The Thirty-Second International Florida
Artificial Intelligence Research Society Conference (FLAIRS-32)

FLAIRS-32 Poster Abstracts

Roman Barták, **Keith Brawner**

Editors

Incorporating System-Level Objectives into Recommender Systems

Himan Abdollahpouri (University of Colorado Boulder, USA)

One of the most essential parts of any recommender system is personalization how acceptable the recommendations are from the user's perspective. However, in many real-world applications, there are multiple objectives often from multiple stakeholders that need to be incorporated into the recommendation generation. In this work, we define the problem of multi-stakeholder recommendation and we focus on finding algorithms for a special case where the recommender system itself is also a stakeholder. We also define different types of system-level objectives and find algorithmic solutions for each of them such that similar problems can be solved by the same class of algorithms. Finally, we will explore the idea of incremental incorporation of system-level objectives into recommender systems to tackle the existing problems with the optimization techniques which only look for optimizing the individual users' lists rather than looking at the whole picture of system performance over time.

How Can Robots Successfully Elicit Help from Humans

Layla Alarifi, Rajaa Rahil, Marius Silaghi
(Florida Institute of Technology, USA)

With autonomous robots being exposed to unstructured environments, they inevitably run across cases in which they stand unable to overcome their limitations, such as removing objects in their path, or location uncertainty. These limitations can be overcome when robots obtain help from humans. We are investigating how robots can effectively interact and request help from human. We conducted an experiment during which passer-by were asked by a NAO robot in different manners, such as indirect, polite, and friendly, and combining different interaction modalities, like speech and gestures. We test a scenario where the robot needs a door to be opened by a human, so the robot could complete its other tasks. The robot is too short to reach the door handle himself. In the polite request, the robot may say "Can you open the door please?", while in the indirect request, the robot may say "I cannot open the door and it is blocking my way". In the friendly way, the robot asks with less formal tones like "You seem taller than me. Would you open the door for me?". The robot can point its hand to the door that it needs opened. Later humans were interviewed as to why they did or did not take the robot seriously. The polite interaction manner was significantly more efficient. We show how to factor the situational awareness effects (whether participants realize the nature of the experiment) in the analysis. The proposed evaluation procedure allows identifying promising mechanisms for such human-robot interactions.

A Group Decision Making Emoji Wall

Badria Alfurhood, Marius Silaghi
(Florida Institute of Technology, USA)

While most group decision processes entail uncertainty, current electronic group decision support systems are deficient in features mitigating corresponding challenges. Appropriate handling can enhance decision makers' abilities for reaching better decisions. We develop the concept of "Group Decision Making Emoji Wall" (GDMEW), an online group e-debate & decide tool that enhances distributed group decision making. Such tool offers the novel feature of conveying a summary of voting results with attached icons that reveal the decision makers' levels of certainty and the most cited decisions' justifications. We propose to relay the group decision stages graphically by displaying only relevant ongoing decisions' results. This feature enables voters to monitor the state of the deliberation process on a resolution. For instance, if a person votes, but is (Still thinking, Undecided, Not Sure, Almost Decided, Decided), GDMEW displays icons to show that level of decision confidence.

Simplicity and openness help to engage more stakeholders in group debates for taking an action. However, there exists a controversy on the effects of different levels of openness. We focus on studying partial openness where in order to reduce framing effects, users can view previous votes with a summary of justifications only after communicating their own thoughts and insights. Thus, GDMEW gives participants the ability to express their level of certainty of going with or against a given resolution and to be aggregated in its analysis. Preliminary experiments are being performed with a current implementation based on case studies and interviews.

A Case-Based Reasoning Application for Guiding Low-Back Pain Patients

Kerstin Bach, Ilya Ashikhmin,
Tomasz Szczepanski, Agnar Aamodt
(Norwegian University of Science and Technology, Norway)

In this work we present the concept and implementation of the selfBACK case-based reasoning (CBR) system in order to guide low back pain (LBP) patients through self-management. selfBACK is an EU project focusing on the self-management of LBP through a digital intervention via a mobile phone application. The Global Burden of Disease study has shown that LBP is the most significant contributor to disability in Europe. Self-management in the form of physical activity and exercise is crucial to improve prognosis; however, adherence to a self-management plan is challenging without feedback and reinforcement.

To create and follow up on individual self-management plans of patients we have developed a CBR system that learns from positive plans and adapts them to fit to users' achievements and goals. As a patient's condition will change over time the CBR system handles time-series data of different resolutions to find most similar cases that fit the situation and comply with clinical guidelines.

Our work focuses on the case structure and similarity modeling developed for selfBACK as well as we show the evolution of cases throughout a 6-week pilot phase of the selfBACK system with patients. A special emphasis is given to the similarity-based case retrieval, candidate selection and their adaptation to the current user. From March 2019, the selfBACK system will be tested in a randomized controlled trial over a period of nine months testing the effectiveness of the intervention.

A System Using Deep Learning and Fuzzy Logic to Detect Fake YELP Reviews

Jun Bai, James Buckley (University of Dayton, USA)

With the prevalence of online searching, looking up online reviews of businesses, such as restaurants, hotel and other services, is a major factor in people's decision making. However, fake reviews cause the sentiment analysis of a corpus of reviews to be clouded. This research uses the YELP data set that is publicly available on the internet. We connect both review content, user information and business information to optimize the fake review detection accuracy. In terms of our solution, we use both feature extraction and machine learning-based classification to detect fake reviews. In feature extraction, features are extracted from reviews using term frequency and frequency-inverted document frequency. We extract features from users' information like number of reviews, review length, number of fake reviews generated, review date, and review helpfulness. We also extract features from businesses like star number, number of reviews, number of associated fake reviews, and type of business. All of the features will be input to different machine learning models (RNN, CNN, LSTM etc.) and results compared to determine which machine learning model has the best performance. We then use fuzzy logic to classify the Yelp dataset into 4 groups: fake review, very likely fake review, less likely fake review and true review. This fuzzy classification allows the user to know when a review is real or fake, and to look at those reviews in between to determine their authenticity.

