Program Book

National Aeronautics and
Space Administration

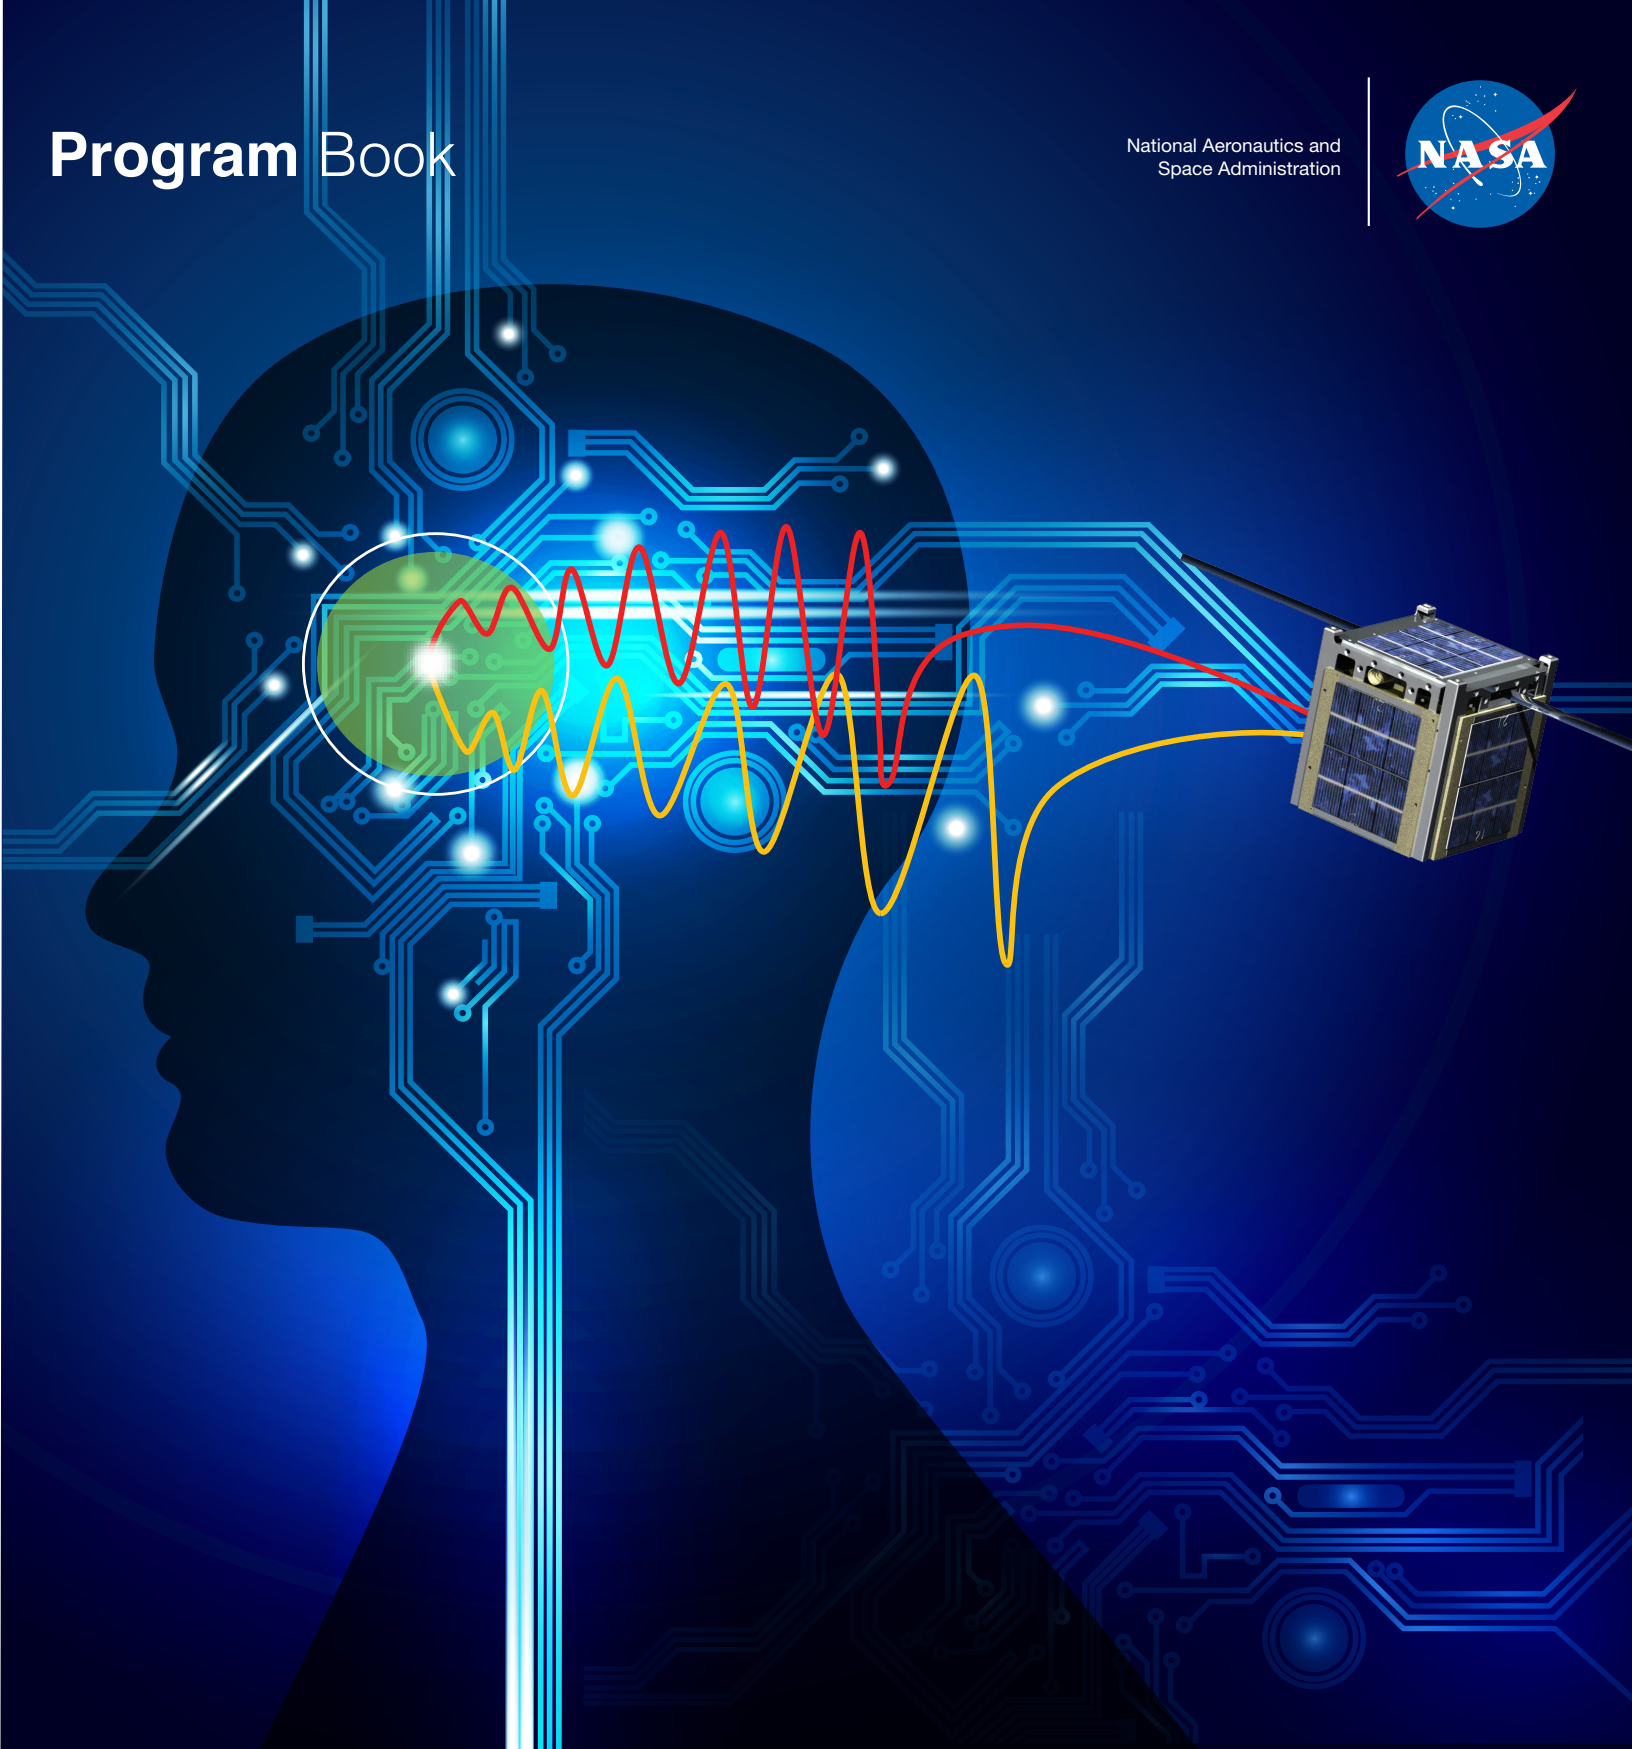

NASA Goddard Workshop on
Artificial Intelligence

Goddard Space Flight Center | November 27-29, 2018

www.nasa.gov

NASA Goddard Artificial Intelligence Workshop

Artificial Intelligence (AI) is a collection of advanced technologies that allows machines to think and act, both humanly and rationally, through sensing, comprehending, acting and learning. AI's foundations lie at the intersection of several traditional fields – Philosophy, Mathematics, Economics, Neuroscience, Psychology and Computer Science. Although the inception of AI started in the 1950's, it has recently made a strong comeback in all aspects of society and all over the world; this is mainly due to the timely combination of increased data volumes, advanced and mature algorithms, and improvements in computing power and storage. Current AI applications include big data analytics, robotics, intelligent sensing, assisted decision making, and speech recognition just to name a few.

As stated in the latest NSF Statement on AI for American Industry, "The effects of AI will be profound. To stay competitive, all companies will, to some extent, have to become AI companies." Compared to Industry and Academia, NASA and Goddard have specific challenges as well as resources that are particularly adapted to the use of AI.

Contents

Agenda 3
Keynote Speakers 7
Invited Speakers.....11
Panel Members 16
Short Talks.....27
Posters 48

Workshop **Agenda**

DAY 1 TUESDAY, NOVEMBER 27, 2018

7:30 am – 9:00 am **ARRIVAL**
Goddard Badging for non-NASA Attendees at the Visiting Center
Workshop Check-in for All Outside of Building 8 Auditorium

9:00 am – 10:30 am **OPENING CEREMONY**
Jacqueline Le Moigne, Chair – General Introduction to the Workshop
Christopher Scolese, Director, NASA Goddard – Welcome and Speaker Introduction

KEYNOTE: David Gunning, DARPA
DARPA's Explainable Artificial Intelligence (XAI) Program

10:30 am – 10:45 am **BREAK**

10:45 am – 12:25 pm **PANEL: AI at NASA**
(Session Chair: Jacqueline Le Moigne)
Dan Crichton/JPL – Data Science at JPL: Integrating Data Analytics into the Full Data Lifecycle
James Ecker/LaRC – Deep Learning for Neuro-visualization and Continuous Control in Autonomous Systems
Nikunj Oza/ARC – Artificial Intelligence in the NASA Ames Intelligent Systems Division
Brian Roberts/GSFC – AI & Computer Vision for Satellite Servicing at NASA Goddard
Brian Thomas/HQ – Elements of an AI/ML Architecture for NASA

12:25 pm – 1:30 pm **LUNCH**

1:30 pm – 2:30 pm **KEYNOTE: Kirk Borne**, Booz Allen (Introduction: Barbara Thompson)
AI at NASA: From Data to Insights to Actionable Intelligence

2:30 pm – 4:40 pm **SESSION 1**
Invited Talk: **Bart Paulhamus,** APL (Introduction: Ron Zellar)
Intelligent Systems Research at JHU/APL

SHORT TALKS: *(Session Chairs: Ioana Rus and Dave Batchelor)*
3:00 – D. Sekora, AI's Missing Real-World Connection, and Its Essential and Multifaceted Roles
3:10 – J. Nanda, Explainable Machine Learning for Aviation Safety Assurance
3:20 – A. Deane, A Cognitive Processing Enhanced Smart Interface Framework For Situational Awareness
3:40 – W.R. Huang, Data Poisoning Attacks Can Compromise Machine Learning Systems

Break: 3:50 pm to 4:00 pm

Workshop **Agenda**

(Session Chairs: Alinda Mashiku and Manohar Deshpande)
4:00 – B. Dean, Deep Multi-Layer Networks for Optical Wavefront Sensing and Control
4:10 – S.R. Alimo, Machine Learning Approaches for General Satellite Maneuvers
4:20 – A. Mashiku, Supervised-machine Learning for Intelligent Collision Avoidance Decision-making and Sensor Tasking
4:30 – J. Krishnan, SEVA-OIE: Open Information Extractor for the Systems Engineering Virtual Assistant (SEVA)

4:40 pm – 5:45 pm *Introduction Breakout Sessions: Burcu Kosar and Jacqueline Le Moigne*
BREAKOUT SESSION: AI for NASA Science Applications

DAY 2 WEDNESDAY, NOVEMBER 28, 2018

8:30 am – 9:40 am *Introduction:* ***Christyl Johnson,*** *NASA Goddard*
KEYNOTE: William Buzz Roberts, NGA
Real World Artificial Intelligence, Automation and Augmentation – Geospatial Intelligence Successes, Challenges and Way Forward

9:40 am – 11:00 am **PANEL: AI in Academia**
(Session Chair: Grey Nearing)
Cynthia Matuszek/UMBC – Learning Grounded Language For and From Interaction
Ray Ptucha/RIT – Deep Learning on Graph Data
Dinesh Manocha/UMD – Autonomy and AI Research at UMD

11:00 am – 11:30 am **BREAK**

11:15 am – 12:15 pm **KEYNOTE: Henry Kautz,** NSF
(Introduction: Jacqueline Le Moigne)
Artificial Intelligence: Everything Old is New Again

12:15 pm – 12:45 pm **BREAK – Grab Lunch**

12:45 pm – 1:15 pm **Brown Bag Lunch with Lika Guhathakurta**, NASA ARC *(Introduction: Michael Kirk)*
The Frontier Development Lab (FDL): Applied Artificial Intelligence for Science and Exploration

1:15 pm – 1:30 pm **BREAK**

Workshop **Agenda**

1:30 pm – 2:40 pm

SESSION 2

Invited Talk: **Tom Goldstein,** UMD *(Introduction: Nargess Memarsdeghi)*

Multi-Scale Neural Networks for Image Processing

SHORT TALKS: *(Session Chairs: Barbara Thompson and Ryan McGranaghan)*

2:00 – Z.liu, Improving NASA Earth Science Data and Information Access Through Natural Language Processing Based Data Analysis and Visualization

2:10 – M. Reiss, Improvements On Coronal Hole Detection Using Supervised Classification

2:20 – K. Tran, X-Net: Bimodal Feature Representation Learning in Satellite Imagery

2:30 – S. Sabogal, Hybrid Semantic Image Segmentation using Deep Learning for On-board Space Processing

2:40 pm – 2:50 pm **BREAK**

2:50 pm – 4:20 pm

SESSION 3

Invited Talk: **Victor Pankratius,** MIT *(Introduction: Sujay Kumar)*

Towards Deriving Theories from Data: Frontiers for Model Inference in Astro-&Geophysics

SHORT TALKS: *(Session Chairs: Craig Pelissier and Troy Ames)*

3:20 – C. Keller, Atmospheric Chemistry Modeling using Machine Learning

3:30 – J. Kouatchou, Implementation of Gaussian Processes in an Hydrological Model

3:40 – D. Josyula, Autonomous Seasonality Adaptation

3:50 – N. Thomas, Machine Learning in Global Scale Classification of Mangrove Forests From remotely sensed imagery

4:00 – T. Maxwell, Machine Learning in the Earth Data Analytic Services (EDAS) Framework

4:10 – M. Halem, RNN/LSTM Ensemble Data Assimilation for the Lorenz Chaotic Models

4:20 pm – 5:10 pm **BREAKOUT SESSION: AI for NASA Engineering Applications**

5:10pm – 5:30 pm **BREAK** (and Poster Setup)

5:30 pm – 6:30 pm **POSTER SESSION**

DAY 3 THURSDAY, NOVEMBER 29, 2018

8:30 am – 9:40 am

Introduction: ***Peter Hughes,*** *NASA Goddard*

KEYNOTE: Vikash Mansinghka, MIT

Probabilistic Programming and Artificial Intelligence

Workshop **Agenda**

9:40 am – 11:25 am **PANEL: AI in Industry** *(Session Chair: Ron Zellar)*
John Hebeler/Lockheed Martin –Determining Normal (and Abnormal) using Deep Learning
Graham Katz/IBM – Watson Intelligent Advisors: Discovery and Conversational Technology for Now and the Future
Jon Neff/Aerospace – Overview of Aerospace Corporation AI Initiatives
Susie Adams/Microsoft – TBD
Larry Brown/NVIDIA – GPU Accelerated High Performance Data Analytics for Federal Applications

11:25 am – 11:40 am **BREAK**

11:40 am – 12:30 pm **SESSION 4**
BREAKOUT SESSION: AI for Intelligent Mission Autonomy

12:30 pm – 1:30 pm **LUNCH**

1:30 pm – 3:10 pm **SESSION 5**
Invited Talk: John Calhoun, Amazon AWS *(Introduction: **Craig Pelissier**)*
Improving Time to Science Using AWS Machine Learning

SHORT TALKS: *(Session Chairs: Nargess Memarsadeghi and Jorge Pinzon)*
2:00 – H. Amiri, Spaced Repetition for Training Artificial Neural Networks
2:10 – T. Yuan, "Application of a Deep U-Net to Automatic Detection of Ship-Tracks Multispectral Images from both Polar-Orbiting and Geostationary Satellites
2:20 – R. McGranaghan, "Ushering in a New Frontier in Geospace Through Data Science
2:30 – R. Attié, Tracking Optical Flows for Better Data Mining on Solar Images
2:40 – D.Hall, Deep Learning Applied to Satellite Data Processing
2:50 – S. Sharma, Data-driven Modeling, Prediction and Predictability: The Complex Systems Framework
3:00 – A. Annex, Automated Stratigraphic Mapping using Convolution Neural Networks

3:10 pm – 3:30 pm **GENERAL DISCUSSION** - CONCLUSIONS and ADJOURN

SPECIAL TUTORIAL *(Organizer: Craig Pelissier)*
Thursday, November 29
4:00 pm – 5:00 pm Python Anaconda Machine Learning Tutorial

Keynote **Speakers**

Kirk Borne

Booz Allen

kirk.borne@gmail.com

AI at NASA: From Data to Insights to Actionable Intelligence

Abstract: This keynote presentation will probe the power of data to inform, predict, and automate next-best actions in a scientific exploration enterprise. Data from sensors are the input to intelligent systems -- in operational environments, streaming data provide the fuel to drive discovery, insights, and decision support through AI. AI is more than artificial intelligence -- it is accelerated, actionable, adaptable, amplified, assisted, and augmented intelligence. Example use cases will be presented for these types of implementations. AI is also applied intelligence, for cases ranging from safety to security to systems to scientific discovery. Methods, relevant opportunities, and some typical algorithms will be presented. Emphasis will be given to machine learning techniques, which are algorithms that learn from experience. In other words, these are algorithms that discover patterns in data and then learn a model that helps to recognize those patterns again or to identify emerging new patterns. Examples will be given from space weather, remote sensing, planetary exploration, and planetary protection. Some of the top trends in AI, machine learning, and data science will be presented also -- first, these will be introduced in a general way, followed by their potential applications in NASA science and engineering projects. Finally, some speculations will be offered on concepts like AI and 4-D printing, AI on a chip, and blockchain for explainable AI (XAI).

Bio: Dr. Kirk Borne is the Principal Data Scientist and an Executive Advisor at global technology and consulting firm Booz Allen Hamilton based in McLean, Virginia. In those roles, he focuses on applications of data science, data management, machine learning, AI (machine intelligence), modeling, and simulation across a wide variety of disciplines. He also provides training and mentoring to multi-disciplinary teams of scientists, modelers, and data scientists. In addition, he consults with numerous external organizations, industries, agencies, and partners in the use of large data repositories and machine learning for discovery, decision support, and innovation. Previously, he was Professor of Astrophysics and Computational Science at George Mason University for 12 years where he did research, taught, and advised students in data science. Prior to that, Kirk spent nearly 20 years supporting data systems activities on NASA space science programs, which included a period as NASA's Data Archive Project Scientist for the Hubble Space Telescope and 10 years as a contract manager in NASA's Space Science Data Operations Office at the Goddard Space Flight Center. Dr. Borne has a B.S. degree Summa Cum Laude in Physics from LSU, and a Ph.D. in Astronomy from Caltech. In 2016 he was elected Fellow of the International Astrostatistics Association for his lifelong contributions to big data research in astronomy. In addition to hundreds of invited talks worldwide, he has also presented keynote presentations at many

dozens of data science and analytics conferences globally. He is an active contributor on social media, where he has been named consistently among the top worldwide influencers in big data and data science since 2013. You can follow him on Twitter at @KirkDBorne.

David Gunning
DARPA
david.gunning@darpa.mil

DARPA's Explainable Artificial Intelligence (XAI) Program

Abstract: The goal of XAI is to create a suite of new or modified machine learning techniques that produce explainable models that, when combined with effective explanation techniques, enable end users to understand, appropriately trust, and effectively manage the emerging generation of AI systems. Dramatic success in machine learning has led to an explosion of new AI capabilities. These systems will offer tremendous benefits, but their inability to explain their actions to human users will limit the effectiveness of these systems. There is an inherent tension between machine learning performance (predictive accuracy) and explainability; often the highest performing methods (e.g., deep learning) are the least explainable, and the most explainable (e.g., decision trees) are less accurate. The program is funding a variety of machine learning techniques to provide future developers with a range of design options covering the performance versus explainability trade space. XAI is focusing these developments on addressing challenge problems in two areas: (1) machine learning problems to classify events of interest in heterogeneous, multimedia data, and (2) machine learning problems to construct decision policies for a simulated autonomous system.

Bio: David Gunning is DARPA program manager in the Information Innovation Office (I2O). Dave has an over 30 years of experience in the development of artificial intelligence (AI) technology. At DARPA, Dave manages the Explainable AI (XAI) and the Communicating with Computers (CwC) programs. This is Dave's 3rd tour as a DARPA PM. Previously, Dave managed the Personalized Assistant that Learns (PAL) project that produced Siri and the Command Post of the Future (CPoF) project that was adopted by the US Army as their Command and Control system for use in Iraq and Afghanistan. In between DARPA tours, Dave was a Program Director for Data Analytics at the Palo Alto Research Center (PARC), a Senior Research Manager at Vulcan Inc., SVP of SET Corp., VP of Cycorp, and a Senior Scientist in the Air Force Research Labs. Dave holds a M.S. in Computer Science from Stanford University, a M.S. in Cognitive Psychology from the University of Dayton, and a B.S. in Psychology from Otterbein College.

Keynote **Speakers**

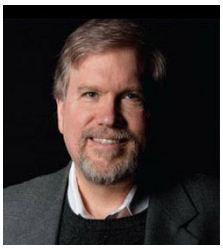

Henry Kautz

NSF, Division Director, Information & Intelligent Systems (CISE/IIS)

hkautz@nsf.gov

Artificial Intelligence: Everything Old is New Again

Abstract: The history of AI is often erroneously summarized as, "AI failed until deep learning came along". In fact, the period between the Dartmouth Workshop of 1956 and the deep learning revolution of 2009 witnessed a steady stream of advances in fundamental models and methods for cognition and commonsense reasoning, in addition to the work on artificial neural networks (ANNs) that culminated in deep learning. Applications of much of the non-ANN oriented research, however, was limited by the inadequacy of methods for grounding symbols in real-world phenomena – that is, perception. Now that deep learning has made machine perception relatively robust, "old" AI research on cognition and commonsense is ripe for revisiting. The next AI revolution will be the creation of AI architectures that synthesize ANNs and symbolic methods.

Bio: Henry Kautz is serving as Division Director for Information & Intelligent Systems (IIS) at the National Science Foundation. He is a Professor in the Department of Computer Science and was the founding director of the Goergen Institute for Data Science and at the University of Rochester. He been a department head at AT&T Bell Labs in Murray Hill, NJ, and a full professor at the University of Washington, Seattle. In 2010, he was elected President of the Association for Advancement of Artificial Intelligence (AAAI), and in 2016 was elected Chair of the American Association for the Advancement of Science (AAAS) Section on Information, Computing, and Communication. His research in artificial intelligence, pervasive computing, and healthcare applications has led him to be honored as a Fellow of the AAAS, Fellow of the AAAI, and Fellow of the Association for Computing Machinery.

