Question 1. AI Use Case Identifying Information				Question 3. Summary	Question 4. Lifecycle Stage		Question 5. (Optional) AI Techniques	Question 6. (Optional) Data Approach				Question 7. (Optional) Technical Solution					Question 8. (Optional) Information System	Question 9. Use Case Releasability	
1A. AI use case name	1B. Agency with AI use case	Optional Note Field: To clarify 1B or to provide additional information (e.g. Agency co-development)	1C. Office with AI use case	Provide a short summary (200 words max) of what the AI does. Include a high-level description of system inputs and outputs.	4A. What stage of production is the AI in?	4B. Additional comments related to lifecycle stage.	What specific AI techniques were used?	6A. Did/does the agency train this AI use case?	6B. If yes - Where did/does the training data originate?	6C. Is the training data, the validation data, and/or test data included in the enterprise data inventory?	6D. If data is publicly available, provide link.	7A. Does the agency have access to the code associate with this AI use case?	7B. If yes, is the code included in the agency source code inventory (e.g. Code.gov)?	7C. If the source code is publicly available, provide the link.	7D. Is the agency able to conduct ongoing testing on the code?	7E. Is the agency able to monitor and/or audit performance?	Provide the name of the Information System (e.g. FISMA system name) associated with the AI use case.	9A. Should this use case be withheld from the public inventory? If yes, the use case will only be shared in the internal government inventory.	9B. If the answer to 9A is yes, provide an explanation (this explanation will be included in the internal government inventory).
Accelerated discovery in TESS full frame images using machine learning	National Aeronautics and Space Administration		Goddard Space Flight Center	This work uses convolutional neural networks to accelerate discovery in TESS (Transiting Exoplanet Survey Satellite) data in the form of full frame images. Techiques include recognizing variances in light curves from stars and identifying patterns which might indicate planets orbitine them.	Prototyped for limited data environment	See PI		Yes	NASA	prototype, still in development	https://archive.stsci.edu/tess/bulk_downloads/bulk_downloads_ffi-tp-lc-dv.html	prototype, still in development	N/A	N/A	N/A	N/A	N/A	No	
Autoresolver/Tailored Arrival Manager	National Aeronautics and Space Administration		Ames Research Center	The Autoresolver system is a tool for autonomous air traffic management. It is designed to perform many of the tasks that air traffic controllers have historically performed including maintaining separation between aircraft, sequencing and scheduling aircraft across locations in space, and avoiding airspace volumes like weather systems and restricted airspace. The tool relies on search/optimization algorithms to select safe and efficient routes for aircraft, aircraft physics models to predict future aircraft states, and probabilistic models of the accuracy of these trajectory predictions to inform resolution selection.	In production: more than 1 year		Convolutional Neural Network Probabilistic (bayesian), Bayesian Network, Uncertainty	Yes				Yes		Yes				No	N/A
ExoMiner discovery of ExoPlanets via data from the Kepler spacecraft	National Aeronautics and Space Administration		Ames Research Center	Machine learning, statistics, and human SME interpretation combined to validate over 300 newly-discovered exoplanets (and growing). The machine learning models were trained via robust human-labelled data, specifically the patterns in star brightness introduced when a planet traverses between a star and NASA sensors. These models were trained, tested and validated, and have now been put into operational use with excellent highly accurate results.	In production: more than 1 year		Machine learning pattern recognition trained with human-labelled data	Yes	Agency Generated								No	N/A
Explainable and robust deep semi-supervised model for multi-class anomaly detection in flight data	National Aeronautics and Space Administration		Ames Research Center	This model is a semi-supervised deep learning based anomaly detection for aircraft flight data. It is designed to work when a small subset of data is reviewed and labeled by experts. The most useful realm is where the size of labeled data is small, so that any supervised learning approach won't reach optimum performance. It is also trained with custom-designed objective function so the learned features are explainable for down-stream tasks such as active learning (data selection for future labeling) and it is tested to be robust to adversarial perturbation. It is currently tested and validated in finding anomaly detection in flight's operational quality assurance data from commerical aircraft.	In production: more than 1 year		Deep Neural Networks	Yes	Agency Generated	Yes	https://c3.ndc.nasa.gov/dashlink/projects/	No					No	N/A
IrrMapper Computer Vision	National Aeronautics and Space Administration		Ames Research Center	High frequency and spatially explicit irrigated land maps are important for understanding the patterns and impacts of consumptive water use by agriculture. This work created an extensive database of land cover across 50,000 irrigated fields, 38,000 dryland fields, and 500,000 square KM of uncultivated lands. The team used 28 years of LandSAT data to create a land-type classifier that provides 97.8% accurate results in classifying irrigated vs. un-irrigated land.	In production: more than 1 year		Big data ML classifiers				https://www.mdpi.com/2072-4292/12/14/2328						No	N/A
Unsupervised anomaly detection in flight data with deep variational autoencoders	National Aeronautics and Space Administration		Ames Research Center	This model is an unsupervised deep learning based anomaly detection for aircraft flight data based on variational autoencoders with convolutional architecture. The model is designed to find anomalies in multivariate time-series and can work with heterogeneous data. It is currently tested and validated in finding anomaly detection in flight's operational quality assurance data from commerical aircraft.	In production: more than 1 year		Deep Neural Networks, Variational Inference	Yes	Agency Generated	Yes	https://c3.ndc.nasa.gov/dashlink/projects/	No					No	N/A
Bone Mineral Density Maintenance During Long-Duration Spaceflight	National Aeronautics and Space Administration		Glenn Research Center	Bone mineral density maintenance is important for astronauts on long-duration missions. This work uses probabilistic image segmentation to automatically generate bespoke finite element models of femoral bones from CT scan images in combination with in-the-loop stochastic optimization to determine loading forces and frequencies required to maintain bone mineral density in microgravity. This information can be used by flight physicians to design customized exercise routines for astronauts during spaceflight using on-board resistive exercise devices.	In production: more than 1 year	Awaiting continued funding from HRP to resume production.	Naïve Bayes Classifiers, Stochastic Optimization	Yes	Other	No	Yes	No					No	N/A