Towards Privacy-Sensitive Robots

Rania Bamagain, Marius Silaghi
(Florida Institute of Technology, USA)

As robot presence grows, they spread to most places where humans also live. Robots possess sensors that can constantly watch, hear, process and record data on the surrounding environment. Numerous benefits and services could be provided, but at the same time, various types of privacy issues appear. Privacy concerns are mainly caused by robots that process data remotely. The main factors that impact on the dangers of violating the users' privacy are the robots' appearance, robots' perception, robots' navigation, robots' authentication, robots' warning system, and robots' application.

We investigate usable, trusted, and comfortable techniques to bring security in the context of social robot utilization, to improve users' awareness towards associated privacy risks, and to find trade-offs between privacy loss and utility achieved by AI planners. We also conducted surveys to identify preferred techniques of mitigating violations of privacy that can be based on robots in order to develop privacy-sensitive social robotics. Our results suggest that users are more comfortable with robots that use both adaptive filters and automatic hardware covers that are placed on the robot's camera and microphone, and that work when they are needed for privacy protection or when the robots are off. Moreover, the users prefer connecting the robots with

environment sensors to constrain their movement, to integrate multi-factor authentication, and to install intelligent warning systems on their robots that enhance user awareness toward the robot's capabilities and actions. It is shown that automated planners can integrate the utility brought by such needs with remaining task goals.

Quantifying the Global Race for Artificial Intelligence Using Clustering and Correlations Analysis

Feras Batarseh, Akhila Kotapati
(George Mason University, USA)

In a rapidly-changing digital world, multiple countries are shifting their focus to investing in Artificial Intelligence (AI). The competition is fiercer than ever, but it is expected to get much more tangible in upcoming years. AI is already proving to having massive returns, locally and globally. Therefore, data (the major driving force of the recent AI surge) became more important, and data-based approaches are already accomplishing multiple victories in many domains such as healthcare, agriculture, warfare, sports, education, and others. In this work, data-sets on this new global race are collected from different sources. The datasets include country wise information relevant to their potential to producing more data, performing more research, developing intelligent applications, and competing in a global race for AI. Example columns include: Population, Number of Internet-less Citizens, Percentage of Network Penetration, Country Gross Domestic Product, Literacy Rates, Research and Development Expenditures, Number of AI companies, and other metrics. All these attributes provide pointers on interest and investment of each country into AI.

K-means Clustering techniques are developed to group the countries according to the mentioned variables. A correlation formula is also developed to identify the most important variables, and an AI-coefficient for each country is defined. The results are deemed close to what is expected, for example, countries such as China and the US are currently leading this race. Many other countries, however, are on the racing behind, such as: Canada, Germany, Vietnam, Brazil, and the UK. The study also showed how certain coalitions (European countries, or countries of Southeast Asia for instance) can potentially tilt the race to their benefit. This work is ongoing, and the aim is to improve on the comprehensive and indicative AI coefficient that can be used for evaluating and measuring the different dimensions of this inevitable global race.

Context-Driven Predictions through Data Imputation and Inference

Feras Batarseh, Ajay Kulkarni
(George Mason University, USA)

Deploying data-driven methods across different domains have assisted in providing intelligent solutions to persisting problems, for example: improving the quality of service (within healthcare), policy making (at the government), defining customer behavioral trends (within commerce), predicting bugs and errors (in software development), and forecasting sports events results (for team sports, athletes scouting, and marketing). However, throughout these deployments, serious show-stopper problems are still unresolved, such as: the lack of domain context in data-sets, blind spots in data collection, hidden biases, and data in-completeness. The major challenge found throughout the process is the quality of the data (outliers, bias, and in-completeness). As Niels Bohr famously stated: “Prediction is very difficult, especially if it's about the future”. The challenge exacerbates however, when the future prediction is an outlier (such as predicting a winner of a sports tournament, or an event such as an economic recession). The data science lifecycle trains and scores predictions based on historical data; the work presented in this poster paper, however, aims to re-define the lifecycle and introduce tangible improvements to its outcomes. Two studies are used for evaluation (in sports predictions). The studies illustrate the need for better methods in detecting bias, outliers’ management, and mitigating data incompleteness through context and other statistical measures. A new Context-driven Data Science Lifecycle (C-DSL) is introduced and tested; and the results are recorded and presented.

Modeling Social Engineering Risk Using Attitudes, Actions, and Intentions Reflected in Language Use

Adam Dalton (IHMC, USA), Alan Zemal, Amirreza Masoumzadeh (SUNY Albany, USA), Archna Bhatia, Bonnie Dorr, Brodie Mather (IHMC, USA), Bryanna Hebenstreit (SUNY Albany, USA), Ehab Al-Shaer (UNCC, USA), Ellisa Khoja, Esteban Castillo Juarez (SUNY Albany, USA), Larry Bunch, Marcus Vlahovic (IHMC, USA), Peng Liu (SUNY Albany, USA), Peter Pirolli (IHMC, USA), Rushabh Shah (SUNY Albany, USA), Sabina Cartacio (IHMC, USA), Samira Shaikh, Sashank Santhanam (UNCC, USA), Sreekar Dhaduvai, Tomek Strzalkowski, Younes Karimi (SUNY Albany, USA).

Social engineering attacks are a significant cybersecurity threat putting individuals and organizations at risk. Detection techniques based on metadata have been used to block such attacks, but early detection success is minimal. Natural language processing and computational sociolinguistics

techniques can provide a means for detecting and countering such attacks. We adopt the view that actions and intentions of adversaries, which we call "asks," in social engineering attacks can be gleaned from the adversaries' language use during communication with targets. Timely recognition and understanding of these asks can facilitate deployment of proper safeguards, thus enabling simultaneous engagement of the adversaries in ways that compel them to divulge information about themselves. Attitudes that the adversaries attempt to induce in their targets for easy compliance can also be identified proactively to compute social engineering risk. These identified attitudes can be used to help the target in making the right decisions. Moreover, the techniques used for detection can also be used to induce attitudes in the adversaries. We present a model of social engineering risk that uses knowledge of actions, attitudes induced, and intentions embedded in the dialogue between an adversary and a target user. Informed by such a model, we are building a platform for Personalized AutoNomous Agents Countering Social Engineering Attacks (PANACEA). The platform informs users about social engineering risk situations and can automatically generate responses to engage the adversary and protect the user.

