FLAIRS-31 Poster Abstracts

Vasile Rus, Keith Brawner

Editors

A Survey about Machine Learning Algorithms in E-Commerce

Alla Abdella
(Florida Institute of Technology)

This survey paper covers two e-commerce applications, marketing selling enhancement, and riskiness fortification. In online marketing, e-commerce sites like Amazon, eBay, etc., utilize data mining algorithms profusely to assist its customers in finding what their purchase needs. With the growth of the e-commerce domain, credit card usage has become a common phenomenon. This gives a chance for fraudulent to commit fraud. This paper presents two topics - recommendation systems and fraud detection - and shows how sophisticated machine learning approaches can enhance both. In particular, it presents a survey about collaborative methods as part of an attempt to improve recommendation systems. It then surveys support vector machines as an attempt to enhance credit card fraud detection. Also, this paper elaborates on how the aforementioned algorithms enrich both users' online buying experience and e-commerce sites' safety.

Value-Aware Recommendation with Multiple Stakeholders

Himan Abdollahpouri, Robin Burke, Bamshad Mobasher
(DePaul University)

The main goal of recommender systems typically is to provide a personalized information access for users. That is, showing products to each user, tailored based on their individual taste and preferences. In many real-world applications, however, that is not the only consideration. An online grocery, for example, has thousands of products from hundreds of brands each of which may have a different commission rate with the platform owner for recommending their products to the user. In this situation, the recommender system must not only take users' preferences into account, it also should try to maximize the system revenue by considering the value of a recommendation. Moreover, it also has to, in some cases, meet item providers' (brands or other owners) preferences regarding which group of users to target. In this work, we introduce a system design for the value-aware multi-stakeholder recommendation that is able to generate recommendations by taking into account item profitability and brand preferences as well as users' preferences. Our system design is general enough that could also be used for other domains which have multiple stakeholders involved in delivering or receiving the recommendations such as online retail stores, apartment sharing businesses, ride sharing systems etc.

Framing Impacts and Avoidance Techniques for Group Decision Support Systems

Badria Alfurhood, Marius Silaghi
(Florida Institute of Technology)

Framing refers to untoward ways in which questions or alternatives are presented to decision makers for deliberation purposes. Framing is an attack method used to manipulate people into preferring a particular alternative. Deliberative framing threatens the integrity of group decision support systems (GDSSs) and has negative impacts on ultimate decisions. Various studies have shown that, unfortunately, GDSSs are vulnerable to framing, as early posted information can sway many later arriving participants. However, framing avoidance or reduction is possible. Framing impacts are studied to get insights for providing mechanisms to avoid negative effects within GDSSs. The GDSSs structure and mechanisms can be designed in ways that reduces framing effects. GDSSs can be designed to incorporate resistance to framing attacks. In this research, we address the GDSSs framing behavior, the techniques people use to frame choices or arguments, factors that influence framing in GDSSs, and the design practices and recommendations that eliminate or lessen the negative framing effects. We investigate several approaches to designing GDSS tools for providing resistance to framing, namely

based on: intrinsic intelligence via input constraints, collaboration filtering, and intelligent rendering interfaces. Mainly, the effort is focused on proposing tools for experiments, identifying GDSSs features that reduce framing, and measuring the framing effects in GDSSs. De-biasing framing in GDSSs has big potential impact on society as people rely heavily on social networks and related technology for gathering information to be used in deliberation and voting purposes, as demonstrated by recent elections.

Can a Computer Learn from a
Natural Conversation with Humans?

Awrad Mohammed Ali and Avelino Gonzalez
(University of Central Florida)

One goal of artificial intelligence is to develop a conversational agent that can communicate efficiently with the user without many restrictions or predetermined patterns. Furthermore, it should be an agent that can learn from an open-ended natural conversation as do humans. However, existing agents require inputs with a clear description or specific words in a specific order to trigger the learning process. This makes the procedure tedious and dependent mostly on the human. In this work, we present a new conversational learning approach that aims to learn from an extended natural dialogue with a human. We focus our discussion on how a computer agent can benefit from using a meaningful representation for its knowledge base, such as a modified version of a semantic network that links sentences that have a high similarity measure rather than representing the nodes with words only. This approach makes linking the acquired knowledge to existing information more efficient and accurate. This work also discusses the benefit of using majority voting among multiple existing classification algorithms to decide whether a provided input should be learned or only considered a chatting statement and therefore neglected. We also propose using a second level of classification based on the word level when there is no unanimous agreement between the classifiers on the label of the user input (relevant information or chatting statements). This step improves the classification process and allows additional information to be captured and learned. This research makes strides toward more natural and robust conversational learning systems.