Toward a deep-learning-based cloud-masking framework for enhanced monitoring of water resources	National Aeronautics and Space Administration		Goddard Space Flight Center	To produce reliable aquatic science products using multispectral satellite imagery, preprocessing steps are often necessary to remove clouds and cloud shadows. These contaminations can lead to highly inaccurate aquatic reflectance spectra and downstream products that may misinform decision-making processes (e.g., issuing advisory). As of yet, various image processing workflows apply simple thresholding approaches; however, they lack fidelity in dealing with thin clouds or cloud shadows or when seamless multi-mission products are desired. Machine learning approaches have shown promise in cloud-masking tasks but require extensive training sets, which are primarily available over terrestrial targets. This effort aims to create a framework for a robust cloud-masking of existing satellite images over inland and coastal waters. To exclude cloud and cloud shadows, a domain adaptation-based deep learning model for semantic segmentation of satellite images is developed. Hand-labeled training samples are generated from Landsat-9/OLI and Sentinel-2/MSI imagery to cover a broad latitudinal band across the Americas.	In production: more than 1 year	See PI	Domain adaptation deep learning	No		No	No		No	No			No	
IMERG Hydro-GV at the catchment scale	National Aeronautics and Space Administration		Goddard Space Flight Center	IMERG (integrated multi-satellite retrievals for GPM (Global Precipitation Measurement)) Hydro-GV at the catchment scale. The Integrated Multi-satellitE Retrievals for GPM (IMERG) algorithm combines information from the GPM satellite constellation to estimate precipitation over the majority of the Earth's surface. Techniques include support vector machine ML models.	In production: more than 1 year												No	
Machine Learning Approaches for Predicting Phytoplankton Community Composition from Ocean Color	National Aeronautics and Space Administration		Goddard Space Flight Center	This work uses Bayesian ML approaches to predict phytoplankton community composition (PCC) from in situ and satellite measurements of ocean color. PCC reveals fundamental insight into oceanic carbon cycling and ecosystem functioning, making its determination vital for understanding a changing ocean. The models are trained using state of the art measurements of PCC (e.g., imaging flow cytometry and molecular assays), which sets it apart from traditional phytoplankton osement based models.	In production: more than 1 year	See PI	Support-vector machine (SVM)										No	N/A
AEGIS: Autonomous Exploration for Gathering Increased Science	National Aeronautics and Space Administration		Jet Propulsion Laboratory	AEGIS enables intelligent targeting and data acquisition by planetary rovers. It uses computer vision techniques to identify targets (e.g., rocks) in wide angles images of the rover's surrounding terrain. If targets are found that match scientists specifications, they are then measured autonomously using remote sensing instruments. AEGIS was first used on the MER Mission. It is currently in use on the MSL Mission to acquire data for the ChemCam instrument. It is planned for use in Spring of 2022 on the M2020 Mission to acquire data for the SuperCam instrument.	In production: more than 1 year		Artificial Neural Network, Dirichlet-multinomial process comouter vision	No				Yes					No	
ASPEN Mission Planner	National Aeronautics and Space Administration		Jet Propulsion Laboratory	Based on AI techniques, ASPEN is a modular, reconfigurable application framework which is capable of supporting a wide variety of planning and scheduling applications. ASPEN provides a set of reusable software components that implement the elements commonly found in complex planning/scheduling systems, including: an expressive modeling language, a resource management system, a temporal reasoning system, and a graphical interface. ASPEN has been used for many space missions including: Modified Antarctic Mapping Mission, Orbital Express, Earth Observing One, and ESA's Rosetta Orbiter.	In production: more than 1 year	Used by multiple satellites and sensors for multiple years	Stochastic Local Search scheduler taking into account numerous criteria and variables					Yes					No	
CLASP Scientific Instrument Planning & Scheduling	National Aeronautics and Space Administration		Jet Propulsion Laboratory	The Compressed Large-scale Activity Scheduling and Planning (CLASP) project is a long-range scheduler for space-based or aerial instruments that can be modelled as pushbrooms – 1D line sensors dragged across the surface of the body being observed. It addresses the problem of choosing the orientation and on/off times of a pushbroom instrument or collection of pushbroom instruments such that the schedule covers as many target points as possible, but without oversubscribing memory and energy. Orientation and time of observation is derived from geometric computations that CLASP performs using the SPICE ephemeris toolkit. CLASP allows mission planning teams to start with a baseline mission concept and simulate the mission's science return using models of science observations, spacecraft operations, downlink, and spacecraft trajectory. This analysis can then be folded back into many aspects of mission design – including trajectory, spacecraft design, operations concept, and downlink concept. The long planning horizons allow this analysis to span an entire mission. Actively in use for optimized scheduling for the ECOSTRESS mission (study of water needs for plant areas), EMIT mission (mineralogy of arid dusty regions), OCO-3 (atmospheric CO2) and more.	In production: more than 1 year	Used by multiple satellites and sensors for multiple years	Squeaky wheel optimization scheduler taking into account numerous criteria and variables	No				Yes	No				No	N/A
Enhanced AutoNav for Perseverance Rover on Mars	National Aeronautics and Space Administration		Jet Propulsion Laboratory	AutoNav on the Perseverance Rover autonomously plans a safe path based on stereo navigation camera images, based on multiple technologies including a tree search for decision making, Dijkstra algorithm for global path planning, stereo processing for 3D terrain reconstruction, and Approximate Clearance Evaluation (ACE) for safety checks. It is deployed on Perseverance rover and being used for autonomous driving on Mars.	In production: more than 1 year	Deployed on Mars.	Tree search, path planning, stereo processing, approximate kinematic solver	No	Other	No	Note: this is NOT ML-based	Yes					No	
Mars2020 Rover (Perseverance)	National Aeronautics and Space Administration		Jet Propulsion Laboratory	Research, experiments, and engineering to empower future rovers with onboard autonomy; planning, scheduling & execution; path planning; onboard science; image processing; terrain classification; fault diagnosis; and location estimation. This is a multi-faceted effort and includes experimentation and demonstrations on-site at JPL's simulated mars navigation yard. The M2020 onboard scheduler incrementally constructs a feasible schedule by iterating through activities in priority-first order. When considering each activity it computes the valid time intervals for placement, taking into account preheating, maintenance heating, and wake/sleep of the rover as required. After an activity is placed (other than a preheat/maintenance or wake/sleep), the activity is never reconsidered by the scheduler for deletion or moving. Therefore the scheduler can be considered non backtracking, and only searches in the sense that it computes valid timeline intervals for legal activity placement. Meta Search: Because the onboard scheduler will be invoked many times in a given sol (Martian Day) with a range of possible contexts (due to execution variations), its non backtracking nature leaves its vulnerable to brittleness. In order to mitigate this potential brittleness, the Copilot systems perform a monte carlo based stochastic analysis to set meta parameters of the scheduler - primarily activity priority but also potentially preferred time and temporal constraints.	In production: more than 1 year		Automated scheduling, Meta Search, Autonomous planning, image processing, terrain classification, fault diagnosis					Yes					No	N/A