Extending Group Recommender Systems: Group Assignments in Massively Multiplayer Online Video Games Based on the Users Behavior

Arman Dehpanah, Jonathan Gemmell (DePaul University, USA)

Recommender systems assist users to discover items from a vast and complex information space. These systems usually recommend items to individuals based on their interests and tastes. However, there are many cases, such as watching a movie or playing music in a gym, which may involve a group of users. In these cases, group recommenders suggest items to users while considering the preferences of each group member. Group recommender systems may include simpler problems such as choosing a group of songs to play for a group of users, to more complicated problems such as assigning players to groups in Massively Multiplayer Online video games. In this work, we focus on the gaming industry and particularly online video games where there are many behavioral and psychological aspects to playing these games. We investigate players behavioral features such as skills, tastes, expectations, and needs for enriching their profiles to generate group recommendations. Combining these features with traditional user/item features, we build a group recommender system which assigns players to groups. Finally, we propose that integrating user behavioral features with traditional group recommenders enhances user engagement in games and as a consequence, increases user satisfaction.

The Effect of Stochastic Approximations on a Gossip-Based Sub-Gradient Solver for Linear SVMs

Haimonti Dutta (University at Buffalo, USA)

In the era of big data, an important weapon in a machine learning researcher's arsenal is a scalable Support Vector Machine (SVM) algorithm. SVMs are extensively used for solving classification problems. Traditional algorithms for learning SVMs often scale super linearly with training set size which becomes infeasible very quickly for large data sets. In recent years, scalable algorithms have been designed which study the primal or dual formulations of the problem. This often suggests a way to decompose the problem and facilitate development of distributed algorithms. In this work, we first present a distributed algorithm for learning linear Support Vector Machines in the primal form for binary classification called Gossip-bAseD sub-GradiEnT (GADGET) SVM. The algorithm is designed such that it can be executed locally on nodes of a distributed system; each node processes its local homogeneously partitioned data and learns a primal SVM model; it then gossips with random neighbors about the classifier learnt and uses this information to update the model. The algorithm attains its scalability by using Stochastic Gradient Descent (SGD), in which each successive iteration is determined by appropriately scaling gradient estimates to the prior iterate. We study several state-of-the-art variants of SGD including mini-batch methods, dynamic and adaptive sampling methods, gradient aggregation schemes, and block proximal gradient methods and observe the effect of these stochastic approximations on the GADGET SVM algorithm. Empirical results are presented on several standard datasets used in literature.

Games, Auctions and Consensus-Based Machine Learning

Haimonti Dutta (University at Buffalo, USA)

The evolution of large and complex collections of digital data has necessitated the development of scalable machine learning algorithms. A special class of these scalable algorithms are consensus-based learning algorithms in which learning algorithms are implemented on a set of networked compute nodes (also called agents). Each node is able to communicate with its neighbors. The nodes want to collaboratively minimize a function that involves global information, while having access to only partial information. The underlying communication protocols used in consensus-based learning algorithms is gossip. These gossip protocols

involve periodic, pairwise interprocess interactions of bounded size which allow the nodes to update their local information based on interactions with neighbors. Often neighbors are chosen at random. In this work, we show that communication between nodes can be modeled as an auction where-in the information at a node can be auctioned to n-neighbors. On winning the bid, the bidder pays a price and is allowed to exchange information with the auctioneer. Symmetric and asymmetric first price auctions will be studied in the context of gossip-based communication protocols and their impact on machine learning algorithms (for e.g. distributed Support Vector Machine or Bregman distance algorithms) will be reported.

Software Agent Architecture for Speeding Up Super Singular Isogeny-Based Diffie-Hellman

Wesam Eid, Timothy Atkinson, Marius Silaghi
(Florida Institute of Technology, USA)

Super singular Isogeny-based Diffie-Hellman (SIDH) key exchange protocol has recently gained an increasing attraction and is considered as a good candidate of the post-quantum algorithms family. As of today, the best-known algorithms against the SIDH protocol have an exponential time complexity for both classical and quantum attackers. The primary downside of isogeny-based cryptography is that it is currently a few orders of magnitude slower than other post-quantum cryptography candidates in both hardware and software. Although the SIDH protocol has been slower than other quantum-resistant schemes, it does feature smaller keys, smaller signatures, forward secrecy, and well-defined mathematical basis. Consequently, some recent works have focused on devising strategies to reduce the runtime cost of the SIDH protocol. In this research we introduce an efficient approach for calculating elliptic curve operations by applying a distribution based on an agent software architecture. The distribution technique enables more opportunities for optimizing computations, parallelizing effort and exacting an important speed-up.
Several achieved communication structures were observed to restructure computation steps in ways leading to significant speed-ups. In addition, the results enable us to apply the obtained protocols to an efficient hardware implementation of SIDH key exchange protocol on an FPGA, which is subject to future evaluation. In conclusion, an agent-based software architecture was proposed for managing and enhancing post-quantum elliptic curve cryptosystems, achieving a distribution helping to maximize parallelization of effort and to reduce computation costs, as well as strategies for computation reorganization that can generalize the efficiency improvement.