Models and Inference Techniques
for Diagnosis of Embedded Components

Timothy Atkinson, Marius Silaghi
(Florida Institute of Technology)

In largescale software projects, it is not commonly known how to make the product perfectly secure. Often software is first written with the sunny day scenario; and then each rainy-day scenario the software/system engineers can think of is explored and tested until the company feels that they have a product they can sell. Anyone who has shipped a piece of software knows that once the software hits the real world, new scenarios they never thought of could happen occur. One approach is to perform more analysis prior to release. An alternative technique which we propose, is to have the system engineers provide a parallel formal model as part of the development process. Then as the software runs, the model can alert the user if there is ever a significant discrepancy between the model and the running software, helping the user know that there is a problem as early as possible. We are experimenting with using statistical models to determine if a trained model is different from the current observed model using logical and probabilistic inference techniques.

Bridging the Gap between Artificial
and Spiking Neural Networks

Sylvain Chartier
(University of Ottawa)

Since the development of Spiking Neural Networks (SNN), it seems that they should be the next successor of artificial neural networks. SNN is based on temporal aspect of spiking dynamics instead of the mere presence or absence of spikes; or mean firing rates. Therefore, they should be used in more applications than standard artificial neural networks in machine learning. However, training SNN to encode desired behaviours reveals to be more challenging than expected. Usually, researchers will used spike timing dependent plasticity to reinforce synapses which is difficult to implement in a network using distributed representation. On the other hand, Recurrent Associative Memories (RAMs) can store information easily in such representation but lack the dynamical aspect of spikes. Few studies were able to use recurrent associative memory to display spiking behaviour and when they do they simply replace the output function with a spiking one. This has a consequence of still having two separate models (spiking and non-spiking). In this study we propose to use an output function that allows both spiking and non-spiking behaviour within the same model. Results show that the resulting model can adapt to context and display essential RAM properties (e.g. patterns storage, noise tolerance) while being able to output spike responses. Therefore, the proposed model opens the door to bring the best to both worlds and allow a wider range of applications in machine learning.

Improving User Acceptability of Recommendations through Opinion Mining

Arman Dehpanah, Jonathan Gemmell
(DePaul University)

Recommender systems have become increasingly popular in many different fields to enhance the user experience, however some users, despite receiving personalized recommendations, end up reading hundreds of reviews of an item in order to make an informed buying decision. In this research, we use Word2Vec to analyze the text reviews and discover meaningful features. Word2Vec includes two implementations of a shallow neural network; Continuous Bag of Words (CBOW) and Skip-gram. We rely on CBOW for the purpose of word representation. CBOW is a neural network that predicts a target word given the surrounding context words in a sentence. Representations obtained from CBOW are used as input for clustering. For instance, when analyzing comments containing words such as "cried", "end", "hero" and "died" for movies, using CBOW we can derive a topic such as "Sad Ending" for those movies and form a cluster containing movies with sad endings. After forming different clusters, each cluster is assigned an importance and sentiment score to determine their ranks. Finally, we propose that presenting these topics along with the recommendations offers a more satisfying experience for the user.

Performance Evaluation of a Real-Time Clustering Algorithm

Gabriel Ferrer
(Hendrix College)

"Clustering algorithms are unsupervised learning algorithms that employ a distance metric to organize their training inputs into clusters. The classification of a previously unseen input is determined by the closest distance from that input to a cluster's reference value. A real-time clustering algorithm updates the reference values for its clusters in a constant amount of time for each training input. We are currently assessing the performance characteristics of the recently introduced real-time clustering algorithm Bounded Self-Organizing Clusters (BSOC).