MLNav (Machine Learning Navigation)	National Aeronautics and Space Administration		Jet Propulsion Laboratory	Accelerates path planning of rovers and other types of vehicles through ML-based heuristics, while guaranteeing safety through conventional, model-based collision checking. Integrated with M2020's Enhanced AutoNav (ENav) and tested with the real terrain data from Marc on ENav simulator. Uses U-net for the ML-based heuristics, trained by simulation-generated terrain data.	In production: more than 1 year		Convolutional Neural Networrk (CNN), U-Net CNN	Yes	Agency Generated	No		Yes					No	
Perseverance Rover on Mars	National Aeronautics and Space Administration		Jet Propulsion Laboratory	3D machine vision via dual cameras to inform convolutional neural networks for rover navigation path planning. Physics / momentum - based Terrain Relative Navigation (TRN). Enhanced AutoNavigation (Enav) plans feasible paths using multiple techniques, to include random forests. Approximate Clearance Evaluation (ACE) assesses obstacles to determine if the suspension of the rover can clear them (drive over them) or needs to route around them. Additional, more complex AIML techniques are being planned and tested for future rovers at JPL. TRN was also used to provide precision landing for the Perseverance Rover in the entry, descent, and landing process to get the rover to the surface of Mars.	In production: more than 1 year	Deployed on Mars	Convolutional Neural Networks, Random Forests,Terrain Relative Navigation	Yes	Agency Generated	No							No	N/A
Planetary Cave Rovers	National Aeronautics and Space Administration		Jet Propulsion Laboratory	In this project, we study multi-rover coordination techniques to allow vehicles to autonomously explore the unknown environments of caves on Mars and the Moon. Since the ultimate goal is to return science data, data routing becomes a key design point. To analyze our different techniques, we developed a simulation framework to easily compare different mission configurations. We developed a simulation framework to easily run different configurations for mission concepts, which provides diagnostic output to evaluate performance. It is implemented in the Robot Operating System (ROS) and allows configuration of asset parameters, environmental parameters, and mission concept parameters, including the number of assets, asset specifications (such as the communication model, navigation model, exploration strategy, and science instruments), and the cave model. Diagnostic output includes an interactive visual playback of the scenario, activity timeline and distribution, paths travelled by the assets, and energy usage distribution.	In production: more than 1 year	Sim environment is in place to experiment with, develop, and test AI-enabled systems to enable autonomous cave exploration	On-board autonomy, multi-rover coordination, planning, scheduling and execution					Yes					No	N/A
SCOTI (Scientific Captioning of Terrain Images)	National Aeronautics and Space Administration		Jet Propulsion Laboratory	SCOTI (Scientific Captioning of Terrain Images) automatically generates natural language explanations of geological images taken by rovers. It uses "ashw-attend-tell" model consisting of CNN (Convolutional Neural Network) and LSTM (Long Short Term Memory), trained by scientist-generated labels on MSL images. SCOTI provides onboard data summarization tha would help the ground operation to selectively downlink high priority data under data bandwidth constraint.	In production: more than 1 year		CNN, LSTM	Yes	Agency Generated	No		Yes					No	
SPOC (Soil Property and Object Classification)	National Aeronautics and Space Administration		Jet Propulsion Laboratory	Using a convolutional neural network (CNN), SPOC (Soil Property and Object Classification) takes rover images and classifies the terrain type (e.g., sand, soil) from visual appearance. This ability enables rover to drive more safely. It is trained by labeled images from MER (Mars Exploration Rover), MSL (Mars Science Laboratory), and Mars 2020 rovers, annotated by tens of thousands of citizen scientists through the AI4Mars project. SPOC deployed on MSL's ground operation system and onboard test on M2020 is being considered.	In production: more than 1 year	Deployed on MSL ground ops; being considered for onboard demonstration on M2020	CNN	Yes	Other	No	https://data.nasa.gov/Space-Science/Al	Yes	No				No	

TRN (Terrain Relative Navigation)	National Aeronautics and Space Administration		Jet Propulsion Laboratory	Terrain Relative Navigation (TRN) estimates position during Mars landing by automatically matching landmarks identified in descent images to a map generated from orbital imagery. The position estimate is used to a select a safe and reachable landing site in a region with many large hazards. TRN was used successfully by the Mars 2020 mission during its landing on February 18th, 2021.	In production: more than 1 year	TRN is planned for use in the next NASA mars lander mission.	Computer vision and state estimation.			Yes						No		
Volcano SensorWeb	National Aeronautics and Space Administration		Jet Propulsion Laboratory	The Sensor Web Project uses a network of sensors linked by software and the internet to an autonomous satellite observation response capability. This system of systems is designed with a flexible, modular, architecture to facilitate expansion in sensors, customization of trigger conditions, and customization of responses. This system has been used to implement a global surveillance program to study volcanos. We have also run sensorweb tests to study flooding, cryosphere events, and atmospheric phenomena. Specifically, in our application, we use low resolution, high coverage sensors to trigger observations by high resolution instruments. Note that there are many other rationales to network sensors into a sensorweb. For example automated response might enable observation using complementary instruments such as imaging radar, infra-red, visible, etc. Or automated response might be used to apply more assets to increase the frequency of observation to improve the temporal resolution of available data. Our sensorweb project is being used to monitor the Earth's 50 most active volcanos. We have also run sensorweb experiments to monitor flooding, wildfires, and cryospheric events (snowfall and melt, lake freezing and thawing, sea ice formation and breakup.)	In production: more than 1 year		Automatic Data Interpretation, Planning and Scheduling, Multi-Sensor Fusion			Yes						No		N/A
Machine Learning for RFID (Radio Frequency Identification) tag localization to support logistics	National Aeronautics and Space Administration		Johnson Space Center	Currently have two production machine learning approaches to tackle RFID (Radio Frequency Identification) tag localization in the highly reflective environment imposed by the International Space Station. First use case, REALMRFC, is a random forest classifier model with feature engineering performed by an RFID localization expert. The second use case is P-RFIDNet, a neural network with a ResNet50 backbone. In continued work, we have leveraged transfer learning to show how P-RFIDNet can be generalized to new RFID environments with limited training data. We benchmark P-RFIDNet and REALMRFC using data from the RFID Enabled Autonomous Logistics Management (REALM) and using truth derived from the Inventory Management System (IMS).	In production: more than 1 year	Derivitive work are in various stages of development.	Deep Neural Network, transfer learning, date image generation, image classification, unsupervised learning, semi-supervised learning, decision tree classifiers, K-nearest neighbors classifiers, statistical methods	Yes	Agency Generated	No	Yes	No	Yes	Yes	No	Data from this use case has potentially derived sensive information.
