium Integration and Engineering Company, and STEMBoard.

References

Schmidhuber, J. (2015). Deep learning in neural networks: An overview. Neural Networks, 61, 85-117.

Angelova, A., Krizhevsky, A., Vanhoucke, V., Ogale, A., & Ferguson, D. (2015) Real-Time Pedestrian Detection With Deep Network Cascades.

Teixeira, T., Dublon, G., & Savvides, A. (2010). A survey of human-sensing: Methods for detecting presence, count, location, track, and identity. ACM Computing Surveys, 5, 427-450.

Fincher, Sally, Marian Petre, and Martyn Clark. 2001. Computer science project work: principles and pragmatics. Springer Science & Business Media, 2001.

Lopatto, D. 2003. The essential features of undergraduate research. Council on Undergraduate Research Quarterly, 23, 139–142.

Erkan, A.; Newmark, S.; and Ommen, N. 2009. Exposure to research through replication of research: a case in complex networks. In Proceedings of the 14th annual ACM SIGCSE conference on Innovation and technology in computer science education (ITiCSE '09). ACM, New York, NY, USA, 114-118.

Qiang Zhu; Yeh, M.-C.; Kwang-Ting Cheng; Avidan, S. 2006. "Fast Human Detection Using a Cascade of Histograms of Oriented Gradients," in Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on , vol.2, no., pp.1491-1498, 2006.

Zivkovic, Z. and Van der Heijden, F. 2006. Efficient Adaptive Density Estimation per Image Pixel for the Task ofBackground Subtraction, Pattern Recognition Letters, vol. 27, no. 7, pages 773-780, 2006.

Dodge S., Robert Weibel R., and Lautenschütz A. 2008. Towards a taxonomy of movement patterns. Information Visualization 7, 3 (June 2008), 240-252.