Institute for
Artificial Intelligence
Franklin College of Arts and Sciences
UNIVERSITY OF GEORGIA

AI Everywhere (All at Once)

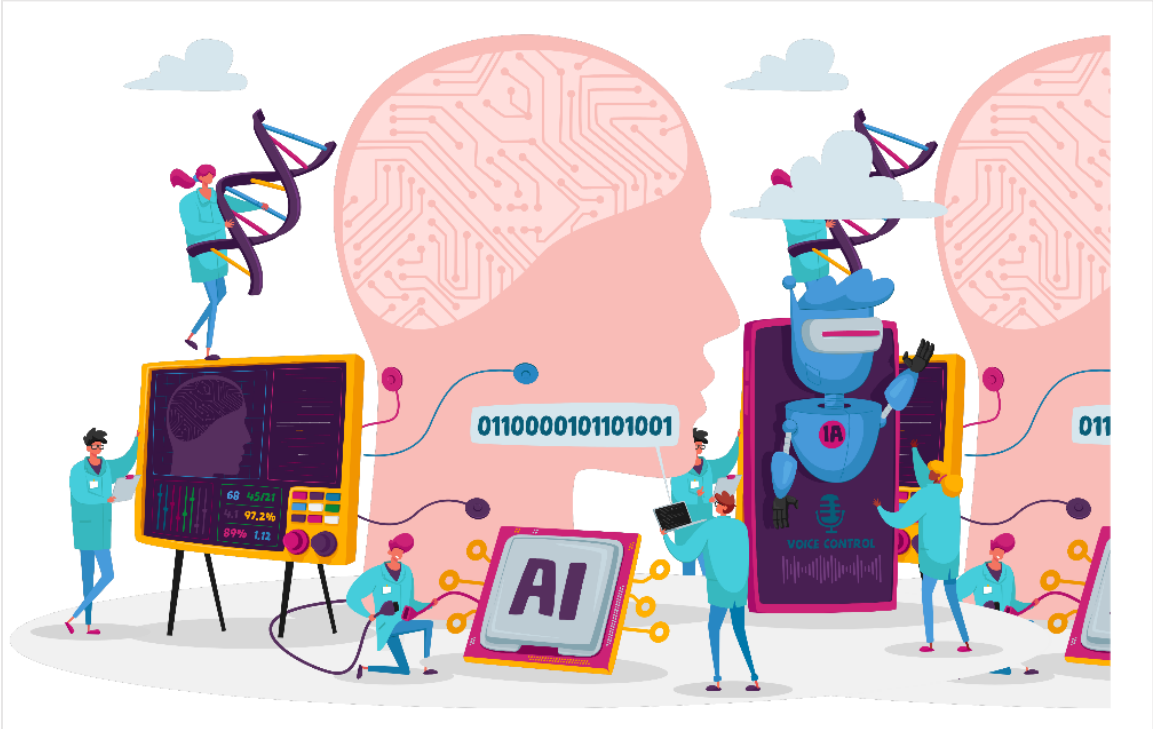

AI Research Day
Fall 2022

Contents

AI Research Day: Schedule

2:00 PM	Welcome
2:15 PM	Keynote
3:15 PM	Lightning Talks (Session I)
3:45 PM	Break
4:00 PM	Panel Discussion
5:00 PM	Lightning Talks (Session II)
5:40 PM	Poster Session

Organizing Committee

- **Frederick Maier** (Committee Chair; Institute for AI)
- **Prashant Doshi** (School of Computing)
- **Evette Dunbar** (Event Planning—Institute for AI)
- **Neal Outland** (Department of Psychology)
- **Ramviyas Parasuraman** (School of Computing)
- **Khaled Rasheed** (School of Computing)
- **Carolina Alves de Lima Salge** (Department of Management Information Systems, Terry College of Business)

Keynote: Peter Stone

@2:15 PM

Director of Texas Robotics, UT Austin; Associate Chair of the Department of Computer Science, UT Austin; Truchard Foundation Chair in Computer Science, UT Austin; Executive Director, Sony AI America

Outracing Champion Gran Turismo Drivers with Deep Reinforcement Learning

Many potential applications of artificial intelligence involve making real-time decisions in physical systems while interacting with humans. Automobile racing represents an extreme example of these conditions; drivers must execute complex tactical manoeuvres to pass or block opponents while operating their vehicles at their traction limits. Racing simulations, such as the PlayStation game Gran Turismo, faithfully reproduce the non-linear control challenges of real race cars while also encapsulating the complex multi-agent interactions. Here we describe how we trained agents for Gran Turismo that can compete with the world's best e-sports drivers. We combine state-of-the-art, model-free, deep reinforcement learning algorithms with mixed-scenario training to learn an integrated control policy that combines exceptional speed with impressive tactics. In addition, we construct a reward function that enables the agent to be competitive while adhering to racing's important, but under-specified, sportsmanship rules. We demonstrate the capabilities of our agent, Gran Turismo Sophy, by winning a head-to-head competition against four of the world's best Gran Turismo drivers. By describing how we trained championship-level racers, we demonstrate the possibilities and challenges of using these techniques to control complex dynamical systems in domains where agents must respect imprecisely defined human norms.