When it was originally presented, the classification performance of BSOC was compared to a real-time formulation of the Self-Organizing Map (SOM). We find this comparison flawed. SOM clusters are deliberately moved close to the reference values of their neighbors, while BSOC tries to maximize distances between the reference values of all clusters. Because of this, we believe that k-means, which also tries to maximize these distances, is a more appropriate clustering performance benchmark.

The goal of this study is to determine the relative effectiveness of BSOC and k-means on a range of clustering problems. We are using the k-means++ initialization algorithm. So far, we have found that the mean clustering error is very similar when both are applied to the color quantization problem. We next plan to evaluate performance on the problem of extracting convolution kernels from input images. Our final comparison will be for clustering video images, comparing k-means++ with BSOC at several different checkpoints in the image sequence."

Hybrid Goal Selection and Planning in a Goal Reasoning Agent Using Partially Specified Preferences

Michael Floyd, Mark Roberts, David Aha
(Knexus Research, Naval Research Laboratory)

Goal Reasoning agents are not restricted to pursuing a static set of predefined goals but can instead reason about their goals and, if necessary, dynamically modify the set of goals that they will pursue. For a solitary agent, goal selection is guided by the agent's own internal motivations. However, an agent that is a member of a team also needs to consider its teammates' preferences when selecting goals. In this work, we propose an online approach to estimate the utility of goals based on a teammate's partially specified preferences and use the estimated goal utilities to guide goal selection. Specifically, we consider the situation where an autonomous agent is teamed with a single operator in a supervisor-supervisee relationship. At the start of a mission, the operator will provide an initial set of goals and, optionally, partially defined preferences for how the mission should be performed. After the initial interaction, the agent acts autonomously to pursue the provided goals. While acting autonomously, the agent may encounter situations that present opportunities for new goals and must determine if those goals align with the operator's preferences. Estimated goal utilities are used during a hybrid goal selection and planning process to select a subset of goals for the agent to pursue. We report evidence from an empirical study which demonstrates that our approach outperforms several baselines in scenarios drawn from a simulated human-agent teaming domain.

eSense 2.0: Modeling Multi-Agent Biomimetic Predation with Multi-Layered Reinforcement Learning

D. Michael Franklin, Derek Martin
(Kennesaw State University)

Accurately modeling predator-prey relationships in a biomimetic was can be a daunting task. Applying machine learning to solve such a modeled task is even more difficult. What is needed is a system that uses multiple simple models to accurately represent complex biological multi-

agent behavior. To do so, eSense 2.0 is being introduced. Building on the success of the eSense BioMimetic modeling done in (Franklin and Martin 2016), eSense 2.0 expands the modeling to include a stronger predator / prey relationship. eSense provides a powerful yet simplistic reinforcement learning algorithm that employs model-based behavior across multiple learning layers. These independent layers split the learning objectives across multiple layers, avoiding the learning-confusion common in many multi-agent systems. The new eSense 2.0 increases the number of layers and the amount of separation between agents so that the behaviors for each agent can be more highly customized and adds specific additional layers for behavior-only learning. In other words, each agent now has multiple layers to model each aspect of their behavior (e.g., obstacle avoidance, prey observation, prey seeking, etc.). This new abstraction of breaking various agent behaviors into multiple levels furthers speeds learning and clarifies the objectives the agent is considering. This builds on the general goal of eSense (splitting out multiple agents into their own levels) because now the agent's behaviors are also split into multiple layers. The learning is now more expressive, faster, and less noisy. This paper presents this new multi-level learning system for multi-agent systems and experimentally confirms its performance.

firstGlimpse: Learning How to Learn through Observation via Memory Modeling with Reinforcement Learning

D. Michael Franklin, Ryan Kessler
(Kennesaw State University)

While there are many systems extant that take well-known models and refine how to learn with them, there are very few that deal with the initial stages of raw model formation. How do we, as human being, learn? Can we teach machines to 'learn to learn'? In this paper, the authors introduce just such a system. firstGlimpse is a hierarchical memory-modeling technique that can learn from scratch via observation. It uses reinforcement learning to develop its models and reinforce which of these models is viable, which are most likely, etc. This hierarchical approach to memory modeling develops from symbols to words to phrases to sentences, thus creating a modeling system for general learning that can be built from the ground up via observation. As a new symbol is observed, it is recorded. Likewise, that symbols connections with other symbols is also recorded in another layer of the memory. Next, these connected symbols are viewed as words, and the words are observed as being connected to other words. In this manner, phrases are added. This continues upwards to higher levels of understanding. firstGlimpse uses the game of Set to show, starting from scratch, that it can learn to play by

observing humans playing the game and recording its observations within its hierarchical memory. The experiments will show that this method is successful both in observing the links of the memory graph and in the results of successfully playing the game.

