

CSEE Research Review – Talk Abstracts

Friday, May 4, 2012

Department of Computer Science and Electrical Engineering
University of Maryland, Baltimore County (UMBC)

Session I

Marie desJardins, MAPLE Laboratory

Using Sample Distributions to Accurately Calibrate Model Confidence

Probability estimates produced by standard machine learning classifiers are typically not highly accurate when used as a measure of predictive confidence. Previous work has applied one-dimensional (1D) regression methods (Platt scaling and isotonic regression) to calibrate these confidence values in order to improve their accuracy. However, these past methods do not take into account the distribution of the training or test data.

In this talk, I will present methods for measuring the sample density and sample purity of test instances, and show how two-dimensional isotonic regression can be used to calibrate model confidence using classifier probabilities in combination with either of these sample measures. I will present empirical results for two classifiers (J48 and naive Bayes) on 13 benchmark data sets, showing that sample-based confidence calibration results in more accurate confidence values than either uncalibrated confidence or confidence calibrated with 1D methods. I will also show visualization methods that can be useful in understanding both the uncertainty associated with model predictions, and how this uncertainty is distributed across the instance space.

Karuna P Joshi (Advisor: Yelena Yesha), ebiquity - *Honorable mention PhD research*

Automating Cloud Services Lifecycle through Semantic Technologies

Managing virtualized services efficiently over the cloud is an open challenge. Traditional models of software development are not appropriate for the cloud computing domain, where software (and other) services are acquired on demand. In this paper, we describe a new integrated methodology for the lifecycle of IT services delivered on the cloud, and demonstrate how it can be used to represent and reason about services and service requirements and so automate service acquisition and consumption from the cloud. We have divided the IT service lifecycle into five phases of requirements, discovery, negotiation, composition, and consumption. We detail each phase and describe the ontologies that we have developed to represent the concepts and relationships for each phase. To show how this lifecycle can automate the usage of cloud services, we describe a cloud storage prototype that we have developed. This methodology complements previous work on ontologies for service descriptions in that it is focused on supporting negotiation for the particulars of a service and going beyond simple matchmaking.

James MacGlashan (Advisor: Marie desJardins), MAPLE Laboratory - *Honorable mention PhD research*

Option-Based Multi-Source Policy Transfer in Object-Oriented MDPs

We present a novel approach, called Option-based Policy Transfer (OPT), that solves the multi-source policy transfer problem in reinforcement learning by casting it as an option learning problem. When learning a new task, OPT constructs options for similar previous tasks and then uses modified option-learning methods to solve the transfer problem. We show how policies for such transfer options can be automatically constructed, even when the source and target tasks have different state spaces and features, by exploiting the state representation of object-oriented MDPs. Empirical evidence in multiple domains is presented to demonstrate the effectiveness of this approach.

Session II

Gymama Slaughter, Bioelectronics Laboratory

Innovative Optical and Power-Harvesting Microsystems for Medicine

This talk will present two ongoing research efforts in the Bioelectronics Laboratory to fabricate and characterize fiber-optic evanescent wave sensors and self-powering microsystems to aid in the diagnosis and treatment of diseases. The first part of the talk will summarize our efforts to develop new strategies for monitoring volatile organic compounds in exhaled breath. First we will describe a flexible polymer-clad fiber for sensing chemicals that permeate the cladding material. Experiments with these fabricated devices demonstrate that they can provide significant improvements in the detection of various chemicals and chemical combinations. The second part of the talk will focus on "self-powering microsystems." A "self-powering microsystem" is a device that converts available inertia energy into electrical energy. These devices enable a human, via the use of surgically implanted electrode array and associated computer decoding algorithms, to harvest electrical energy to control other implantable devices by pure inertia energy alone. We will focus on a novel miniature glucose biofuel cell electrodes configuration that autonomously harvests electrical power within subcutaneous tissue so as to monitor and maintain optimal blood glucose level. Demonstrations of the device in glucose solutions and our ongoing efforts to produce miniaturized implantable versions of these devices will be reviewed.