A Programming Environment for Robotics Education

Gabriel Ferrer (Hendrix College, USA)

We have created and analyzed a system for educational robot programming that presents natural abstractions for robot programming without the use of a programming language. The system is implemented entirely aboard each robot. No auxiliary computing device is necessary for programming. The system enables creation of reactive robot behaviors. Symbolic flags are defined in terms of sensor values. Both Boolean flags and fuzzy flags are available. Actions are defined in terms of motor power levels. Again, both on-off actions and fuzzy actions are available. Each program is a finite state machine. Each node in the machine contains a transition table. Each row in a table contains a flag, an action, and the node whose table is to be used on the next iteration. The first flag that evaluates to "true" determines the action to be executed by the robot, as well as its successor node and accompanying transition table. Available sensors include sonars and several vision-derived inputs, such as color thresholding, color blob location, and neural network image classification. Each robot has two motors, a plastic chassis, three sonar sensors, and a webcam-equipped Android tablet. During the Fall 2018 semester, 24 students used the system. Students give a mildly positive overall assessment of their ability to express their intended behavior using the system. The flag abstraction, the fuzzy logic subsystem, and integrated computer vision were especially popular. Student feedback identified key areas of improvement, including better visualization of the finite state machine and more sophisticated Boolean logic for the flags.

Improving Peak-Shaving Algorithms with Strategic Multi-Agent Multi-Team Reinforcement Learning

D. Michael Franklin (Kennesaw State University, USA)

Energy systems and thermodynamic interactions are intricate and complex. As a result, many extant studies have used many simplifications or generalizations that do not accurately reflect the nature of this complex system. In particular, most HVAC systems are modeled as a single unit, or several large units, rather than as a hierarchical composite (e.g., as a floor rather than as a collection of disparate rooms). The net result of this is that the simulations are too generic to perform meaningful analysis, machine learning, or integrated simulation. We propose using a multi-agent multi-team strategic simulations framework called SiMAMT to better define, model, simulate, and learn the HVAC environment. SiMAMT allows us to create distinct models for each type of room, hierarchically aggregate them

into units (like floors, or sections), and then into larger sets (like buildings or a campus), and then perform a simulation that interacts with each sub-element individually, the teams of sub-elements collectively, and the entire set in aggregation. Further, and most importantly, we additionally model another 'team' within the simulation framework – the users of the systems. Again, each individual is modeled distinctly, aggregated into sub-sets, then collected into large sets. Each user, or agent, is performing on their own but with respect to the larger team goals. This provides a simulation that has a much higher model fidelity and more applicable results that match the real-world.

Addressing Learning Gaps in Astrophysics through Interactive 3D Gravimetric Simulations

D. Michael Franklin (Kennesaw State University, USA),
Andrew Tran (Kennesaw Mountain High School, USA)

Educators in the sciences at the secondary and undergraduate level, specifically for astrophysics or astronomy, do not have ways to effectively model celestial scenarios or phenomenon that are being taught. Also, many of the current methods of visualization—such as commonly used texts and pictures in PowerPoints—are not stimulating enough, which leave students uninterested and less likely to be willing to participate in the activity or lesson. This could result in a general lack of knowledge of astronomy and science in general, with false assumptions being made in certain concepts, since they are being taught unclearly. However, with the rise of modern technology, especially in computer software, various fields in the natural sciences have implemented 3D modeling in order to represent certain physical phenomenon. Companies that utilize scientific expertise develop 3D models and simulations that are based on physical laws during the prototyping stages of their projects. We wish to address these learning gaps by connecting the educational concepts via an interactive, 3D-modeled solar system that will allow for direct-learning through mechanical interaction with the high-fidelity gravimetric model.

A Grounded Theory of the Perception of Recommender Systems

Mohammed Muheeb Faizan Ghori, Jonathan Gemmell
(DePaul University, USA)

Recommender systems have become a cornerstone of modern internet applications, helping users discover new and interesting items. As recommender systems have become increasingly sophisticated, so too has the users' understanding of how these systems work. Users may recognize that these systems use demographic data, explicit or implicit online behavior and contextual data to generate personalized recommendations. Similarly, they may also presume the motives of the system such as customer satisfaction, profit maximization or brand promotions. Based on this understanding, we argue that users possess a cognitive model of how recommender systems work. Consequently, this model may impact how users interact with the system. In this research, we propose a grounded model that describes the working of a recommender system from the users' perspective. Grounded models encompass abstract, detailed and novel concepts that are not influenced by any pre-existing theories. An open-ended survey was given to 100 internet users probing their understanding of how and why online systems produce recommendations. We performed a rigorous data analysis using the data collected to develop a theoretical account of the users' concept of the recommender system. The results were manually annotated by four domain experts and validated for agreement. To the best of our knowledge, this is the first comprehensive study of how users comprehend recommender systems. This analysis will lay the groundwork for the design of recommender systems that will account for the users' understanding of the recommender and its influence on their interactions.

Spectral Unmixing with Semi-Supervised Autoencoders

Jeffrey Graves, Thomas Blum, Ramakrishnan Kannan
(Oak Ridge National Laboratory, USA)

In imaging spectroscopy, special hardware is used to capture hundreds of wavelength bands across the electromagnetic spectrum for every pixel in an image. These spatially coregistered images can then be used to identify objects or materials. Despite having a high spectral resolution, many spectral images have a low spatial resolution. This, along with the effects of multiple scattering, can make it difficult to analyze the composition of a spectral image. Spectral unmixing is the problem of extracting components (endmembers) and their associated spatial locations (abundance maps) from an image using the observed spectra at every pixel. Autoencoders, a type of artificial neural network trained to reconstruct their inputs, have been applied to the problem of blind-source spectral unmixing with relatively good success. However, it is not uncommon for an autoencoder to produce redundant endmembers, uninterruptible abundance maps, and results that do meet physical constraints. To address these issues, we explore the use of semi-supervised autoencoders for spectral unmixing. For many unmixing problems, a domain expert can often provide guidance as to the material most likely present in a scene. Given a collection of reference spectra, we can train an autoencoder to simultaneously learn the provided endmembers and reconstruct the given image. This enables us to guide the model towards

producing physically meaningful endmembers while still allowing the autoencoder to capture nonlinearies in the data and produce unsupervised abundance maps. We have explored the use of autoencoders with semi-supervised training on datasets obtained using electron microscopy as well as satellite imagery.