Lightning Talks (Session I)

@3:15 PM

Gengchen Mai
Assistant Professor, Department of Geography

Mohamad Kazem Shirani Faradonbeh
Assistant Professor in Data Science, Department of Statistics

Katherine Ireland
Interim Head of the DigiLab, Digital Humanities

Camila Lívio
Ph.D. Candidate in Romance Languages

Panel Discussion

@4:00 PM

Prashant Doshi (moderator)
Professor, School of Computing

Anna Abraham
E. Paul Torrance Professor, Department of Educational Psychology; Director, Torrance Center for Creativity & Talent Development

Aaron Schecter
Assistant Professor of Management Information Systems, Terry College of Business

Wenzhan Song
Georgia Power Mickey A. Brown Professor, School of Electrical and Computer Engineering; Director of the Center for Cyber-Physical Systems

Neal Outland
Assistant Professor (Industrial-Organizational), Department of Psychology

Lightning Talks (Session II)

@5:00 PM

Sergio Bernardes
Associate Director, Center for Geospatial Research; Director, Disruptive Geospatial Technologies Lab

Jeremy Davis
Assistant Professor, Department of Philosophy

Eric Miller
Small Satellite Research Lab; MSAI Student

John Gibbs
Associate Professor, Department of Theatre and Film Studies; President & CEO, Artimatic, Inc.

Research Poster Session

@5:40

Poster #1

Mapping the glycosyltransferase fold landscape using interpretable deep learning

Zhongliang Zhou, Rahil Taujale, Wayland Yeung, Kelley W. Moremen, Sheng Li and Natarajan Kannan
School of Computing

Glycosyltransferases play fundamental roles in nearly all cellular processes through the biosynthesis of complex carbohydrates and glycosylation of diverse protein and small molecule substrates. The extensive structural and functional diversification of GTs presents a major challenge in mapping the relationships connecting sequence, structure, fold, and function using traditional bioinformatics approaches. We present a CNN-attention model that leverages secondary structure representations to provide GT fold prediction with high accuracy. The model learns to distinguish secondary structure features without primary sequence alignment and is highly interpretable. It delineates sequence and structural features characteristic of individual fold types while classifying them into distinct clusters that group evolutionarily divergent families based on shared structural features. We further extend our model to classify unknown or variants of known GT folds. By identifying families that are likely to adopt novel folds, our studies expand the GT fold landscape and prioritize targets for future structural studies.

Poster #2

Anytime Learning of Sum-Product Networks

Swaraj Pawar[1] and Prashant Doshi[1] (Presented by: Hannah Tawashy[2])
[1]School of Computing, [2]Institute for Artificial Intelligence

Prominent algorithms for learning sum-product networks (SPN) focus on learning models from data that deliver good modeling performance without regard to the size of the learned network. Consequently, the learned networks can get very large, which negatively impacts inference time. We introduce anytime algorithms for learning SPNs. These algorithms generate intermediate but provably valid models whose performance progressively improves as more time and computational resources are allocated to the learning. They flexibly trade off good model performance with reduced learning time, offering the benefit that SPNs of small sizes (but with reduced likelihoods) can be learned quickly. We comprehensively evaluate the anytime algorithms on two testbeds and demonstrate that the network performance improves with time and reflects the expected performance profile of an anytime algorithm. We expect these anytime algorithms to become the default learning techniques for SPNs given their clear benefit over classical batch learning techniques.

Poster #3

Applying AI in STEM Education

Xiaoming Zhai, Holly Amerman, Omodele Fakayode, Xinyu He, Sahrish Panjwani, Ehsan Latif
AI for STEM Education Research Lab

AI4STEM Education Research Lab employs cutting-edge technologies to tackle challenges in STEM education. Quality STEM education is needed to address a number of issues in industry and society in the coming decades. However, teachers face challenges due to a lack of resources (e.g., time, technical knowledge, physical resources) and the complexity of instruction. Moreover, a deep and nuanced understanding of student knowledge is needed in order to address learning and preparation for the future. Advanced technologies such as machine learning have advantages in capturing essential information for teachers in real-time and supporting teachers in teaching practices. By using AI and machine learning in STEM education, our research has found (1) high levels of human-machine agreement of scoring of student-drawn models; (2) cognitive differences for middle school students in argumentation practices; (3) student computational thinking patterns when engaging in maker activities; and (4) teachers' instructional decision processes when using Automatic Reports (AutoRs).