NASA OCIO STI Concept Tagging Service	National Aeronautics and Space Administration		Langley Research Center	An API (application program interface) for exposing topic models created with the STI (Scientific & Technical Information) concept training repository.	In production: more than 1 year		NLP	Yes			Yes		https://github.com/nasa/concept-tagging-api		No	N/A	
Watson Explorer	National Aeronautics and Space Administration		Langley Research Center	Langley Research Center has used IBM Watson Explorer (WEX) Natural Language Processing for over 8 years to analyze numerous text-based data sets. The capability provides cognitive exploration and content analysis and is used to explore and analyze structured and unstructured, internal, external and public content to uncover trends and patterns. The team has applied WEX to data sets up to 600,000 documents to cluster concepts, conduct author analysis, and explore the data via powerful faceted search	In production: more than 1 year		Natural language processing									No		
iNtelligent Stennis Gas House Technology (INSIGHT)	National Aeronautics and Space Administration		Stennis Space Center	INSIGHT is an operational system that performs autonomous Integrated System Health Management (ISHM) and autonomous operations of the Nitrogen System of the High Pressure Gas Facility at NASA Stennis Space Center. It is an application implemented using the NASA Platform for Autonomous Systems (NPAS) described in this document as AI Use Case Name: NASA Platform for Autonomous Systems (NPAS). The inputs are (1) descriptions of system structure and behavior including schematics, (2) description of operations, (3) autonomy strategies to mitigate off-nominal behavior, (4) strategies for ISHM (anomaly detection, diagnostics, prognostics, comprehensive awareness.	In production: more than 1 year	INSIGHT is currently at a base autonomy capability, and the expectation is to continue its evolution to achieve high levels of autonomy capability. INSIGHT represents a validation of the AI technologies of NPAS for operational systems that have to meet safety criticality criteria and requirements. INSIGHT has been classified as Class C Safety Critical.	INSIGHT employs the AI technologies provided by NPAS, described in this document as AI Use Case Name: NASA Platform for Autonomous Systems (NPAS).	No			Yes	No		Yes	Yes	No		
NASA Platform for Autonomous Systems (NPAS)	National Aeronautics and Space Administration		Stennis Space Center	The NASA Platform for Autonomous Systems (NPAS) enables implementation of "thinking" systems, and in particular of "thinking" autonomous systems. A broad range of systems can be made to display "thinking" autonomous behavior; including fluid, mechanical, electrical, networks, and computer. Additional types of systems can be easily included. AI behavior, "thinking," is grounded in a comprehensive representation of the system (comparable to SysML model descriptions that include health management and autonomy behaviors as well as schematic level descriptions), behavior/function models (physics-based, heuristic, rule-based - probabilistic and neural-network models can also be incorporated), Failure Modes and Effects Analysis (FMEA) with generic/re-usable libraries. NPAS systems/applications incorporate autonomy strategies to deal with off-nominal cases, and strategies for fault management. NPAS includes infrastructure for autonomous operations (task definition, planning, scheduling, and execution) embodying specific concepts of operations. NPAS supports implementation of hierarchical distributed autonomous systems and operations.	In production: more than 1 year	NPAS has been used to implement autonomous operations of the Nitrogen System of the High Pressure Gas Facility (HPGF) at Stennis Space Center. NPAS is currently being used to implement autonomous satellite capability under the AES Polaris project "Autonomous Satellite Technology for Resilient Applications (ASTRA)," in collaboration with commercial partner Sidus Space. NPAS was used to prototype hierarchical distributed autonomous operations of a Vehicle System Manager, multiple module managers, and system managers; according to the Gateway autonomy design. This was a risk reduction collaboration with commercial partner Lockheed Martin.	On-board "thinking" by the system rather than on-board application of "thinking" that is done off-line by teams of humans. The AI technologies embody a transfer of the thought and process capabilities of humans to the system itself. Model-based on-board "thinking" using models of existing knowledge (e.g. physics models, heuristic models, etc.), processes (e.g. engineering and science processes, strategies for health management and autonomy), Failure Modes and Effects Analysis (FMEA) with generic cause-effect tree representation and resolution for failure instances occurring in real-time, goal-based task management including task definition, planning, scheduling, and execution . In general, thinking that involves perception, analysis, evaluation, reasoning, decision making, task and mission execution.	No			Yes	No	Not publicly available	Yes	Yes	No		
High-Performance Quantum-Classical Hybrid Deep Generative Modeling Parameterized by Energy-based Models for Flight-Operations Anomaly Detection	National Aeronautics and Space Administration		Ames Research Center	Our project conducts high-performance scalable and explainable machine learning for flight-operations anomaly detection, with contributions from classical computing (enhanced performance, reduced cost) and quantum computing (encoding of quantum correlations, quantum-resource estimates). Our deep-learning model takes time series of 19 flight metrics collected by flight recorders of commercial aircraft as input and predicts operational and safety-relevant anomalies during the take-off and landing phases of flight.	In production: less than 6 months		Deep learning, generative modeling, semi-supervised learning, model explainability, active learning, quantum regularization	Yes	Agency Generated	Yes	https://c1.ndc.nasa.gov/dashlink/projects/81/					No		N/A
Integrated System for Autonomous and Adaptive Caretaking (ISAAC)	National Aeronautics and Space Administration	Teamwork with Johnson Space Center	Ames Research Center	Human exploration infrastructure in deep space, such as the planned "Gateway", will be uncrewed for extended periods. During these periods NASA needs autonomous systems that can perform Fault Detection, Isolation, and Recovery (FDIR) as well as perform routine maintenance and logistics operations through high latency communication to ground controllers. The Integrated System for Autonomous and Adaptive Caretaking (ISAAC) project directly addresses this key technical need.	In production: less than 6 months		Semantic classification, Automated planning, Automated change and anomoly detection	Yes	Agency Generated			Yes				No		
Cognitive Communications	National Aeronautics and Space Administration		Glenn Research Center	The Cognitive Communications project will design, build, and demonstrate a cognitive communication system prototype called Cognitive Engine – 1 (CE-1). CE-1 will seamlessly operate with NASA and commercial networks, eliminate the need for operators to forecast communication service schedules, adapt to changing environmental conditions without human interaction, and maximize data throughput. CE-1 is a ground-based prototype cognitive system capable of handling communication tasks for emulated spacecraft, ground-based network assets, and mission operation centers. CE-1 represents NASA's first step towards enabling autonomous spacecraft communications. The project will assemble a CE-1 prototype by incrementally combining AI/ML and communication technologies that deliver automated, reactive, predictive, and finally cognitive functionality.	In production: less than 6 months	See PI	Reinforcement Learning, Deep Learning, Generative Adversarial Networks, Spiking Neural Networks	Yes	Agency Generated	No	No		Data and Code to be made available upon p	Yes	Yes		No		N/A