Data Mining Approach to Estimate Field Popularity from the US College Scorecard Data

Fazel Keshtkar, Shiromani Neerudu, Md Suruz Miah*
(St. John's University, *Bradley University)

The College Scorecard is launched by the U.S. Department of Education in September 2015 that provides information about different colleges. It has a rich resource of information intended to help students and parents to attend colleges in USA. The data set includes information from 1996 through 2015 for all undergraduate degree granting institutions of higher education, and is updated each year. In this paper we aim to explore The College Scorecard Data information based on college affordability and values. Our goal is to extract patterns in these data that can provide students to make better decisions to which college to attend in United State. College Scorecard data is an online tool created by U.S. Department of Education that launched in 2015. The Data contains a variety of information about admissions, student body, cost, department, completion and earnings by institution. We used unsupervised learning, clustering analysis, PCA, and data visualization approaches to explore, underlying the structure of data, and pattern extraction. The model shows the percent of total enrollment for each field ranked highest to lowest. The Business Program was enrolled 18.23% and second highest enrollment was Health Program with 12.83%. The analysis will use the number of enrolled undergraduates and the percent of degrees awarded in each field to estimate the total number of enrolled undergraduates for each program. The PCA component is finding schools with students who did well on the SAT score, complete college at a high rate, and from high income.

Compensating for Rating Distribution through Percentile Transformation

Masoud Mansoury, Robin Burke, Bamshad Mobasher
(DePaul University)

"Recommender systems use information from user profiles to generate personalized recommendations. User profiles are either implicitly inferred by the system through contextual information, or explicitly provided by users. In the latter case, users are asked to rate different items based on their preferences and may have individual differences in how they use explicit rating scales: some users may tend to rate higher, while some users may tend to rate lower. It is a well-known phenomenon in recommender systems that ex-

plicit user ratings are biased toward high ratings, and the users differ significantly in their distributions. This is usually compensated for through mean-adjustment in rating normalization or the inclusion of a user bias term in factorization models.

In this work, we propose a rating transformation model that compensates for skew in the rating distribution by converting ratings into percentile values as a pre-processing step before recommendation generation. This transformation flattens the rating distribution to reduce its skew, which better compensates for differences in rating distributions and improves recommendation performance. Also, each percentile value associated with an item reflects its rank among all of the items that the user has rated. Thus, the percentile captures an item's position within a user's profile better than the raw rating value or a mean-adjusted version of the same, and therefore compensates for differences in users' overall rating behavior. A comprehensive set of experiments is performed, showing improved ranking performance for the percentile technique across a variety of state-of-the-art recommendation algorithms in four real-world data sets."

Genetic Algorithms on Tensor Network Contraction Order Finding

Reamonn Norat, Annie Wu, David Anekstein,
Jonathan Jakes-Schauer, Pawel Wocjan
(University of Central Florida)

We analyze the use of genetic algorithms on the problem of tensor network contraction order finding, which can be reduced to a permutation problem. The tensor network can be represented as a connected graph with nodes corresponding to tensors. Each edge incident to a node corresponds to a particular index and each weight corresponds to the number of values assumable by that index. Tensor network contraction involves a series of edge contractions; contractions occur until there is only one node remaining in the network. Each contraction has an associated cost, and the sum of these costs represent the approximate computational cost of contracting the entire tensor network. This total cost can vary greatly depending on the edge contraction order. In practice, the current state of the art is guaranteed to find an optimal solution, but the time required will scale exponentially with network size. We investigate the use of a genetic algorithm and compare it against the state of the art approach to tensor network contraction ordering. Within the genetic algorithm, we examine the performance tradeoffs among three different representation strategies.