Case Study Using GitHub for AI Programming Assignments

William H. Hooper (Belmont University, USA)

Programming assignments in Artificial Intelligence courses are particularly difficult to distribute and collect because they are built on a large framework of common code. Importing the code and its dependencies is a problem with many possible solutions. Student submissions that work on their own machines often fail when tested by the grader, often for reasons not clear to either party.

This poster outlines the process used to submit, collect, and test student submissions for an Artificial Intelligence course offered from August to December of 2016, and again from August to December of 2018. Many of the assignments were based on Russell and Norvig, Artificial Intelligence: A Modern Approach, 3rd Edition (AIMA). AIMA presupposes an agent framework that subsumes most of the algorithms presented. The AIMA Online Code Repository, a GitHub project launched in 2016, implements this framework in various programming languages. This repository is the starting point for thousands of student assignments every year.

The essential tools in this process were git (a tool for code management on individual machines) and GitHub (a free website for posting git repositories). Advantages of GitHub over more traditional Course Management System and email submission processes are described, as are limitations of the GitHub-based submission process. Some of these limitations are addressed by GitHub Classroom, another free website, which will likely be used in future iterations of the course.

Using Computational Linguistics to Identify Aggression Patterns in Social Media Data

Sayef Iqbal, Fazel Keshtkar (St. John's University, USA),
Md Suruz Miah (Bradley University, USA)

Aggression such as cyberbullying and hate speech are rising in social media networks which is drawing attention in the research community to detect and investigate the use of such languages. Aggression can be expressed in various forms and devastate victims and affect them for life. Families and social media users prefer a safer platform to interact with each other. Hence, detection and prevention of aggression and hatred over internet is important. In this paper, we extract different features from our social media dataset and perform supervised machine learning methods to understand which model produces the best results and finally we perform some experiments and results. Our dataset comprised of 11,999 Facebook comments in English. We annotated the comments into two categories: Non-aggressive and Aggressive. We extracted TF-IDF, N-gram features, Part-of-speech tags, Sentiment Intensity score and psycholinguistic features from LIWC. We also applied SVM and Random Forest supervised machine learning methods and achieved an F-Score of 0.67 for our final model which involves all the features explained above. We also performed experiments on the features to understand aggression patterns in dataset. Our research shows that the top aggressive comments were related to religion and politics. For part of speech, we found adjective, verb and noun to be the most frequently occurring and associated to aggression. On the other hand, LIWC features demonstrated most psycholinguistic patterns to be about positive, negative and anger in emotion categories.

SPLAIN: Augmenting Cybersecurity Warnings with Reasons and Data

Vera A. Kazakova, Jena D. Hwang, Bonnie J. Dorr, Yorick Wilks
(Institute for Human and Machine Cognition, USA),
J. Blake Gage, Alex Memory, Mark A. Clark (Leidos, USA).

Comprehensible and reliable forecasting systems are needed if cyber threats are to be recognized and cyber-attacks avoided. Prior systems provide little information beyond threat levels and confidence values. We design Simplified Plaintext Language for Actionable and Informative Narratives (SPLAIN), a natural language generator that converts warning data into human-readable prose. SPLAIN endeavors to produce clear and informative outputs with actionable cyber-threat information backed by details regarding input data and system functionality. SPLAIN's inputs are individual sensor-induced forecasting signals coupled with an overall warning produced by a fusion module. SPLAIN queries each signal to obtain information regarding contributing sensors and data signals. The collected information is combined into a single coherent narrative in English that conveys the overall warning – alongside forecasting, sensing, and data elements – to be reviewed by the user at their desired level of detail. SPLAIN employs a hierarchical template-based approach to produce warnings with consistent structure and vocabulary. Each SPLAIN output can be further expanded to reveal the underlying explanation for the overall warning. We conclude that: (1) explainable outputs require that component designers provide the specifics of how or why (not just what) behind cyber warnings; (2) a simple set of structured templates can be developed for generating warnings and explanations with consistent sentence

structure and vocabulary; and (3) direct causal links between inputs and outputs are not always identifiable within Machine Learning approaches, necessitating that some explanations describe general methodology only (e.g., model and training data).

Using Computational Linguistics to Explore Patterns and Features to Identify Disease Outbreak Signals in Social Network Data

Fazel Keshtkar, Neelesh Rastogi (St. John's University, USA)

Social network data, such as twitter has turned out to be an excellent medium for providing public health surveillance and predict the rise of infectious diseases and finally predict the disease outbreak. Based on this factual finding, we present an Infectious Disease Analysis Software (IDAS), an online web-based system, for detecting and visualizing potential disease outbreaks among a certain community that appears as contextual information within Tweets. The system currently analyzes tweet queries for a specific disease (Measles) from a selected geographical area. We conducted classification algorithm on relevance topics, then visualize the outcomes using street map using geo-location information. We used content-based knowledge in our dataset to observe patterns and features selected via various natural language processing tasks conducted over our collected tweets database from an historical disease outbreak timeframe. This content knowledge includes both information on common language used for infectious diseases terminology and analysis on outbreak spread information as a time series data. The system currently consists of six main stages: point of view tagging (i.e., first, second- and third-person references within tweets), topic classification, disease classification, sentiment polarity, and geo-location visualization. Evaluation of the system is conducted by applying training and testing our model against a gold standard corpus of disease outbreak. The results show that system performed well with high accuracy. This led us to observe the detailed understanding of various patterns among tweets related to infectious disease outbreaks.

Towards Parsing Unscoped Episodic Logical Forms with a Cache Transition Parser

Gene Louis Kim (University of Rochester, USA)

Unscoped Episodic Logical Form (ULF) is an underspecified semantic representation resolving the predicate-argument structure within a logical type system. Episodic Logic (EL) is a logical representation that closely matches the expressivity and surface form of natural language and enables inference based on the model theory. ULF resolves the semantic type structure in EL while leaving scoped operators, word sense, and anaphora unresolved. With this, ULF aims to be an accurately recoverable, initial representation towards EL while supporting context-sensitive inferences through the type structure. A small but growing number of annotated examples of ULF are available enabling the deployment of statistical parsing to ULF. We present preliminary results on parsing ULF from English with only naive modifications to a cache transition AMR parser. This parser transforms the problem of mapping a sequence of words to a graph into a problem of mapping a sequence of words to a sequence of actions in a transition system. As ULFs can be syntactically mapped to AMR-like structures, transforming the original AMR parser to a ULF parser consists of building a word-to-ULF-symbol aligner and removing AMR-specific features. Using gold alignments of words to ULF symbols, the parser achieves an EL-smatch score of 0.738, with a training set of only 725 sentences. We also discuss on-going work to refine to the parser by exploiting ULF-specific information such as restrictions in the logical type system.