Poster #4

VRDEO: Closing the Learning Loop With Extended Reality

Brook Bowers, Kyle Johnsen
School of Computing

Evaluating learning outcomes today typically relies on evaluating static artifacts from which learning is inferred. However, the static nature of the artifacts hides the information of the learning process itself. Understanding the learning process relies on capturing the dynamic elements of learning. This presents immediate problems: Where does learning occur? Not only in the classroom, but in study halls and private dwellings. To alleviate these pains, we built VRDEO, a multi-platform, mobile, networked, general research tool that enables simultaneous viewing of video content and generation of written artifacts across multiple devices. VRDEO enables us to flexibly deploy hardware, either headsets or tablets, into laboratories, classrooms, or to be sent with students in order to more fully capture the learning process.

Poster #5

Efficient Algorithms for Learning to Control Bandits with Unobserved Contexts

Hongju Park, Mohamad Kazem Shirani Faradonbeh,
Department of Statistics, University of Georgia

Contextual bandits are canonical models for sequential decision-making under uncertainty in environments with time-varying components. In this setting, the expected reward of each bandit arm consists of the inner product of an unknown parameter with the context vector of that arm. The classical bandit settings heavily rely on assuming that the contexts are fully observed, while the study of the richer model of imperfectly observed contextual bandits is immature. This work considers Greedy reinforcement learning policies that take action as if the current estimates of the parameter and of the unobserved contexts coincide with the corresponding true values. We establish that the non-asymptotic worst-case regret grows poly-logarithmically with the time horizon and the failure probability, while it scales linearly with the number of arms. Numerical analysis showcasing the above efficiency of Greedy policies is also provided.

Poster #6

Application of Deep Learning to the Classification of Palatalized [t] in UK English

Austin Brailey-Jones, Sara E. Miller, Margaret Renwick
Department of Linguistics, University of Georgia, USA

Feed forward neural networks can classify acoustic data and can serve as a possible alternative to forced alignment which can overlook phonetic variation. Here we explore the application of deep learning in the classification of phonetic variation. Model training data was taken from the Audio British National Corpus and annotated to produce gold standard tokens. MFCC features were used to represent palatalization of [t] to [t͡ʃ] in naturalistic speech across word boundaries preceding /ju/, "you". The model performs at 90% accuracy compared to gold judgments. In testing with novel data, the model achieves 70% - 100% success.

Poster #7

An anchor-free CenterTrack network for cotton seedling and flower detection and counting

Chenjiao Tan (Chenjiao.Tan@uga.edu), Changying Li (cyli@uga.edu), Dongjian He, Huaibo Song
Bio-Sensing and Instrumentation Lab, College of Engineering, Institute of AI, University of Georgia

Accurately counting the number of plants and plant organs in natural environments is essential for breeders and growers. Cotton plant count, especially in the seedling stage, can help breeders to select genotypes with a high emergence rate and growers to provide information on the necessity of replanting. Cotton plant organ count, especially the flower count, can provide the yields in advance. In order to avoid repeated counting in consecutive frames of videos collected in natural environments, CenterTrack, a deep convolutional neural network-based tracking method, was investigated for cotton seedling and flower counting. The network is extended from a customized CenterNet, which is an anchor-free object detector. CenterTrack predicts the detections of the current frame and displacements of detections between the previous frame and the current frame, which are used to associate the same object in consecutive frames. The modified CenterNet detector achieved high accuracy on both seedling and flower datasets. Experimental results showed that seedling and flower counts with optimized hyperparameters highly correlated with those of manual counts (=0.98 and =0.95). The anchor-free deep convolution neural network-based tracking method provides an automatic tracking and counting in video frames, which will significantly benefit plant breeding and crop management.

Poster #8

Sharing Autonomy of Exploration and Exploitation via Control Interface

Aiman Munir, Ramviyas Parasuraman
Heterogeneous Robotics (HeRo) Research Lab, School of Computing

Shared autonomy is a control paradigm that refers to the adaptation of a robot's autonomy level in dynamic environments while taking human intentions and status into account at the same time. Here, the autonomy level can be changed based on internal/external information and human input. However, there are no clear guidelines and studies that help understand "when" should a robot adapt its autonomy level to different functionalities. Therefore, in this paper, we create a framework that helps to improve the human-robot control interface by allowing humans to adapt to the robots' autonomy level as well as to create a study design to gather insights into human's preference to switch autonomy levels based on the current situation. We create two high-level strategies - Exploration to gather more data and Exploitation to make use of current data - for a search and rescue task. These two strategies can be achieved with human inputs or autonomous algorithms. We intend to understand the human preferences to the autonomy levels (and "when" they want to switch) to these two strategies. The analysis is expected to provide insights into designing shared

autonomy schemes and algorithms to consider human preferences in adaptively using autonomy levels of certain high-level strategies.