Graph Neural Networks for Airfoil Performance Prediction	National Aeronautics and Space Administration		Glenn Research Center	We are investigating the use of Graph Convolutional Neural Networks to learn relationship between airfoil coordinates and predict the performance for aerodynamics analysis. Inputs include the shape of the airfoil and outputs are the coefficient of lift, drag, and moment. The impact is we have a new type of neural network architecture that we can potentially use for other projects.	In production: less than 6 months	Software safety is currently holding up the project from being released.	Graph Neural Networks, Spline Convolution	Yes	Agency Generated		https://nasa-public-data.s3.amazonaws.com/pioclbd_utilities/airfoil-learning-dataset.zip				Some services run on Goddard Commercial Cloud AWS	No		N/A
Inverse Design of Materials	National Aeronautics and Space Administration		Glenn Research Center	Discovering new materials is typically a mix of art and science, with timelines to create and robustly test a new material mix / manufacturing method ranging from ten to twenty years. This project seeks to enable rapid discovery, optimization, qualification and deployment of fit-for-purpose materials. The team is using supervised machine learning techniques to develop ML models of material structure based on 100,000+ NASA images, to include using ML to optimize materials recipes based on desired materials properties. ML models are also acting as surrogate models, with 100 to 10,000 times speedup over traditional physics simulations already realized for select materials. Inputs include materials images, metadata, and SME-based labelling. Outputs include recipes and approaches for new materials custom-tailored to applications with an 4x speedup for the overall materials discovery / design lifecycle, and potential 10x throughput for the same cycle based on parallitizing discovery of multiple materials at once.	In production: less than 6 months	This project is in the process of iterative development between data scientist and SME. It is beyond "planned," but not "in production" yet.	Machine learning image recognition, Supervised ML techniques	Yes	Agency Generated	Yes	Yes	No	Data and Code to be made available upon p	Yes	Yes	This system is running on cloud and on-premise systems behind the NASA firewall	No		N/A
Physics Based Neural Networks for Fluid Dynamics	National Aeronautics and Space Administration		Glenn Research Center	Physics based neural networks are a new field. They can be used to predict CFD (computational fluid dynamics) results. One of the main advantages is that they do not need a mesh, they can solve the navier stokes using randomly simulated points. This makes it advantageous in cavity flow, narrow passages, turbine blade tip geometries, just to name a few. Since this is a new field, it is important that we advance the science to improve accuracy.	In production: less than 6 months		Physics Based Neural Networks with Deep Neural Network architecture	Yes	Agency Generated							No		N/A
Probe Calibration using Neural Networks	National Aeronautics and Space Administration		Glenn Research Center	Sensor / probe calibration for wind tunnel testing has historically relied on maps and surfaces. Interpolation is also difficult because of the multi-dimensionality of the calibration data. Calibration data includes the angles that the probe can experience in pitch and yaw, in addition to the different test conditions. Machine learning can be used to better estimate the multi-dimensionality of the dataset resulting in an accurate pitch and yaw angles based on what the probe is experiencing in the test section.	In production: less than 6 months		Deep Neural networks	Yes	Agency Generated							No		N/A
Using a ML regression model for processing hyperspectral observations over optically complex aquatic ecosystems	National Aeronautics and Space Administration		Goddard Space Flight Center	Coastal and freshwater ecosystems are among the most productive ecosystems. They are highly vulnerable to the changing climate and to anthropogenic activities while supporting human lives through the services they provide. For over two decades, satellite systems with ocean color (OC) capabilities have produced valuable records of observations for studying the carbon and biogeochemical cycling in these optically complex environments. In preparation for NASA's Surface Biology and Geology mission, this project aims at developing and demonstrating a processing pipeline which utilizes Mixture Density Networks (MDNs) to enable high-quality retrievals of aquatic products from imaging spectrometry data. Our retrieval scheme constitutes two steps. First, we will conduct an atmospheric correction to retrieve hyperspectral remote sensing reflectance using an MDN model trained with in situ Rrs and simulated top-of-atmosphere reflectance. Second, we will utilize the derived Rrs products to simultaneously retrieve in-water properties. Our framework will be implemented for, and validated with, Hyperspectral Imager for Coastal Ocean (HICO) images using the AERONET-OC data and in situ data acquired at well-monitored U.S. coastal estuaries and lakes. To further extend our validation opportunities, we will also demonstrate our retrieval scheme for the Italian Space Agency's existing hyperspectral mission (PRISMA).	In production: less than 6 months	See PI	A class of NNs referred to as Mixture Density Networks (MDNs)	No		No	No	https://github.com/STREAM-RS	Yes	Yes	The assembled training dataset has been submitted to Nature's Scientific Data for	No		
High Resolution Earth and Planetary Atmospheric Predictions using Machine Learning	National Aeronautics and Space Administration		Goddard Space Flight Center	We are developing a machine learning (ML) tool that can predict high-resolution in situ atmospheric conditions using relatively lower resolution remote data e.g., from an orbiting spacecraft. Our ML tool could be used to map and track atmospheric cycling on Earth and planetary bodies, not only as a fundamental science tool, but also as a mechanism for tracking planetary weather from orbit.	In production: less than 6 months		Autoencoders, convolutional neural network, UMAP	Yes	Agency Generated							No		N/A
Autonomous Marine Vehicles (Single, Multiple)	National Aeronautics and Space Administration		Jet Propulsion Laboratory	Due to the communication paradigm associated with operating an underwater submersible on an Ocean World, the vehicle must be able to act autonomously when achieving scientific goals. One such goal is the study of hydrothermal venting. Evidence for hydrothermal activity has been found on one Ocean World, Enceladus. On Earth, these geological phenomena harbor unique ecosystems and are potentially critical to the origin of life. Similar vents on Ocean Worlds could be the best chance at extra-terrestrial life in our Solar System. We focus on performing autonomous science, specifically the localization of features of interest - such as hydrothermal venting - with limited to no human interaction. A field program to Karasik Seamount in the Arctic Ocean was completed in Fall 2016 to study and understand the human-in-the-loop approach to the localizing hydrothermal venting. In 2017/2018 an autonomous nested search method for hydrothermal venting was developed and tested in simulation using a hydrothermal plume dispersion model developed by Woods Hole Oceanographic Institution	In production: less than 6 months	Testing and experimentation with a variety of techniques is underway	Planning and Execution, Nested Search, Anomaly Detection									No		N/A

Hybrid On-Board and Ground-Based Processing of Massive Sensor Data (HyspIRI IPM)	National Aeronautics and Space Administration		Jet Propulsion Laboratory	Future space missions will enable unprecedented monitoring of the Earth's environment and will generate immense volumes of science data. Getting this data to ground communications stations, through science processing, and delivered to end users is a tremendous challenge. On the ground, the spacecraft's orbit is projected, and automated mission-planning tools determine which onboard-processing mode the spacecraft should use. The orbit determines the type of terrain that the spacecraft would be overflying—land, ice, coast, or ocean, for instance. Each terrain mask implies a set of requested modes and priorities. For example, when a spacecraft overflies polar or mountainous regions, producing snow and ice coverage maps can provide valuable science data. The science team can adjust these priorities on the basis of additional information (such as external knowledge of an active volcano, a flooded area, an active wildfire, or a harmful algal bloom). The mission-planning tool accepts all these requests and priorities, then determines which onboard-processing algorithms will be active by selecting the highest-priority requests that fit within the onboard CPU resources, band-processing limitations, and downlink bandwidth.In the intelligent onboard processing concept, HyspIRI's onboard processing algorithms would consist of expert-derived decision tree classifiers, machine-learned classifiers such as SVM classifiers and regressions, classification and regression trees (CART), Bayesian maximum-likelihood classifiers, spectral angle mappers, and direct implementations of spectral band indices and science products	In production: less than 6 months	In prototyping and experimentation	decision tree classifiers, classification & regression trees, Bayesian maximum likelihood classifiers, spectral angle mappers, support vector machine classifiers				Yes							No		N/A