Merging in Weighted Voting Games: The Case of $k > 2$ Strategic Agents

Ramoni Lasisi (Virginia Military Institute, USA)

Weighted Voting Games (WVGs) are important in multiagent systems because of their usage in automated decision-making. WVGs are not immune from the vulnerability of manipulation by players referred to as being strategic. Previous works have considered merging in WVGs using the well-known Shapley-Shubik and Banzhaf indices to compute agent's power. Upper and lower bounds on the extent of power that manipulators may gain exist for the simple case of when there are k=2 strategic agents in the games. The bounds on the more complicated case of when there are $k > 2$ strategic agents, until now, have remained open for the two indices. We resolve one of these open problems by providing two nontrivial bounds to characterize the effects of merging in WVGs using the Shapley-Shubik index.

Agent-Based Learning Platform for Introductory Artificial Intelligence

Ramoni Lasisi, Robert Dupont (Virginia Military Institute, USA)

We developed ALP4AI, an Agent-based Learning Platform for Introductory Artificial Intelligence. ALP4AI is a graphical-based tool that is suitable for teaching introductory AI, places emphasis on hands-on learning, and provides for visualization of results. The tool we have developed is suitable for solving problems in both the state space search and local search problem domains. ALP4AI is designed to be a simple tool with little learning curve. Thus, students do not need several hours of study to understand the details of how to

use the simulator. ALP4AI is entirely written in Java and makes provision for several functionalities that students can use to model different AI problems and develop solutions.

The functionalities are based on these themes:

Obstacle: No obstacle or obstacles are present in an environment

Agent: Single agent or multi-agent environment

State: micro or macro state modeling of an environment

Goal: Single goal or multi-goal are present in an environment

Themes can be combined in different ways to implement and test AI algorithms. An environment may be defined to contain obstacles, be multi-agent, uses macro state, and be multi-goal. ALP4AI framework is situated in a two-dimensional grid representing environments that agents need to explore. Agents are given the task of locating goals. Many problems in introductory AI can be modeled to use this environment. Students can also conduct and report results of experiments. ALP4AI makes provisions for experiments to evaluate performance of suggested students' solutions. The setup of each experiment includes random placement of agent(s) and goal(s) in the environments.

Intelligent Text Extraction and Summarization for an Improved Community Initiative

Xingbang Liu, Janyl Jumadinova (Allegheny College, USA)

Natural language processing and machine learning algorithms have been shown to be effective in a variety of applications. In this work, we present an intelligent system that was used to process textual information, generate knowledge, and automatically summarize key findings of the My Meadville community statements. My Meadville is a non-profit organization with the goal of highlighting the positive work that is being done in the city of Meadville, PA and bringing upfront the enhancements and improvements that can be made. Over a period of two years, My Meadville conducted a large number of interviews with the residents of Meadville during the community events and transcribed these interviews into textual data files. We developed an automated system that processes these community statements, finds important keywords, and then produces a summary of the key excerpts from all data. Our key contributions lie in the automated workflow that was created utilizing text processing and keyword extraction techniques, and in the text summarization algorithm that utilizes a deep learning framework. Automatic text summarization is the process tasked with creating a succinct and descriptive summary by automatically extracting the most important points of the original text. Our findings have been used by My Meadville to create community value statements, highlight relevant community assets and to develop an action plan based on the concerns and areas of improvement identified by the community members.

Trust-Aware Recommender Systems: Fairness Analysis

Masoud Mansoury, Bamshad Mobasher (DePaul University, USA), Robin Burke (University of Colorado, Boulder, USA)

Fairness has become an important criterion for evaluating recommender systems' performance which refers to the fact that how a recommendation method equally considers both protected and unprotected groups in generating recommendation lists. The aim of this research is to examine the impact of trust-aware recommender systems on generating fair recommendations for protected/unprotected groups. We define protected/unprotected groups based on users' gender and we explore the possibility of improving the quality of recommendations in terms of being fair toward both groups when we use trust ratings as input for recommendation algorithms instead of original ratings. We use Yelp dataset for showing the validity of our hypothesis. Yelp has all required features for this research except users' gender. To resolve this issue, we use an existing API (e.g., https://gender-api.com) to derive users' gender from users' name. Using a comprehensive set of experiments with several real-world data sets, we show that trust-aware recommender systems improves the equity of recommendation lists while it maintains the quality of recommendations.

Vikings' Challenge: A Game to Support Students in Reading Comprehension

Diego Pereira, Fabiana Lorenzi (Universidade Luterana do Brasil, Brasil), Eliseo Reategyu (Universidade Federal do Rio Grande do Sul, Brasil), Andre Peres, Evandro Miletto (IFRS, Brasil)

This research presents the Vikings' Challenge Game: a tool that uses text mining to identify the most relevant words in the text and challenges to trigger reflexive processes and make students think harder about what they read. The Vikings' Challenge game invite the student to read a text in a Viking world. Then the student gets into a room where seven challenges are presented. These challenges are related to the context presented in the text and they are defined by the teacher. The text mining process helps the teacher in the task of creating the challenges previously. For this research, we used the tool Sobek, that analyzes a text and provides a graphical representation of its most relevant terms and relationships. It is able to extract information from text and to represent this information in a graphical view that exhibits the main terms extracted (as nodes) and the relationship between them (as links between nodes). This graphical representation makes it easier for the teacher to understand the text main ideas and how they are connected. Sobek has been

developed using a mining algorithm based on the n-simple distance graph model, and later modified to improve the relevance of the information extracted and to mine larger texts and better display, the information extracted. The student has to overcome all the challenges to get rewards and deliver all the rewards to Odin at Valhalla. The game aims to balance entertainment and education.