Poster #9

Scheduling Multi-Tenant AI applications on Heterogeneous Edge Devices

Ting Jiang, Jacob Stein, Jianwei Hao, Lakshmish Ramaswamy, In Kee Kim
School of Computing

We observed significant differences in DNN inference throughput when the models were run on different resource types. For example, for Mobile Net-V1/V2, EdgeTPUs produced the highest throughput compared to EdgeGPUs or CPUs. However, when processing large models like Inception-V3, the inference throughput on EdgeTPUs could be lower than on EdgeGPUs. Moreover, DNN inference throughput can also be changed by the choice of DL frameworks. That shows the difficulty in managing the performance of AI/DNN executions on heterogeneous edge devices. To address the resource heterogeneity on edge devices, we try to develop novel scheduling techniques using the performance statistics and profiling results to meet the performance goal, such as accuracy, latency, throughput, and power consumption.

Poster #10

Efficient Power Management for In-situ, Energy-harvesting Platform

Jianwei Hao, Emmanuel Oni, Connor Caddell, In Kee Kim, Lakshmish Ramaswamy
School of Computing, University of Georgia, Athens, Georgia

AWeSOMSense (A Wetland Soil Organic Matter Sensor) is a project to transform salt marsh SOM sensing with high-quality SOM data at a geographically dense spatial pattern. This work is for developing an in-situ sensing platform for AWeSOMSense. Energy harvesting can be leveraged to store abundant energy from solar irradiance for future use. However, the limited capacity of the battery is an issue. It may be insufficient to support sensor operations for a long period of cloudy or overcast weather. To overcome the above limitation, it is important to dynamically adapt and optimize the sensing operations by considering the available energy budget and future energy gain to maintain the sensing data quality and support a long operation time of the sensors.

Poster #11

RGB-D Based Instantaneous Evaluation of Safety in Human-Robot Interaction Settings

Pranav Pandey, Ramviyas N. Parasuraman, Prashant Doshi
HeRo Lab and THINC Lab, School of Computing

Due to extensive studies and technological advancements, including the International Standard ISO 10218 and ISO/TS 15066:2016 for collaborative robots are easy and the safety of humans is guaranteed. In this paper, our objective is to provide real-time safety levels for humans in an overlapping workspace for collaborative tasks from an external agent's perspective. The external agent's perspective is essential for any Human-Robot collaboration setting because of the fact that it is not necessary that the human is always in the viewable range of the camera view especially when we are using vision-based safety systems. We use vision-based distance and velocity measures used in our quantitative framework to give the safety level for a specific interaction type.

Poster #12

Decision-making in Open Multi-Agent Systems

Gayathri Anil
Institute for Artificial Intelligence

In the real world, decision-making in multiagent systems tends to be open as the composition of the system can change over time due to several endogenous and exogenous factors. To deal with Openness in multi-agent environments, agents are required to not only reason about other agents to come up with the best response, but also reason the openness that their system can be subjected to. With this work, we try to establish an approach to tackle Agent and Task Openness in Multi-agent Systems using Planning and Learning based approaches.

Poster #13

Enhancing Sustainable Weed Management in Turf Grass Through Deep Learning and Proximal Sensing

Chintan B. Maniyar[1,3], Dhiraj Srivastava[2,3], Sathish Samiappan3 , Joby Czarnecki[3] , Amy Wilber[4]
[1]Department of Geography, University of Georgia, [2]School of Plant and Environmental Sciences, Virginia Tech, [3]Geosystems Research Institute, Mississippi State University, [4]Plant and Soil Sciences, Mississippi State University

Weeds are a serious issue in turf in the United States. Real-time detection of weeds is crucial for developing a strategy for effective weed management. Spot spraying in turf is one of the most common ways for dealing with weed growth, but proves to be manually extensive and time consuming. This study attempts to automate the process of detecting weeds and optimize consequent preventive measures. A custom dataset of consisting of 4 different invasive weed species is curated using more than 6000 proximal pictures from a greenhouse and turf fields. Deep neural networks pretrained for classification and object detection based on ResNet-50 and YOLO are assessed for weed detection in turf grass across diverse weather conditions. Results show that both ResNet-50 and YOLO-based models achieve more than 95% accuracy in detecting the weeds in the turf. Inference speeds indicate that YOLO-based models are faster than the ResNet-50 models (YOLO: 6.75ms; ResNet-50: 1257.5 ms) and hence are more suitable to be deployed on a ground control vehicle for near-real time mapping of weeds in a turf.