William Buzz Roberts

NGA

William.T.Roberts@nga.mil

keith.l.keller@nasa.gov

Real World Artificial Intelligence, Automation, and Augmentation – Geospatial Intelligence Successes, Challenges and Way Forward

Abstract: With the growth in commercially available geospatial information and remote sensing, as well as the democratization of science and technologies to make sense of it, characterizing our changing planet remains a daunting task. The National Geospatial Intelligence Agency's (NGA) Artificial Intelligence,

Keynote **Speakers**

Automation and Augmentation (AAA) strategy, plans and actions are focused on maximizing the potential information and knowledge from this growing corps of data and reducing the data to decision timelines to meet time sensitive missions like humanitarian assistance, disaster response and safety of navigation. Successes, challenges and realities are forming NGA's way forward on its AAA activities. Informing researchers, practitioners and decision makers with real world results, gaps and challenges is critical to NGA's mission success. The speaker will share success stories on how AAA is currently being applied to augment a variety of NGA missions towards geospatial-intelligence as the speed of human understanding.

Bio: Mr. William T. "Buzz" Roberts currently serves as the Analytic Automation POD Lead within the Research Directorate of the National Geospatial-Intelligence Agency (NGA). He's responsible for leading, coordinating, and developing state-of-the-art analysis methods, processes, skills and technologies to advance geospatial intelligence analysis, functional capabilities, operations, personnel training and skills, technologies and architectures of the National System for Geospatial-Intelligence (NSG) and Allied System for Geospatial-Intelligence (ASG). Prior to this, Mr. Roberts served as the inaugural Director, Artificial Intelligence, Automation and Augmentation for NGA.

Mr. Roberts' more than 36 years of service in the Intelligence and GEOINT community began in March of 1982 upon entering the United States Air Force as an Imagery Analyst. After his Air Force retirement, he immediately began his civilian career, with NGA's legacy organization the National Imagery and Mapping Agency (NIMA), in July of 2004.

Throughout his career he's held numerous military and civilian positions in a variety of capacities at tactical, operational and theater command headquarters levels, as well as service staffs, national agency and national senior decision maker levels. His numerous positions within the Intelligence Community and Department of Defense Intelligence, Surveillance and Reconnaissance (ISR) community spanned intelligence analysis, collection management, operations, planning and programming, training, resource and human capital management and advanced science and technologies activities. His service includes 16 years of overseas assignments. His lengthy operational and intelligence expertise is highlighted by a distinguished and proven career of operations excellence and transformation and innovative improvements in intelligence, decision, warfighting, acquisition, training and science and technology operations and capabilities.

Mr. Roberts is a 2016 graduate of the National War College, with a Master of Science in National Security Studies. He also is a graduate of University of Maryland University College with Master of Science in Technology Management, Master of Business Administration and Bachelor of Science in Computer and Information Science degrees. He is Defense Acquisition Program Management and Science and Technology Program Management, level 3 certified.

He is a native of Royal Oak, Michigan and is an active community member and enjoys participating in various academic, sporting and outdoor activities. He is married to the former Ms. Patricia Ann Rich. They have a blended family of six adult children; Kasey, Samira, Jesse, Daniel, Andrew and Heather and live in Fauquier County, Virginia.

Invited **Speakers**

James A. Bednar
Senior Technical Consultant, Anaconda, Inc.
jbednar@anaconda.com

Python Tools for Geoscience Research Using Machine Learning and AI

Abstract: Python is a very popular choice for machine-learning and artificial-intelligence applications, in part because of its

- wide range of available ML/AI algorithms
- support for easily transforming data into a usable form
- tools for big-data visualization and analysis
- support for automating complex and computationally intensive data processing

However, it can be difficult to determine the best way to approach any specific problem in such a general framework. Earth-science and climate data also present specific difficulties that can make NASA-related data workflows awkward, error-prone, time consuming, and limited in data size and complexity.

The EarthML project (EarthML.pyviz.org) is a new joint initiative between NASA Goddard and Anaconda, Inc. to make it simpler to apply machine-learning and related techniques to satellite imagery, climate measurements, and other data sources used by NASA.

EarthML consists of:

1. Best-practice examples of using open-source Python libraries in complete data pipelines including data preparation, visualization, and analysis (EarthML.PyViz.org and EarthSim.PyViz.org).
2. Improvements to underlying libraries for scalable visualization (PyViz.org), scalable data processing (Dask.pydata.org), and multidimensional data cataloging and retrieval (Intake.readthedocs.io and XArray.pydata.org).
3. Documentation and training materials to make it easier to get started with each of these libraries (EarthML.PyViz.org and EarthSim.PyViz.org).
4. We will demonstrate these tools with examples of satellite image segmentation via spectral clustering, tagging data on maps as ML training examples, and predicting climate variables of interest from local or satellite measurements.

Bio: Jim Bednar leads the PyViz group at Anaconda, working with commercial and government clients to improve Python software for visualizing and analyzing large and complex datasets. Dr. Bednar holds an M.A. and a Ph.D. in Computer Science from the University of Texas, along with degrees in Electrical Engineering and Philosophy. He has published more than 50 papers and books about the visual system and about software development. Dr. Bednar manages the open source Python projects PyViz, Panel, Datashader, HoloViews, GeoViews, Param, Colorcet, and ImaGen. Before Continuum, Dr. Bednar was a lecturer and researcher in computational neuroscience at the University of Edinburgh, Scotland, as well as a software and hardware engineer at National Instruments.

Invited **Speakers**

John Calhoun

Amazon AWS

jcalho@amazon.com

Improving "Time to Science" Using AWS Machine Learning

Abstract: In this presentation we will talk about how machine learning is being used in science to discover new results and how AWS can accelerate those machine learning workloads.

Bio: John Cahoun is a solutions architect on Amazon Web Services (AWS) public-sector partners team who specializes in machine learning. Before AWS John was a mathematician who did research and education around machine learning. Now at AWS he works closely with a wide variety of customers, helping them meet their missions by adopting and using machine learning.

Tom Goldstein

UMD

tomg@cs.umd.edu

Multi-scale Neural Networks for Image Processing

Abstract: We present Stacked U-Nets, a simple neural net architecture that combines the information globalization properties of multigrid solvers with the expressive power of neural nets. Stacked U-Nets are very effective for image processing problems, and achieve state of the art performance on image segmentation using relatively few parameters.

Bio: Tom Goldstein is an Assistant Professor at the University of Maryland Department of Computer Science. His research lies at the intersection of optimization and distributed computing, and targets applications in machine learning computer vision. Tom designs intelligent systems for a wide range of platforms. This includes powerful cluster/cloud computing environments for machine learning and computer vision, in addition to resource limited integrated circuits and FPGAs for real-time signal processing. Tom's research takes an integrative approach that considers algorithms and hardware to build practical high performance systems, and theoretical studies to understand how these systems work. Before joining the faculty at Maryland, Tom completed his PhD in Mathematics at UCLA, and was a research scientist at Rice University and Stanford University. Tom has been the recipient of several awards, including SIAM's DiPrima Prize, a DARPA Young Faculty Award, and a Sloan Fellowship.

Invited **Speakers**

Madhulika (Lika) Guhathakurta
ARC (SETI)
mguhathakurta@gmail.com

Applied Artificial Intelligence for Science & Exploration

Abstract: The recent advances in Artificial Intelligence (AI) capabilities are particularly relevant to NASA Heliophysics because there is growing evidence that AI techniques can improve our ability to model, understand and predict solar activity using the petabytes of space weather data already within NASA archives. This represents a strategic opportunity, since the need to improve our understanding of space weather is not only mandated by directives such as the National Space Weather Action Plan and the Presidential Executive Order for Coordinating Efforts to Prepare the Nation for Space Weather Events, but also because space weather is a critical consideration for astronaut safety as NASA moves forward with the Space Policy Directive to leave LEO and return to the Moon.

The Frontier Development Lab (FDL) is an AI research accelerator that was established in 2016 to apply emerging AI technologies to space science challenges which are central to NASA's mission priorities. FDL is a partnership between NASA Ames Research Center and the SETI Institute, with corporate sponsors that include IBM, Intel, NVidia, Google, Lockheed, Autodesk, Xprize, Space Resources Luxembourg, as well as USC and other organizations. The goal of FDL is to apply leading edge Artificial Intelligence and Machine Learning (AI/ML) tools to space challenges that impact space exploration and development, and even humanity.

Bio: LAs a NASA astrophysicist, Dr. Madhulika Guhathakurta (also known as Lika) has had the opportunity to work as a scientist, mission designer, instrument builder, directing and managing science programs and teacher and spokesperson for NASA's mission and vision in the Heliophysics Division. Occasionally, she performs all of these roles in a single day. Before joining NASA Headquarters in December of 1998, her career has focused on studying the importance of the scientific exploration of space in particular understanding the Sun as a star and its influence on the planet Earth, with research focus on understanding the magneto hydrodynamics of the Sun's outermost layer, the solar corona. She has been a Co-Investigator on five Spartan 201 missions on aboard space shuttles (STS-56, STS-64, STS-69, STS-87, STS-95) to study the solar corona in white-light and UV radiation and nine eclipse expeditions. She has led the Living with a Star Program for the past 15 years whose goal is to understand and ultimately predict solar variability and its diverse effects on Earth, human technology and astronauts in space, also known as "Space Weather". She has led missions such as STEREO, SDO, Van Allen Probes, Solar Orbiter Collaboration, Parker Solar Probe and others. She initiated Space Weather to be a permanent agenda item at UNCOPUOS in 2013 and started International Living with a Star (ILWS) in 2003. She has partnered with the American Museum of Natural History in New York and NASM in DC to produce full dome planetarium and 3D IMAX shows that are being exhibited internationally and used by teachers to excite the next generation of space scientists and helped create graduate level textbooks and train the next generation in heliophysics.. She was also the lead scientist for the 2017 Eclipse and presently on detail to NASA Ames Research Center exploring concepts for new initiatives.

Invited **Speakers**

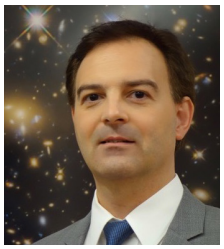

Victor Pankratius

MIT

pankrat@mit.edu

Towards Deriving Theories from Data: Frontiers for Model Inference in Astro-&Geophysic

Abstract: Machine learning techniques that address classification and feature detection problems have gained popularity in data science. Inference techniques, by contrast, are receiving much less attention even though they have significant untapped potential for scientific insight generation. Advancing these techniques beyond state-of-the-art has the potential to enhance the scalability of computer-aided discovery and reveal missing links between data-derived features and their explainability through our current body of knowledge. This talk outlines early explorations on deriving geophysical and astrophysical models from data by leveraging techniques involving symbolic computation and genetic programming. This context will also highlight the benefits of infusing domain knowledge into artificial intelligence methods and provide illustrations from case studies in volcanology and exoplanet search. Support for this work is provided by NASA AIST80NSSC17K0125 and NSF ACI1442997.

Bio: Victor Pankratius leads the Data Science in Astro-&Geoinformatics group at MIT where he is currently affiliated with the MIT Kavli Institute for Astrophysics and Space Research. He also serves a principal investigator in the NASA AIST program. His research advances data science through novel methods involving domain-aware artificial intelligence, scalable parallel computing, and software engineering for artificial intelligence systems. Contact him at pankrat@mit.edu and victorpankratius.com.

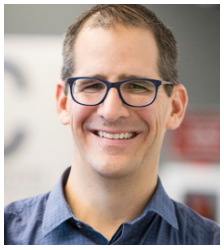

Bart Paulhamus

APL

Bart.Paulhamus@jhuapl.edu

Intelligent Systems Research at JHU/APL

Abstract: Established in 2016, the Intelligent Systems Center (ISC) acts as a focal point for research and development in intelligent systems at the Johns Hopkins University Applied Physics Laboratory (JHU/APL). In addition to fostering new internal and external partnerships, the ISC is home to a cross-disciplinary research group spanning artificial intelligence, robotics and applied neuroscience. The ISC leverages APL's broad expertise across defense, intelligence, homeland protection, space exploration, and health care to fundamentally advance the employment of intelligent systems for our Nation's critical

challenges. This talk will provide an overview of intelligent systems research at JHU/APL before diving into specific artificial intelligence research currently being investigated within the ISC.

Bio: Bart Paulhamus currently runs the Intelligent Systems & Physical Sciences Branch within the Research and Exploratory Development Department at Johns Hopkins University's Applied Physics Laboratory (JHU/APL). One of the goals of the branch is to tackle the challenges underlying intelligent systems through a multidisciplinary approach that combines machine learning, neuroscience, and robotics. His technical background is in machine learning, and he has over 15 years of experience applying advanced analytics to technical challenges within the Department of Defense, the Intelligence Community, and the Department of Homeland Security. In 2016, he co-led the establishment of JHU/APL's new Intelligent Systems Center. He joined JHU/APL in 2000 after receiving M.S. and B.S. degrees in Computer Science and Engineering from Penn State University. On the weekends, you can find Bart exploring the outdoors with his wife and four adventurous kids.

Panel **Members**

Susie Adams

Microsoft

Susie.Adams@microsoft.com

Abstract: ABSTRACT TO COME

Bio: Susie Adams is the Chief Technology Officer for Microsoft Federal where she oversees Microsoft Federal's cloud and security initiatives. Susie joined Microsoft in 1999 and has held multiple leadership positions as Director of the Reston Microsoft Technology Center, serving as the CTO, Federal Civilian Business and now as the CTO, Microsoft Federal. Prior to joining Microsoft, she spent 16 years in the Government Consulting arena in a variety of management and leadership roles as a practice manager and software developer. Susie was named as a Fed100 award winner, is a frequent speaker at industry events on Cloud computing, security and Big Data and has authored several books on the topics of software integration and web development. Susie is a graduate of George Mason University where she received a BS in Information Systems.

Larry Brown

NVIDIA

larryb@nvidia.com

GPU Accelerated High Performance Data Analytics for Federal Applications

Abstract: Dr.Brown will provide a quick overview of the role GPUs and high-performance computing play in current AI and data analytics techniques and give example use cases from the Federal ecosystem.

Bio: Dr. Larry Brown is a Sr. Solutions Architect with NVIDIA, and manager of the Federal Solutions Architecture team. His team helps customers design and deploy GPU accelerated workflows in high performance computing, data analytics and deep learning. Larry has 20 years of experience designing, implementing and supporting a variety of advanced sofware and hardware systems for defense and national security applications. Prior to joining NVIDIA he worked for some technology start-ups, as well as some large defense systems companies including General Dynamics, Textron and Booz Allen. Larry has designed electro-optical systems for head-mounted displays and training

simulators, developed GIS applications for multi-touch displays, and developed sofware on UGV and UAV programs. He has a Ph.D. from the Johns Hopkins University in the area of Vision Science and a graduate certificate in Sofware Engineering from the University of Colorado.

Dan Crichton
JPL
Daniel.J.Crichton@jpl.nasa.gov

Data Science at JPL: Integrating Data Analytics into the Full Data Lifecycle

Abstract: JPL and NASA have unlocked unprecedented amounts of scientific knowledge through exploration of our solar system, universe and planet Earth. The robotic spacecraft that JPL builds to support this scientific research have generated enormous amounts of data that have also challenged the traditional approaches to capturing, managing, analyzing and ultimately gaining insight from the data. Newer architectures and methodologies are needed to consider the entire observing system, from spacecraft to archive, with integrated data-driven discovery approaches enabling data exploitation both onboard and on the ground. Through a joint partnership in Data Science, JPL and Caltech are working to innovate new architectures, methodologies, and technologies that are not only enabling new approaches for space and Earth science, but now benefiting other disciplines such as biomedicine. This talk will address progress in applying data science to projects at JPL realizing a data-driven strategy across the mission-science data lifecycle.

Bio: Daniel Crichton is a program manager, principal investigator, and principal computer scientist at NASA's Jet Propulsion Laboratory, which he joined in 1995. He is the leader of the Center for Data Science and Technology, a joint center formed with Caltech, focusing on the research, development and implementation of data intensive systems for science and missions. Mr. Crichton has program management appointments to multiple JPL program offices for data science and data intensive system projects working in earth science, solar system exploration, and technology for NASA and non-NASA sponsors. He is one of the founding members of the International Planetary Data Alliance (IPDA), an organization committed to building compatible international planetary science data archives. He also serves as principal investigator of several distributed data system projects for NASA and the NIH, including the Informatics Center for the NCI Early Detection Research Network (EDRN). He architected an open source software framework, Apache OODT, for the management, distribution, and analysis of massive scientific data. This won runner-up for NASA Software of the Year 2003 and is now hosted by the Apache Software Foundation as a top-level project, NASA's first. Mr. Crichton recently served on the U.S. National Research Council Committee on the Analysis of Massive Data, which released its report on big data analytics. He has published over 100 papers on software and information architectures, distributed systems, and scientific data management and analysis, as well as five book chapters.

Panel **Members**

James E. Ecker

LaRC

james.e.ecker@nasa.gov

Deep Learning for Neuro-visualization and Continuous Control in Autonomous Systems

Abstract: This talk will give an overview of work taking place at NASA Langley investigating both deep unsupervised learning methods to generate low-dimensional mental representations of an object without real sensory data and deep reinforcement learning for continuous control and task learning in autonomous systems.

Bio: Mr. Ecker is a Computer Scientist in the NASA Langley Office of Chief Information Officer's Data Science group. He conducts research in Artificial Intelligence and Machine Learning, specializing in Reinforcement and Deep Learning. He is currently serving as a principal investigator in a multi-group initiative which seeks to explore methods to quantify trust and achieve effective behavioral explainability in autonomous systems. Mr. Ecker holds a Bachelor of Science in Computer Science from Florida Southern College and a Master of Science in Computer Science from the Georgia Institute of Technology.

John Hebeler

Lockheed Martin

john.w.hebeler@lmco.com

Determining Normal (and Abnormal) using Deep Learning

Abstract: Although normal is an ambiguous term, many operations, such as operational platforms (e.g. ship), electricity use, computer operations, and network traffic, depend on an understanding of normal conditions. Recognizing abnormalities can raise a concern before it becomes critical and help identify unusual situations worthy of investigation, even ones formally unrecognized. Fortunately, data exists across these environments from the many "sensors" that exist. Outlined methods enable machine learning to form a predictive fabric. The fabric is constructed using two deep learning methods; self-organizing maps (SOM) create the envelopes of normal states and recurrent neural networks (RNN) predict sequences. The two combine to form a high performance prediction model to determine normal/abnormal conditions. The methods "evolve" the predictive models to incorporate "normal" changes over time using model comparisons. Additionally, the methods reveal the sensitivity of the various data sources in contributing to normal determinations. Both deep learning methods depend on a powerful architecture based on GPUs.

Panel **Members**

Bio: John Hebeler, PhD is a Lockheed Martin Fellow focusing on large data set analysis via machine learning. Formally, he led a five-year program to analyze large, diverse data streams to form complex policy determinations in an event-driven architecture. John holds three patents and is the coauthor of two technical books and multiple journal articles on networking, data semantics, and machine learning. He also teaches graduate technology courses for University of Maryland. John holds a BS in Electrical Engineering, an MBA, and a PhD in Information Systems. In his free time, he's an avid tennis player, beer brewer, and amateur audiophile, usually not simultaneously.

Graham Katz

IBM

gkatz@us.ibm.com

Watson Intelligent Advisors: Discovery and Conversational Technology for Now and the Future

Abstract: In collaboration with NASA, IBM Watson has developed prototype intelligent assistants based on it's award-winning Watson question answering system, exploiting natural language understanding technology to provide the information needed to aid experts in problem solving, event response and technical research. These prototypes provide insight into the needs, requirements and challenges that there are for intelligent advisor systems based on exploiting information in technical documents. In this brief talk we will review lessons learned from these engagements that have fueled IBM's efforts to advance the state of the art, unlocking more of the information in documents and using embodied technologies to move conversational interfaces beyond the keyboard and microphone, expanding the role of context.

Bio: Graham Katz received his PhD in Linguistics and Cognitive Science (AI) from the University of Rochester in 1995. He has held research and teaching positions at the Universities of Tübingen and Osnabrück and at Stanford University and headed the Computational Linguistics program at Georgetown University until 2012. He joined IBM Watson in 2015 where he is now a Senior Natural Language Processing Analyst with the Watson Public Sector Delivery team. At IBM his work has focused on the Watson Discovery Adviser offering and the IBM Cloud language processing services, including Watson Conversation, Watson Discovery, Watson Knowledge Studio, Watson Natural Language Understanding and Watson Speech to Text.

Panel **Members**

Dinesh Manocha

UMD

dm@cs.umd.edu

Autonomy and AI Research at UMD

Abstract: There has been considerable interest in AI technologies and their applications. This has been facilitated by recent developments in machine learning, availability of large datasets, and high performance commodity processors. I will give a short history of AI research at University of Maryland and give an overview of current projects in computer vision, robotics, natural language processing, machine learning, game theory, and natural language processing.

Bio: Dinesh Manocha is currently the Paul Chrisman Iribe Chair of Computer Science and Electrical & Computer Engineering at the University of Maryland at College Park. He has published more than 480 papers in the leading conferences and journals in computer graphics, robotics, computational geometry, databases, multimedia, high performance computing and symbolic computation, and received 16 best paper and time of test awards. He is a co-inventor of 9 patents, several of which have been licensed to industry. He has won many awards including NSF Career Award, ONR Young Investigator Award, Sloan Fellowship, IBM Fellowship, SIGMOD IndySort Winner, Honda Research Award, UNC Hettleman Prize, etc. He is a Fellow of ACM, AAAS, AAAI, and IEEE, and received Distinguished Alumni Award from Indian Institute of Technology, Delhi.

Vikash Mansinghka

MIT

vkm@mit.edu

probcomp.ea@gmail.com

abrower@mit.edu

Paradigms of Probabilistic Programming

Abstract: Is it possible to develop general-purpose software for knowledge representation, learning, and reasoning, that can be used to combine prior knowledge with empirical learning, reduce the development time and cost for state-of-the-art AI research prototypes, and solve real-world modeling and inference problems in any domain where AI techniques can be applied? Over the last ten years, the emerging field of probabilistic programming has begun to map out a rigorous approach to achieving this ambitious goal. Probabilistic programming languages are now under active development at universities such as MIT, Stanford, and Berkeley, and by teams at companies such as Google, Microsoft, Uber, Amazon, and Facebook.

Panel **Members**

This talk will illustrate key principles and paradigms of probabilistic programming using open-source probabilistic programming systems recently prototyped at MIT. Specifically, it will present probabilistic programs in Gen, which builds on TensorFlow and Julia; Metaprob, which builds on the Clojure variant of Lisp; and BayesDB, which builds on relational databases and SQL-like query engines. Example applications will be drawn from machine perception and scientific data analysis, specifically:

1. Inferring 3D structure from 2D images, by using combinations of deep learning and Monte Carlo inference to invert generative models based on graphics software. Example applications include inferring 3D models for human faces, human bodies, and the geometry of radially symmetric objects. The key technology enabling these applications is probabilistic programming languages with programmable inference.
2. Learning probabilistic programs that can simulate synthetic biology lab experiments, and serve as the basis for machine-assisted data analysis. In these problems, data is too scarce for neural network techniques to be sufficient on their own. The key technology enabling these applications is new techniques for Bayesian learning of probabilistic programs from data.

Gen is a probabilistic programming language with fast programmable inference, embedded in Julia and integrated with TensorFlow. Gen allows users to concisely express generative models and perform inference using combinations of state-of-the-art sequential Monte Carlo, gradient-based optimization, and deep learning techniques.

BayesDB is a probabilistic programming platform for machine-assisted data science, which lets users without statistics training query probabilistic implications of their data using an SQL-like language. BayesDB automatically builds models by synthesizing probabilistic programs from data, using new Bayesian model discovery techniques.

Applications will be drawn from machine perception, such as inferring 3D models of human bodies from single images, and machine-assisted analysis of scientific data.