The Game of Chicken and Bitcoin Trading

Marius Silaghi, Badria Alfurhood, Timothy Atkinson
(Florida Institute of Technology)

Modeling large-scale social phenomena related to Bitcoin trading and processing using theoretical games has the potential to predict and even control outcomes. Efforts have also been in the past directed to the modeling of other significant world problems such as nuclear war, Ponzi schemes, and the theory of money. The Bitcoin phenomenon has overtaken the world attention through an unusual combination of idealism and utilitarian interest. As with other recent crypto currencies, Bitcoin trade arguably adds intrinsic value by supporting anonymous transfers between idealistic or criminal entities. Entities banned from selling their coal or other energetic resources, can convert them to anonymous value on an open and still largely unregulated international market. There exists also a significant speculative side of this trade, which alone would amount to a zero-sum game, if not for the price fluctuations and high energetic costs of transactions. This is superposed with the already studied game of the miners and traders, into a complex environment influenced by available information. We propose an agent-based simulation method based on iterations where the result of each long-term simulation is used to refine the pay-off matrix of each type of participant, to capture the impact of media and information on the game structure. Initial simulation inputs are currently approximated with "reasonable values" but can be refined with corresponding social studies. The result of our simulations, while currently heavily dependent on approximations for initial participant structure and pay-off matrix values and dynamics, can help improve predictability and stability in the market.

Fairness-Aware Recommendation Systems

Nasim Sonboli, Robin Burke, Farzad Eskandanian
(DePaul University)

The basic literacy is a complex phase, composed of several stages that require dedication to the conclusion of the process. Nowadays, with the evolution of electronic devices and with the advent of data mining technology, respectively, the teaching and monitoring of the stages of literacy can be facilitated by the introduction of its principles into the educational process. Based on this, this paper proposes the presentation of a learning object, consisting mainly of a mobile application and a follow-up system, named Object of study of Literacy, Objeto de Estudo de Letramento (OEL), for teaching and monitoring the initial stages of the process of basic literacy.

Comparing General and Domain-Specific LSA Classifiers in the Context of Virtual Internships

Zachari Swiecki,[1] Vasile Rus,[2] Zhiqiang Cai,[2]
Dipesh Gautam,[2] David Williamson Shaffer[1]
([1]University of Wisconsin-Madison, [2]University of Memphis)

Virtual internships are simulations that give students the opportunity to experience realistic professional practices and problem solving in an online environment. As part of the virtual internship experience, students submit short written reports of their work, which are assessed by human raters. To implement virtual internships at scale, automated assessment methods are needed to reduce this human effort. Our prior work trained classifiers on student data to automatically identify domain-relevant concepts in student reports and assess their quality. However, educators can now develop virtual internships for new domains using authoring tools. In such cases, student data does not initially exist for classifier training. In other words, to include automated assessment models in their new internships, educators need a way to bootstrap these models before students participate in the internships and before any student data is available to train classifiers. To address this problem, we developed a method for generating classifiers based on limited information provided by educators during their authoring process, including concept definitions and exemplar reports. Our classifiers rely on a latent semantic analysis (LSA) approach using a domain-general corpus. Prior results showed that classifiers developed using this corpus had mixed performance. In this work, we compare the performance of the domain-general LSA classifiers to LSA classifiers developed using domain-specific corpora of varying document and dimension sizes, and we investigate how concept and exemplar features relate to performance in each case.

Cooperation Protocols for Ad-Hoc Robot Teams Composition in Labyrinth Exploration

Muntaser Syed, Marius Silaghi, Rajaa Rahil,
Sam Kellar, Shakre Elmane
(Florida Institute of Technology)

One's cohesive team of robots may participate in disaster rescue operations where other teams from other organizations/countries are also ready and willing to help. Such robotic teams may not be specifically prepared to work together, but they can apply ad-hoc collaboration mechanisms if appropriate communication interfaces and decision-making mechanisms are defined and supported. Significant effort has already been laid into the design of ontologies and their standardization. Further, physical communication interfaces have also emerged that are widely supported, such as wireless and Bluetooth. The challenge is to now come out with collaboration mechanisms that can exploit such opportunities. Protocols can be standardized to support a flexible engagement in such decision-making. Flexibility here is understood from the perspective of the depth and amount of details of the involvement. We take as departing point the advancement in Distributed Constraint Reasoning, which converged to a set of flexible protocols based on the assumption of privacy of constraints. This privacy assumption inspired the definition of protocols leaving agents the flexibility of selecting their degree of collaboration within limits that maintain guarantees concerning decision-making processes. Protocols are developed for guaranteeing that the robots advance in their labyrinth exploration tasks while respecting physical constraints and pursuing progress, exchanging map information, avoiding collisions, and reducing overlapped exploration.