Poster #14

Cyberinfrastructure to Monitor Soil Organic Matter in Salt Marshes of the Georgia Coast

Rajneesh Sharma[1], Jianwei Hao[2], Deepak Mishra[1], In Kee Kim[2], S. Sonny Kim[3], Lakshmish Ramaswamy[2], and Lori Sutter[4]
[1]Department of Geography, University of Georgia, Athens, Georgia; [2]School of Computing, University of Georgia, Athens, Georgia; [3]College of Engineering, University of Georgia, Athens, Georgia; [4]Biology & Marine Biology, University of North Carolina Wilmington

Wetlands serve serval ecological services such as nutrient absorption, floodwater reduction, storm attenuation, and essential fish industries. Moreover, they are great sequesters of carbon and are widely recognized for their terrestrial carbon stock. Wetlands cover 4% - 6% of the globe and store almost 20% - 25% of global soil organic matter (SOM). However, studies suggest worrying losses of wetlands in North America, 53% in the United States, 16% in Canada, and 62% in Mexico. Once drained it takes almost 500 years to stock the same amounts of SOM back before reaching equilibrium after relocation. With increasing human activities and increasing greenhouse gases it's essential to monitor SOM in wetlands. However, mapping SOM in remote areas is expensive and time-consuming. For an economical and time-effective way to monitor SOM in wetlands, we showcase a way to remotely sense SOM in coastal wetlands. We are utilizing soil sensors to collect essential soil information that is used to predict SOM in coastal wetlands. We are developing data collection infrastructure, the soil sensors such as pH, salinity, soil temperature, electrical conductivity,

redox potential, and soil color mounted over a platform in coastal wetlands. The data collection infrastructure would make data cost and time effective. This would make timely data viable compared to conventional soil sample collection. Proper monitoring of SOM would conserve soil carbon and prevent it to escape into the atmosphere and reduce the concentration of greenhouse gases in the air.

Poster #15

Long-distance linguistic dependencies in Chinese and English brains

Donald Dunagan[1], Maximin Coavoux[2], Shulin Zhang[1], Shohini Bhattasali[3], Jixing Li[4], Jonathan Brennan[5], John Hale[1]
University of Georgia[1], Université Grenoble Alpes[2], University of Toronto Scarborough[3], City University of Hong Kong[4], University of Michigan[5]

Words can occur arbitrarily far away from where they contribute their meaning in a sentence. Two examples are WH-questions (WHQs), which begin with a WH-word like 'what' and object-extracted relative clauses (ORCs), in which a noun is modified by a sentence-like grammatical unit. While these long-distance dependencies have been extensively studied, never before have their brain bases been examined from a multi-lingual, naturalistic perspective. This study fills this gap by analyzing WHQs and ORCs in fMRI data collected while 35 Chinese participants (15 females) and 49 English participants (30 females) listen to translation-equivalent stories. These languages exhibit radical typological differences in word order in these constructions. It remains unknown whether the brain basis for comprehension in these languages is similar or different. Separate general linear model analyses were performed and voxel-level intersections were calculated between the results to identify common regions of selectively increased activation during the comprehension of these linguistic constructions. Further Bayesian region of interest analyses probed whether common increases were truly similar. We found remarkable cross-linguistic commonality for both constructions. WHQs were associated with increased activation in the left middle and superior temporal lobe, left temporoparietal junction, left inferior frontal gyrus, and bilateral medial frontal lobe. ORCs were associated with increased activation in the left middle temporal lobe, left inferior frontal gyrus, bilateral angular gyrus, bilateral posterior cingulate, bilateral precuneus, and left medial frontal lobe. These results support the hypothesis that, regardless of form, the brain bases of higher-level language processing are uniform across languages.

Poster #16

Aerial Imagery and Deep Learning to Characterize Season-long Flowering Response in Cotton

Daniel Petti, Jeevan Adhikari, Changying “Charlie” Li, Andrew H Paterson
Bio-Sensing and Instrumentation Lab, University of Georgia

Many perennial plants make important contributions to agroeconomies and agroecosystems, but often have complex architecture and/or long duration of flowering that hinders measurement and selection. Iteratively tracking productivity over a long flowering/fruiting season may permit the identification of genetic factors conferring different reproductive strategies that might be successful in different environments. In cotton, we apply aerial imagery and deep learning methods to novel genetic stocks, identifying genetic factors influencing the duration and rate of fruiting. With transfer learning of the deep learning models, this approach could be applied to other crops as well.