Bio: VVikash Mansinghka is a research scientist at MIT, where he leads the Probabilistic Computing Project. Vikash holds S.B. degrees in Mathematics and in Computer Science from MIT, as well as an M.Eng. in Computer Science and a PhD in Computation. He also held graduate fellowships from the National Science Foundation and MIT's Lincoln Laboratory. His PhD dissertation on natively probabilistic computation won the MIT George M. Sprowls dissertation award in computer science, and his research on the Picture probabilistic programming language won an award at CVPR. He served on DARPA's Information Science and Technology advisory board from 2010-2012, and currently serves on the editorial boards for the Journal of Machine Learning Research and the journal Statistics and Computation. He was an advisor to Google DeepMind and has co-founded two AI-related startups, one acquired and one currently operational.

Panel **Members**

Cynthia Matuszek

UMBC

cmat@umbc.edu

Learning Grounded Language For and From Interaction

Abstract: As robots become more powerful, capable, and autonomous, they are moving from controlled industrial settings to human-centric spaces such as medical environments, workplaces, and homes. As physical agents, robots have the potential to collaborate with people in entirely new categories of tasks that require intelligence. Before that can happen, robots must be able to interact gracefully with people and the noisy, unpredictable world they occupy, an undertaking that requires insight from multiple areas of AI. Useful robots will need to be flexible in dynamic environments with evolving tasks, meaning they must learn from and communicate effectively with people. In this talk, I will describe current research in our lab on combining natural language and robotics to build robots people can use in human environments, and some ways in which academia is key to building a diverse, informed AI research community.

Bio: Cynthia Matuszek is an assistant professor of computer science and electrical engineering at the University of Maryland, Baltimore County. Her research focuses on robots' acquisition of grounded language, in which robots learn to understand how language relates to the real, physical world. She has developed algorithms and approaches that make it possible for robots to learn about their environment and how to follow instructions from interactions with non-technical end users. She received her Ph.D. in computer science and engineering from the University of Washington in 2014, where she was co-advised by Drs. Dieter Fox and Luke Zettlemoyer in robotics and natural language processing. She has published in robotics, machine learning, artificial intelligence, natural language processing, and human-robot interaction venues. Dr. Matuszek was recently named one of IEEE's bi-annual "10 to watch in AI."

Jon Neff

Aerospace

jon.m.neff@aero.org

Overview of Aerospace Corporation AI Initiatives

Abstract: The Aerospace Corporation has a number of initiatives involving applications of artificial intelligence and neural networks. Our work includes detection of near Earth objects from telescope

imagery, robust intelligent ecosystems for weather prediction, prediction of spacecraft anomalies from telemetry, nondestructive evaluation of additive manufacturing, identification of chemical species in hyperspectral imagery and autonomous docking. This presentation provides an overview of these activities.

Bio: Jon Neff is a data scientist and system engineer in the Civil Systems Group Artificial Intelligence Team at The Aerospace Corporation. He has worked on development and operations of several NASA robotic space missions and also has experience in software startups and medical devices. Most recently, he was Director Risk Analytics at Visa. He has a Ph.D. in aerospace engineering from the University of Texas at Austin and an MBA from Pepperdine University.

Nikunj Oza
ARC
nikunj.c.oza@nasa.gov

Artificial Intelligence in the NASA Ames Intelligent Systems Division

Abstract: The Intelligent Systems Division at NASA Ames provides leadership in information technology for NASA by conducting mission-driven, user-centered computational sciences research, developing and demonstrating innovative technologies, and transferring these new capabilities for utilization in support of NASA missions. Much of this research and development involves Artificial Intelligence (AI), ranging from Machine Learning to Information Retrieval to Robotics and Automated Planning and Scheduling. The 2017 Ames Machine Learning Workshop discussed some of this work and looked to future developments. This talk will discuss the range of work just described.

Bio: Nikunj Oza is the leader of the Data Sciences Group at NASA Ames Research Center. He also leads a NASA project team, which applies data mining to aviation safety and operations problems. Dr. Oza¹s 50+ research papers represent his research interests, which include data mining, machine learning, ensemble learning, anomaly detection, and their applications to Aeronautics and Earth Science. He received the Arch T. Colwell Award for co-authoring one of the five most innovative technical papers selected from 3300+ SAE technical papers in 2005. His data mining team received the 2018 NASA Honor Award and the 2010 NASA Aeronautics Research Mission Directorate Associate Administrator¹s Award. He is an Associate Editor for the peer-reviewed journal Information Fusion (Elsevier) and has served as organizer, senior program committee member, and program committee member of several data mining and machine learning conferences. He received his B.S. in Mathematics with Computer Science from MIT in 1994, and M.S. (in 1998) and Ph.D. (in 2001) in Computer Science from the University of California at Berkeley.

Panel **Members**

Ray Ptucha
RIT
rwpeec@rit.edu

Deep Learning on Graph Data

Abstract: Deep learning methodologies using Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) have profoundly impacted the pattern recognition community. For example, the Machine Intelligence Laboratory at RIT has recently developed CNN and RNN architectures to solve tasks in the fields of remote sensing, facial recognition, video summarization, robot localization, and handwriting recognition. CNN/RNN techniques are predicated on structured inputs such as dense gridded tensors, lending themselves to a simple, but restrictive set of filtering and pooling operations. Unfortunately, most of the world's problems are unstructured and the analogy of CNN/RNN techniques to heterogeneous graphs with varying vertices, neighbors and edge weights is not straight forward. To address this gap, filtering techniques called Graph-CNNs are introduced. Graph-CNNs operate on generic unstructured data, lending itself to a simple, but non-restrictive set of filtering and pooling operations on both homogeneous and heterogeneous data. Several examples will be given to demonstrate the usefulness of the proposed methods on protein structure, fMRI, and point cloud data, as well as the fusion of unstructured inputs like LiDAR with structured data such as computer vision.

Bio: Raymond Ptucha is an Assistant Professor in Computer Engineering and Director of the Machine Intelligence Laboratory at Rochester Institute of Technology. His research specializes in machine learning, computer vision, and robotics. Ray was a research scientist with Eastman Kodak Company where he worked on computational imaging algorithms and was awarded 31 U.S. patents with another 19 applications on file. He graduated from SUNY/Buffalo with a B.S. in Computer Science and a B.S. in Electrical Engineering. He earned a M.S. in Image Science from RIT. He earned a Ph.D. in Computer Science from RIT in 2013. Ray was awarded an NSF Graduate Research Fellowship in 2010 and his Ph.D. research earned the 2014 Best RIT Doctoral Dissertation Award. Ray is a passionate supporter of STEM education and is an active member of his local IEEE chapter and FIRST robotics organizations.

Panel **Members**

Brian Roberts

GSFC

brian.roberts@nasa.gov

AI, Autonomous Systems, and Computer Vision for Satellite Servicing

Abstract: Ground simulations are a key part of enabling on-orbit robotic satellite servicing since the necessary technologies have never been tested in space. Autonomous systems are used for rendezvousing with client satellites in need of servicing and for robotic system operations. The presentation will describe the ways in which NASA is maturing these autonomous systems on the ground and on the International Space Station, for their future incorporation in on-orbit servicing, assembly, and manufacturing missions.

Bio: Brian Roberts is the Robotic Technologist in the Satellite Servicing Projects Division (https://sspd.gsfc.nasa.gov/) at NASA's Goddard Space Flight Center. His team is maturing the robotic technology needed to perform satellite servicing in space as well as developing the capability to simulate the dynamics of a robotic system interacting with space objects and using industrial robotic platforms to simulate motion of space objects. Before coming to Goddard, Brian spent 6 years as a research engineer at the University of Maryland. There he work on teams that developed and tested various robotic systems ranging from those designed to service satellites and fly on the shuttle, to those that can put themselves together and take themselves apart in space, to those that assist physical therapists working with shoulder rehabilitation patients, to those that autonomously find and sample life at the bottom of the ocean. Most of his time was spent coordinating the design, assembly, testing, and operation of the systems and conducting much of the testing underwater in the lab's Neutral Buoyancy Research Facility.

Brian earned a Bachelor of Science in Aerospace Engineering from Case Western Reserve University in Cleveland, Ohio and completed a Master of Science in the same field at the University of Maryland where he also completed coursework in Fire Protection Engineering.

Brian Thomas

HQ

brian.a.thomas@nasa.gov

Elements of an AI/ML Architecture for NASA

Abstract: I will discuss important issues which are preventing NASA from doing a better job at leveraging Machine Learning (ML) and discuss possible solutions.

Keynote **Speakers**

Bio: Brian Thomas is the NASA Office of the Chief Information Officer (OCIO) Data Scientist. He has 25+ years of experience supporting or leading scientific research, data analysis and scientific programming. This experience spans diverse environments at a number of institutions and facilities including helping guide the agency data science strategy and mission for NASA (NASA/HQ), developing scientific data archives (National Optical Astronomy Observatory; NOAO, the Astronomical Data Center; NASA/GSFC), and developing and working in operations at a scientific satellite operations center (XTE SOF; NASA/GSFC).

Mr. Thomas is fluent in most of the common technologies and methodologies needed for effective scientific and information systems research and development including various programming languages (Java/J2EE, Python, Perl, Ruby/Rails, C/C++), solving research problems with scientific and/or big data, working with semantic technologies, machine learning, natural language processing, XML/RDF/JSON, webservice development, service oriented architectures, microservices, scientific/data visualization, continuous/agile/testdriven software development, database modeling, database architecting, database application design and administration (relational, XML & NoSQL databases). He has demonstrated software development management/leadership skills and has effectively led distributed software development and testing teams.

His research experience is diverse and includes multi-wavelength data analysis in astronomy (X-ray/optical) and development of novel algorithms/platforms/applications for scientific software and research. He has extensive experience writing white papers, posters and giving oral presentations. He has participated in and helped organize national and international meetings/forums.

Short **Talks**

By Order of Presentation:

AI's Missing Real-World Connection, and Its Essential and Multifaceted Roles

David Sekora
djsekora@cs.umd.edu
University of Maryland, College Park

Co-Authors: Timothy C. Clausner, Justin Brody, Don Perlis

For an intelligent system, knowledge about itself and how it fits into the world is important1, but has been under-explored in the AI literature. While ad hoc solutions might solve any particular instance of a lack of self-awareness leading to undesirable behavior, they won't be able to tackle a multitude of others. Too many things can go wrong for engineers to foresee them all. Instead, a system should assess when things are going badly and make appropriate autonomous responses.

There is an underlying issue here: an artificial agent's knowledge is generally regarded as residing within the agent, an internal encoding of properties and behaviors apparent in the world.
Yet, that knowledge is (largely) about aspects of the real world external to the agent. An open question is, how is the internal connected to the external (Perlis 1991 and 1994; Clausner 2005)?

This appearance-reality distinction has historically interested linguists, philosophers and psychologists, not AI researchers. We claim that a rich internal-external connection is essential for progress toward human-level automated intelligence, including anomaly detection, general learning, scene interpretation, and explanation.

Recent work (Perlis, 2016) has detailed just how critical this distinction can be. Empirical studies have contributed new understanding that how people imagine the world, affects their performance (Clausner, 2009; Palmer et al., 2009). Work in AI (Perlis et al, 2017) uses the time situated active logic formalism and its implementation, ALMA, to study temporal forms of connection between a robot and its surroundings. Unlike traditional temporal-logic approaches to AI reasoning, ALMA has a physical clock-based connection with real time, unlike traditional temporal-logic approaches to AI reasoning. A related project demonstrated the importance of self-representation when a robot was unable to function in conversation until it was outfitted with the ability to detect that it was hearing itself as it was speaking (Brody et al, 2015).

We are beginning to apply these ideas to physical movement. Our approach includes an internal world-and-self simulator with a naive physics engine, alongside the addition of spatial dimensions to active logic. The spatial connection to the world, through the body, is the key – a system that understands that its actions affect the physical world will be able to better comprehend the consequences of those actions. Our aim is to connect spatial reasoning to actual evolving movement in space, analogous to how ALMA

connects temporal reasoning to actual evolving time.

While this work may not cleanly fit into only one of the seven workshop categories, we believe it to be an important topic that underlies and cuts across many of them, especially learning and explainability.

[1] See (Sekora et al 2017) for a broad attempt to set out requirements for an artificial mind in a practical setting.

References:

Brody, J., Perlis, D., and Shamwell, J.. "Who's Talking—Efference Copy and a Robot's Sense of Agency." 2015 AAAI Fall Symposium Series. 2015.
Clausner, T.C. (2005). Image schema paradoxes: Implications for cognitive semantics. In B. Hampe & in collaboartion with J. E. Grady (Eds.), From Perception to Meaning: Image Schemas in Cognitive Linguistics. (pp. 93-110). Berlin; New York: Mouton de Gruyter.
Clausner, T.C., Palmer, E., Brown, C., Bates, C., & Kellman, P.J. (2009). Conceptualization ofperceptual cues guided by metaphor, perception, and imagined perspective. Psychonomic Society: 50th Annual Meeting, 14, 102. Boston, MA.
Palmer, E.M, Brown, C.M., Bates, C.F., Kellman, P.J. & Clausner, T.C. (2009). Perceptual affordances and imagined viewpoints modulate visual search in air traffic control displays. Proceedings of the 53rd Annual Meeting of the Human Factors and Ergonomics Society. 53(17).
Perlis, D (1991). Putting ones foot in ones head – part I: Why. Noûs 25 (4):435-55.
Perlis, D (1994). Putting ones foot in ones head – part II: How. In Eric Dietrich (ed.), Thinking Computers and Virtual Persons. Academic Press. pp. 435-455.
Perlis, D. (2016). Five dimensions of reasoning in the wild. Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, pp. 4152-4156.
Perlis, D., Brody, J., Kraus, S., Miller, M. (2017). The internal reasoning of robots. Proceedings of the Thirteenth International Symposium on Commonsense Reasoning.
Sekora D., Barham, S., Brody, J., Perlis, D. (2017). Anatomy of a Task: Towards a Tentative Taxonomy of the Mind. AAAI Fall Symposium on a Standard Model of the Mind.

Explainable Machine Learning for Aviation Safety Assurance

Jyotirmaya Nanda
jnanda@i-a-i.com
Intelligent Automation, Inc.

Co-authors: Alexander Grushin, Jyotirmaya Nanda, Ankit Tyagi, David Miller, Joshua Gluck and Nikunj C. Oza

During the past decade, the field of machine learning has been revolutionized by the use of deep and recurrent neural networks. While these biologically-inspired models are often capable of making decisions with an unprecedented level of accuracy, they have also faced increasing criticism for the inability to explain their decisions, which is problematic in a number of applications. For example, in the domain

of aviation safety assurance, it is of significant interest to reduce the likelihood of degraded states, i.e., deviations from normal operations that may (in some cases) have operational significance. To achieve this, it is important to understand why these degraded states arise, i.e., to identify their precursors, which can then aid in degraded state prevention. While neural networks are now being used to predict degraded states before they occur, an understanding of precursors requires a machine learning methodology that is explainable. To address this gap, we have developed a method where a "black box" neural network model B is converted into a "white box" model W that approximates B's decisions, but makes them explainable/interpretable. The latter model W is created by generating sample points on B's decision boundary, and then approximating this decision boundary via a set of hyperplanes in the input space. Each decision boundary point is found by perturbing some input via gradient descent until the network changes its decision (prediction), where the gradient of the neural network's output is computed with respect to the input. A local hyperplane is then defined for each decision boundary point, and attempts to separate nearby inputs for which a degraded state is predicted from those for which such a prediction is not made. We apply the above explainable machine learning procedure to identify precursors to degraded states, in a three-step process. First, we use a statistical procedure to detect degraded states, by identifying significant deviations in aircraft state variables. Second, we train the recurrent long short-term memory neural network (LSTM) to predict that a degraded state will occur in the future, based on data that was observed in the past and present. Third, we convert the LSTM network into a "white box" model, and analyze this model's hyperplane coefficients to explain how it makes the predictions; these explanations are then formulated as precursors. We have found that (a) the original "black box" LSTM-based model can accurately predict degraded states in advance of their occurrence, (b) the corresponding "white box" model achieves comparable prediction performance, and therefore provides a useful approximation of the "black box" model, and (c) the "white box" model helps to elucidate precursors to the degraded states.

Area of Interest: 7 (Explainable Artificial Intelligence (XAI))

A Cognitive Processing Enhanced Smart Interface Framework For Situational Awareness

Anil Deane
deane@umd.edu
University of Maryland

Co-Author: Renuka Ellis

We present a multi-(data)source, multi-(science) domain framework for situational awareness that includes visual analytics, image recognition, natural language and other cognitive processing elements to enhance the user experience (UX). The framework provides user configurable components that displays information and adds elements based on user commands via natural language interaction (NLI). A UX assumption is that as the number of data types and collections becomes large a point and click solution becomes unwieldy and for some applications a smart interface is needed.

Processing of the NLI utilizes AI and is done via a combination of local and cloud processing resources. Local master code mediates the bursting to the cloud for the actual cognitive API calls. We demonstrate using well known providers of cognitive services and show how the implementation is agnostic to the particular choice of cognitive services provider, although factor into performance which we measure. We present an abstract framework in the form of a reference architecture, and then describe specific implementations using Microsoft Azure, IBM Watson and Amazon Alexa that deploys as a thin client, web-browser, and mobile. The example implementation use Grafana to query, visualize and issue alerts.

To conclude we discuss the limitations, extensions and real-world application considerations of the implementation. Among the extensions are Kafka based analytics model processing and Elasticsearch. Example domains are robotics and IOT sensor networks of interest to NASA.

Data Poisoning Attacks Can Compromise Machine Learning Systems

W. Ronny Huang
wronnyhuang@gmail.com
Ernst & Young and University of Maryland

Data poisoning is an attack on machine learning models wherein the attacker adds examples to the training set to manipulate the behavior of the model at test time. This paper explores poisoning attacks on neural nets. The proposed attacks use “clean-labels”; they don’t require the attacker to have any control over the labeling of training data. They are also targeted; they control the behavior of the classifier on a specific test instance without degrading overall classifier performance. For example, a politically motivated employee could add a seemingly innocuous image (that is properly labeled) to a training set for an Arctic ice satellite vision system, and control the predicted ice mass at test time (to, e.g., falsify findings related to climate change). Because the attacker does not need to control the labeling function, poisons can be injected into the training set and easily bypass human auditors. We present an optimization-based method for crafting poisons, and show that just one single poison image can control classifier behavior when transfer learning is used. For full end-to-end training, we present a “watermarking” strategy that makes poisoning reliable using multiple (~50) poisoned training instances. We demonstrate our method by generating poisoned frog images from the CIFAR dataset and using them to manipulate image classifiers.

Deep Multi-Layer Networks for Optical Wavefront Sensing and Control

Bruce Dean
bruce.dean@nasa.gov
NASA GSFC Code 551

We discuss an application of deep multi-layer neural networks to optical wavefront sensing and control. A comparison is given to how the problem is traditionally solved and then how several new improvements

and insights can be made through the use of multi-layer convolutional networks. We also present how candidate networks can be optimized using hyper-paramater tuning, genetic optimization, and then finally, we discuss the application of reinforcement learning where a self-directed agent determines an optimal control strategy.

Machine Learning Approaches for General Satellite Maneuvers

Shahrouz Ryan Alimo
sralimo@jpl.nasa.gov
Jet Propulsion Laboratory, California Institute of Technology

Satellite swarms will be an integral tool for many future space-based scientific endeavors, including space-based interferometry. These endeavors will require the ability to precisely reconfigure satellites swarms to specifications generated on the fly according to scientific priorities. Crucial to reconfiguring swarms is the ability to maneuver individual satellites between prespecified start and end positions in a fixed amount of time. However, there is no consensus solution on how to do these single satellite maneuvers fuel-efficiently in the general case.

Consider the position of a single satellite in the swarm's Local Vertical/Local Horizontal (LVLH) reference frame relative to the swarm's reference orbit. For specific single satellite transition scenarios, such as transferring the satellite between LVLH passive reference orbits (PROs) with minimal thrust expenditure in the Clohessy-Wiltshire (CW) linearized orbit dynamics system, provably optimal algorithms have been developed. But these provably optimal algorithms fail to generalize in a multitude of ways. They fail to generalize to arbitrary start and end states in the LVLH frame, they fail to generalize to more complex dynamical models such as those including J2 and drag terms, they cannot handle more detailed models of thrusters and satellite pose, and cannot account for more complex optimization criteria which balance the needs of scientific instruments with fuel efficiency and timeliness. We develop two machine learning techniques to tackle this general problem of maneuvering a satellite between the specified start and end dates in a fixed time frame. Firstly, DeltaDOGS, a Delaunay triangulation-based black-box optimization technique, is applied to finding efficient three-burst maneuver solutions. As a black-box optimization technique, DeltaDOGS simulates different maneuvers to evaluate their effectiveness. DeltaDOGs, in particular, uses an interpolation function and explicit exploitation/exploration trade-off parameter to quickly home in on efficient maneuvers using a minimal number of simulations/computational resources, while still exhibiting good exploration over all possible maneuvers. Secondly, we train neural networks to produce efficient three and four-burst maneuver solutions. Neural networks front all simulation executions during a training period, allowing them to act quickly at runtime with minimal computational resources, making them candidates for being run onboard satellite computers. Also, as neural networks have been shown to be highly effective high-dimensional function approximators, this technique should be extensible to continuous control schemes. It is shown that both the DeltaDOGS and neural network approaches produce seemingly good solutions for transfers between general start and end states. It is then shown that for a specific subset of these transfers - namely transfers

between PROs under CW dynamics - both techniques achieve high levels of efficiency compared to the known optimum in this setting. This indicates that the solutions produced by these techniques in the general case are also likely to be efficient. Lastly, it is noted that neither DeltaDOGS nor neural networks are reliant on the choice of dynamical model or optimality criteria, suggesting that they should be extensible to more complicated and realistic scenarios without significant loss of efficiency.

Supervised-machine Learning for Intelligent Collision Avoidance Decision-Making and Sensor Tasking

Alinda Mashiku
@alinda.k.mashiku@nasa.gov
NASA Goddard Space Flight Center

Collision avoidance (CA) usually depends on the collision probability (Pc) based on relative position and velocity uncertainties. When the Pc passes a certain predetermined threshold, a decision on whether to maneuver or not is made. This however may neglect the further knowledge that is present about the system, such as the knowledge of how orbits and uncertainties are generated from the data, what the sensor schedules are etc. For example, at times, based on orbit uncertainty and data updates, a maximum Pc peak is passed that quickly abates. The objective of this project is to construct a neuro-fuzzy logic based decision making and a sensor tasking system that is based on intelligently selected parameters beyond the Pc values that takes the comprehensive knowledge of data gain, data processing, orbit generation and object dynamics into account. Fuzzy logic is a form of many values of logic that span from 0 to 1, in contrast to Boolean logic that has values of either 0 or 1. The project consists of two main interlaced tasks: a) the selection and robust simulation environment of the parameters that are relevant for the problem at hand and b) the design of a decision-making tool that uses supervised and unsupervised machine learning. The selected parameters include classical simplifying concepts, assuming Gaussian uncertainty for Pc, and using the means of the objects' relative positional distributions for miss-distance. In addition, statistically more comprehensive parameters, such as the Mahalanobis distance and Kullback-Leiber divergence are considered. Close conjunctions happen in a physical space, but are complicated by the fact that the exact object shapes, position and velocity realizations are not available, rather only their statistical representations. Furthermore, given the dynamics of the objects are highly non-linear; the divergence from linear approximations is most prominent over longer propagation time scales. A set of representative test cases are formed and simulated in a high fidelity environment to create the data points that enter the machine learning parameter space. Whereas the physical influence of certain parameters is relatively clear, it is unclear, which evaluation criteria or combination of criteria are most reliable and robust for the prediction of collisions and for automated decision making. This wide parameter set is used to create the Fuzzy Membership Function (FMF) inputs, an indicator function that provides the degree of truth of a set membership. We will construct fuzzy rules to operate the overall architecture based on clustering algorithms for supervised learning using neural networks. The neural network will train previous or simulated data of both the ground and space components to provide us with the correct mapping for the subset of the wide parameter set FMF inputs and outputs. The FMF outputs would provide us with a desired sensor tasking that enables quality measurements for improved uncertainty quantification along with a decision strategy on whether to maneuver or not.