Plan Communication to Robots

Rajaa Rahil, Layla Alarifi, Marius Silaghi
(Florida Institute of Technology, USA)

The robots still benefit from some human skills and advice. For the problem of giving a robot directions for navigation in indoor environments, we identify which out of a set of combination of speech, gestures, and drawing mechanisms are the most comfortable, easier to learn, and least error-prone for human users. Three different methods: a Speaking method, a Gesturing and Speaking method, and a Drawing method, were tested for guiding a NAO robot to a desired place, assuming that the robot has no priori map of the building, and even no information about the human faces. Our experiment consists of having participants ask the robot to accomplish a complex indoor navigation task, with each of the above methods. In one experiment, no additional details have been provided about the proposed methods to help the participants, except for a video showing the problem in neuter human terms. In other tests the participants were also given restrictions with respect to robot capabilities to only use speech, gestures, or vision. In the third set of tests, the participants were also given examples of efficient communication with corresponding methods. The task was decomposed, with points assigned for each component. The methods have been investigated and evaluated best on the average task success in simulated plan execution. The preliminary results suggest that the Drawing method leads to the highest level of task accomplishment, and it is also the fastest communication method. Participants feel that speech was more comfortable, despite accomplishing less tasks, and taking longer.

Algorithmic Composition of Classical Music through Data Analysis

Tom Richmond, Ryan Strelow,
Caitlin Harvey, Jeremy Iverson, Imad Rahal
(The College of Saint Benedict and Saint John's University, USA)

The desire to teach computers how to compose music has been a topic of special interest in Computer Science since the 1950's, with roots of computer-less algorithmic composition dating back to Mozart himself. One limitation of algorithmically composing music is the difficulty of eliminating the human intervention required to achieve a homogeneous composition. In this poster, we present a music composition system that examines data of musical scores stored in **kern files to produce "original" music based on the desired era of traditional Western classical music ranging from Medieval to Modern. A piece of classical music is comprised of a finite number of musical notes each containing a pitch and a duration. A musical interval is the distance between the pitches of any two successive notes within the piece. The decision was made to focus upon the frequency with which certain musical intervals occur within the pieces of music of a given era. A musical measure is a meaningful collection of related musical notes within a piece. To generate music that is more aurally pleasing, the system generates sequences of successive notes based upon the intervals between previous notes and newly generated ones. It does so while incorporating aspects of rhythm – such as note durations, rests, and patterns within music measures for a given era – into the generation process resulting in the composition of musical pieces, authored entirely by a computer with no human intervention, that incorporate both pitch and rhythm from a desired classical music era.

An Application of Correspondence Analysis to Word Association Mining Questions

Seon-Hi Shin (Korea Research Institute for School Education, South Korea), Daeryong Seo (Pearson, USA), Moongyu Park (Oakland University, USA)

Mining associations of words from text is an important branch of text data mining and analysis. Among others, paradigmatic and syntagmatic relations are known to capture fundamental relations between units in arbitrary sequences. There have been reported ways to capture such relations between the occurrences of two words. The current study proposes an application of the Correspondence Analysis (CA) scheme to these word association mining questions. CA is a statistical technique to deal with categorical data in the form of numerical frequencies. It not only offers a simple graphical display which permits more rapid interpretation and understanding of the data, but also provides pairwise correlations of elements within and between the analyzed variables. When the CA is applied to word association questions, it can quantify the correlation between the occurrences of two words, and thus one can discover the strongest paradigmatic as well as syntagmatic relations from a collection of documents. For the current study, the authors use 114 written essays of male and female students from two European countries who responded the following questions: "Some people say that in our modern world, dominated by science and technology and industrialization, there is no longer a place for dreaming and imagination. What is your opinion?" and "The prison system is outdated. No civilized society should punish its criminals: it should rehabilitate them". The

implication of the current study is that one can discover word relations by utilizing a well-known existing statistical approach instead of needing to develop a new semantic model.

Effectiveness of Latent Class Analysis Techniques to Handle Count Data with Excessive Zeros in Topic Analysis

Seon-Hi Shin (Korea Research Institute for School Education, South Korea),
Daeryong Seo (Pearson, USA),
Moongyu Park (Oakland University, USA)

Discovering topics discussed in text data is a crucial task for machine learning and scoring engine in the field of educational assessment and/or language testing. This pilot study explored the usefulness of the statistical modeling techniques, Latent Class Analysis (LCA) as a toolkit for topic analysis. LCA offers a flexible framework to handle count data with excessive zeros. Under the framework of LCA, the conditional probability of a word within a topic is permitted to follow a zero-inflated Poisson (ZIP) distribution. The ZIP distribution is proposed to deal with count data with an excessive number of zeros, which is typically the case of text analysis data. The ZIP distribution combines two different distributions to formulate if the behavior ever occurs and how frequently the behavior occurs if it occurs. The results of the pilot study were promising. The text data consisted of forty-eight abstracts of the research papers published in two academic journals from two distinct disciplines. The LCA model could group on average 85% of abstracts into correct topics. When the text length was longer, which was the case of the journal Machine Learning, the accuracy reached 91%. Regarding the task to identify words that indicated topics substantively, the LCA model performed in the topic associated with the journal Nursing relatively well while it showed a somewhat ambiguous result in the other topic. Future research is to compare the LCA model to probabilistic LSA (pLSA) and/or Latent Dirichlet Allocation (LDA) models to test the effectiveness of the LCA approach.