Neural network accelerated radiative transfer modeling	National Aeronautics and Space Administration		Jet Propulsion Laboratory	Started in FY21 the AI behind the Neural network accelerated radiative transfer modeling is intended to enhance efforts in the Earth Science domain. Specifically, JPL constructed a flexible radiative transfer model (RTM) that combines physics-based models and artificial neural networks with the intent of providing fast radiative transfer modeling for global imaging spectroscopy missions, as well as large-scale airborne campaigns (ABoVE, Western Diversity Time Series, FIREX-AQ, etc.)	In production: less than 6 months		Neural networks	Yes	Agency Generated			Yes							No		
Intelligent Contingency Management	National Aeronautics and Space Administration		Langley Research Center	Adapt and train AI algorithms to contribute to an autonomous vehicle mission manager for Advanced Air Mobility (Cargo, Air Taxis). At a high level, the AI must recognize contingency flight conditions and react appropriately to return the aircraft to safe flight status. The project has three main objectives: 1. Explore machine learning for intelligent contingency management, with a focus on assessing/projecting vehicle capability and maintaining nominal performance via reinforcement learning. 2. Develop vehicle intelligent contingency management system architecture at a functional level and validate against a specific Unmanned Air Mobility (UAM)-class vehicle. 3. Incorporate (1) and (2) into an evolving toolset for an autonomous vehicle. Inputs include sensors for aircraft state and the surrounding environment (natural and man-made). Outputs include recognition of off-nominal conditions (contingencies) and mission executon strategy adjustments.	In production: less than 6 months	This project is in the process of iterative adaptation, development, and training within a complex, detailed modeling and simulation environment. It is in the applied R&D stage (under development); this means that it is not "planned," but not yet in production operations either.	Reinforcement Learning, Knowledge Reasoning, Probabilistic Inference, Generative Adversarial Networks, Physics-Informed Neural Networks	Yes	Agency Generated	Yes		Yes	No	Data and Code will be made available upon project completion	Yes	Yes	This is running in a firewalled NASA Aeronautics modeling and simulation system.	No		N/A
Lessons Learned Bot (LLB)	National Aeronautics and Space Administration		Langley Research Center	In near real-time, the Lessons Learned Bot, or LLB, brings lessons learned (LL) documents to users through a Microsoft Excel add-in application locally installed to search for LL content relevant to the text within the selected Excel cell. The application will encompass a corpus of documents, a trained Machine Learning (ML) model, built-in ML tools to train user's documents, and an easy-to-use user interface to allow for the streamlined discovery of LL content. Today, NASA's LL are online and searchable via keywords. Nevertheless, users often face a challenge to find lessons relevant to their issues. Applying the advancement in Natural Language Processing (NLP) ML algorithm, the LLB can find and rank LL records relevant to text in the user's selected Excel cells, containing just a few words or entire paragraphs of text. Results are displayed to the user in their existing Excel workflow. The LLB's installation package comes with a pre-trained NASA LL dataset and a NASA Scientific and Technical Information (STI) dataset, as well as on-demand training tools allowing the user to apply the LLB search algorithm to their own discipline specific datasets.Additionally, we also have an API version of this software that can be called from any application within the Agency firewall.	In production: less than 6 months	Currently in production and recieves updates/security patches	Doc2Vec NLP Machine Learning Approach	Yes	Agency Generated	Yes	https://ntrs.nasa.gov/ https://llis.nasa.gov/	Yes	No		Yes	Yes	llb.larc.nasa.gov	No		N/A
Pedestrian Safety Corridors for Drone Test Range	National Aeronautics and Space Administration		Langley Research Center	NASA Langley Research Center (LaRC) is actively experimenting with Unmanned Aerial Systems (UAS - Drones and surrounding systems) to include command, control, coordination and safety mechanisms. LaRC is expanding an on-site UAS test range, to include areas where people walk, drive, etc. This project leverages the parking advisor image recognition project and applies it to detecting pedestrian traffic to supplement statistical assessment of human-heavy and human trie traffic areas with near-real time human-presance-detection. Inputs include camera signals and hand-labelled training data. Outputs include maps indicating density of human pedestrian traffic.	In production: less than 6 months	This project is in the process of iterative development between data scientist and SME. It is beyond "planned," but not "in production" yet.	Image recognition, Recurrent convolutional neural networks, Perspective-based grid generation system to account for parallax and similar phenomena	Yes	Agency Generated	No		Yes	No	Code to be made available upon project completion; imagery data is inappropriate to share however due to physical security concerns.	Yes	Yes	This system is running on a firewalled on-premesis server in a local data center	No		N/A
Storm Prediction via Above Anvil Cirrus Plume Image Recognition	National Aeronautics and Space Administration		Langley Research Center	Above-Anvil Cirrus Plumes (AACPs) are a precursor to destructive hailstorms. This project seeks to train machine learning image recognition techniques to identify AACPs in satellite imagery to eventually contribute to an early storm warning system. The system uses image recognition, Li-Nets, and Long-Term-Short-Term memory systems. Inputs include satellite imagery and human labeling of AACPs. Outputs include machine-identified AACP instances.	In production: less than 6 months	This project is in the process of iterative development between data scientist and SME. It is beyond "planned," but not "in production" yet.	Image recognition, Li-Nets, LSTM Neural Networks and more	Yes	Agency Generated	Yes		Yes	No	Data and Code will be made available upon project completion	Yes	Yes	This system is running on a firewalled cloud R&D environment	No		N/A
Superimposition and Collation of Images with Non-English Translations (SCI-NET)	National Aeronautics and Space Administration	In teamwork with Marshall Space Flight Center	Langley Research Center	NASA seeks to leverage global aerospace knowledge, and many key concepts lie in historical technical literature in foreign languages. This project leverages ever-improving language translation capabilities from industry to translate Russian and other languages to English, and then overlay the translated text over the original document, maintaining synchronization with figures, diagrams, graphs, formulas, etc. Inputs include foreign language technical documents as well as SME interpretation of jargon. The primary output is the translated text overlaid over the other document elements from the original document.	In production: less than 6 months	This project is in the process of iterative development between data scientist and SME. It is beyond "planned," but not "in production" yet.	Foreign language translation and document element recognition systems	Yes	Other	Yes		Yes	No	From Google Cloud Platform AIML services	Yes	Yes	These services run on Google Cloud Platform; NASA's instance is part of a firewalled GovCloud instance.	No		N/A