Short **Talks**

SEVA-OIE: Open Information Extractor for the Systems Engineering Virtual Assistant (SEVA)

Jitin Krishnan
jkrishn2@gmu.edu
George Mason University

This project aims at implementing the Natural Language Processing (NLP) component of the Systems Engineering Virtual Assistant (SEVA). SEVA, designed at GSFC, is a novel attempt towards a trustable, human-in-the-loop, interactive, and personal software system designed to assist a Systems Engineer in their daily work environment through complex information management and high-level query-answering to augment their problem-solving abilities. The foundational knowledge of SEVA is built on two types of knowledge bases: 1) Domain Specific Common Sense and 2) Project Specific Information. We attempt to extract high accuracy relation triples/n-tuples to populate the knowledge bases. A new Information Extraction (IE) paradigm called Open Information Extraction (OIE) such as Stanford Open IE, Standalone Open IE by AI2, ClausIE, and CSD-IE help in identifying wide range of domain-independent relations as compared to the traditional IEs. However, this makes the extractions extremely noisy and erroneous which can be improved by knowing more about the domain language. We study the Systems Engineering domain, analyze a set of open information extraction (OIE) systems, and design an integrated, heuristic, and semi-supervised information extractor that extract relation triples/n-tuples from Systems Engineering text. This information is then refined to produce high accuracy tuples to construct a knowledge graph/ontology for a Question Answering system.

Improving NASA Earth science data and information access through natural language processing based data analysis and visualization

Zhong Liu
zhong.liu@nasa.gov
CSISS, George Mason University and NASA GES DISC

LCo-Authors: Zhong Liu[1,2], Shimei Pan[3], Jianwu Wang[3], Angela Li[2], and David Meyer[2]
1. CSISS, George Mason University; 2. NASA GES DISC; 3. Department of Information Systems, UMBC.
This presentation addresses the challenge listed in the announcement: Customize intelligent user interfaces, including visual analytics and natural language processing and No. 6 in the areas of interest. (An oral presentation is preferred) Abstract: NASA Earth science data collected from satellites, model assimilation, airborne missions, and field campaigns are large, complex and evolving. Such characteristics pose great challenges for end users (e.g., Earth science and applied science users, students, citizen scientists) particularly for those who are unfamiliar with NASA's EOSDIS system and thus unable to access and utilize datasets effectively.

For example, a novice user may simply ask for what is the total rainfall for a flooding event in my county yesterday? For an experienced user (e.g., algorithm developer), a question can be, how did my rainfall product perform, compared to ground observations, during a flooding event? Nonetheless, with rapid

information technology development such as natural language processing, it is possible to develop simplified Web interfaces and back-end processing components to handle such questions and deliver answers in terms of data or graphic results directly to users. In this presentation, we describe main challenges for end users with different levels of expertise in accessing and utilizing NASA Earth science data. Surveys reveal that most non-professional users normally do not want to download and handle raw data as well as conduct heavy-duty data processing tasks. Often they just want some simple graphics or data for various purposes. To them, simple and intuitive user interfaces are sufficient because complicated ones can be difficult and time consuming to learn. Professionals also want such interfaces to answer many questions from datasets.

Over the years, the NASA Goddard Earth Sciences Data and Information Services Center, or GES DISC, has developed various online tools to facilitate NASA remote sensing data access. The GES DISC Interactive Online Visualization and Analysis Infrastructure (Giovanni) is an example that provides access to over 1900 variables from satellite measurements and data assimilation projects without downloading data and software. However, the GUI-based interface in Giovanni is still quite different and complicated, compared to that of Web search or a search box. The latter, based on natural language, is much simpler and easy to use. Now the question is, with natural language processing, can we design a system to process a scientific question typed in by a user? It is evident that major search engines are working toward this direction, namely, delivering information, not just Web site links that users use to find information or data, but progress is slow in handling scientific queries. In this presentation, we describe our plan to develop such prototype. The workflow is, 1) extract needed information (e.g., variables, spatial and temporal information, processing methods, etc.) from the input, 2) process the data in the backend, and 3) deliver the results (data or graphics) to the user. Some existing backend software packages can be reused in this project. The team consists of researchers from the Department of Information Systems at UMBC who have extensive research experiences in natural language processing and system design as well as members at GES DISC who are experts on satellite remote sensing products and processing.

Improvements on coronal hole detection using supervised classification

Martin Reiss
martin.reiss@nasa.gov
NASA Goddard Space Flight Center

The application of machine learning continues to be an important driver of innovation in space weather research. In this study, we present the use of machine learning algorithms to improve the automated detection of coronal holes in NASA's Solar Dynamics Observatory (SDO) images of the Sun. It is well known that the evolution of coronal hole boundaries plays an important role in the origin of the slow wind. However, it is difficult to study coronal hole regions in a consistent manner because they evolve in space and time, and are easily confused with quiet-Sun regions and filaments. Here we use solar images from the Fe XII 19.3 nm channel from the SDO/AIA instrument and apply the Spatial Possibilistic Clustering Algorithm (SPoCA) to prepare training sets of manually labelled 'coronal hole' and 'filament'

regions over the years 2011 to 2016. We propose characteristic coronal hole attributes and use them as input for different classification algorithms such as Support Vector Machine, Linear Support Vector Machine, Decision Tree, and k-Nearest Neighbours. We systematically quantify their performance in terms of True Skill Statistics and find that the best results are obtained with cost-sensitive learning and Support Vector Machine classifiers. We conclude that a comprehensive classification rule may help improve the performance of other community models and outline future directions.

X-Net: Bimodal Feature Representation Learning in Satellite Imagery

Kenneth Tran
kvtran2@ncsu.edu
North Carolina State University

In September 2017, IARPA released a data challenge called the Functional Map of the World (fMoW) that introduced the largest remote sensing data set to date. This entailed the classification of annotated satellite image patches with their given bounding boxes. IARPA's goal was to re-assess the successes of deep learning on traditional images and explore its extension to remote sensing data. A unique property of the fMoW is that each data point is bimodal, including both the visible color image (RGB bands) at a high resolution and its associated multispectral image at a lower resolution. The top three contestants of the competition only used visible color imagery in their solutions. In our work, we conjecture that fusing the two modalities will enhance classification performance. We propose using a multimodal autoencoder to learn a joint representation between the bimodal data. We adapt notions from recent deep learning architectures, such as U-Net and densely connected CNNs, to design X-Net. In U-Net, each stage of the autoencoder consists of two convolutional layers followed by a down/up-sampling operation. This autoencoder has a "U" shape, where the first half of the "U" (split vertically) is the encoder and the second half is the decoder. Adopting this architecture to our multimodal autoencoder yields X-Net. To that end, we concatenate the activations from the previous encoding state from each modality into the input to the last encoding. We refer to it as the fusion stage. We proceed in the decoding layers by simply separating the joint representation back into the two separate modalities. In implementing the decoders, we follow the U-Net policy of making the stages symmetric. To gain perfect symmetry, we chose to use convolutional layers with larger strides instead of pooling in the encoder. This resulting fully convolutional neural network learns filter weights for the convolutional down-sampling in the encoder as well as transposed convolutional up-sampling in the decoder. With the proper padding, it also ensures that all the dimensions are symmetric throughout the autoencoder. Following densely connected CNNs, we propagate forward activations of each encoder stage in the network to the fusion stage to help the network retain information from previous layers. These activations undergo a convolution that make their dimensionality match that of the fusion layer input. They are then concatenated to the input to the fusion stage. Lastly, we replace the last transposed convolution with a pixel shuffler, a popular up-sampling method in super resolution applications. This results in higher quality reconstructed images, which is beneficial for training an autoencoder. Our method has outperformed the original multimodal autoencoder and bidirectional DNN (biDNN), achieving an PSNR of 30.384 between the original image

and reconstructed RGB image. Comparatively, the multimodal autoencoder and biDNN architectures achieve a PSNR of 28.086 and 27.947, respectively. Since the training is unsupervised, we can leverage the massive amount of unlabeled remote sensing data to learn an effective representation. This encoder can then act as a pre-trained feature extractor for downstream classification tasks.

Areas of Interest: 2. NN and Deep Learning (DL), 4. Computer Vision and Image Processing.

Hybrid Semantic Image Segmentation using Deep Learning for On-board Space Processing

Sabastian Sabogal
sebastian.sabogal@chrec.org
University of Pittsburgh

Due to ongoing innovations in sensor technology and autonomous spacecraft operations, on-board space processing continues to be outpaced by the computational demands required for future missions. The application of deep-learning concepts for on-board processing can enable spacecraft to efficiently process large volumes of raw sensor-data into scientific knowledge or actionable data to overcome limitations in downlink communication. On-board artificial intelligence can also enable spacecraft to formulate decisions for critical operations. Spacecraft often employ radiation-hardened (RadHard) space processors, which tend to be generations behind their commercial-off-the-shelf (COTS) counterparts in terms of performance and energy-efficiency. COTS hybrid system-on-chip (SoC) devices combine multiple, distinct computing architectures to attain and combine the advantages of each. FPGA-based hybrid SoCs (e.g., Xilinx Zynq) combine fixed-logic CPUs (e.g., dual-core ARM Cortex-A9) with a reconfigurable-logic FPGA fabric (e.g., Artix-7) and provide numerous architectural advantages that make them well-suited to address the escalating demands for high-performance, on-board processing. Recent advancements in deep learning present new opportunities for enhanced scientific methods, autonomous spacecraft operations, and intelligent applications for spaceflight missions. Semantic image segmentation is an advanced deep-learning algorithm based on convolutional neural networks (CNNs) that learns to infer dense labels for every pixel of an image at the original, spatial resolution. One example is the SegNet model, which introduces an encoder-decoder network architecture. In this model, objects are classified in the encoder stages, and feature maps are upsampled in the decoder stages. Semantic image segmentation has numerous space-science and defense applications, from semantic labeling of Earth observations for insights about our changing planet, to monitoring natural disasters for damage control, to gathering intelligence for national defense and security. Despite the high applicability of semantic image segmentation, these advanced deep-learning algorithms are computationally prohibitive on traditional RadHard processors.

To enable semantic image segmentation for on-board space computers, researchers at the NSF Center for Space, High-performance, and Resilient Computing (SHREC) at the University of Pittsburgh, in collaboration with NASA Goddard Space Flight Center, developed a hybrid (hardware/software codesigned) framework for accelerating CNN inferencing on hybrid SoCs. Our framework is scalable and

parameterizable to accommodate various mission applications. The framework architecture includes a scatter-gather DMA (SGDMA) and reconfigurable CNN accelerator (ReCoN) tuned for the highly parallel dataflow of the SegNet model. The SGDMA provides high-throughput streams of multi-dimensional feature maps to the ReCoN accelerator for high-performance parallel processing. When evaluated on Xilinx ZC706 (Zynq SoC) and Xilinx ZCU102 (Zynq UltraScale+ MPSoC) development boards, our hybrid approach achieves a performance and energy-efficiency improvement of up to two orders of magnitude versus the software-only baseline. Due to significant performance speedup and reduced energy consumption, our hybrid approach can be an enabling technology for applying semantic image segmentation and other CNN applications in future spaceflight missions.

Atmospheric Chemistry Modeling using Machine Learning

Christoph Keller
christoph.a.keller@nasa.gov
NASA Goddard Global Modeling and Assimilation Office (GMAO) and
Universities Space Research Association (USRA)

Modeling of the chemical composition of the atmosphere is central for a wide range of applications, including air quality, stratospheric ozone loss, and climate. Since 2017, the NASA Global Modeling and Assimilation Office (GMAO) produces global air quality forecasts at unprecedented horizontal resolution of 25 km using the Global Earth Observing System Model (GEOS-5) with GEOS-Chem chemistry. However, the computational cost associated with atmospheric chemistry modeling is large due to the large number of chemical species and the stiffness of the underlying set of ordinary differential equations. This severely limits the usability of current atmospheric chemistry models such as GEOS-Chem for applications such as ultra-high-resolution air quality modeling or chemical data assimilation. Instead of solving the differential equations of atmospheric chemistry numerically, we have been exploring alternative approaches based on machine learning to predict the rate of change of chemically reactive trace gases. We will present results from model simulations using random forests and neural networks, and discuss some of the advantages and limitations of both approaches. We find that random forests (30 trees, 10,000 leaves) offer better predictability than neural networks trained on the same data set. However, the performance of the neural network can be improved significantly by reducing the dimensionality of the system through analysis of the cross-correlations between chemical species. Further improvements can be achieved through mass balance considerations and by accounting for error correlations. Both machine learning approaches show many of the characteristics of the GEOS-5 full chemistry reference simulation and have the potential to be orders of magnitude faster. For instance, the evaluation of a single tree in the random forest is 1000 times faster than the reference chemical solver. The machine learning methods are also much more parallelizable than conventional numerical solvers and naturally lend themselves to new soft/hardware environments such as GPUs and cloud-based computing. There are a wide range of applications for an atmospheric chemistry module based on machine learning, e.g. for use in air quality forecasts, sub-grid parameterization, chemical data assimilation, or inverse modeling.

Short **Talks**

Implementation of Gaussian Processes in an Hydrological Model

Jules Kouatchou
Jules.Kouatchou@nasa.gov
NASA Goddard Space Flight Center

Co-Authors: Craig Pelissier, Grey Stephen Nearing, Daniel Duffy, Christa D Peters-Lidard and James V Geiger

The NASA Land Information Systems model (LIS) is a driver for high-performance terrestrial hydrology modeling. It enables modeling of several state-of-the-art hydrology models including NOAH, NOAH-MP, CLM, VIC, CABLE, Sacramento/SNOW17, and SUMMA. The land surface models can be validated using ground-based observations collected at hundreds of locations around the globe. The data products from LIS are used in many important applications including flood prediction, weather forecasting, and agriculture to name a few. With a large amount of observational data being collected, data-driven modeling offers a potential to improve the accuracy of current land-surface models. In particular, we show the potential to add a machine-learning model to the time-step integration of land surface models in LIS. A data-driven model is trained on the deviation between observations and a terrestrial hydrology model (e.g., NOAH), and then used at every timestep of the NOAH run to reduce simulation errors. In effect, this is a form of real-time data assimilation, and can also be combined with more common types of (e.g., Bayesian, variational, ensemble) data assimilation. By using the terrestrial hydrology model as the mean field or mean function for the data-driven model, simulations revert to the hydrology model in cases that are dissimilar to past observation data used for training. In this work, Gaussian Processes (GPs) are used to construct the data-driven model; these are trained and added to the LIS simulation code at the timestep level. The model is trained using data from 40 AmeriFlux tower sites with a two-year split-record time period for calibration and validation. In addition a leave-one-out strategy is used to test spatial extrapolation - here, one tower site is reserved for validation and all others are used for training. The GP model is run within LIS and compared with the NOAH model, and the results are presented and discussed.

Autonomous seasonality adaptation

Darsana Josyula
darsana@cs.umd.edu
Bowie State University

Co-Authors: Marvin Conn and Kenneth M'Bale; Anthony Herron and Hubert Kyeremateng Boateng.

Artificial agents are often trained on the specific conditions under which they will operate before they are deployed for autonomous or semi-autonomous operations. However, agents that need to operate in unknown terrains autonomously cannot be pre-trained for all the possible contingencies that may arise. Still, for successful autonomous operations, they should have the ability to detect anomalies, proactively avoid potential failures and recover from any failures that may occur. What goes into providing this ability is the topic of this research. The successful functioning of an autonomous agent depends heavily on the

efficiency of the agent's three components – anomaly detection, failure avoidance and failure recovery. Anomaly detection involves monitoring a data stream (observations from sensors, time-series of goals, time-series of behaviors, time-series of performance measures etc.) and detecting when the observed values deviate from what is typically normal for the agent. Failure avoidance involves detecting patterns in the data stream to predict when a failure may occur. Failure recovery involves identifying the correct sequence of behaviors that will bring the agent back to safe operating mode. Thus, all three components mandate effective processing of vast amounts of data. With memories and processors becoming smaller and faster while requiring lesser power to operate, it has become feasible to store vast amounts of data and process them quickly. Therefore, there is increased interest in collecting and organizing large amounts of data in order to detect patterns in the data and to particularly identify deviations from normal, in real-time. This research focuses on processing large amounts of data in real-time for autonomous anomaly detection, failure avoidance and failure recovery in seasonal environments. In seasonal environments, agents can exploit the lessons learned in a previous similar season to improve its performance when the season repeats. If the current season has never been experienced before, then it is an anomaly that requires special attention for learning further about how to behave in that season. For instance, a cognitive radar system may choose its transmission parameters based on what worked best in the previous season that had similar spectrum signals. The failed transmission attempts in a previous season need not be repeated if the cognitive system can recognize the similarity of the current season. Also, the learning that occurred during a failure can be reused for failure recovery when a similar failure occurs. Thus, the system can proactively avoid repeating failures and recover from failures that occur. This research addresses the following questions. i. How can an agent detect if the agent has experienced the current season before? ii. How can an agent predict its next season? iii. How can an agent avoid relearning its behavior when a season repeats? In this ongoing work, a reinforcement learner is used to learn appropriate behavior to maximize performance rewards of an agent. The data processing element Kasai organizes the incoming stream of data into seasonal patterns and predicts the next season. The Q-tables associated with known seasons are brought into play whenever a season repeats. When Kasai detects a new season, the exploration factor of the learner is adjusted to accelerate learning in the new season. Since the reinforcement learner and Kasai operate in real-time, the agent can autonomously adapt to changing seasons.

Machine learning in global scale classification of mangrove forests from remotely sensed imagery

Nathan Thomas
nathan.m.thomas@nasa.gov
ESSIC University of Maryland and NASA Goddard Space Flight Center

Co-Authors: Pete Bunting[3], Richard Lucas[3], Andy Hardy[3], Ake, Rosenqvist[4], Lola Fatoyinbo[2]

2. Biospheric Sciences, NASA Goddard Space Flight Center, MD, USA;
3. DGES, Aberystwyth University, Wales, UK;
4. solo Earth Observation (soloEO), Tokyo, Japan.

Machine learning has provided powerful tools for the classification of remotely sensed data from the advent of spaceborne image acquisition. Increased access to greater computer processing power accompanied by progressively large data volumes, has led to a rise in the number of available algorithms for image classification. These have demonstrated increased capability in achieving greater accuracies than their predecessors. Developed alongside technology in distributed computing that is now available (High Performance Computing), a set of tools is now available to the Earth science community to understand Earth and its processes at unprecedented geographical scales. Here we present an application driven use of machine learning to map the global extent of mangrove forests, a crucially important yet critically threatened ecosystem. We used a Random Forest[1] classifier within a geographic object-based image analysis (GEOBIA) framework, utilizing a range of open source software and tools. Spaceborne Phased Array L-band Synthetic Aperture Radar (ALOS PALSAR) and Landsat imagery was classified at 16 study site locations covering large geographical areas, to develop a globally applicable method. ALOS PALSAR and Landsat image stacks were segmented using the Shepherd et al algorithm[2] within the Remote Sensing and GIS Library (RSGISLib) python module[3]. Each image object represented the mean pixel values contained within it, for each band of each input image. Training data for three classes of 1) mangrove, 2) water and 3) other were generated and the Scikit-learn[4] python module implementation of Random Forests with 1000 estimators was used. The results yielded a comprehensive mangrove cover map for each of the 16 study sites, mapping a total of 2,529,760 ha and provided significantly improved results over existing products. Accuracy assessments yielded a map accuracy in excess of 90%. A modified version of the method was subsequently applied to 1500 ALOS PALSAR and 1800 Landsat composites to map 14,112,300 ha (2.26E+08 pixels) of mangrove, generating a new up-to-date baseline of global extent. Extremely randomized trees were used to classify mangroves across 128 subsets, utilizing a total training size of 12.8 million samples. This utilized a high-performance computing environment benefitting from the use of 100 cores. We outline a method of using machine learning within an open source GEOBIA framework for achieving global scale science products, facilitated by the python scripting language. The use of non-parametric machine learning algorithms provided numerous advantage over traditional algorithms used within the field. 1. Breiman, L. (2001). Random forests. Machine learning, 45(1), 5-32. 2. Clewley et al. 2014. A Python-Based Open Source System for Geographic Object-Based Image Analysis (GEOBIA) Utilizing Raster Attribute Tables. Remote Sensing. Volume 6, Pages 6111-6135 3. Bunting et al. 2014. The Remote Sensing and GIS Software Library (RSGISLib), Computers & Geosciences. Volume 62, Pages 216-226 4. Pedregosa et al. 2011. Scikit-learn: Machine Learning in Python, Journal of Machine Learning Research. 12, pp. 2825-2830.

Machine Learning in the Earth Data Analytic Services (EDAS) Framework

Thomas Maxwell
thomas.maxwell@nasa.gov
NASA Center for Climate Simulation (NCCS)

Faced with unprecedented growth in earth data volume and demand, NASA has developed the Earth Data Analytic Services (EDAS) framework, a high performance big data analytics and machine learning

framework. This framework enables scientists to execute data processing workflows combining common analysis and forecast operations close to the massive data stores at NASA. The data is accessed in standard (NetCDF, HDF, etc.) formats in a POSIX file system and processed using vetted tools of earth data science, e.g. ESMF, CDAT, NCO, Keras, Tensorflow, etc. EDAS utilizes high performance parallel data access, a custom distributed array framework, and a streaming parallel in-memory workflow for efficiently processing huge datasets within limited memory spaces with interactive response times. EDAS services are accessed via a WPS API being developed in collaboration with the ESGF Compute Working Team to support server-side analytics for ESGF. The API can be accessed using direct web service calls, a Python script, a Unix-like shell client, or a JavaScript-based web application. New analytic operations can be developed in Python, Java, or Scala (with support for other languages planned). Client packages in Python, Java/Scala, or JavaScript contain everything needed to build and submit EDAS requests. Capabilities The EDAS architecture brings together the tools, data storage, and high-performance computing required for timely analysis of large-scale data sets, where the data resides, to ultimately produce societal benefits. It is currently deployed at NASA in support of the Collaborative REAnalysis Technical Environment (CREATE) project, which centralizes numerous global reanalysis datasets onto a single advanced data analytics platform. This service enables decision makers to compare multiple reanalysis datasets and investigate trends, variability, and anomalies in earth system dynamics around the globe. EDAS services include configurable high performance neural network learning modules designed to operate on the products of EDAS workflows. As a science technology driver we have explored the capabilities of these services for long-range forecasting of the interannual variation of important regional scale seasonal cycles. Neural networks were trained to forecast All-India Summer Monsoon Rainfall (AISMR) one year in advance using (as input) the top 8-64 principal components of the global surface temperature and 200 hPa geopotential height fields from NASA's MERRA2 and NOAA's Twentieth Century Reanalyses. The promising results from these investigations illustrate the power of easily accessible machine learning services coupled to huge repositories of earth science data.

RNN/LSTM Ensemble Data Assimilation for the Lorenz Chaotic Models

Milton Halem
halem@umbc.edu
University of Maryland Baltimore County (UMBC)

Co-Authors: H. Vashistha and M. Halem

Data assimilation (DA) is the process of updating model forecasts (priors) with information from the observations of complete or incomplete state variables. The goal is to produce an improved model state (posteriors), which better represents the dynamical system. Previous work on using Artificial Neural Networks (ANN) for the purpose of data assimilation mostly used Feed Forward Back Propagation (FFBP) networks. There has been generally less use of Recurrent Neural Networks (RNN) for DA. In this talk, we propose to compare Long Short-Term Memory (LSTM) networks, a type of RNN, with FFBP data assimilation for the Lorenz-63 and Lorenz-96 models as well as Ensemble Kalman Filters.

Both of the Lorenz models are known to be chaotic in nature and while Lorenz-63 has only 3 state variables, the Lorenz-96 system of equations can represent as much or more than 36 variables. We implemented an OSSE experiment to prepare the model data for these networks. We present the results of the LSTM data assimilation for the Lorenz 63 and 96 models. We show that LSTM networks can produce results as good as the Ensemble Kalman Filter algorithm. Our LSTM network was trained against the Data Assimilation results of an EnKF for both of these models. LSTM networks were chosen since these networks avoid the problem of vanishing or exploding gradients, which vanilla RNN implementations can suffer.