Early Dropout Prediction with Neural Co-Embeddings

Milagro Teruel, Laura Alonso Alemany
(Universidad Nacional de Córdoba, Argentina)

We present an approach to dropout prediction in Massive Online Courses (MOOCS) that relies on a neural model of student behavior. The evaluation is focused on predicting the dropout weekly, using partial information, to simulate a more realistic scenario. We propose to obtain a joint representation (a co-embedding) of students and course components with a recurrent neural network (RNN) trained with logs of student activity. A joint representation is more adequate than disjoint representations because they elicit insights on the interactions between students and contents. Such insights are useful for early prediction, when less information is available. This approach does not require manual labeling of the data, which makes it less prone to theoretical bias, more portable and less costly to develop. Results indicate that a joint embedding improves the performance for datasets with less students.

An Evidence-Based, Pragmatic Approach to Argument Mining

Milagro Teruel, Cristian Cardellino, Johanna Frau
(Universidad Nacional de Córdoba, Argentina),
Serena Villata (INRIA, France),
Laura Alonso Alemany
(Universidad Nacional de Córdoba, Argentina)

Argument Mining is a growing area of work within Natural Language Processing, because it has a number of applications, ranging from targeted information retrieval for e-debate to building argument-aware summaries. However, the performance of automated systems is still far from satisfactory. The complexity of argumentation makes it difficult for humans to create rule-based systems or annotated corpora. It also makes it very difficult for automated systems to infer models that solve the problem adequately.

In this poster we deal with the complexity of argument mining by analyzing errors in automated systems and disagreements between human annotators. We then propose how to reduce such errors, in various ways:

a) Changing the representation of argumentation to match what human annotators identify more clearly, based on inter-annotator agreement. We have found that annotators tend to agree on what a claim is, but disagree on the scope and type of justifications.

b) Applying it on a dataset that is regular enough for machine learning approaches to infer a model. We are focusing on sentences from the European Court of Justice, which have the advantage of being parallel in various languages.

c) Targeting a concrete application as a way to evaluate. We will evaluate the results of annotation and automated analysis as a reading aid. This application is concrete enough for a clear evaluation and yet general enough to assess the goodness of a system in the argument mining task.

We show that our proposal is general enough to be of use, or at least as an inspiration, for applications to other domains.

Exploring Bias Through Model Explanations in NLP Tasks

Khonzodakhon Umarova, Eni Mustafaraj
(Wellesley College, USA)

Recently, decision-making by machine learning algorithms in areas like recidivism risk-assessment, predictive policing, and automated hiring has been criticized because the learned models reinforce biases already present in society and exacerbate unfair treatment of members of certain groups. In addition, most of these models such as Deep Neural Networks are very complex and often operate as "black-boxes." The need for fairness on one hand and lack of transparency from these models on the other inspired a research movement towards Fairness, Accountability, and Transparency in AI and machine learning. One of the concepts in the context of this research is that of model explainability, which is concerned with informing stakeholders of a model's functionality, decisions, and rationale. Given the novelty of this research, the goal of our poster is to provide an overview of explanation methods in NLP and examine whether these methods are effective in detecting gender bias in texts. Current NLP research for model explanations in classification tasks includes explanation methods such as Gradient-based explanations, Layer-wise relevance propagation, DeepLIFT (Deep Learning Important FeaTures), etc. However, such research is conducted in the context of datasets that do not take into account various biases. Therefore, we look into existing explanation methods on datasets that potentially exhibit naturally-occurring gender bias and determine whether these methods are able to capture and explain bias in the model. In particular, we will look into media coverage of women running for political office, job application and hiring decisions for women, and misogyny in social media.

Towards Improving Semantic Parsing Using Statistical Word Sense Disambiguation

Siddharth Vashishtha, Ritwik Bose, James Allen
(University of Rochester, USA)

The TRIPS semantic parser is a bottom-up chart parser which uses a handcrafted ontology, lexicon, and grammar to produce logical forms (LFs) for sentences. Incorrect Word Sense Disambiguation (WSD) decisions made early in the parsing process can lead to large errors in the final LF. Errors compound as the parser attempts to adequately satisfy selectional type restrictions for semantic roles or attempts to fit predicates to incorrect type signatures. Additionally, in order to maintain a manageable state-space, the parser prunes off lower-likelihood constituents from the chart using a variety of heuristics. The aim of this work is to add a prior distribution over the initial constituents on the chart to the heuristics already present in the parser. This will allow us to keep higher likelihood senses on the chart longer without sacrificing the ability of the parser to honor type restrictions. We perform independent WordNet-based WSD using existing WSD systems and convert the sense-likelihoods to distributions over TRIPS ontology types. This approach allows us to maintain the high level of fine-grained control afforded by hand-crafted rules while benefiting from statistical word sense knowledge. This poster shows a qualitative and quantitative analysis of our initial experiments. We examine approaches for incorporating the statistical scores into the parser. We observe that in cases where statistical WSD does not respect role restrictions, the parser does. In the future, we hope to leverage this fact to also correct errors in statistical WSD by applying sentence level selectional semantic restrictions.

Underwater Robotic Unit with Smart Sensing Technology

Elisia Wright (Allegheny College, USA),
David Boughton (Pennsylvania Sea Grant, USA),
Janyl Jumadinova (Allegheny College, USA)

Water quality testing is important as it can identify possible contaminants in the water. The current methods for water quality testing use individual stationary sensors to obtain random readings for testing each parameter, such as temperature, pH, conductivity and dissolved oxygen, among others, separately. In this work, we have designed a sensor unit comprised of multiple sensors for water quality testing that are able to collect data simultaneously and autonomously. We also built an underwater robot that is able to travel to different depths in the water. The sensor unit is waterproofed and attached to the underwater robot that can be used to collect more encompassing data from different levels of the water column when sampling water quality. The robotic unit with the waterproof sensors is able to collect data underwater for several hours at a time and transmit it to the analytics software autonomously. The software is built with the coordination from the environmental water experts to automate data analysis and utilizes machine learning algorithms to aggregate, learn and predict trends in the collected data. Sensor calibrations and initial testing of the underwater robot has been conducted in the pool, with more extensive experiments to be run in the waters of northwest Pennsylvania, including local French Creek and Lake Erie. The results of this project, including the collected data and its analysis, will be shared with other researchers and will be used to help environmental science students understand variations in water quality measurements.