Poster #17

Multi-robot Collaboration for Localization and Map Exploration

Ehsan Latif, Ramviyas Parasurman
Heterogeneous Robotics (HeRo) Research Lab, School of Computing

Simple robots can collaborate to carry out sophisticated and more competent tasks in multi-robotic systems (MRS). Some of the fundamental processes of an MRS include localization and map exploration. To obtain high accuracy for relative localization and increased mapping efficiency during an exploration activity, our research focuses on

overcoming established hurdles, such as those related to communication and processing. Numerous attempts have been made at model-based multi-robot localization for relative localization. Still, few have articulated the multi-robot collaborative localization as a graph problem that may be resolved using graph optimization techniques. We propose a Distributed Graph Optimization approach for Relative Localization termed DGORL. Three main parts comprise the DGORL we provide for solving this issue: constructing (connectivity) graphs, expanding via the transition model, and optimizing relative poses. Regarding multi-robot autonomous exploration strategies, the robots must preserve an internal map of their surroundings, which they would use localization and planning methods to steer.

On the other hand, learning-based techniques improve exploratory abilities as the robot engages with its surroundings. However, due to their substantial sample complexity, they take longer to converge to the optimal solution, have effective communication and update costs, and are consequently difficult to deploy in real-world scenarios. To satisfy the demand for quick convergence, we propose in this study a novel cooperative Q-learning-based method called CQLite. CQLite implements a coverage-weighted reward function while taking into account minimal communication costs. We theoretically evaluate our proposed approaches to ensure practicality and better performance for collaborative localization and map exploration.

Additionally, We experimentally confirm our methodology using CQLite to simulate several robots exploring an indoor map. DGORL provided 23% higher localization accuracy than Terrain Relative Localization. Furthermore, with more than a twofold reduction in computation and communication, our proposed CQLite significantly outperformed cutting-edge multi-robot exploration techniques (like the Rapidly Exploring Random Trees (RRT) and the Deep Reinforcement Learning (DRL) based approaches) in terms of mapping performance (coverage and distance traveled).

Poster #18

The Emerging Roles and Characteristics of AI to Foster Teacher and AI Collaboration from the Pedagogy Perspective: A Systematic Review

Lehong Shi & Ikseon Choi
Learning, Design, and Technology

As an emerging technology, artificial intelligence (AI) has become among the most promising technologies to significantly change teaching, learning, and assessment and open a new educational horizon. However, AI advancement raises critical questions about AI potential, such as (1) how the emergence of AI technology differs from traditional technologies in assisting teaching and improving learning. (2) will AI integration in education change teachers' professional roles or take over teachers dramatically? and (3) what would be the emerging relationships between teachers and AI? To explore those questions, we first investigated AI's pedagogical roles and characteristics by systematically reviewing 40 peer-reviewed articles about AI integration in education published from 2010 to 2022. We extracted 19 critical functions of AI and classified them into six pedagogical roles of AI, including (a) AI instructor, (b) learning partner and scaffolder, (c) teacher supporter, (d) learning monitor and analyst, (e) affection carer, and (f) communicator and interaction facilitator. We also discovered ten characteristics of AI, from which we built a four-characteristic AI profile model (i.e., adaptivity, AI-human interactivity, autonomy, and automaticity). Based on the discovered AI pedagogical roles and characteristics, we then recommended teachers to consider three perspectives when they construct their professional roles, including (1) AI is a new player in classrooms; (2) AI forms a new pedagogical partnership with teachers; and (3) teachers should have shared educational paradigms.

Poster #19

Is AI Bad for Aesthetic Life?

Sam Bennett
Department of Philosophy

Art and aesthetic objects (i.e., objects that possess an aesthetic quality, like beauty or elegance) are an important dimension of human life. After the so-called "fourth industrial revolution," AI is now "threaded into large tracts of everyday life" (Elliot 2022: 4), including our aesthetic activities. It is therefore worthwhile to ask whether the inclusion of AI is likely to help or impede flourishing in aesthetic life. I approach this question by distinguishing three kinds of aesthetic activities. First is production (where an individual produces art, or changes a non-art object such that it has an aesthetic, such as re-arranging a bedroom to make it more beautiful); second is reception (where an individual responds, for example by appreciating, either to art or an object with aesthetic features); and third is mediation (where an individual does something that enables activities of either production or reception, such as the activity of choosing to read a certain novel). Using these distinctions, my research responds to three arguments within the literature, each of which claims that AI impedes aesthetic flourishing in one of the three domains of activity. The first argument is Kelly's (2019) historicity argument against AI in aesthetic production; the second is Melchionne's (2017) autonomy argument against AI in aesthetic mediation; and the third is Turk's (2021) aesthetic filter bubble argument against AI in aesthetic reception. While agreeing with and bolstering the arguments of Kelly (2019) and Melchionne (2017), I disagree with Turk's (2021) argument, which fails to appreciate the common operational goals of recommender systems.