Spaced Repetition for Training Artificial Neural Networks

Hadi Amiri
amirieb@gmail.com
Harvard University

We present a novel approach for training artificial neural networks. Our approach is inspired by broad evidence in psychology that shows human learners can learn efficiently and effectively by increasing intervals of time between subsequent reviews of previously learned materials (spaced repetition). We develop an efficient and effective spaced repetition-based algorithm to train neural models. The core part of our algorithm is a cognitively-motivated scheduler according to which training instances and their "reviews" (use for training) are spaced over time. Our algorithm uses only 34-50% of data per epoch, is 2.9-4.8 times faster than standard training, and outperforms competing state-of-the-art baselines. Our results show that standard training spends most of the computation on instances that are properly handled and can be ignored. Our code is available at scholar.harvard.edu/hadi/RbF/ We present two spaced repetition-based algorithms: a modified version of the Leitner system developed for humans and our Repeat before Forgetting (RbF) model respectively: Leitner system: this approach assumes n queues {q_0, q_1,..., q_{n-1}}. Its scheduler initially places all training instances in the first queue, q_0. During training, if an instance from q_i is correctly classified by the downstream learner (here an a neural network), the instance will be "promoted" to q_{i+1}, otherwise it will be "demoted" to the first queue, q_0. Leitner scheduler uses instances in q_i for training at every 2^i iterations. Therefore, instance in lower queues (difficult instances) are reviewed more frequently than those in higher queues (easy instances). The overhead imposed on training by the Leitner system is O(|current_batch|) at every epoch for moving instances between queues. Repeat before forgetting: The challenge in developing training models is to accurately estimate the time by which a training instance should be "reviewed" (used for training) before it is "forgotten" (misclassified) by the network. However, a heuristic scheduler such as Leitner system is sub-optimal as its hard review schedules (i.e. only 2^i-iteration delays) may lead to early or late reviews. We develop flexible schedulers that utilize "recall indicators" (namely item difficulty, delay since last review, and network strength) to lengthen or shorten review intervals with respect to individual training instances and downstream learners (neural networks) [1,2]. In particular, we propose using density kernel functions as schedulers to estimate the latest epoch in which a given training instance can be

recalled (correctly classified) by the network. Our kernels more confidently delay less difficult training instances in stronger networks. Our approach considerably outperforms Leitner scheduler and standard training across several tasks including image categorization and text classification. We also show the application of the above framework in spotting spurious instances (those with potentially wrong labels) in datasets. References [1] Hadi Amiri, Timothy A. Miller, Guergana Savova. Repeat before Forgetting: Spaced Repetition for Efficient and Effective Training of Neural Networks, EMNLP'17. [2] Hadi Amiri, Timothy A. Miller, Guergana Savova. Spotting Spurious Data with Neural Networks, NAACL'18.
Area: 2. Neural Networks and Deep Learning (DL).

Application of a Deep U-Net to Automatic Detection of Ship-Tracks Multispectral Images from both Polar-Orbiting and Geostationary Satellites

Tianle Yuan
tianle.yuan@nasa.gov
Goddard Space Flight Center and University of Maryland

Ship-tracks are an iconic demonstration of aerosol-cloud interactions. They have been observed in satellite images for more than 60 years. Many studies have resulted from analyses of remote sensing, in-situ, and modeling results. Satellite data are instrumental in discovering and understanding of ship-tracks. However, it is hugely time consuming to identify ship-tracks within satellite multispectral image with manual processing. Whole careers have been spent on manually picking out and analyzing ship-tracks. Here we apply a deep (16-layer) U-Net architecture to train a neural network algorithm to automatically find ship-tracks in the satellite data. The training set includes hundreds of ship-tracks with varying lengths from MODIS data. Native MODIS images come at 2030pixels by 1350pixels, which are way too large for deep NN to train. We adopt a strategy of breaking up native images into small (64X64 and 128X128) clips and train the NN on these sub-images and recombine them once the training is done. The risk is whether the structural information in the large images can be fully recovered with training only on smaller regional scales. We experimented with various setting of the model and training methods to test if this break-and-recombine strategy is viable . Our results show that this strategy works excellent. The recombined ship-track detection algorithm can not only pick out objects that is in the training data, but on many occasions find objects that are not noticed by human eyes. With this algorithm, we can now automatically find ship-tracks at a global scale during both daytime and nighttime. Initial test run on a few months of MODIS data, on the order of terabytes, show that our algorithm can map out ship-tracks occurrences with high accuracy. The occurrence frequency map agrees with similar maps from the commercial shipping industry very nicely. This unprecedented capability of finding ship-tracks at a global scale enables us to analyze radiance data at a large scale and to analyze ship-track properties with sample sizes of orders magnitude larger than what is currently available. For example, the current most comprehensive dataset for ship-tracks includes less than 10,000 samples. Our test run already finds more than a million. More importantly, our algorithm works directly on radiance data from multispectral instruments. This means we can apply it to geostationary as well as older polar orbiting sensors like AVHRR, which expands the possibility of studying ship-tracks. For example, we are

applying it to the new GOES-16 and -17 data to study the temporal evolution of individual ship-tracks. We can also combine the detection algorithm with multiple satellite sensors to dive deeper into the physical processes that are responsible for ship-tracks. Initial results are encouraging. We characterize the global distribution of ship-tracks and its variability on monthly to seasonal time scales. A significant portion (up to 5%) of the low clouds are affected by instantaneous ship-tracks. We also obtain results regarding the distribution of ship-track properties such as their length, width, and overall size (in terms of number of pixels) for different regions and seasons. Based on millions of ship-tracks we are able to find, it becomes possible to start comprehensively addressing the long-standing science question: how significant are ship-tracks in modulating low cloud properties.

Ushering in a New Frontier in Geospace Through Data Science

Ryan McGranaghan
ryan.mcgranaghan@jpl.nasa.gov
NASA Jet Propulsion Laboratory

The hallmarks of the geospace system, which extends from the near Earth space environment, through the magnetosphere and interplanetary space, to the Sun, are variability and complexity. To unravel the critical variabilities and complexities and to evolve beyond current approaches to understand geospace, new data-driven approaches and data analysis technologies are required. These data‐driven methods are taking on new importance in light of the shifting data landscape of the geospace system. The space physics community faces both an exciting opportunity and an important imperative to create a new frontier built at the intersection of traditional approaches and state-of-the-art data driven sciences and technologies. This talk will first discuss the meaning of data science in the context of geospace, and then reveal efforts from a Jet Propulsion Laboratory Data Science Working Group pilot project to leverage data science innovation to utilize a powerful data set for geospace science – Global Navigation Satellite Systems (GNSS) signals. We take advantage of a large volume of GNSS data, increasingly sophisticated tools for data-driven discovery, and machine learning algorithms covering a spectrum of complexity to develop novel predictive models of space weather-driven disruptions to GNSS signals (i.e., scintillation). We find that machine learning approaches significantly outperform current predictive capabilities, which, at high-latitudes, only consist of climatology and persistence. Using a robust metric known as the Total Skill Score (TSS), we position our results as a benchmark upon which to evaluate future predictive models, following a similar approach developed by the solar flare prediction community [Barnes et al., 2016]. Finally, this talk will be targeted to spark discussion of how data science provides a common paradigm to bridge disciplines, leading to the opening of new interdisciplinary research vistas and future funding opportunities for geospace science.

Reference: G. Barnes et al., (2016), A comparison of flare forecasting methods. I. Results from the ‘All Clear’ workshop, The Astrophysical Journal, 829(2), 89.
Workshop Areas of Interest addressed: 1, 2, 6, and 7.

Short **Talks**

Tracking optical flows for better data mining on solar images

Raphael Attie
raphael.attie@nasa.gov
NASA GSFC

Motions of plasma at the surface of the sun are known to the solar physics community thanks to the provision of an ever growing amount of image series. This community has developed various image processing techniques to track and characterize these motions, offering an interesting approach to solve the problems posed by optical flows: "Local Correlation Tracking" was the first automated technique to track horizontal photospheric motions in the 80's, and " Balltracking" is the most recent that automates the tracking of photospheric flows to data mine terabytes of images. A similar problem has been at the core of the research & development of robotics and autonomous vehicles that need curated and actionable information about what is moving around them. This has spawned another set of very advanced solutions, some of them making use of artificial intelligence. Due to the different nature of this industry, there have been little or no overlap with the solar physics community who is nonetheless facing similar problems. I will present the latest flow tracking techniques that exist in each of these realms: (i) the latest advances in tracking plasma flows in the solar atmosphere, and (ii) through a demo of how self-driving cars and drones efficiently measure optical flows, I will explain how they analyze the physics of their surroundings. Finally, I will show where are the interesting overlaps that can help us perfect our knowledge in both areas.

Deep Learning Applied to Satellite Data Processing

David Hall
dhall@nvidia.com
NVIDIA Corporation

NVIDIA and NOAA are collaborating to research the application of deep learning techniques for processing satellite observation data. With the recent launch of new satellites like the Geostationary Operational Environmental Satellite (GOES)-16, the data volume has increased by orders of magnitude and can be difficult to process in a timely manner using traditional methods. Deep learning has the potential to significantly improve both the processing speed and scientific accuracy of results. Through our research, we are using Convolution Neural Networks (CNN) to identify regions of interest (ROI) from satellite observations. These areas include cyclones, both tropical and extratropical, cyclogenesis, and eventually convection initiation. We are also exploring the use of conditional Generative Adversarial Networks (cGANs) to facilitation the translation of satellite observations to model variables for data assimilation in the Global Forecast System (GFS) model. This presentation will provide an introduction to our research efforts into the application of deep learning, the tradeoffs on computing and accuracy when designing neural networks, the challenges of data preparation and model training, and where we see these applications heading into the future.

Short **Talks**

Data-driven Modeling, Prediction and Predictability: The Complex Systems Framework

Surja Sharma
ssh@umd.edu
University of Maryland

The data-driven modeling of dynamics, spatial complexity and fluctuations in large-scale open systems in nature using observational data have yielded many significant advances. This approach is based on the dynamical systems theory which provides a framework with proper mathematical foundation for developing models that embody the features of the system inherent in the data, independent of modeling assumptions. This key feature is enabled by the recognition that the dynamics of nonlinear dissipative systems evolves in a phase space with a limited number of degrees of freedom, thus only a small number of variables are essential to model the dynamics. As demonstrated by the well-known Lorenz attractor the complexity of the system arises from the nature of the dynamical trajectories in the limited phase space. The embedding theorem of complex systems enables the modeling of such dynamical behavior from observational data and thus provide predictions. The early applications of this data-driven modeling provided many advances, including the first predictions of space weather, which used data from extensive ground-based and spacecraft-borne measurements of the coupled solar wind – magnetosphere system. Along with predictions, quantifying the predictability has been a long-standing challenge, in particular for extreme events. The recognition that an ensemble of similar initial states will undergo a spreading as they approach extreme situation has led to a new technique for prediction of extreme events. An ensemble transform Kalman filter (ETKF) technique developed for a dynamical system whose phase space is reconstructed from data, with unknown model equations, can be used to investigate the nature of the ensemble spread. In a study of extreme space weather, the ETKF was used to follow the trajectories of an ensemble of states in the phase space reconstructed from the time series data of the auroral elctrojet index AL. The ensemble spread was found to vary proportional to the AL intensity, thus providing a potential precursor to an extreme event. In weather and climate, there has been a growing realization of the need for prediction beyond weather, placing the predictability at intraseasonal and seasonal time scales as a key question that need to be resolved. This question was addressed using the extensive data of Indian monsoon rainfall to develop a dynamical model in the reconstructed phase space. An analysis of the evolution of many initial states showed predictability at intraseasonal time scale, and comparisons with global climate models lead to identification of ways to improve their predictive capability.

Automated stratigraphic mapping using Convolutional Neural Networks

Andrew Annex
annex@jhu.edu
Johns Hopkins University

Co-authors: Kevin W. Lewis

Meter scale sedimentary units in the Arabia Terra region of Mars [1] can be mapped using traditional GIS methods and HiRISE imagery at 1 meter GSD or finer resolution. These layered units appear in a variety of morphologies including buttes and stair-step formations [1]. To derive bed properties, sample lines are drawn in a GIS environment near but not over strata edges with curvature to perform OLS to fit 3d planes to data extracted from corresponding Digital Terrain Models using the sample lines [2]. However, scalability of human based efforts are a limiting factor in fully analyzing the available data. In private industry companies such as Descartes Labs Inc, Mapbox, and Planet Labs Inc advertise (and sometimes open source) the use of neural networks for scalable image classification and segmentation tasks using commercial Earth observing datasets and open source data sets like OpenStreetMap as label sources. Inspired by these approaches, we applied a convolutional neural network of the U-Net architecture [2] for image segmentation to automate mapping of the Arabia Terra layered units. Several models were trained experimenting with different network depths, cost functions, and configurations with and without: batch normalization, dropout, residual connections, and transposed convolution. The data set used for training was the 1 meter GSD orthoimagery grayscale products paired with binary mask representations of the vector GIS product of layers resulting in approximately 1000 image pairs. The data was always split into training and validation subsets, and image augmentation was employed to expand the training data set. Ultimately networks were found to perform well with 3 to 4 convolution and deconvolution steps requiring generally no more than 11 million parameters and the use of the Dice metric loss function to handle the unbalanced classes. Trained models were then used to produce a binary mask over the full spatial extent of a given HiRISE image, with some basic image processing techniques produced vectorized representations of the predicted sampling regions. Over a data set of 11 HiRISE images with a single computer, this process was able to fit bed properties to over 30,000 strata in less than 20 minutes before optimization. We found that the model was able to efficiently learn to segment layers across a variety of lighting conditions, and was able to simultaneously map both Butte formation and broader stair-step terrain. More work however will likely need to be attempted to reduce the amount of false negative mask regions in the training data and introduce more training data with true negative masks. Using a “off-the-shelf” neural network enabled the production a viable trained model quickly with modest resources, and furthermore enabled both a massive improvement in sample size across our study region.

Ref: 1. Edgett, Kenneth S., and Michael C. Malin. “Martian sedimentary rock stratigraphy: Outcrops and interbedded craters of northwest Sinus Meridiani and southwest Arabia Terra.” Geophysical Research Letters 29.24 (2002): 32-1; 2. Lewis, Kevin W., et al. “Quasi-periodic bedding in the sedimentary rock record of Mars.” science 322.5907 (2008): 1532-1535; 3. Ronneberger, Olaf, Philipp Fischer, and Thomas Brox. “U-net: Convolutional networks for biomedical image segmentation.” International Conference on Medical image computing and computer-assisted intervention. Springer, Cham, 2015.
Areas of interest: 2, 3 Presentation

Poster **Session**

(P-2) S. Ryan Alimo
sralimo@jpl.nasa.gov
Jet Propulsion Laboratory, California Institute of Technology

AI for AI: Efficient Hyperparameter Optimization Algorithm for Deep Learning

One important challenge in machine learning is determining how to modify a network architecture to better fit a new dataset. Although there are numerous models available for standard datasets, adopting such models for other datasets is not straightforward as it could likely be improved using hyper-parameter optimization. We propose to extend Delaunay-based Derivative-free Optimization via Global Surrogates (deltaDOGS) for automatically tuning deep neural network hyper-parameters in the case that our dataset is changing. deltaDOGS is a data-driven global optimization scheme which is provably convergent to the global minimum under appropriate assumptions for the unknown objective function. Additionally, deltaDOGS is designed for problems with computationally expensive objective functions that could take several hours to run and have an unknown analytic form and derivative. For the purpose of demonstration, we apply deltaDOGS to a Network Intrusion Detection System (NIDS) to find malicious activity. The user behavior analysis involves training an unsupervised network which considers that user behavior can change both from computer to computer and during a user's session. We show that by using deltaDOGS with our NIDS we can find suspicious network activity. Moreover, the results indicate that our approach is efficient and can therefore significantly save time and minimize damage caused by hackers.

(P-1) Shahrouz Ryan Alimo
sralimo@jpl.nasa.gov
Jet Propulsion Laboratory, California Institute of Technology

A Deep Learning Approach for Pose Estimation of Non-Cooperative Spacecraft

In order to support close-proximity operations such as rendezvous, final approach and docking, station-keeping, in missions of On-Orbit Servicing (OOS), Active Debris Removal (ADR) or Formation Flying (FF), several approaches can be adopted for onboard, real-time spacecraft pose determination, depending on whether the target spacecraft is cooperative, uncooperative, known or unknown. If the target spacecraft is uncooperative (i.e., not equipped with any communication link or artificial marker), and the chaser spacecraft has limited onboard power resources, a vision-based strategy can be adopted for the pose estimation of the target by making use of a low power passive electro-optical sensor (monocular or binocular cameras) onboard the chaser. The relative position and attitude of the target with respect to the camera can be estimated with onboard processing of the raw images of the target. However, determining the pose of a spacecraft through real-time image processing with limited computational power is a very challenging problem. Moreover, poor illumination conditions and high image contrasts can provide misleading features which further complicates accurate pose estimation. In this work, we propose a deep learning approach for vision-based pose estimation. We constructed a monocular

image pose estimation model for satellites by extending VGG19 initialized with the weights obtained by training on the ImageNet dataset. Rather than updating all of the weights in our model, only the layers appended to the VGG19 model were trained. This was done to reduce the training time without significantly reducing the accuracy under the assumption that low-level features are roughly equal in both images from ImageNet and images of satellites. Reducing the training time in this manner allowed us to quickly create and evaluate several similar models. Furthermore, we constructed both models that output continuous labels representing the pose, and the models that predicted the pose as one of many possible classes. The continuous models predicted seven values, three of which represented the translation vector, and four representing the rotation quaternion. The classification models also produced multiple outputs: one for the translation class and another for the rotation class. By splitting the classification problem into separate classification problems, we were able to have more instances per class and thus train our models more accurately.

(P-2) – Shahrouz Ryan Alimo
sralimo@jpl.nasa.gov
Jet Propulsion Laboratory, California Institute of Technology

Promises of Convex Optimization for Training Polynomial Neural Networks

The presence of spurious local minima and saddle points limits the application of local search algorithms in machine learning. In this work, we investigate the use of global optimization methods for training polynomial neural networks (PNNs). We formulate training tasks as polynomial optimization and employ a recently developed convex relaxation, namely parabolic relaxation, to solve the resulting non-convex problems. While the state-of-the-art methods for polynomial optimization such as cone programming relaxations are computationally prohibitive, we demonstrate promising theoretical results for parabolic relaxation. The proposed approach is guaranteed to produce a perfect fit as long as the given data-set is large enough. Moreover, we show that using parabolic relaxation, sufficiently small training errors can be reduced to zero if perfect fitting is feasible. As a proof of concept, we demonstrate our approach on illustrative examples of shallow PNNs.

(P-3) – Troy Ames
troy.j.ames@nasa.gov
NASA Goddard Space Flight Center

Deep Learning for Super-Resolution of Time Series Satellite Images

This presentation will explore how machine learning techniques, when applied to super-resolution (SR), can compensate for relatively low resolution sensors on resource constrained environments, such a SmallSats and CubeSats. Higher resolution satellite imaging data is often required to understand the science or process being observed. However, resource constraints (size, cost, power, limited transmission bandwidth, etc.) may limit the sensor capability. Since the spatial resolution needs will

continue to increase as quickly as new sensor technology, the trade-offs are not likely to change for future missions. The field of machine learning, in particular deep learning, has the potential to mitigate trade-offs and drive significant advances in the processing of all types of science data. Recently, deep learning techniques achieved promising results for producing super-resolution of hyper-spectral images; this presentation will focus on approaches for super-resolution of time series satellite hyper-spectral images and compare the results to previous methods.

(P-4) – Richard Barry
richard.k.barry@nasa.gov
NASA Goddard Space Flight Center

Rapid Machine lEarned Triage (RAMjET): a GPU code to extract and categorize photometric lightcurves from wide-field astronomical images

We describe a new approach, based on machine learning, to categorize astronomical time- ordered photometric measurements (lightcurves) as they are accumulated. We discuss a method to encode lightcurves into a form that is readily usable for the training of a support vector machine (SVM). Categorization efficiency and accuracy together with performance results in both a graphics processing unit (GPU) and multi-threaded environment will be contrasted. Finally, we discuss implications for the Transiting Exoplanet Survey Satellite (TESS), the Wide-Field Infrared Survey Satellite (WFIRST) and for ground observing programs such as the Microlensing Observations in Astrophysics (MOA) collaboration.

(P-5) – James Bednar
jbednar@anaconda.com
Anaconda, Inc.

EarthML.PyViz.org: Machine-learning workflows in Python

Many different computer languages can be used for artificial intelligence and machine learning. Python is a very popular choice because of the wide range of open-source ML/AI libraries available for it, its support for easily transforming data into a usable form, and its tools for big-data visualization and analysis. However, it can be difficult to determine the best way to approach any specific problem in such a general framework, and earth-science and climate data present specific difficulties that can make NASA-related data workflows awkward, error-prone, time consuming, and limited in data size and complexity. The EarthML project (http://earthml.pyviz.org) is a new joint initiative between NASA Goddard and Anaconda, Inc. to make it simpler to apply machine-learning and related techniques to satellite imagery, climate measurements, and other data sources used by NASA. EarthML consists of: 1. Best-practice examples of using Python libraries in complete data workflows including data preparation, visualization, and analysis (EarthML.PyViz.org and EarthSim.PyViz.org). 2. Improvements to underlying libraries for scalable visualization (PyViz.org), scalable data processing (Dask.PyData.org), and multidimensional data cataloging and retrieval (Intake.readthedocs.io and XArray.PyData.org). 3.

Documentation and training materials to make it easier to get started with each of these libraries. We will demonstrate these tools with examples of satellite image segmentation via spectral clustering, tagging data on maps as ML training examples, and predicting climate variables of interest from local or satellite measurements.

(P-23) Combined – Aniket Bera
ab@cs.unc.edu
The University of North Carolina at Chapel Hill, NC, USA
(Presented by Tanmay Vitthal Randhavane)

Real-time Anomaly Detection using Behavior Learning of Pedestrians

There has been a growing interest in developing computational methodologies for simulating and analyzing the movements and behaviors of crowds in real-world videos using artificial intelligence. We present an algorithm for real-time anomaly detection in low to medium density crowd videos. Our approach uses online methods to track each pedestrian and learn the trajectory-level behaviors for each agent by combining non-linear motion models and machine learning. We use AI techniques to compute the trajectory behavior feature for each pedestrian. These trajectory behavior features are used for anomaly detection in terms of pedestrian movement or behaviors. Our approach can be used for interactive surveillance and any crowd videos.

To capture the essence of pedestrian behavior in a crowd, we need to capture both the individual level movement features and also the dynamics of the group or cluster of which it is a part. Our approach computes global movement flows of pedestrians in semi-dense to dense settings. As a result, we compute clusters of pedestrians in a crowd based on their positions, velocity, inter-pedestrian distance, orientations, etc. We initially assign each pedestrian to a separate cluster, one consisting of a single pedestrian.

Modeling and classifying the behavior of different pedestrians using AI in a crowd is an important problem in various domains including psychology, robotics, pedestrian dynamics, and behavior learning. According to Convergence Theory [1], a well-known approach used in sociology and economics, crowd behavior is not a sole product of the crowd itself; rather, it is defined by the individual pedestrians in that crowd. As a result, it is important to accurately predict the behavior of individuals and their interactions with the environment to capture realistic, heterogeneous crowd behaviors.
We address the problem of classifying the behaviors of different pedestrians in a crowd video based on their movement patterns and use these patterns for crowd behavior prediction. We present a novel learning algorithm to classify pedestrian behaviors based on their movement patterns. We extract the trajectory of each pedestrian in a video and use a combination of machine learning and pedestrian dynamics techniques to compute the local and global characteristics at interactive rates. The local characteristics include the time-varying motion model that is used to compute the personality traits. We

also present new statistical algorithms to learn high-level characteristics and global movement patterns. Our approach extends the method presented in [5]. The global dynamics consist of factors that govern pedestrians' trajectory behaviors in a group or crowd, i.e., the factors that influence a pedestrian's overall movement or flow. We primarily use our trajectory-based behavior characteristics for local anomaly detection. Formally, we represent these dynamic characteristics for each pedestrian with a vector-valued function, whose initial value is determined by the motion function (more details in the paper). At every few steps, we compute the new behavior features for each pedestrian. We group similar features and find the most common behavior patterns, which correspond to the movement flow clusters. We use recently observed behavior features to learn the time-varying movement flow. We combine these characteristics with Eysenck's 3-factor PEN model [6] and characterize the personality into six weighted behavior classes: aggressive, assertive, shy, active, tense, and impulsive. We also use individual personalities to predict the state of the crowd under different environmental scenarios.

References:
[1] Ralph H. Turner and Lewis M. Killian. Collective Behavior. Prentice Hall, 1987
[2] J. Funge, X. Tu, and D. Terzopoulos. Cognitive modeling: knowledge, reasoning and planning for intelligent characters. In SIGGRAPH, pages 29–38. ACM Press, 1999
[3] Julio Godoy, Ioannis Karamouzas, Stephen J Guy, and Maria Gini. Moving in a crowd: Safe and efficient navigation among heterogeneous agents. In Proc. Int. Joint Conf. on Artificial Intelligence, 2016.
[4] F. Durupinar, N. Pelechano, J.M. Allbeck, U. Gu andd ¨ u andkbay, and N.I. Badler. How the ¨ocean personality model affects the perception of crowds. Computer Graphics and Applications, IEEE, 31(3):22 – 31, may-june 2011
[5] Sujeong Kim, Stephen J Guy, Dinesh Manocha, and Ming C Lin. Interactive simulation of dynamic crowd behaviors using general adaptation syndrome theory. In Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games, 2012.
[6] H.J. Eysenck and M.W. Eysenck. Personality and individual differences: A natural science approach. Plenum Press New York, 1985.
[7] L. Kratz and K. Nishino. Anomaly detection in extremely crowded scenes using spatio-temporal motion pattern models. In Computer Vision and Pattern Recognition. IEEE Conference on, pages 1446–1453. IEEE, 2009.