Anupam Joshi, ebiquity

Policies and Context for Situationally-Aware Security Systems

Han Dong (Advisor: Shujia Zhou) - *Best research by MS student*

Cross-Platform OpenCL Code and Performance Portability for CPU and GPU Architectures Investigated with a Climate and Weather Physics Model

Current multi- and many-core computing typically involves multi-core Central Processing Units (CPU) and many-core Graphical Processing Units (GPU) whose architectures are distinctly different. To keep longevity of application codes, it is highly desirable to have a programming paradigm to support these current and future architectures. Open Computing Language (OpenCL) is created to address this problem. While the current implementations of OpenCL compiler provide the capability to compile and run on the architectures above, most of the current researches investigate the performance of GPU's as a compute device. In this paper we will investigate the portability of OpenCL across CPU and GPU architectures in terms of code and performance via representative climate and weather physics model, NASA's GEOS-5 solar radiation model, SOLAR. An OpenCL implementation portable between CPU's and GPU's has been obtained with significant performance improvement in some CPU's and GPU's. We found that OpenCL's vector-oriented programming paradigm assists compilers with implicit vectorization and consequently significant performance gains were achieved.

Session III

Chein-I Chang, Remote Sensing Signal and Image Processing Laboratory

Subpixel Target Detection in Hyperspectral Imaging

Hyperspectral imaging has expanded capability of multispectral imaging in many prospects. Specifically, it changes many ways how algorithms are designed from traditional spatial-domain (literal) analysis to spectral-domain (non-literal) analysis. This is mainly due to the fact that many unknown substances which cannot be identified by visual inspection or prior knowledge can now be uncovered by high spectral resolution provided by hyperspectral imagery. It is particularly evidential in endmember extraction, spectral unmixing, anomaly detection and subpixel/mixed pixel analysis where target signatures are on longer pure but rather mixed by other substances or background signatures within a single pixel. In this talk we will focus on subpixel target detection.

Matthew Anderson (Advisor: Tulay Adali), Machine Learning for Signal Processing Laboratory

An Effective Decoupling Method for Matrix Optimization and its Application to the ICA Problem

Matrix optimization of cost functions is a common problem. Construction of methods that enable each row or column to be individually optimized, i.e., decoupled, are desirable for a number of reasons. With proper decoupling, the convergence characteristics such as local stability can be improved. Decoupling can enable density matching in applications such as independent component analysis (ICA). Lastly, efficient Newton algorithms become tractable after decoupling. The most common method for decoupling rows is to reduce the optimization space to orthogonal matrices. Such restrictions can degrade performance. We present a decoupling procedure that uses standard vector optimization procedures while still admitting nonorthogonal solutions. We utilize the decoupling procedure to develop a new decoupled ICA algorithm that uses Newton optimization enabling superior performance when the sample size is limited.

Josiah Dykstra (Advisor: Alan Sherman), Cyber Defense Lab - *Best research by PhD student*

Acquiring Forensic Evidence from Infrastructure-as-a-Service Cloud Computing: Exploring and Evaluating Tools, Trust, and Techniques

We expose and explore technical and trust issues that arise in acquiring forensic evidence from infrastructure-as-a-service cloud computing and analyze some strategies for addressing these challenges. First, we create a model to show the layers of trust required in the cloud. Second, we present the overarching context for a cloud forensic exam and analyze choices available to an examiner. Third, we provide for the first time an evaluation of popular forensic acquisition tools including Guidance EnCase and AccessData Forensic Toolkit, and show that they can successfully return volatile and non-volatile data from the cloud. We explain, however, that with those techniques judge and jury must accept a great deal of trust in the authenticity and integrity of the data from many layers of the cloud model. In addition, we explore four other solutions for acquisition—Trusted Platform Modules, the management plane, forensics as a service, and legal solutions, which assume less trust but require more cooperation from the cloud service provider. Our work lays a foundation for future development of new acquisition methods for the cloud that will be trustworthy and forensically sound. Our work also helps forensic examiners, law enforcement, and the court evaluate confidence in evidence from the cloud.