Poster #20

Optimal Plant Part Segmentation using 3D Neural Architecture Search

Farah Saeed[1,3], Changying "Charlie" Li[1,2,3], Tianming Liu[1,3]
[1]School of Computing, [2]College of Engineering, University of Georgia[3]

The automatic, and accurate plant phenotyping plays important role to improve the crop yield through enabling efficient plant analysis and plant breeding studies. The 3d deep learning allows automatic segmentation of plant parts from point cloud data. However, the network architecture is designed manually and performance is limited to prior experience. The aim of this study is to search for optimal 3d deep networks to perform the plant part segmentation. We perform the 3d neural architecture search by training a super network composed of candidate networks. Using the trained super network, the evolutionary searching is used to search for top performing architecture. The results demonstrate the searched architecture outperforms manually designed architectures by attaining mean IoU and accuracy of more than 90% and 96%, respectively. The searched architecture achieves more than 83% class-wise IoU for all main stem, branches, and boll class. This plant part segmentation method shows promising results and holds potential to be utilized by plant breeders for enhancing the production quality.

Poster #21

Baseline Models for Representing Mitochondrial Dynamics

Rachel Mattson§, Delaney Ott‡, Victoria Schneider||, Meekail Zain‡, Blessing Ojeme, Marcus Hill‡, Neelima Pulagam‡, Mojtaba Fazli‡, Andrew Durden‡, Russell Karls¶, Frederick D Quinn¶, Shannon Quinn‡***
‡ Department of Computer Science, University of Georgia | § Institute of Artificial Intelligence, University of Georgia | || College of Engineering, University of Georgia | ¶ Department of Infectious Diseases, University of Georgia | ***Department of Cellular Biology, University of Georgia

The physical characteristics of mitochondria, such as size, shape and quantity, are reflected in their spatial distributions within cells. Certain diseases, such as tuberculosis, are known to correspond with drastic structural changes in mitochondrial protein distribution. Such shifts are observable under fluorescent microscopy but their identification requires the meticulous efforts of trained clinicians. The Organelle Networks (OrNet) project seeks to automatically characterize the dynamics of mitochondria and organelles with similar spatial patterns. The OrNet model and its associated task are unique, with no baseline models to compare against. Thus, determining the quality of such a representation directly is difficult.

To evaluate the utility of OrNet's graph-based approach, we build baseline models which do not leverage such structure. We first attempt to evaluate the representative power of local protein interactions by applying a CNN on a fixed region of interest (ROI) taken across the first and last frames. Secondly, we attempt to capture global information by applying a similar CNN to the entirety of the two frames. We compare these alternative representations against OrNet's representation space using simple classification models. The results of these baselines will clarify the importance of local, global, and temporal features in quantifying protein dynamics and help motivate the direction of future OrNet developments.

Poster #22

The Limitations and Opportunities of Fine-tuned Language Models on Oral, Abstract Datasets

Rex Vanhorn
Institute for Artificial Intelligence

Language models have significantly improved over the last few years, to the point where they can generate unique, logical responses to all kinds of questions. Often, with fine-tuning, these models can also generate factual responses to very specific questions, in a desired format or style. Little research, however, has been conducted on the performance and abilities of model models when tuned on spoken data, in highly nuanced and abstract domains. This research seeks to determine the OpenAI's GPT-3 and AI2's MACAW models abilities and limitations to answering questions in an abstract domain, specifically to determine if answers can be influenced through fine-tuning, and the general relationship between the amount of data provided, and its influence on the language model's output.

Poster #23

Perceptions of Artificial Intelligence Systems

Brandon Y. Kang & Neal Outland, Ph.D
Department of Psychology

Artificial intelligence is rapidly reshaping the way work happens. Increasingly AI is conceptualized as autonomous, representing an independent employee or a teammate in cases of collaboration. The rapid expansion of AI faces barriers in utilization, predominantly in terms of the trust. Currently, models of trust assume isomorphism of trust models from human subjects. This study explorers directly if those models are applicable. Authors present evidence that human models of trust are applicable in conceptualizing and measuring trust in AI.