(P-24) – Aniket Bera
ab@cs.unc.edu
The University of North Carolina at Chapel Hill, NC, USA
(Presented by Tanmay Vitthal Randhavane)

Automatically Learning Driver Behaviors for Safe Autonomous Vehicle Navigation

We present an AI-driven autonomous driving planning algorithm that takes into account neighboring drivers' behaviors and achieves safer and more efficient navigation. Our approach leverages the advantages of a large-scale data-driven mapping that is used to characterize the behavior of other

drivers on the road. Our formulation also takes into account pedestrians and cyclists and uses psychology-based models to perform safe navigation. Our method includes a novel set of features that can be easily extracted from car trajectories. We derive a data-driven mapping between these features and six driver behaviors using an elaborate web-based user study. We also compute a summarized score indicating a level of awareness that is needed while driving next to other vehicles. Prior work in transportation research [1,2] often characterizes drivers using their levels of aggressiveness and carefulness. Several works in modeling pedestrian trajectories [3] and navigation [4] algorithms have applied psychological theory to capture human behavior. We conducted a comprehensive user survey to establish a mapping between five features and six different driving behaviors (Aggressive, Reckless, Threatening, Careful, Cautious, and Timid), involving 200 participants to identify driver behaviors from videos rendered from the Interstate 80 Freeway Dataset [5]. We perform least absolute shrinkage and selection operator (LASSO) analysis on six driving behaviors and four attention metrics which does regularization and feature selection by eliminating weak subsets of features. Our studies show that there could be one latent variable that is negatively correlated with aggressiveness and positively correlated with carefulness. We further verify these results by analyzing the correlation of the Principal Components with the amount of awareness that the users indicated they would pay to the targeted car. Additionally, we perform data augmentation to that the dataset has a sufficiently broad spectrum of driving behaviors corresponding to lane changes, fast-moving cars, passing cars, etc. To improve numerical stability during the regression analysis, the data linearly using the 5th and the 95th percentile samples to minimize the effects of extreme values. We enhance our existing autonomous driving algorithm [6] to navigate according to the neighboring drivers' behavior and demonstrate its benefits in terms of safer real-time navigation while driving next to aggressive or dangerous drivers. Our driving algorithm is based on a data-driven vehicle dynamics model and optimization-based maneuver planning, which generates a set of favorable trajectories from among a set of possible candidates, and performs selection among this set of trajectories using optimization. It can handle dynamic lane-changes and different traffic conditions. We demonstrate our benefits over previous methods: safer behavior in avoiding dangerous neighboring drivers, pedestrians and cyclists, and efficient navigation around careful drivers. With advancement in computer vision, one can expect more trajectory data in urban environments would be made available to the autonomous driving research community. In the future, we would like to apply our approach to analyzing and developing different AI driven navigation strategies that adapt to new situations and local driving styles.

Ref: [1] G. A. M. Meiring and H. C. Myburgh, "A review of intelligent driving style analysis systems and related artificial intelligence algorithms," Sensors, vol. 15, no. 12, pp. 30 653–30 682, 2015; [2] Z.-X. Feng, J. Liu, Y.-Y. Li, and W.-H. Zhang, "Selected model and sensitivity analysis of aggressive driving behavior," Zhongguo Gonglu Xuebao (China Journal of Highway and Transport), vol. 25, no. 2; [3] S. J. Guy, S. Kim, M. C. Lin, and D. Manocha, "Simulating heterogeneous crowd behaviors using personality trait theory," in Proceedings of the ACM SIGGRAPH/Eurographics symposium on computer animation, 2011, pp.43–52; [4] A. Bera, T. Randhavane, and D. Manocha, "Aggressive, tense, or shy? Identifying personality traits from crowd videos," in Proceedings of the Twenty Sixth International Joint Conference

on Artificial Intelligence, IJCAI-17, 2017; [5] J. Halkia and J. Colyar, "Interstate 80 freeway dataset,"Federal Highway Administration, U.S. Department of Transportation, 2006. [6] A. Best, S. Narang, L. Pasqualin, D. Barber, and D. Manocha, "AutonoVI: Autonomous vehicle planning with dynamic maneuvers and traffic constraints," arXiv:1703.08561, 2017.

Areas of interest: 4,3,1

(P-6) – Ricardo Martin Campos
riwave@gmail.com
Dept.of Atmospheric & Oceanic Science, University of Maryland

Nonlinear Wave Ensemble Averaging using Neural Networks

Co-Authors: Vladimir Krasnopolsky2, Jose-Henrique Alves3, Stephen G. Penny1,2;;1Dept.of Atmospheric & Oceanic Science / University of Maryland; 2EMC/NCEP / NOAA Center for Weather and Climate Prediction; 3SRG/EMC/NCEP / NOAA Center for Weather and Climate Prediction.

Artificial neural networks (ANNs) applied to nonlinear ensemble averaging are developed to improve wave forecasts. This is an approach that expands the conservative arithmetic ensemble mean (EM) from the NCEP Global Wave Ensemble Forecast System (GWES) to a nonlinear mapping that better captures the differences among the ensemble members and reduces the systematic and scatter errors of the forecasts. The goal is to improve the long-range predictability of significant wave height (Hs), peak wave period (Tp), and 10-m wind speed (U10). Several ANNs with different architectures and growing complexity have been tested and a detailed assessment of GWES has been performed. The first experiment was based on ANN training at two locations (Atlantic and Pacific Oceans) using NDBC buoys. A second experiment was conducted in the Gulf of Mexico, using a spatial approach. Results show that a small number of neurons are sufficient to reduce the bias, while 35 to 50 neurons are optimum to reduce both the scatter and systematic errors. The correlation coefficient for forecast Day 10 was increased from 0.39 to 0.61 for U10, from 0.50 to 0.76 for Hs, and from 0.38 to 0.63 for Tp. We are currently running the third experiment using altimeter data to train the ANN for the whole globe.

(P-28) – Kamal Choudhary
kamal.choudary@nist.gov
National Institute of Standards and Technology, Gaithersburg, MD, USA
(Presented by Francesca Tavazza)

Machine learning with force-field inspired descriptors for materials: fast screening and mapping energy landscape

We present a complete set of chemo-structural descriptors to significantly extend the applicability of machine-learning (ML) in material screening and mapping energy landscape for multicomponent systems. These new descriptors allow differentiating between structural prototypes, which is not possible using

the commonly used chemical-only descriptors. Specifically, we demonstrate that the combination of pairwise radial, nearest neighbor, bond-angle, dihedral-angle and core-charge distributions plays an important role in predicting formation energies, bandgaps, static refractive indices, magnetic properties, and modulus of elasticity for 3D materials as well as exfoliation energies of two-dimensional (2D) layered materials. The training data consists of 24549 bulk and 616 monolayer materials taken from JARVIS-DFT database. We obtained very accurate ML models using gradient boosting algorithm. Then we use the trained models to discover exfoliable 2D-layered materials satisfying specific property requirements. Additionally, we integrate our formation energy ML model with a genetic algorithm for structure search to verify if the ML model reproduces the DFT convex hull. This verification establishes a more stringent evaluation metric for the ML model than what commonly used in data sciences. Our learnt model is publicly available on the web (https://www.ctcms.nist.gov/jarvisml) property predictions of generalized materials.

(P-7) – Daniel da Silva
daniel.e.dasilva@nasa.gov
NASA Goddard Space Flight Center
and Trident Vantage Systems, JHU

Classifying Spacecraft Magnetospheric Region with Ambient Plasma Sensing and Support Vector Machines

Link to poster: https://docs.google.com/presentation/d/1awfDS26G3dqe8jG6j3m_35IXNiSaYB086hIgnks07rA/edit?usp=sharing Abstract Earth's magnetosphere can be partitioned into separate regions with unique properties. The Magnetospheric Multiscale Mission (MMS) houses an in-situ plasma sensor, the Fast Plasma Instrument (FPI) in high earth orbit. Flight calibration of this instrument utilizes data from different regions in provide a dynamic set of calibration inputs. Calibration can be automated if we can have the ability classify flight region between: magnetosphere, magnetosheath, and solar wind. In this poster we outline, explain, and interpret a support vector machine model that successful classifies between these regions using ambient plasma sensor data. As a bonus, we verify the weights of the trained model agree with prior knowledge of the physics.

(P-8) – Richard Garnett
garnetri@mcmaster.ca
McMaster University

GammaNet: a Convolution Neural Network for Event Classification on Space-borne Time Projection Chambers

The Advanced Energetic Pair Telescope (AdEPT) mission uses a gaseous Time Projection Chamber to fully characterize gamma-ray induced pair production events in low earth orbit. The data rate for AdEPT is expected to be on the order of 16 Gbps due to the high resolution and imaging rate of the TPC, 400 x

400 x 400 μm3 over an 8 m3 active volume and 20 images per second. The rate of background events is expected to be 104 times higher than pair production events, mostly due to cosmic ray protons. These two parameters set the requirement for an on board background rejection technique. The use of an anti-coincidence detector for background rejection is precluded by the imaging rate and the event rate of cosmic ray protons, which would result in near 100% dead time. This produces the need for a new event classification technique. GammaNet proposes to be the solution to this problem by using machine vision techniques to classify images produced by the AdEPT TPC in real time on board the satellite. Machine vision is a subset of artificial intelligence that uses convolutional neural networks, which need to be trained on large sets of images that have been classified beforehand. To generate this image set for GammaNet, a simulation of the AdEPT TPC was developed in the Geant4 Monte Carlo toolkit. The simulation is constructed with a 253 x 253 x 253 mm3 volume of Argon and CS2 gas at 1.5 atm to simulate a reduced AdEPT TPC volume. For each event in the simulation the ionization electrons generated in the gas are projected onto the XZ and YZ planes, which are used for classification in GammaNet. These two projections are used to reduce the probability that electron and positron tracks overlap in both images used by GammaNet. The simulation images are generated in two sets, positive and negative. The positive set contains images with 1-2 pair production events overlaid on 1-5 cosmic ray proton events for background. The negative set contains only the 1-5 cosmic ray proton events. Electronic noise was added to each image from a normal distribution with a standard deviation of 2 and mean of 0. The pair production events in the simulation were generated from mono-energetic photons from 5-250 MeV, and the cosmic ray protons were generated randomly with energies following the spectrum expected at a 550 km orbit. Both particle sources were generated on a spherical shell sampled inward with isotropy. From training on these simulation images, it was found that GammaNet can obtain a 99.6% background rejection rate while maintaining 40-50% sensitivity to pair production events over the energy range simulated.

Area of interest: 4 – Computer Vision and Image Processing.

(P-9) – Evana Gizzi
Evana13G@gmail.com
Tufts University

Machine Learning for Intelligent Collision Avoidance

Conjunction Assessment is an area of study that looks into ways evaluate and handle situations where satellites may be in risk of colliding. It is important that an accurate determination is made, as collisions are costly and dangerous. The main method for collision avoidance is to calculate a probability of collision (Pc) value and a corresponding threshold, and chose to maneuver in situations where the Pc value is above that threshold. We investigated using alternative parameters in a machine learning construct to help inform the collision avoidance decision. Specifically, we investigated how Fuzzy Inference Systems (FIS) could process those parameters into collision decision output, in either a standalone way, or in a way that could enhance the Pc decision. Fuzzy Inference Systems are able to capture partial membership of variables into different sets, and generate output that exists in a continuous space,

versus discrete classifications, which is characteristic of classic logical systems and of many traditional machine learning constructs. We hypothesized that this would be beneficial in resolving the collision decision ambiguities that exist at parameter sets that are within the neighborhood of one another. We first used k-means clustering and SVM classification methods to determine whether the parameters correlated to ground truth classifications, and found a higher and more promising performance metric in SVM classification. We looked at the output's from both trained and untrained FIS's using 5 parameters (Miss Distance, Mahalanobis Distance, Bhattacharyya Distance, Kullback-Leibler Distance, and Orbit Angle), and found that for any FIS with n inputs, an n-dimensional decision space is generated as a "best fit" to the data used in training, or the functionality defined within untrainable systems. By virtue of the "best fit" nature, the decision space is not able separate data into classes which are defined based on feature extraction. Thus, we concluded that use of FIS's may not be the best construct for classification, but rather, is better suited for model fitting problems. In future work, we would like to explore how the system behavior would change with using a Monte Carlo Pc value as a ground truth, and also how a Neural Network may help inform and/or reinforce our findings.

(P-10) – Mircea Grecu
mircea.grecu-1@nasa.gov
Goddard Earth Sciences Technology and Research (GESTAR), Morgan State University and NASA Goddard Space Flight Center

k-NN machine learning methodology to improve the precipitation estimates derived from combined TRMM and GPM radar and radiometer observations

The complex variability of Particle Size Distributions (PSDs) is one of the main sources of uncertainties in the precipitation estimates derived from satellite radar observations such as those collected by the NASA Tropical Rainfall Measuring Mission (TRMM) and the Global Precipitation Measurement Mission (GPM). To mitigate PSD uncertainties, advanced precipitation estimation algorithms make use of the fact that satellite radar observations are attenuated due to precipitation and that the Path Integrated Attenuation (PIA) can be estimated from the analysis of the surface return. Multiple studies showed that the PIA is crucial not only in the attenuation correction process, but also in the quantification of the precipitation rates associated with the attenuation corrected radar observations. One unique and remarkable feature of the GPM mission is the Dual-Frequency Precipitation Radar (DPR). The DPR enables the derivation of PIA estimates significantly more accurate than those derived from single frequency satellite radar observations. To extend the benefit of accurate GPM dual-frequency PIA estimates to TRMM and GPM single frequency radar observations, a k-Nearest Neighbors (k-NN) approach is developed to estimate the PIA from coincident radiometer observations. Specifically, a large dataset of over ocean GPM Microwave Imager (GMI) radiometer observations and high-quality near-nadir dual-frequency PIA estimates is constructed and used to develop a k-NN estimation procedure. It is found that PIA estimates derived from the GMI observations using the k-NN procedure are significantly more accurate than those derived from single frequency radar observations. It is therefore possible to use the k-NN procedure in the derivation of PIA estimates in the outer portion of the DPR swath where only single frequency radar

observations are available. The developed procedure is also applicable to TRMM Microwave Imager (observations) for the derivation of PIA estimates usable in the TRMM Precipitation Radar (PR) algorithm. Category 1, i.e. general Machine Learning.

(P-11) – Bob Harberts
robert.harberts@nasa.gov
Global Science and Technology (GST)

Machine learning in the science data enterprise

Data intensive science coupled with rapid technological change restlessly prompt challenges to manage unprecedented volumes of data efficiently, swiftly, and within budgets. "Big data" and "Cloud" technology represent examples of contemporary transformations embraced by data management organizations. Additional challenges linked with increasing rates of growth in science data acquisition, data product generations, and access by fast-paced electronically connected user communities have applicable NASA science enterprise contexts.

The NASA science enterprise represents a context of contexts ranging from missions to spacecraft, instruments to operations, data processing and management to data distribution and use. Contextualization helps sort out unique challenges as well as shared opportunities for advances in machine intelligence technology. Accomplishing this in a dynamic distributed environment of scientists, observatories, multiple sources of data, and complex infrastructures poses a significant challenge to data management operations. Clearer contexts for new kinds of smart automation are essential.

Smarter automation incorporating machine-learning techniques promises to upgrade system capabilities with a transforming effect for data management. Astutely developed and embedded machine–learning functions and services such as telemetry analysis, log history understanding, science data understanding, data transformations, and self-management describes a technological path of increasingly sophisticated automation.

Contextualized research identifies realistic opportunities, specifications, and integration pathways for embedding smart automation across the science data enterprise. A first research step characterizes current implementations and current planned evolution for selected contexts (e.g. spacecraft telemetry, operations, science data production, archiving, and distribution) to contrast with the next step of relating appropriate smart automation and machine learning techniques. The third step assembles a contextualized architecture for embedded machine learning capabilities that enable smart feedback analysis, autonomous applications, self-improving operational capabilities, imminent fault prediction and root-cause analysis, and algorithms for collaborative processes expected of a more adaptive interactive digital science enterprise.

(P-12) – Yuning He
yuning.he@nasa.gov
NASA Ames Research Center

Statistical Learning and Modeling of Safety Boundaries for AI-based Aerospace Applications

Complex aerospace systems operating in a changing and potentially unknown environment are subject to operational envelopes and safety boundaries, which separate safe from unsafe behavior. For example, the operational flight envelope is governed by numerous parameters and determines if the aircraft is operating safely or getting into a dangerous stall condition. In traditional systems, such envelopes and boundaries have been analyzed and are present in manual and trainings to help guide the human operator or pilot. However, the task of always safely operating the system requires many years of human experience. Cyberphysical systems that are to be operated autonomously under the auspices of an Artificial Intelligence (AI) are expected to do so safely even in changing, unknown, or hostile environments. Typical examples include unmanned aerial systems (UAS) and autonomous cars. For safe operations, it must be ensured that the AI or the adaptive controller keeps the system always within the safety boundaries. Such boundaries must be modeled and characterized in order to be able provide safety guarantees during Verification and Validation (V&V) and system analysis. “The UAS must keep a minimal horizontal distance of 100ft to the restricted zone” is such a typical property, which needs to be asserted during V&V or continously checked during its operations. High dimensional state spaces and non-linear systems make a direct boundary analysis impossible because brute-force testing does not scale for such complex parameter spaces. In addition, boundaries are time-varying due to the system dynamics, interactions with the environment and adapttation through learning. Time-varying boundaries add yet another layer of complexity to the analysis process rendering traditional testing impractical. We present a statistical technique for the modeling of safety boundaries in high-dimensional state spaces. Active learning based upon a boundary-aware metric allows us to substantially reduce the number of required data points to model the boundaries. We use DynaTrees, a particle-filter based algorithm for the efficient handling of high dimensional dynamic spaces and time series. For proper analysis and review by the human expert, safety boundaries must be characterized and presented in way suitable for human understanding. Our tool chracterizes safety boundaries as parameterized geometric shapes like planes, polygons, or spheres and uses a Bayesian approach to automatically estimate optimal shape parameters. We will illustrate our approach with selected case studies: Even a simple adaptive controller like the NASA Intelligent Flight Control System (IFCS) features non-linear boundaries that change over time as the adaptation progresses. A complex aircraft collision avoidance system like ACAS-X provides suggestions for evasive maneuvers to the on-board AI or autopilot. We present results of experiments on modeling time-series safety boundaries that characterize regions where dangerous near mid-air collisions can happen.

(P-13) – Mikel Holcomb
mikel.holcomb@mail.wvu.edu
West Virginia University

Using Machine Learning and other Analytics to Optimize Material Properties

Our long-term goal is to develop, test and improve algorithms for the optimization of several material properties (including magnetization, absorption and conductivity) of different materials. For example, magnets have a variety of current space applications, from satellites to landing devices (including magnetic torquers for orientation with Earth's magnetic field, magnetic motors to generate power which can be less susceptible to power outages, and magnetic spectrometers to study antimatter and dark matter) and more ideas have been proposed that may improve spacecraft launch, assist formation flying, and speed up long space travel times. Additionally, the enhancement of various properties can lead to better energy efficiency, lower the limit of device size and reduce the necessary amount of material needed for a device to function properly (saving time and manufacturing/development/material costs/launch weight). The material LaxSr1-xMnO3 (LSMO) is an excellent model and our starting system, which has been proposed for use in many applications, such as spintronics, magnetic tunnel junctions for computing and sensing and solid oxide fuel cells. The opportunity and challenge in this and similar materials is that its properties are affected by adjusting even a single parameter (e.g. the amount of oxygen) out of all parameters available (including the growth temperature, termination and smoothness of the starting material, etc.). Within the LSMO system (as we can also observe in other similar model systems), we have strongly correlated properties, i.e. slight tweaks in one property can have major effects on material properties—a defining feature in devices. What is needed is a more holistic approach, in which all of these properties are measured and compared either statistically, or empirically by an expert. Our approach may allow researchers to determine the dominating factors and system limitations which may lead to enhancements in terms of magnetization strength or other desired properties, which stands in the way of current use of this system and others. Not only could this project revolutionize the progress and scientific understanding of LSMO, but the process could also be generalized to any material system. The work is expected to ultimately illuminate not only the physics and chemistry behind the factors that contribute to stronger properties, but also improve the engineering strategies to capitalize on the properties of other complex materials. We are working to develop a relational network of sample parameters to magnetic and other properties and a prediction tool for improved magnetization. Most current machine learning approaches in materials physics focus on theoretical calculations, whereas we are working from experiments. We grow and characterize samples from predictions to enhance the variability and predictability of the database analysis to build upon and analyze the work already performed on over 100 of these samples, partially sponsored by the NASA WV Space Grant Consortium, NSF (DMR-1608656) and DOE (DE-SC0016176), along with including work from publications. We use mathematical modeling and operations research methodologies to find the optimal values for input parameters with the objective of maximizing the desired material characteristics.