Dynamic Lightning Prediction Algorithm for lightning safety	National Aeronautics and Space Administration		Marshall Space Flight Center	Uses a long short-term memory model and convolutional neural network to predict lightning occurrence and extent in for a short term 5-15 minute period.	In production: less than 6 months		Neural network long-short-term memory models and convolutional neural networks.									Local ST11 machine	No		
Sinatra	National Aeronautics and Space Administration		Marshall Space Flight Center	A flexible software framework that analyzes big data and detects anomalies. Allows for integrating ML with minimal background in deep learning and software development.	In production: less than 6 months		Framework for big data analysis and anomaly detection										No		
Deep Learning for asteroid damages prediction	National Aeronautics and Space Administration		Ames Research Center	This project is focused on prediction of ground damages of different asteroid impact scenarios based on deep neural network. The training data used comes from millions of numerical simulations based on the physics - the neural networks aim at providing a lower fidelity but fast prediction of the damages.	In production: less than 1 year		Deep Neural networks	Yes	Agency Generated	Yes		Yes	No	https://github.com/gregorychomette/ml-atap.git				No		N/A
Guided Materials Discovery	National Aeronautics and Space Administration	ARC-TI	Ames Research Center	Guided materials discovery combines predictive theory with experimental validation. Machine learning methods allow to combine electronic-structure calculations with thermodynamic and kinetic modeling; such multi-scale methods are used to predict dependence of material properties on composition and structure on multiple scales. The composition-structure-property relations are used in materials design: they help to improve and optimize properties by compositional adjustment and thermo-mechanical treatment of material samples. We provide theoretical guidance for our experimental partners, and we develop methods and software for making reliable predictions.	In production: less than 1 year		Optimal choice of ML method	Yes		No		Yes	No		Yes	Yes	NAS	No		N/A
Machine Learning Airport Surface Model: Airport Configuration Prediction	National Aeronautics and Space Administration		Ames Research Center	The ML-airport-configuration software is developed to provide a reference implementation to serve as a research example how to train and register Machine Learning (ML) models intended for predicting airport configurations. The software is designed to point to databases which are not provided as part of the software release and thus this software is only intended to serve as an example of best practices. The software is built in python and leverages open-source libraries kedro, scikitlearn, MLFlow, and others. The software provides examples how to build three distinct pipelines for data query and save, data engineering, and data science. These pipelines enable scalable, repeatable, and maintainable development of ML models.	In production: less than 1 year		XGBoost, kedro, scikit-learn, MLFlow, pipelines	Yes				Yes		https://github.com/nasa/ML-airport-conf	Yes	Yes		No		N/A
Machine Learning Airport Surface Model: Arrival Runway Prediction	National Aeronautics and Space Administration		Ames Research Center	The ML-airport-arrival-runway software is developed to provide a reference implementation to serve as a research example how to train and register Machine Learning (ML) models intended for predicting arrival runway assignments. The software is designed to point to databases which are not provided as part of the software release and thus this software is only intended to serve as an example of best practices. The software is built in python and leverages open-source libraries kedro, scikitlearn, MLFlow, and others. The software provides examples how to build three distinct pipelines for data query and save, data engineering, and data science. These pipelines enable scalable, repeatable, and maintainable development of ML models.	In production: less than 1 year		XGBoost, kedro, scikit-learn, MLFlow, pipelines	Yes				Yes		https://github.com/nasa/ML-airport-arriv	Yes	Yes		No		N/A
Machine Learning Airport Surface Model: Departure Runway Prediction	National Aeronautics and Space Administration		Ames Research Center	The ML-airport-departure-runway software is developed to provide a reference implementation to serve as a research example how to train and register Machine Learning (ML) models intended for predicting departure runway assignments. The software is designed to point to databases which are not provided as part of the software release and thus this software is only intended to serve as an example of best practices. The software is built in python and leverages open-source libraries kedro, scikitlearn, MLFlow, and others. The software provides examples how to build three distinct pipelines for data query and save, data engineering, and data science. These pipelines enable scalable, repeatable, and maintainable development of ML models.	In production: less than 1 year		XGBoost, kedro, scikit-learn, MLFlow, pipelines	Yes				Yes		https://github.com/nasa/ML-airport-depa	Yes	Yes		No		N/A
Machine Learning Airport Surface Model: Estimated ON Time Prediction	National Aeronautics and Space Administration		Ames Research Center	The ML-airport-estimated-ON software is developed to provide a reference implementation to serve as a research example how to train and register Machine Learning (ML) models intended for predicting landing time. The software is designed to point to databases which are not provided as part of the software release and thus this software is only intended to serve as an example of best practices. The software is built in python and leverages open-source libraries kedro, scikitlearn, MLFlow, and others. The software provides examples how to build three distinct pipelines for data query and save, data engineering, and data science. These pipelines enable scalable, repeatable, and maintainable development of ML models.	In production: less than 1 year		kedro, scikit-learn, MLFlow, pipelines	Yes				Yes			Yes	Yes		No		N/A
Machine Learning Airport Surface Model: Taxi-in Prediction	National Aeronautics and Space Administration		Ames Research Center	The ML-airport-taxi-in software is developed to provide a reference implementation to serve as a research example how to train and register Machine Learning (ML) models intended for four distinct use cases: 1) unimpeded AMA taxi in, 2) unimpeded ramp taxi in, 3) impeded AMA taxi in, and 4) impeded ramp taxi in. The software is designed to point to databases which are not provided as part of the software release and thus this software is only intended to serve as an example of best practices. The software is built in python and leverages open-source libraries kedro, scikitlearn, MLFlow, and others. The software provides examples of how to build three distinct pipelines for data query and save, data engineering, and data science. These pipelines enable scalable, repeatable, and maintainable development of ML models.	In production: less than 1 year		XGBoost, kedro, scikit-learn, MLFlow, pipelines	Yes				Yes		GitHub - nasa/ML-airport-estimated-ON https://github.com/nasa/ML-airport-taxi-	Yes	Yes		No		N/A
Machine Learning Airport Surface Model: Taxi-out Prediction	National Aeronautics and Space Administration		Ames Research Center	The ML-airport-taxi-out software is developed to provide a reference implementation to serve as a research example how to train and register Machine Learning (ML) models intended for predicting impeded and unimpeded taxi out duration. The software is designed to point to databases which are not provided as part of the software release and thus this software is only intended to serve as an example of best practices. The software is built in python and leverages open-source libraries kedro, scikitlearn, MLFlow, and others. The software provides examples how to build three distinct pipelines for data query and save, data engineering, and data science for each model. These pipelines enable scalable, repeatable, and maintainable development of ML models.	In production: less than 1 year		XGBoost, kedro, scikit-learn, MLFlow, pipelines	Yes				Yes		https://github.com/nasa/ML-airport-taxi-	Yes	Yes		No		N/A