Recognizing and Exemplifying Gender Bias in Online Articles

Khonzodakhon Umarova, Eni Mustafaraj
(Wellesley College)

There is so much information online, with millions of new websites being created continuously. The notion of credible sources is being lost, as users rely on search engines and social networks for encountering news or fulfilling their information needs. This has lead to an increase in the spread of misinformation, which takes many forms and disguises, from pseudoscience and conspiracy theories, to fake news and propaganda. At the same time, numerous studies have shown that young people lack skills necessary to critically assess and evaluate online information that is consumed so readily. Middle-, high-school and even college students struggle to identify biases in the information sources; differentiate between fact and opinion based texts; or recognize promotional pieces. In order to develop web/news literacy skills and educate young people, we imagine augmenting search results or web pages with particular signals that will raise awareness and invite users to reason about the credibility of the source. One such signal can refer to a possible bias in the source. Examples of biases are: ideological, cultural, or social. How to learn to recognize biases automatically at web scale? We explore the feasibility of a new approach, Word2vec, which computes vector representations of words and has been shown to be effective in determining words that are semantically related to each other. We we will use it to identify examples of gender bias in articles, not only to recognize such instances, but to extract concrete phrases that can be used to educate through examples.

Detecting Vehicular Patterns Using a Graph-Based Approach

Sirisha Velampalli, Lenin Mookiah, William Eberle
(University of Hyderabad Campus, CA Technologies,
Tennessee Tech University)

Extracting knowledge from heterogeneous datasets is a complicated but useful task. The Visual Analytics Science and Technology (VAST) 2017 challenge, deals with unusual activity at a wildlife preserve. It is identified that in this preserve, there are a number of nesting pairs of Rose-Crested Blue Pipit, a popular local, singing bird with attractive plumage. In the VAST competition, one of the challenges is to discover vehicular traffic patterns for understanding the reasons behind a decrease in the number of nesting pairs. One approach to solve this problem is a graph-based approach. Graph-based approaches enable one to handle rich contextual data and provide a deeper understanding of data due to the ability to discover patterns in databases that are not easily found using traditional query or statistical tools. In this work, we present a graph-based approach that analyzes the data for structural patterns in the data. Our approach first reports the normative patterns in the data, and then discovers any anomalous patterns. After analyzing the patterns, we are able to hypothesize some of the reasons that include: (a) most vehicles in the park are 2-axle cars, emitting more pollution than other types of vehicles; (b) vehicular traffic is at its peak during the months of July and August - the prime breeding months for these birds; and (c) traffic is at its peak in the mornings, which effects the atmospheric conditions surrounding when the birds sing – effecting the migratory patterns of the birds.

An Analysis of WordNet's Coverage of Personality Disorder Terms Using A Personality Disorder Corpus

Morgan Wixted, Amanda Hicks, Mary Cate Espinosa
(University of Florida)

The healthcare industry as a whole is turning to natural language processing for fast solutions to complex problems. Personality disorders, a subset of the mental health field, has limited and underdeveloped computational resources that could benefit researchers if developed. WordNet is becoming an increasingly popular resource in the mental health and natural language processing industries. With the addition of personality disorder characteristics into WordNet, there could be an enhanced version of WordNet that is designed to be more helpful to its various users, including researchers and the general public. In this project we analyze the coverage of various mental health terms in WordNet gathered though corpora of six different personality disorders: antisocial personality disorder, borderline personality disorder, histrionic personality disorder, narcissistic personality disorder, paranoid personality disorder, and schizoid personality disorder. We then discuss some key considerations and next steps for adding personality disorder terms and their meanings to WordNet.