(P-15) – Vincent Houston
vincent.e.houston@nasa.gov
NASA Langley Research Center

Machine Learning Algorithms To Improve Model Accuracy and Latency, and Human-Autonomy Teaming

Vincent E. Houston NASA Langley Research Center, Hampton, VA, 23681 Bryan A. Barrows2 Virginia Polytechnic Institute & University, Blacksburg, VA 24061 Walter J. Manuel3 Purdue University, West Lafayette, IN 47907 Lisa R. Le Vie4 NASA Langley Research Center, Hampton, VA, 23681

Improving the approach to the development of algorithms that allow machine agents to act independently within a circumscribed set of goals has garnered more traction, over recent years, than just a phase of observation, but has progressed rapidly into the necessary stage of small to moderate scale implementation. This is evident when operating in those areas that were once considered for human tasking only, particularly the area of aviation. To introduce an autonomous agents within the cockpit of an aircraft remains a very detailed and complex endeavor; one that requires the arduous verification and validation of the agent's ability to perform its task in a complex connected national airspace system (NAS), where in 2025 daily aircraft operations are predicted to number 95,000 twice that recorded in 2016. In addition, these machine agents are required to work in concert with humans and other manned aircraft. This is accomplished with the agent(s) achieving a task(s) that help reduce pilot workload, without the degradation of aviation safety, reliability, costs, or operational efficiency. Therefore, it is vital that the proper performance of any cockpit related task given to this burgeoning technology be vetted thoroughly to facilitate trust between human and machine, and help further human autonomy teaming (HAT). As a step toward these goals, a modified increasingly autonomous system (IAS) has been developed that optimizes the identification, distribution, and rendering of "relevant" air traffic onto the ownship's navigation display for use in a NAS Net-Centric environment. This paper describes the second phase of our work in this effort and the necessary V&V to prove out its ability as an autonomous cockpit agent. I. Nomenclature ASTRAO = Autonomous System Technologies for Resilient Airspace Operations D.A.T.A = Dynamic Air Traffic Application H.A.T = Human Autonomy Teaming IAS = Increasingly Autonomous Systems MATIMAL = Machine Learning Algorithms To Improve Model Accuracy and Latency ML = Machine Learning MVL = Multi-View Learning NextGen = Next Generation Air Transportation System NAS = National Airspace System SASO = Safe Autonomous Systems Operations TDM = Traffic Data Manager V&V = Verification and Validation II. Introduction Air traffic operations within the NAS have continued to increase at an annual rate of 3.8 percent since 2009, which translates into ≈ 96,000 operations daily by 2025.1 This statistic alone represents a challenge for aviation and begs the question: How can flight safety be maintained with the influx of aircraft operations on the rise at this pace? NASA's Safe Autonomous Systems Operations (SASO) project has as one of its tasks, the development of autonomous system technology that would allow human pilots the ability to navigate within this, anticipated, heavily populated airspace by formulating an IAS that identifies relevant air traffic with respect to an ownship's positioning. This concept is fundamental to an 'intelligent party-line' where the human can be kept informed properly, without being bombarded by information, during computer-to-computer interactions

in a NextGen Autonomic Airspace Architecture.5 SASO's initial salvo in this technological direction took place during Phase One of the Autonomous System Technologies for Resilient Airspace Operations (ASTRAO) research.2 Phase Two of ASTRAO, Machine Learning Algorithms To Improve Model Accuracy and Latency (MATIMAL), is focused on implementing algorithm enhancements to the first phase machine learning used in the Traffic Data Manager (TDM), as well as implementing a robust way of validating and verifying the prediction accuracy of TDM and its enhanced algorithms, while fostering HAT. This last element is a fundamental technical challenge in creating autonomous cockpit agents that have the required levels of functionality and trust of human participants in this environment. Phase One ML algorithms, using state data from surrounding air traffic, offered a relevancy prediction for each traffic element with respect to the position of an ownship. There were several algorithms tested during the ASTRAO TDM prototype phase, incorporating standard supervised learning techniques such as Bootstrap-Aggregated Decision Trees, K-Nearest Neighbor, Naïve Bayes, and Support Vector Machine. MATLAB's TreeBagger ensemble method demonstrated the most promise in predicting the relevancy of air traffic. It provided the TDM prediction model in the categories of not-relevant, maybe-relevant, and relevant air traffic with accuracies of 88.7%, 48.5%, and 46.7% respectively, and an overall accuracy of 61.3%.2 Although the prediction model from the prototype phase of ASTRAO produced promising results, they were skewed heavily toward non-relevant air traffic. It was determined that this outcome was due to the fact that the subject matter experts (SMEs) had labeled the majority of the air traffic within the traffic scenarios presented as non-relevant; it was the SMEs' relevancy determinations that biased the TDM prediction model during training in producing inferences on air traffic relevance. The consensus direction for Phase Two work was twofold: develop more sophisticated ensemble method algorithms than were used previously and develop a verification and validation process that would incorporate more of a model relevancy prediction confirmation by the SMEs. III. Results and Discussion Algorithm improvements for Phase Two focused on a Multi-View learning (MVL) approach toward data analysis and machine learning. MVL addresses issues related to the high dimensionality of data; in the case of ASTRAO Phase One, it was non-relevant air traffic. Utilizing only the data labels/features that are pertinent, MVL views a data set as a sub-table of training data with respect to a subset of the data labels/features. The elimination of irrelevant data features performed by MVL is useful toward the reduction in dimensions and addressing the issue of having insufficient training examples than even a reduced set of pertinent features. During MATIMAL, the performance marker highlights are classification accuracy, disagreement among the classifiers, latency, and the reduction of attributes/features.3 The V&V of TDM's prediction model, during MATIMAL, took on a different approach than what was instituted during the first phase of ASTRAO. To generate and promote a symbiotic HAT environment, it was decided that the system (the human pilots) would provide a confirmation of the model (the TDM MVL) air traffic relevancy predictions. Given a set of aircraft state data, the MVL algorithm within TDM would render its inferences as to relevant air traffic with respect to ownship position, for which the SMEs would either confirm or deny the model's prediction. It was concluded that this approach would facilitate a more unified effort between human and machine rather than the competitive one that grew from the comparison approach. More importantly, the model has the ability to learn immediately from the human's confirmation of relevance during this V&V process. During Phase Two testing, various classifiers were explored, one-

vs-one binary classification, one-vs-all binary classification, ensemble learning, and multi-view learning. MVL had the greatest relevant true prediction accuracy, with 88% correct relevant predictions, and 78% non-relevant prediction accuracy. These results show a significant improvement over the TreeBagger algorithm utilized during the Phase One development of TDM. The binary classifiers did not meet the desired accuracies and correctness for prediction. Similarly, the trained ensemble classifier, without varying feature subsets, provided notably lower accuracies and did not generalize well for each of the ensemble models. A challenge that presented itself during both phases of TDM development was ensuring that the algorithms did not show signs of over-fitting data and therefore generating excessive false positive predictions. MVL allowed the model to be adjusted such that the false positives were minimized, and provided insight for each algorithms' performance level according to the supplied data labels/features.4 IV. Proposed Paper In the proposed paper, details of the ML algorithm development will be presented. The objective of MATIMAL is to improve model prediction accuracy for classifying relevant data and to generate HAT through V&V. In addition, a high priority is to produce an improved definition and understanding of the "maybe relevant" data prediction boundaries.2 This paper will describe our work in data subsampling to assess the impact of equal size datasets for each of the relevancy classes. The Phase One K-NN, Naïve Bayes, Support Vector Machines, and TreeBagger were applied using the equal-size data bins for each class. Although the models resulted in having more balanced prediction accuracies than the first generation models, they did not meet our minimum threshold requirements for use. After further exploration of the data through feature selection and feature analysis, it was found that model combinations of the classifiers tended to perform greater than all of the elementary standalone models.2 4 In the proposed paper, the fundamental technical challenge in creating autonomous cockpit agents and replicating the expertise of human pilots is explored by V&V and examined as it pertains to levels of functionality and trust of human participants using autonomous cockpit agents. IV. References [1] FAA Aerospace Forecast Fiscal Years 2009 – 2025 https://www.faa.gov/data_research/aviation/aerospace_forecasts/2009-2025/media/2009%20Forecast%20Doc.pdf [2] Houston, V, E., Le Vie, L, R., "Autonomous System Technologies for Resilient Airspace Operations", June 2017 [3] Kumar, V., Minz, S., "Multi-view ensemble learning: an optimal feature set partitioning for high-dimensional data classification", September 2015 [4] Barrow, B, A., Manuel, W, J., Houston, V, E., Le Vie, L, R., "Development of Multi-View Ensemble Machine Learning for Predicting Relevant Airspace Traffic" [5] NASA Aeronautics Airspace Operations and Safety Program Discussion with the ARTR http://sites.nationalacademies.org/cs/groups/depssite/documents/webpage/deps_152342.pdf

(P-16) – Bert Huang
bhuang@vt.edu
Virginia Tech

Integrating Machine Learning to Improve Optimal Estimation of Atmospheric Composition

Bert Huang (bhuang@vt.edu), Dept. of Computer Science, Virginia Tech;
Chidubem Arachie (achid17@vt.edu), Dept. of Computer Science, Virginia Tech;
Elena Spinei (eslind@vt.edu), Space @ VT, Bradley Dept. of Electrical and Computer Engineering, Virginia Tech;
Natalya Kramarova (Natalya.a.kramarova@nasa.gov), NASA GSFC;
Kris Wargan (Krzysztof.wargan-1@nasa.gov), SSAI/NASA GSFC.

Machine learning is a maturing technology that can improve many scientific analyses. However, it is important to integrate machine learning into scientific analysis in ways that build on existing knowledge, rather than tasking machine learning with rediscovering what is already established. The optimal estimation technique is commonly used in remote sensing to retrieve vertical profiles of atmospheric constitutes. This technique requires a priori information to stabilize the retrieval process. In cases when the measurements provide only limited information, accuracy of the a priori profiles can affect the retrieved quantity. As an example, measurements from ground-based Brewer spectrometers and multi-axis DOAS measurements have limited information about the vertical shape of ozone profiles, and accurate a priori information is needed. Often seasonal ozone climatology is used for this purpose. However, ozone vertical distribution can significantly change within shorter time scales, and seasonal climatology is unable to reflect those changes. In this study, we explore a new approach for predicting ozone profiles using machine learning, improving a critical component in the overall remote sensing process. We train a neural network model with a set of more than 10,000 assimilated ozone profiles and corresponding meteorological parameters. The Modern-Era Retrospective analysis for Research and Applications version 2 (MERRA-2) assimilates satellite limb observations from NASA's Aura Microwave Limb Sounder (MLS) and Ozone Mapping Instrument (OMI). Using MERRA-2 data, we train the neural network to predict vertical ozone profiles over The NASA Goddard Space Flight Center from inputs of meteorological readings and temporal indicators. The neural network uses a convolutional architecture to capture correlations among ozone concentrations in nearby elevations. In preliminary experiments, we train the neural network using MERRA-2 data from January 2005 to December 2008, with eight snapshots spread throughout each day, and we evaluate its ability to predict ozone profiles for all of 2011 and 2012. The resulting model is able to predict more refined ozone profiles than seasonal climatology, reflecting the effects of time-of-day and the increased ozone variability in winter. These results are promising indicators of potential benefits of embedding modern machine learning methods into atmospheric remote sensing. Areas of interest: General Machine Learning (1) to improve the outcomes of science modeling. Presentation preference: short oral presentation

(P-17) – Kendall Johnson
kjohns21@masonlive.gmu.edu
George Mason University, CCMC, and NSBE

R2D2's little Brother

I hope to present a demonstration of topic 1 on general machine learning and how Goddard can utilize its properties while designing satellites and rovers and their long-term operation in space. I will be presenting the robot S.P.O.V. (Solar powered operational vehicle), that I built myself, with the basic functions of motion, detecting distance, detecting light, and the ability to use solar energy, but instead of being discreetly programed to do those tasks we will use machine learning to train the robot to learning its own way of adapting to its environment with the programed extremes. The robot's main components will be an Arduino Uno, motors, ultrasonic sensor, solar panel, LCD display, and photo-resistors. Machine learning is used with these particular sensors because of a satellite's necessity to do the same thing due to a satellite's constant need to avoid space debris and move toward better lite areas to charge itself appropriately. A satellite's ability to essentially automate itself to move using this low grade Artificial Intelligent may soon become standard to its bus and to NASA for many new upcoming challenges and reasons. The property of greater success and longevity for NASA missions are crucial for Goddard and to all of NASA, and the success rate of a machine in space to control itself without the help of man will be the next lap of the space race. Machine learning with the photo-resistors will be used as a light sensor for the robot to interpret the strength and direction of the incoming light and precede toward it. This robot will be trained many times under conditions that allow it to create its own working neural net so that it is able to complete its purpose of avoiding objects and going toward light. I plan to demonstrate training tests to visually show the progress of the coded algorithm go from running into walls to precise eloquent movement, or simply just accomplishing the task of avoiding obstacles and going to where it is most brightly lite to gather solar energy. I will use the allotted 10 mins for oral demonstration as I explain the physically represented properties of machine learning and program the robot as it accomplishes the following tasks, tests, and goals. 1) Show the "learning" aspect of the machine as its fails and succeeds while compiling data. I will likely be using q-learning for the code. We will determine the exercise by the number of times the robot fails at a task before it is successful. The next number that follows that first success, that would be most important to NASA, and is how many times it is successful before it is unsuccessful again. 2) Demonstrate the ability of the newly learned machine with variables and obstacles. - Show that the robot can now follow a light source - Show that the robot can avoid obstacles 3) Show machine learned code vs Standard coded movement - Program the robot with standard obstacle avoidance code then show and discuss the difference.

(P-18) – Erdem Karakoylu
erdem.m.karakoylu@nasa.gov
NASA Goddard Space Flight Center; Ocean Biology Processing Group

Comparing Probabilistic Programming and Neural Network Based Approaches in Information Retrieval from Ocean Colour

For several decades, chlorophyll-a (Chl-a) and inherent optical properties (IOP), key metrics of marine ecological and biogeochemical processes, have been derived from ocean colour using empirical or semi-analytical methods with reasonable accuracy across a wide variety of water types. However, accurate retrieval of these parameters in optically complex waters, where atmospheric correction is challenging, remains elusive. We argue that it is time to modernize our approach to information retrieval from ocean colour in these challenging scenarios. Instead of trying to correct for the atmosphere by estimating and subtracting it, we propose to use signals that include the atmosphere as input to more sophisticated techniques. The increased availability and accessibility of a variety of powerful machine learning approaches makes such an exploration feasible. Here, we compare and contrast two different machine learning paradigms used in research performed as part of the first NASA PACE (Phytoplankton, Aerosol, Cloud and Ocean Ecosystem) Science Team activities to demonstrate the feasibility of deriving Chl-a and IOPs from ocean colour. One approach we tested is neural networks (NN). NN are relatively easy to implement and generally work well as universal approximators. However, their black-box nature makes interpretation difficult and may not generalize well. The other approach we investigated belongs under the umbrella of probabilistic programming (PP). This framework is more challenging to implement and remains computationally intensive, though recent advances in algorithm development have made PP much more applicable to a wider range of problems. And while PP does not always scale well to big datasets its advantages makes it an attractive proposition for small to medium size datasets. Moreover, larger datasets can be tackled with some compromises on inference quality. Advantages of PP include uncertainty estimation ‘out of the box’; PP models are generative, making them amenable to simulations before and after data collection; the ability to include quantified background knowledge in the form of priors, which reduces the probability of overfitting. We specifically focus on comparing machine learning pipeline design, including data preparation and transformation, model performance evaluation and validation, and model selection. While the models, trained and validated both on a synthetic dataset and an in-situ-to-satellite matchup dataset, performed reasonably well with r2, mean percent difference and model-to-true slope values nearing in some instances 0.9, 27% and ~1, respectively, we show that the approach chosen is ultimately determined by dataset characteristics, how the resulting model is to be used post-fitting, and how critical the quantification of uncertainty is, e.g. as part of a subsequent decision process.

Area of Interest 1.

(P-19) – Benjamin Marchant
benjamin.marchant@nasa.gov
Universities Space Research Association (USRA)

Gaussian processes for satellite remote sensing classification: Application to MODIS cloud thermodynamic phase

Discriminating liquid and ice cloud pixel from passive sensors is an important task, since those two type of clouds have very different radiative properties. To continue improving the SW-NIR MODIS cloud phase classification algorithm, new machine learning methods have been tested such as Gaussian process methods. One of the advantage of Gaussian process is to provide uncertainty quantification. To train the model, a dataset has been developed based on collocated data between MODIS and the lidar CALIOP. First results and the pros and cons of Gaussian processes will be discussed. Oral or poster

(P-20) – Daniel Miller
Daniel.J.Miller@nasa.gov
NASA GSFC Ocean Ecology Lab (616)

Ongoing neural network algorithm development for polarimetric observations of clouds and aerosols above clouds

We have developed a multi-layer neural network (NN) based algorithm for the retrieval of cloud microphysical properties (cloud optical depth, cloud droplet effective radius and variance) from airborne multi-angle polarimetric measurements. This feed forward back-propagation multi-layer perceptron network is developed and applied to data from the airborne Research Scanning Polarimeter (RSP) instrument. RSP measures both polarized and total reflectance in the spectral range of 410 to 2260 nm and scans along the flight track obtaining ~150 viewing zenith angles spanning -60˚ to 60”. The neural network architecture training, and input parameters were developed using a synthetic training set and informed by an information content analysis (Segal-Rozenhaimer, 2018 in process). In this study, we present further development of the algorithm, including the correction for output parameters based on comparisons to existing retrieval algorithms using data from the ObseRvations of Aerosols above CLouds and their intEractionS (ORACLES) 2016 and 2017 NASA field campaign. We will also discuss our approach toward extending the lessons from this NN exercise to retrievals of more complicated remote sensing systems - such as simultaneous retrieval of aerosol and cloud properties.

(P-21) – Jaideep Pathak
jpathak@umd.edu
University of Maryland, College Park

Machine Learning Techniques for Analysis of High-Dimensional Chaotic Spatiotemporal Dynamics

We demonstrate the effectiveness of machine learning for analysis and prediction of spatiotemporal

chaotic dynamical systems from data. Using a computationally efficient recurrent neural network called an echo state network or reservoir computer, we show that we can reconstruct the attractor of very high dimensional chaotic dynamical systems with unprecedented fidelity. This reconstruction allows us to determine the ergodic properties (e.g., the spectrum of Lyapunov exponents) of a dynamical system purely from data [1]. We also develop and introduce a computationally parallelized extension of a prediction scheme based on reservoir computing that allows us to obtain model-free predictions of spatiotemporal chaotic flows of arbitrarily large spatial extent and attractor dimension[2]. We obtain outstanding results using machine learning for these difficult tasks where previous methods are either unfeasible or have had limited success. We demonstrate the scalability and computational efficiency of our approach using a toy model often used for testing techniques in weather prediction (Lorenz 1996) and the spatiotemporally chaotic Kuramoto-Sivashinsky partial differential equation. Further, we develop a hybrid approach for forecasting spatiotemporal chaotic systems [3] that combines machine learning with an approximate knowledge-based model. Such a hybrid technique is able to accurately predict for a much longer period of time than either its machine-learning component or its knowledge-based component alone and requires less training data than a purely data driven, machine learning approach. This research could have wide applicability to tasks such as improving geophysical weather prediction models, better algorithms for data assimilation as well as prediction of space weather events such as solar flares.

Ref: [1] Pathak et al. Chaos 27, 121102 (2017); [2] Pathak et al. Phys. Rev. Lett. 120, 024102; [3] Pathak et al. Chaos 28, 041101 (2018).

(P-22) – Eric Pollack
eric.z.pollack@nasa.gov
NASA Goddard Space Flight Center/Wallops Flight Facility

Intelligent Collaborative Constellation

With the turn of the decade, the scope of satellite experimentation is expanding beyond what a single spacecraft can provide. Scientific instruments carry the weight of a mission's purpose, and the vehicle on which they conduct their observation provides context for scientific data generated. The classical interpretation of a satellite mission is restrictive in that a singular spacecraft must handle all aspects of the operation. Missions may require several spacecraft conducting parallel observations using varied instrumentation, some of which may be entirely decoupled from the rest of those in the mission cluster, yet reliant on each other's data to perform science. A monolithic approach may have sufficed to realize the majority of scientific goals in the past, but as orbital and interstellar experimentation vehicles begin to evolve in their roles to take on more complex mission assignments, scientists have shifted toward the pursuit of more ambitious goals, many of which demand the applicative use of Distributive Artificial Intelligence (DAI). A variety of missions calling for the utilization of intelligent swarm-configuration constellations has been conceptualized and proposed, but the complete realization of such concepts has yet to be achieved. A consolidated peer-to-peer communication framework is a prerequisite for the

viability of missions of this classification, and this is the challenge Intelligent Collaborative Constellations (ICC) seeks to address. The context in which collaborative missions are coordinated presents a litany of challenges in the formulation of the overarching mechanisms which would govern the operation of a DAI network. The way in which ICC facilitates peer-to-peer communication is tailored to account for the challenges of space with respect to concurrent data transmission and to the benefit of autonomy. Delivery of shared data and commands from one originating node to specific target nodes for the purpose of collaborative scientific data processing is a core responsibility of ICC. ICC decouples the operation of the network from the relative position of one satellite in a host constellation to another by autonomously discovering the constellation's network topology and propagating data through the network until it reaches its target. The host network's topology is maintained on each node as a series of known reference points from prior transmissions and is updated via the same propagation techniques used to pass telemetry through the network. The topology of the network is utilized to predict the shortest path between two nodes and make dynamic adjustments in accordance with its real-time performance while minimizing power consumption and optimizing transmission latency. This is significant, as the desire for minimal power consumption dictates nodes cannot constantly broadcast updates to their location despite the network's intrinsically dynamic landscape. Not only does this approach account for a potentially transient location in any particular member of that constellation, but it also mitigates the overall susceptibility for network failure due to a malfunction or operational conflict. Furthermore, ICC accounts for one or more satellites in the computed shortest path being unavailable or out of range, duplicated transmissions, planetary eclipse, volume of data being transmitted, and interference from other transmissions.

(P-25) – Maryam Rahnemoonfar
maryam.rahnemoonfar@tamucc.edu
Texas A&M University-Corpus Christi

Intelligent Solutions for Navigating Big Data from the Arctic and Antarctic

Significant resources have been spent in collecting and storing large and heterogeneous datasets collected during expensive Arctic and Antarctic fieldwork. While traditional analyses provide some insight, the complexity, scale, and multidisciplinary nature of the data necessitate advanced intelligent solutions. In recent years, the research community has witnessed advances in artificial intelligence (AI). Recent advances in deep neural networks (DNNs) and massive datasets have facilitated progress in AI tasks such as image classification, object detection, scene recognition, semantic segmentation, and natural language processing. Despite this progress, these algorithms are limited to electro-optical data with large labeled datasets. There is a critical need to develop more advanced machine learning and deep learning algorithms for both visual and non-visual sensors collecting data for real-world scenarios during various polar ice missions. For polar ice applications, data-driven algorithms that can be utilized by domain-experts are increasingly valuable. In this presentation, different AI approaches for detecting ice surface, bottom, and internal layers from Radar imagery will be discussed. Ice loss in Greenland

and Antarctica has accelerated in recent decades. Melting polar ice sheets and mountain glaciers have considerable influence on sea level rise (SLR) and ocean currents; potential floods in coastal regions could put millions of people around the world at risk. The Intergovernmental Panel on Climate Change (IPCC) estimates that sea level could increase by 26–98cm by the end of this century. This large range in predicted SLR can be partially attributed to an incomplete understanding of bed topography and basal conditions in fast-flowing regions of ice sheets in Greenland and Antarctica. Therefore, precise calculation of ice thickness is very important for sea level and flood monitoring. The shape of the landscape hidden beneath the thick ice sheets is a key factor in predicting ice flow and future contribution to SLR in response to a changing climate. Recent large-scale radar surveys of Greenland and Antarctica reveal internal ice layers on a continental scale. Our large-scale dataset enables accurate detection and tracing of these internal layers to illuminate many aspects of ice sheet dynamics, including their history and their response to climate and subglacial forcing. Therefore, it is important to develop fully automatic artificial intelligence techniques for this big data to detect ice surface, internal layers, and sub-glacial topography hidden beneath the thick ice sheets. To the best of our knowledge there is not any fully automatic technique to detect ice internal layers from Radar imagery. In this work, we developed a multi-scale architecture in combination with the wavelet network to detect the internal annual snow accumulation boundaries from large dataset on NASA IceBridge mission. Our multi-scale network takes advantage of different levels of contextual information. Our experimental results in comparison with the ground-truth and state-of the art-results show the effectiveness of our approach.

(Lobby) – Sharad Sharma
ssharma@bowiestate.edu
Bowie State University

Indoor building evacuation application for emergency response using HoloLens

Early hands-on experiences with the Microsoft Hololens augmented/mixed reality device have given promising results for building evacuation applications. A range of use cases are tested, including data visualization and immersive data spaces, in-situ visualization of 3D models and full scale architectural form visualization. Our Hololens application gives a visual representation of a building on campus in 3D space, allowing people to see where exits are in the building. It also gives path to the various exits; shortest path to the exist as well as directions to a safe zone. Our proposed AR application was developed in Unity 3D for Microsoft HoloLens. It is a fast and robust marker detection technique inspired by the use of Vuforia AR library. The application offers users an enhanced evacuation experience by offering enthralling visuals, helping them learn the evacuation path they could use during an emergency situation where evacuation is necessary. The goal of this project is to enhance the evacuation process by ensuring that all building patrons know all of the building exits and how to get to them, which would improve evacuation time and eradicate the injuries and fatalities occurring during indoor crises such as building fires and active shooter events. We have incorporated existing features in the building as markers for the HoloLens application to trigger the floor plan and subsequent location of the person in

the building. This work also describes the system architecture as well as the design and implementation of this AR application to leverage the Microsoft HoloLens for building evacuation purposes. Pilot studies were conducted with the system showing its partial success and demonstrated the effectiveness of the application in an emergency evacuation. Our results also indicate that majority of participants felt that HoloLens application can be used as a substitute for evacuation plans (2D plan) in a building. Usually, the evacuation plans are displayed as a 2D plans in the buildings. Sometimes it becomes difficult for users to visualize a building through a 2D plan. The use of AR application gives the user the flexibility and ability to visualize the building and exits in a 3D space. We believe that AR technologies like HoloLens could be adopted by people for building evacuating during emergencies as it offers enriched experience in navigating large-scale environments.