NextGen Advanced Methods: ATCSCC Webinar Speech2Text and Analysis	National Aeronautics and Space Administration	Federal Aviation Administration	Ames Research Center	The Advanced Methods project explores the use of innovative and emerging technologies to drive post operational analysis of Traffic Management for aircraft. Technologies such as machine learning, (ML), artificial intelligence (AI), and advanced data analytics for use in improving the FAA's traffic flow management. In this specific use case, our aim is to use deep learning to convert live ATCSCC webinar meeting conversation to text, and then apply natural language processing to the converted text data for later analysis and review.	In Production: Less than 1 year		Natural Language Processing, Named Entity Recognition using Transformer Models, Entity Linking, Unsupervised Learning, Unsupervised clustering using TF-IDF/HDBSCAN/Umap, Transformer Model Embeddings, Speech2Text using Deep Nerual Network	Yes	Agency Generated	No	N/A	Yes	No	N/A	Yes	Yes		No		N/A

NextGen Data Analytics: Letters of Agreement	National Aeronautics and Space Administration	Federal Aviation Administration	Ames Research Center	Today, operation constraints are documented via Standard Operating Procedures (SOP) and Letters of Agreement (LOA) and are not made available to the public in a consistent manner. SOPs are specific to an air traffic control facility and specify the procedures necessary for safe operation in the sector. LOA outline agreements establishing procedures and responsibilities between two parties (including crossing restrictions, holding patters, emergency procedure coordination, etc.) The LOA/SOPs are published internally as scanned PDFs and are the responsibility of the facility to maintain. To reduce the manual effort of tagging the documents for ease of reference, there is an opportunity to use modern data analytics and machine learning to produce and disseminate constraints in a standardized manner. Providing LOAs or SOPs to stakeholders will enable flight planners (pilots and vendors) to study or ingest this information and thereby plan flight trajectories that remain consistent with air traffic constraints. It is also fundamental to Next Gen capabilities to share accurate data for purposes of creating new noise abatement procedures; improve NAS information for common situational awareness and alignment to implement new tools to assist in future time-based flow management.	In Production: Less than 1 year		Natural Language Processing, Named Entity Recognition using Transformer Models, Unsupervised Learning, Unsupervised clustering using TF-IDF/HDBSCAN/Umap, Trasnformer Model Embeddings	Yes	Agency Generated	No	N/A	Yes	No	N/A	Yes	Yes		No	N/A
Aero-Engines AI - a machine-learning app for aircraft engine system-performance prediction	National Aeronautics and Space Administration		Glenn Research Center	Aero-Engines AI is a Windows app that deploys machine-learning analytics to predict aircraft engine performance. The app was created using Tkinter, a GUI framework that is built into the standard Python library. Employing Tkinter greatly facilitates the sharing of machine-learning application as an executable file which can be run on Windows machines (without the need to have Python or any library installed). Current version of the app focuses on the performance prediction of conventional turbofans. The app gets user input for a turbofan design, preprocesses the input data, and deploys machine-learning analytics to predict turbofan thrust-specific fuel consumption (TSFC), engine weight, engine diameter, and core size, respectively. The predictive analytics were built by employing supervised deep-learning algorithm to study patterns in an existing open-source database of production and research turbofan engines. They were trained, cross-validated, and tested in Keras, an open-source neural network's API written in Python, with TensorFlow (Google open-source artificial intelligence library) serving as the backend engine. The smooth deployment of the machine-learning analytics using the app shows that Aero-Engines AI is an easy-to-use and a time-saving tool for aircraft engine design-space exploration during the conceptual design stage.	In production: less than 1 year		Deep Neural networks, Keras, Python, Tensorflow	Yes	Other	No	Jane's Aero-Engines (available on Amazon); FAA TCDS datasheet (https://sgl.faa.gov/Regulatory_and_Guidance_Library/rgMakeModel.nsf/Frameset?OpenPage); GRC in-house generated data; engine companies websites	Yes	No		Yes	Yes		No	N/A
Providing visualization tools and streamlining the detection and tracking of wildfire-induced smoke plumes during the Fire Influence on Regional to Global Environments and Air Quality (FIREX-AQ) mission	National Aeronautics and Space Administration		Jet Propulsion Laboratory	Started in FY21 the AI behind the Providing visualization tools and streamlining the detection and tracking of wildfire-induced smoke plumes during the Fire Influence on Regional to Global Environments and Air Quality (FIREX-AQ) mission is intended to enhance efforts in the Earth Science domain. Specifically by providing a hybrid unsupervised/supervised data processing pipeline for data fusion and wildfire/smoke identification with unique classification products from multiple instruments for further structural understanding smoke/fire dynamics.	In production: less than 1 year		Restricted Boltzmann Machine (RBM) and Conditional Generative Adversarial Network (CGAN)	Yes	Agency Generated			Yes						No	
DRAFT Ethical AI Considerations	National Aeronautics and Space Administration		Langley Research Center	Not AI, but guidance to inform AI development. Ethical AI considerations in the form of draft principles (Fair, Explainable & Transparent, Accountable, Secure & Safe, Human-Centric & Societally Beneficial, Scientifically & Technically Robust). Most draft principles were adapted from other organizations' work; the Scientifically & Technically Robust draft principle was amplified for NASA because of the Agency's culture and history of strong scientific & technical work. The guidance is contributing to multiple NASA discussions and debates regarding ethics in AI, to include how existing processes can encourage ethical AI, as well as any gaps where new approaches are needed.	In production: less than 1 year	Contributing to discussion & debate regarding how to approach ethical AI	Ethical AI draft principles										No	N/A
Parking Lot Advisor	National Aeronautics and Space Administration		Langley Research Center	Select parking lots at NASA Langley Research Center were overcrowded and employees needed parking advice to avoid spending time searching for parking. Interns experimented with a variety of mechanisms to provide advice on how full a given parking lot was. The end solution was one or more cameras strategically positioned to cover the most crowded parking lots, with machine learning image recognition used to identify full spaces, empty spaces, and give an estimated count of free parking spaces. The prototype worked well but the project was put on pause with COVID19 quarantine. Inputs included parking lot images and hand labelling of special parking spaces (e.g., handicap, visitor, etc.). Outputs included a count of free parking spots.	In production: less than 1 year	This prototype succeeded; however, it was put on pause when parking space demand plummeting due to quarantine	Image recognition, Recurrent convolutional neural networks	Yes	Agency Generated	Yes		Yes	No	Code could be made available upon project refinement/restart following quarantine. However, the imagery data is inappropriate to share due to physical security concerns.	Yes	Yes	This system is running on a firewalled on-premesis server in a local data center	No	N/A
Guided Materials Discovery	National Aeronautics and Space Administration	ARC-TI	Ames Research Center	Guided materials discovery combines predictive theory with experimental validation. Machine learning methods allow to combine electronic-structure calculations with thermodynamic and kinetic modeling; such multi-scale methods are used to predict dependence of material properties on composition and structure on multiple scales. The composition-structure-property relations are used in materials design: they help to improve and optimize properties by compositional adjustment and thermo-mechanical treatment of material samples. We provide theoretical guidance for our experimental partners, and we develop methods and software for making reliable predictions.	In production: less than 1 year		Optimal choice of ML method	Yes		No		Yes	No		Yes	Yes	NAS	No	