Poster #24

A Qualitative Investigation of Fear and Acceptance of AI at Work

Riley A. Hess, Sierra R. Stryker, Brandon Kang, Jared Richardson, MA, & Neal Outland, Ph.D
Department of Psychology

While organizations rush to implement artificial intelligence (AI)-enabled technologies to support business practices, many workers fear that AI will take over part of or the entirety of their job. This fear of "technological unemployment" induces anxiety in workers and is a hindrance to successful integration of AI into the workplace. The present study takes a data-driven approach to describe why workers do or do not experience fear of technological unemployment. Results of this work will help organizations identify barriers to user trust and successful integration of AI.

Poster #25

Integrating AI-Enabled Hardware Into an Aerial Data Collection Platform

Bryan Ponce, Beni Rodriguez, Jonathan Porter, Malachi Brewer & Sergio Bernardes
Disruptive Geospatial Technologies Laboratory/Center for Geospatial Research (CGR), Department of Geography, The University of Georgia (bryan.ponce@uga.edu; sbernard@uga.edu, digtl.uga.edu)

Poster #26

Integration of an AI-Enabled Drone System: Software and Data Processing Considerations

Porter Squires[1], Malachi Brewer[1], Bryan Ponce[1], Beni Rodriguez[1], Holly Presley-Wright[2], Ali Missaoui[2] & Sergio Bernardes[1]
[1]Disruptive Geospatial Technologies Laboratory/Center for Geospatial Research (CGR), Department of Geography, The University of Georgia; [2]Institute of Plant Breeding and Genomics, The University of Georgia (jonathan.squires@uga.edu; sbernard@uga.edu, http://digtl.uga.edu)

About the AI Institute

In 1984, interested researchers at the University of Georgia formed the *Artificial Intelligence Research Group*. The Group was officially established as the A*rtificial Intelligence Center* in 1995, and it became the *Institute for Artificial Intelligence* in 2008. It is part of the Franklin College of Arts and Sciences.

The mission of the Institute is to encourage and support interdisciplinary research in Artificial Intelligence involving University faculty and staff, and to facilitate the integration of Artificial Intelligence technology by industry and government agencies in the State of Georgia and the Nation.

The Institute is both a research and an instructional unit. It has for many years administered both the AB in Cognitive Science and the MS in Artificial Intelligence degrees, and in 2022 a PhD in Artificial Intelligence was approved by the Board of Regents of the University System of Georgia. The Institute expects its first PhD students to begin Fall 2023. It is among the first such degree programs in the country.

Participation in the AI Institute by UGA faculty is voluntary. There are currently 50 associated AI Faculty Fellows, taken from many departments across multiple colleges and schools at UGA. Some of the units involved include: Computer Science, Philosophy, Linguistics, Psychology, Public Administration & Policy, Political Science, Educational Psychology, Statistics, Theatre & Film Studies, Marine Sciences, Electrical and Computer Engineering, Business Administration, Advertising & Public Relations, English, Geography, Management Information Systems, Epidemiology & Biostatistics, Mathematics and Science Education, Health Promotion & Behavior, Epidemiology & Biostatistics.

More information on the AI Institute can be found at *https://ai.uga.edu*.

AI Faculty Fellows

Jason Anastasopoulos
Drew Abney
Anna Abraham
Budak Arpinar
Yuri Balashov
Pete Bettinger
Suchendra Bhandarkar
Ikseon Choi
Chris Cieszewski
Jeremy Davis
Prashant Doshi
Mark Ebell
Mohamad Kazem Shirani Faradonbeh
Jennifer Gay
John Gibbs
Adam Goodie
John Hale
Bill Hollingsworth
Brian Hopkinson
Pengsheng Ji
Kyle Johnsen
Elena Karahanna
Yuan Ke
In Kee Kim
Jooyoung Kim
Bill Kretzschmar
Jaewoo Lee
Changying Li
Ninghao Liu
Tianming Liu
Guoyu Lu
Ping Ma
Gengchen Mai
Frederick Maier
Aaron Meskin
John Miller
Deepak Mishra
Neal Outland
Ramviyas Parasuraman
Roberto Perdisci
Shannon Quinn
Lakshmish Ramaswamy
Khaled Rasheed
Margaret Renwick
Carolina Salge
Juliet Sekandi
Kimberly Van Orman
Javad Mohammadpour Velni
Sarah Wright
Xiaoming Zhai

AI Research Day: Schedule

2:00 PM	Welcome
2:15 PM	Keynote
3:15 PM	Lightning Talks (Session I)
3:45 PM	Break
4:00 PM	Panel Discussion
5:00 PM	Lightning Talks (Session II)
5:40 PM	Poster Session

Institute for Artificial Intelligence
Boyd GSRC, Room 516
University of Georgia
Athens, GA 30602-7415

Phone: (706) 542-0358
Web: https://ai.uga.edu