(P-26) – Steven Slojkowski
Steven.E.Slojkowski@nasa.gov
Flight Dynamics Facility/Omitron Goddard Space Flight Center

Applications of Machine Learning to Flight Dynamics Operations

Steven Slojkowski, Joel Getchius, Adam Michaels, Daniel Mattern, Arda Aslan, Toni Deboeck, Ryan Patterson, Juan Crenshaw, Travis Schrift, Douglas Ward

The Flight Dynamics Facility (FDF) of the Goddard Space Flight Center (GSFC) performs orbit determination, tracking data evaluation, antenna calibration, launch support, and pre-launch mission analysis for a fleet of missions in low-earth, geosynchronous, lunar, and deep-space orbit regimes. We foresee numerous potential applications of Machine Learning to the flight dynamics domain. Machine Learning is applicable to tracking data anomaly detection and classification, orbit determination quality assurance, orbit estimation filter tuning, optical navigation, novel methods of orbit determination, intelligent and rapid data analysis, and space weather monitoring and prediction, among other possibilities. In this presentation, we introduce our facility and activities. We talk about the data we collect and maintain, and the ways we are seeking to apply machine learning to our data. We introduce our current effort to apply machine learning to tracking data evaluation and present some preliminary results. Supervised methods like Decision Trees, Random Forests, and simple Neural Networks are easily trained to identify “first-order” anomalies like measurement biases, based on classifications assigned by analysts using currently employed tracking data evaluation methods. We show some results in this area. Our effort also seeks to employ unsupervised learning methods to identify deeper patterns in the data, with the hope of identifying new anomaly classes beyond the categories currently employed by human analysts in their routine work. These include more subtle “second-order” anomalies such as time-tag errors, bi-modal residuals, cycle slips, un-modeled or poorly-modeled maneuvers, and station geodetics errors. We show examples of some of these cases and discuss attempts to identify and classify such anomalies, as well as some results from application of unsupervised learning methods to identifying natural clusters of tracking data residuals. The goal of this tracking data analysis effort is a system that works alongside analysts to offer suggestions and assistance for classifying and troubleshooting

tracking data anomalies. In the future, as the system becomes more competent and users gain trust in it, it can be given more responsibility for classifying tracking data, perhaps taking over as the prime analyst with task members reviewing its decisions or stepping in only when it is unable to make a clear judgement. We expect this type of system to have a number of advantages: • More informative and diverse anomaly criteria, • Consistent anomaly reporting, eliminating possible differences in analysts' experience and judgement, • The ability to place confidence limits on classification decisions, which reduces false positives and allows the facility to report only those anomalies for which we have high confidence, • Facilitates consideration of a larger number of metrics when classifying data, • Hopefully reduces the manual labor involved in classifying tracking data. This effort has only recently begun in the FDF and we face a number of challenges going forward: • Our personnel have shallow expertise with machine learning and are confronted with a vast array of literature and methods, • We have a large amount of data that is not yet structured or labeled ideally for machine learning, • The potential feature space is large and needs to be reduced to the most relevant and important metrics, • An unclear path towards data normalization or standardization for methods that require it, • Analyzing and classifying diverse observation types simultaneously.

Categories 5 or 6.

(P-27) – Connor Sprague
connor.sprague@noaa.gov
The Aerospace Corporation

Assured Operations of the Environmental Information Enterprise

U.S. agencies are accomplishing their missions to protect lives and property using complex systems comprised of hardware, software, networks, and human-machine interfaces. Automation of these systems is increasingly required to transfer information between enterprise system components at the speed of need. NOAA environmental systems generate data at rates of >20TB per day that are processed through the data chain and delivered to provide decision support to people serving in roles from air traffic controllers to emergency managers. State-of-health and state-of-performance are currently conducted on select systems within the Environmental Enterprise Value Chain using a range of tools, techniques, processes, and people. There is value in a holistic, enterprise-wide approach to verify performance of systems along the value chain, as well as impacts of degraded performance at any point in the system of systems to the overall mission of protecting lives and property. This paper will discuss prototyping efforts which intake large volume and velocity environmental data streams and observatory telemetry and perform descriptive and prescriptive analytics, while identifying and presenting anomalous behavior using a configurable dashboard. Having learned nominal data stream behaviors, the system detects and predicts future anomalies and downstream impacts. The system operates on three scales: • Real-Time Information Assurance – Monitor spacecraft and ground system telemetry to perform real-time diagnostics of mission performance through the value chain. Uses data analytics to assess anomalies and trends for proactive enterprise system operations management. Perform corrective action to prevent

disruptive degradation or outages impacting the operational mission. • Operational Impact Assessment – Leverage forecast sensitivity to observation systems impacts (FSOI) using data analytics and machine learning to assess and monitor quality of products in response to detected anomalies. • Long-Term Enterprise Performance – Improve NOAA observation systems resilience over time by assessing potential impacts on enterprise integrity to meet operational mission to provide warnings, watches and other environmental information. Using a common system for three scales opens the possibility for exploring and applying mitigation techniques related to anomalies. This paper will discuss secondary data sources and algorithms that hold promise in providing gracefully degraded forecasts when primary data is missing or corrupted. Preliminary work on the application of data analytical techniques to tropical storm track and intensity prediction indicates that machine learning techniques can produce useful predictions of future storm directions based on comparisons with past tracks and environmental conditions, and use a smaller set of inputs with lower likelihood for data outages/corruption compared to typical physical models for storm track. The situation with respect to storm intensity is even more promising, as statistical techniques are competitive with physical models for accuracy; we will discuss some new potential data sources that have become practical with the advent of high-volume data processing.

(P-29) – Barbara Thompson
barbara.j.thompson@nasa.gov
NASA GSFC

Driving Scientific Discovery with Deep Learning and AI at the NASA GSFC Center for HelioAnalytics

What is HelioAnalytics? This is a broad term meant to cover all the ways that we harness advanced statistics, informatics and computer science methods to achieve our science. Our focus is on problems that we can attack with modern methods that we cannot attack otherwise. A keener understanding of how information is derived from data, and how machine learning can be harnessed to accomplish this, will expand the discovery potential for key heliophysics research topics and missions. We report on a new program to integrate modern information science, statistics, and scientific knowledge to advance the fundamental physics of connected sun-heliosphere-geospace system. The Center for HelioAnalytics is an “expert group” at NASA GSFC focusing on topics such as machine learning, neural networks, and data analytics in order to expand the discovery potential for key heliophysics research topics and missions. We define “HelioAnalytics” as a hybrid of Heliophysics + Machine Learning + Statistics + Information Design. Each of these are fields that are well developed in their own right; HelioAnalytics is the cross-disciplinary convergence of communities of physicists, statisticians, and computer scientists. HelioAnalytics is intended to foster research into advanced methodologies for heliophysical research, and to promulgate such methods into the broader community.

(P-30) – Chenxi Wang
chenxi.wang@nasa.gov
ESSIC/GSFC

A new machine learning based cloud phase discrimination algorithm designed for passive infrared satellite sensors

Clouds play critical roles in the Earth's energy budget due to their large coverage and strong radiative effect. Among various important cloud properties, cloud thermodynamic phase is a critical midway to link cloud microphysical properties with cloud optical and radiative properties. In this study, we developed a novel cloud thermodynamic phase algorithm based on a Random Forest (RF) classifier. The training (75%) and validation (25%) datasets are generated using a 3-year coincidental MODIS (onboard both Aqua and Terra) and CATS (onboard the ISS) observations. The "true" cloud thermodynamic phases are provided by a lidar on CATS and inputs are solely from MODIS thermal infrared (IR) observations and surface temperature from reanalysis. An independent 1-year cloud phase dataset from CALIPSO/CloudSat is also used for validation purpose. Our preliminary results show that the RF-based phase algorithm performs much better than current MODIS MOD06 IR Phase 1km product. In the near future, we intend to apply a similar RF-based phase algorithm to daytime clouds with a more complete spectral information (e.g., using IR and shortwave observations).

(P-31) – Jianwu Wang
jianwu@umbc.edu
University of Maryland, Baltimore County

Spatiotemporal Climate Data Causality Analytics: An Analysis of ENSOís Global Impacts

El Niño and the Southern Oscillation (ENSO) is a local phenomenon of the variation in sea surface temperature (SST) and surface air pressure across the equatorial eastern Pacific Ocean. Numerous studies have indicated that ENSO could have determinant impacts on remote weather and climate through the atmospheric "teleconnection" using the conventional correlation-based methods, which however cannot identify cause-and-effect of such linkage and ultimately determine a direction of causality. Lagged linear regression is frequently used to infer causality between climate variables. This method has weaknesses when one or more of the variables have high memory or autocorrelation. Granger causality method, which consists of a lagged autoregression and a lagged multiple linear regression, is suitable to determine the causality relations with high memory data. We use the Vector Auto-Regressive (VAR) model estimation method to find the Granger causality relations between ENSO and some climate variables (land surface temperature, sea level pressure, precipitation and wind). We also use the climate model simulation to double confirm the causality relations between ENSO and climate variables from the observation-based analyses. This work is carried out in order to determine the spatiotemporal causality relationships between ENSO and abnormal events in remote regions, and provide some valuable insights for the prediction of some extreme weather/climate events under different ENSO backgrounds. The observational data we use are the Hadley Centre Sea Ice and Sea Surface Temperature data (HadiSST: 1948-present), the Global Precipitation Climate Project Precipitation (GPCP) data (1979-present) and the NCEP/NCAR reanalysis

data (1948-present). Sensitivity simulations with global climate model forced by different SST patterns are carried out to investigate the responses of atmospheric fields. We use the Community Atmospheric Model (version 5.3, CAM5.3) with the CAM5 standard parameterization schemes. Three sensitivity simulations include: the control run forced by climatological SST; the +2K run forced by climatological SST + 2K at the Nino3.4 region; the -2K run is forced by climatological SST – 2K at the Nino3.4 region. We compare the simulated wind, SLP, precipitation and temperature fields from three simulations. We determine the causality relation between ENSO and land surface air temperature on the global scale using the VAR method. Our experiments show ENSO is a driver of surface temperature anomalies in remote regions such as South America, northwest North America, equatorial South Africa, and northern Australia; while ENSO variation is not caused by surface temperature over land. The global distribution of the maximum lag correlation between ENSO index and surface temperature that ENSO has strong positive relationship with surface temperature in South America and equatorial South Africa. Results of the climate model sensitivity simulations are consistent with the observational-based analyses. In the ENSO warm phase, there are positive anomalies in surface temperature over South America, northwest North America while in the ENSO cold phase there are negative anomalies in surface temperature over these regions. Further analyses with the ENSO and other climate variables (precipitation, SLP, wind) also show consistent conclusions between the observational-based and model-based results.

(P-32) – Jake Wilson
s18jwilson@severnschool.com
Intern at NASA Goddard Space Flight Center

Machine Learning on Solar Eruptive Events: Determining Key Physical Properties from Multi-Instrumental Observations

Coronal mass ejections (CMEs) often leave behind a dark region on the sun known as a coronal dimming. The two largest catalogues of coronal dimmings are the DEMON and JEDI catalogues. JEDI and DEMON use two different methods of detecting dimmings, using two different instruments on the Solar Dynamics Observatory. This project consisted of two parts: 1) Perform an assessment of the preliminary JEDI catalogue results (Version 1.0) to determine detection accuracy and complementary with DEMON (which has already been validated). 2) Understand how the different populations of coronal dimmings as observed in the catalogues can be used to predict future dimming populations. This is done through machine learning applications such as neural networks. We make recommendations for improving the JEDI catalog, such as refining the algorithm of well-documented multi-wavelength events. These recommendations that are being made as our assessment of the JEDI catalog showed that it did not accurately reproduce the results from papers it is based on. Currently, JEDI is triggered by GOES flare times, but only half of DEMON dimmings have a corresponding flare. If we alter JEDI to trigger of the times used by DEMON, we can more understand more about detecting non-flare dimmings using JEDI. We use machine learning to understand the behavior of dimming as detected by DEMON. We use the information derived from DEMON to train a basic neural network that can predict the intensity

of dimming on the sun, with a validation score of 70%. With this neural network, we hope that it can be applied to coronal mass ejections on other stars.
Area of Interest: 2, NN and Deep Learning(DL) Preference: A poster would be preferred over a short oral presentation.

(P-33) – Soni Yatheendradas
soni.yatheendradas@nasa.gov
Earth System Science Interdisciplinary Center, University of Maryland, College Park, MD 20742, USA and Hydrological Sciences Laboratory, NASA Goddard Space Flight Center, Greenbelt, MD 20771

Co-Author: Sujay Kumar Hydrological Sciences Laboratory, NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

Developing fine-scale snow cover fraction estimates using deep learning

Evaluating local effects of snow on land hydrology can be difficult with coarse resolution maps of snow cover fraction (SCF). We use Deep Learning to recreate the fine-resolution target 1-km version of the original MOD10A1 SCF from coarser 5-km MOD10C1 SCF and auxiliary 1-km data. Note that while the MOD10A1 SCF values are continuous, the original MOD10C1 SCF and cloud cover fraction (CCF) layers were respectively created by upscaling from the binary masks of snow and cloud presence. The MOD10C1 thus has complementary information layers of the SCF, CCF and the Confidence Index that must be combined to derive the 1-km target MOD10A1 values. These factors considered together make this problem a mixture of downscaling and regression. In addition, our study considers the utility of different auxiliary data types: static ones like terrain and land cover, and dynamic including satellite products like MOD10A1 snow albedo and MOD11A1 land surface temperature (LST) and outputs from a land surface model. Our initial try used the auxiliary data of dynamic satellite products like MOD10A1 snow albedo and MOD11A1. Training data of 1-3 years for the deep network showed that the percentage of usable data can go to a low of upto about a third to two-thirds of all available data points, depending on whether the MOD11A1 Quality Control flag value is 0 for good-quality data and 1 for unreliable-quality data. For incorporating the surrounding spatial influence typically done in a Convolutional Neural Network (CNN) via a filter, this brought into focus the requirement of either using a partial convolution, or a spatial infilling using partial convolution for obtaining valid values everywhere so. Hence we focused our initial CNN architectures on not considering the influence of spatially surrounding pixels, by using a simple 1X1 spatial filter to consider only valid values. We implemented a 3-layer CNN with this 1X1 filter in Keras/Tensorflow and a custom root mean square error (RMSE) loss function that calculates the SCF RMSE at only the valid-data pixels. To increase the network depth without losing the MOD10C1 information over large number of layers, we also experimented with Residual Networks (ResNets). Our ResNet combines CNN stacks on two paths between the input and target layers: a longer residual path that uses only auxiliary information, and a main path that combines information from the complementary layers of MOD10C1. However, these networks could attain only ~12% SCF RMSE instead of our desired

5% RMSE. To pinpoint the problem, we then implemented a serial stack architecture of regression (5-km) and downscaling (5-to-1 km) components, where the intermediate target layer is a simple spatial averaging of the 1-km SCF. The downscaling component was seen to have minimal RMSE (1.5%), indicating a need to better capture the regression component. In this presentation, we will describe the deep learning infrastructure as well as the key results from this snow cover application. The developed infrastructure is expected to facilitate future machine learning prototypes for a variety of terrestrial modeling applications.

Area of Interest: 2: NN and Deep Learning (DL).

(P-34) – Kiley Yeakel
kiley.yeakel@jhuapl.edu
Johns Hopkins Applied Physics Lab

Automatic Detection of Saturn's Magnetospheric Regions using Cassini Data

Co-Authors: Jon Vandegriff, Don Mitchell, JHU/APL; Doug C. Hamilton, University of Maryland; Caitriona Jackman, University of Southampton; Peter Delamere, University of Alaska Elias Roussos, Max Planck Institute for Solar System Research Artificial intelligence (AI) continues to show promise in the area of space exploration, from early detection of spacecraft systems anomalies to reduction of downlinked data. One key area of focus as we face data sets that are increasingly growing in size and complexity is automated detection of scientifically-relevant events. Much of the current mission framework relies on having a scientist-in-the-loop (SITL) to parse out interesting snippets of data – often a tedious task bordering impossible when facing massive datasets. However, AI has shown promise in the pattern detection techniques SITLs use to parse data streams and thus could be a potential outsource. Here, we present one such example from the Cassini-Huygens mission in detection of Kronian magnetospheric regions. In post-processing of the mission data, scientists labeled various regions near Saturn, viz., magnetosphere, magnetosheath, and solar wind, using data from the Cassini Plasma Spectrometer (CAPS) and magnetometer (MAG) sensors. Simultaneous to the CAPS measurements are near-continuous measurements from the Magnetospheric Imaging Instrument (MIMI). Both MIMI and CAPS provide measurements of in situ ions and electrons, with CAPS covering the lower end of the energy spectrum and MIMI focused on the higher energy levels. MIMI provides measurements via a pairing of instruments – the charge-energy-mass spectrometer (CHEMS) and low energy magnetospheric measurements system (LEMMS). Given that He++ can be assumed to be sourced solely from solar wind, and H2+ and water group ions (O+, OH+, H2O+, and H3O+) are sourced solely from Saturn and its moons, MIMI data can provide an indication of which region the spacecraft is currently residing in. Likewise, measurements of the magnetic field strength and direction from the magnetometer can inform the region identification based on deviations from a modeled dipole. Within the magnetosphere, the magnetic field can be expected to behave much like a modeled dipole but suffers significant variation within the magnetopause and solar wind. Thus, both MAG and MIMI data provide independent

measures of the various plasma regions. Using the SITL-derived region destinations as our "label" and overlapping MIMI, MAG, and region labels as our "training set," we attempt to classify regions of the mission timeseries for the entirety of the mission life. Preliminary results indicate that flux ratios of particular pairings of ion species may be the best indicator of magnetospheric regions.

(P-35) – Tianle Yuan
tianle.yuan@nasa.gov
Goddard Space Flight Center and University of Maryland

Application of a Deep U-Net to Automatic Detection of Ship-Tracks Multispectral Images from both Polar-Orbiting and Geostationary Satellites

Ship-tracks are an iconic demonstration of aerosol-cloud interactions. They have been observed in satellite images for more than 60 years. Many studies have resulted from analyses of remote sensing, in-situ, and modeling results. Satellite data are instrumental in discovering and understanding of ship-tracks. However, it is hugely time consuming to identify ship-tracks within satellite multispectral image with manual processing. Whole careers have been spent on manually picking out and analyzing ship-tracks. Here we apply a deep (16-layer) U-Net architecture to train a neural network algorithm to automatically find ship-tracks in the satellite data. The training set includes hundreds of ship-tracks with varying lengths from MODIS data. Native MODIS images come at 2030pixels by 1350pixels, which are way too large for deep NN to train. We adopt a strategy of breaking up native images into small (64X64 and 128X128) clips and train the NN on these sub-images and recombine them once the training is done. The risk is whether the structural information in the large images can be fully recovered with training only on smaller regional scales. We experimented with various setting of the model and training methods to test if this break-and-recombine strategy is viable . Our results show that this strategy works excellent. The recombined ship-track detection algorithm can not only pick out objects that is in the training data, but on many occasions find objects that are not noticed by human eyes. With this algorithm, we can now automatically find ship-tracks at a global scale during both daytime and nighttime. Initial test run on a few months of MODIS data, on the order of terabytes, show that our algorithm can map out ship-tracks occurrences with high accuracy. The occurrence frequency map agrees with similar maps from the commercial shipping industry very nicely. This unprecedented capability of finding ship-tracks at a global scale enables us to analyze radiance data at a large scale and to analyze ship-track properties with sample sizes of orders magnitude larger than what is currently available. For example, the current most comprehensive dataset for ship-tracks includes less than 10,000 samples. Our test run already finds more than a million. More importantly, our algorithm works directly on radiance data from multispectral instruments. This means we can apply it to geostationary as well as older polar orbiting sensors like AVHRR, which expands the possibility of studying ship-tracks. For example, we are applying it to the new GOES-16 and -17 data to study the temporal evolution of individual ship-tracks. We can also combine the detection algorithm with multiple satellite sensors to dive deeper into the physical processes that are responsible for ship-tracks. Initial results are encouraging. We characterize the global distribution of ship-tracks and its variability on monthly to seasonal time scales. A significant

portion (up to 5%) of the low clouds are affected by instantaneous ship-tracks. We also obtain results regarding the distribution of ship-track properties such as their length, width, and overall size (in terms of number of pixels) for different regions and seasons. Based on millions of ship-tracks we are able to find, it becomes possible to start comprehensively addressing the long-standing science question: how significant are ship-tracks in modulating low cloud properties.

(P-36) – Tianle Yuan
tianle.yuan@nasa.gov
Goddard Space Flight Center and UMBC

Unsupervised classification of low cloud organization using multispectral satellite data and a deep neural network model

Low cloud modeling in current general circulation models (GCMs) has been a challenge for more than three decades. How low cloud would change in the context of climate change, often described as low cloud feedback, contributes the largest uncertainty in future climate projections. It is a pressing issue that needs to be resolved before uncertainty in future climate projection can be significantly lowered. We attack this issue from the observation side to understand low cloud organization using multispectral satellite images. Low clouds can organize themselves into many forms at a scale on the order of 100km, the so-called mesoscale organizations, and different organizations have vastly different impacts on cloud radiative effect both at the surface and the top of the atmosphere. Mesoscale organization forms can either evolve smoothly or switch dramatically in a short time. Neither GCMs nor high resolution regional models can model these constantly occurring organization changes. Here we develop a deep neural network model and train it with MODIS reflectance data in an unsupervised fashion. This is the first step in our approach and the trained model serves as a feature extractor. The trained model extracts 2048 features from the raw reflectance data. A principal component analysis is then applied to the resulting feature vectors to reduce the dimension of the data. Finally, we use a clustering algorithm to classify the PCA processed, lower dimension data. Experimenting with clustering algorithm parameters we settle on a set of parameters that give physically meaningful classifications. With this, we obtain an unsupervised algorithm that can classify low cloud mesoscale organization into physically meaningful categories. The algorithm is based purely on reflectance data and does not require absolute calibration or physical retrievals of cloud parameters. This algorithm can be applied to data from similar sensors such as AVHRR and GOES-series. Resulting cloud organization classification is a powerful tool for scientists to better understanding low cloud behavior using observations.

Poster **Session**

Location	Last Name	Poster Title
P-1	Alimo	A Deep Learning Approach for Pose Estimation of non-cooperative spacecraft
P-2	Alimo	AI for AI: Efficient Hyperparameter Optimization Algorithm for Deep Learning AND Promises of Convex Optimization for Training Polynomial Neural Networks
P-3	Ames	Deep Learning for Super-Resolution of Time Series Satellite Images
P-4	Barry	A new machine learning based cloud phase discrimination algorithm designed for passive infrared satellite sensors
P-5	Bednar	EarthML.PyViz.org: Machine-learning workflows in Python
P-6	Campos	Nonlinear Wave Ensemble Averaging using Neural Networks
P-7	da Silva	Classifying Spacecraft Magnetospheric Region with Ambient Plasma Sensing and Support Vector Machines
P-8	Garnett	GammaNet: a Convolution Neural Network for Event Classification on Space-borne Time Projection Chambers
P-9	Gizzi	Machine Learning for Intelligent Collision Avoidance
P-10	Grecu	k-NN machine learning methodology to improve the precipitation estimates derived from combined TRMM and GPM radar and radiometer observations
P-11	Harberts	Intelligent Archive (IA)
P-12	He	Statistical Learning and Modeling of Safety Boundaries for AI-based Aerospace Applications
P-13	Holcomb	Using Machine Learning and other Analytics to Optimize Material Properties
P-15	Houston	Machine Learning Algorithms To Improve Model Accuracy and Latency, and Human-Autonomy Teaming
P-16	Huang	Integrating Machine Learning to Improve Optimal Estimation of Atmospheric Composition
P-17	Johnson	R2D2's little Brother (final)
P-18	Karakoylu	Comparing Probabilistic Programming and Neural Network Based Approaches in Information Retrieval from Ocean Colour
P-19	Marchant	Gaussian processes for satellite remote sensing classification: Application to MODIS cloud thermodynamic phase
P-20	Miller	Ongoing neural network algorithm development for polarimetric observations of clouds and aerosols above clouds
P-21	Pathak	Machine Learning Techniques for Analysis of High-Dimensional Chaotic Spatiotemporal Dynamics
P-22	Pollack	Intelligent Collaborative Constellation
P-23	Randhavane	Real-time Anomaly Detection using Behavior Learning of Pedestrians
P-24	Randhavane	Automatically Learning Driver Behaviors for Safe Autonomous Vehicle Navigation
P-25	Rahnemoonfar	Intelligent Solutions for Navigating Big Data from the Arctic and Antarctic
P-26	Slojkowski	Applications of Machine Learning to Flight Dynamics Operations
P-27	Sprague	Assured Operations of the Environmental Information Enterprise
P-28	Tavazza	Machine learning with force-field inspired descriptors for materials: fast screening and mapping energy landscape
P-29	Thompson	Driving Scientific Discovery with Deep Learning and AI at the NASA GSFC Center for HelioAnalytics
P-30	Wang	A new machine learning based cloud phase discrimination algorithm designed for passive infrared satellite sensors
P-31	Wang	Spatiotemporal Climate Data Causality Analytics: An Analysis of ENSO's Global Impacts
P-32	Wilson	Machine Learning on Solar Eruptive Events: Determining Key Physical Properties from Multi-Instrumental Observations
P-33	Yatheendradas	Developing fine-scale snow cover fraction estimates using deep learning
P-34	Yeakel	Automatic Detection of Saturn's Magnetospheric Regions using Cassini Data
P-35	Yuan	Application of a deep U-Net to automatic detection of ship-tracks multispectral images from both polar-orbiting and geostationary satellites
P-36	Yuan	Unsupervised classification of low cloud organization using multispectral satellite data and a deep neural network model
Lobby	Sharma	Indoor building evacuation application for emergency response using HoloLens
Lobby	Grubb	Science and Engineering AR/VR
	Sharma	Emergency Evacuation AR/VR

National Aeronautics and
Space Administration

www.nasa.gov