Project Title	Advising Team	Project Description
Cryotechnology for 3D Bioprinting and Isochoric Preservation of Tissues and Organs	Boris Rubinsky [ME], Matthew Powell-Palm [ME], Gideon Ukpai [ME]	Only 10% of patients worldwide in need of an organ receive one, while the window of organ delivery is currently limited to 4-6 hours. The most common way to preserve an organ is to cool it in order to slow its metabolism, but current methods are limited because lower temperatures lead to the formation of intratissue ice crystals that damage the structural integrity and function of the organ, making it a major deterrent to organ viability. In order to address this, our team has developed a two-phase solution of 3D bioprinting and isochoric preservation technology aiming to produce transplantable tissues and extend organ viability time to 24 hours using low temperatures. Tissues printed with a cell-loaded bioink will comprise of cells at low temperatures, increasing the overall survivability rate and allowing for more complex printing. These tissues will then be stored in an isochoric chamber, which utilizes a constant volume, high pressure system to suppress ice formation at low temperatures, effectively slowing the metabolic rate and keeping the organ cold for longer periods of time. Tissues and organs can then be globally distributed to eliminate current time and cost constraints, potentially saving thousands of lives and improving millions globally.
Affordable Electrochromic Glass	Junqiao Wu [MSE]	Electrochromic (EC) coating is one of the ways to alter glass attributes through the application of voltage on conductive layers. However, current manufacturing processes for such glass aren't optimized and required high-end equipment resulting in high price tag for the product. To make it more accessible, this capstone project approaches from an unconventional angle. Instead of inventing new glass products, a better alternative to the sputtering process called chemical bath deposition (CBD) has been utilized for fabrication at Berkeley. To date, with the help of many specially tuned processes, customized equipments and tools, a 10cm by 10cm glass prototype has been successfully developed employing the CBD process proving that the approach is indeed inexpensive and accessible.
Selective Growth of Hafnia thin films	Matthew Sherbourne [MSE]	With the advent of technologies like machine learning, data analytics, virtual reality, and augmented reality, the demand for greater processing power and data storage has significantly increased. Devices are getting smaller, more powerful and increasingly energy-efficient. Team Selective Growth Group is working on exploiting the ferroelectric switching characteristics of the latest high-k gate dielectric like Hafnia thin films. This problem will be addressed by two distinct yet synergistic approach: computational modelling and experimental. Our goal is to combine computation and experiment to drive rapid optimization of enhanced dielectric, piezoelectric, pyroelectric or electrocaloric responses and figures of merit
Innovative Materials Solution for Thermal Infrared Cameras	Kechao Tang [MSE], Junqiao Wu [MSE]	Infrared detection is the critical component for a large market of applications involving night vision, imaging, remote sensing, and smart homes. Bolometers measure the power of incident infrared radiation via the heating of a material with a temperature-sensitive electrical resistance. Our capstone project, DeepRed Technologies, seeks to develop uncooled bolometers with phase change materials for low-cost, high-performance, near-room-temperature infrared detection, and therefore eventually enhance the sensitivity of infrared cameras.
Reducing Temperature Gradients to minimize defects in 3D Parts Using Genetic Algorithms	Tarek Zohdi [ME]	The market size of the 3D printing industry has increased to 7.4 billion dollars in 2018, and users rely on 3D printers to achieve their ideas. However, most of the users seldom know what is the proper printing parameters that will give them the best printing qualities. Our team has worked on giving 3D printers the ability to learn from printing experience and to set the optimal parameters automatically so that users do not need to care about settings of their 3D printers. This new technology will make 3D printers more intelligent and more efficient.
Object Scanning and Replication Using 3D Camera	Gabriel Gomez [ME]	3D printing technology is not only becoming ever more pervasive in design and manufacturing applications, but it is also making an increasing impact to our daily life. A CAD model is usually required as an input, which are often difficult to create. Our project addresses this problem by providing a cheap and fast 3D scanning solution that creates a 3D mesh straight from existing objects, with the application of a depth camera and a scanning setup prototype. Regardless of your background knowledge of CAD, we help embody your design ideas!
Coding-Free Robot Control: Rapid Learning for Automated Processes	Chengtao Wen [Siemens], Gabriel Gomes [ME]	Industrial robot arms are widely used in various fields including manufacturing automation. When these robots need to adapt to a new production task, reprogramming of the control is the key. Our goal is to make this transition process more user-friendly, cost-effective, and adaptive, by enabling a robot arm to mimic human motions without explicit programming. To achieve this, we adopted computer vision object-tracking techniques to automatically record, analyse, and extract useful information from human motion, and execute the desired trajectory with Model Predictive Control.

Optimizing the Design and Manufacturing of Aligned Carbon Fiber	Scott Perkins [Arris Composites], Matthew Shernurne [MSE]	In 1855, a process for inexpensively manufacturing aluminum was developed, and the price dropped from $500 to $40 per kilogram in just 5 years. This process transformed aluminum from a metal reserved for royalty into the ubiquitous material we know today. Arris Composites, a Berkeley-based startup, has developed a process that could do the same for carbon fiber. Currently, manufacturing aligned composites is expensive and reserved for Bugatti's and Boeing Dreamliners. Arris' new manufacturing technique could significantly reduce the price of carbon fiber and make it a viable material for a host of new applications. This project aims to answer some of the questions raised by the development of this technology. Where could this material be used, would it be designed differently and how would it perform relative to current materials?
Bipedal Robotic Locomotion	Koushil Sreenath [ME]	In the field of robotics, legged robot locomotion is one of the hardest problems. We developed a framework for autonomous locomotion. This was accomplished through controllers, path planning, and computer vision.
IoT System-on-Chip for Bluetooth Low Energy	Kristofer Pister [EECS]	Bluetooth technology has developed rapidly during past decades. For IoT (Internet of Things) usage, power consumption and space occupation have become major concerns for Bluetooth technology. Our team aims at developing a highly customized chip that has significantly lower power, and space occupation for IoT application. The design is a robust system that can transmit and receive Bluetooth packets from other devices. In our product, we used mixer-first architecture for RF (radio frequency) transmitter and an open-source RISC-V processor with an out of order core (BOOM), and implemented our system in Chisel (hardware construction language).
Collaborative Radio Networks	Anant Sahai [EECS], John Wawrzynek [EECS]	Wireless communication is used everywhere, yet the system remains static with a fixed allocation of resources. We explore which AI algorithms create more efficient and flexible wireless communication, allowing Radio transmission stations to cooperate with each other based on urgency and need. We created a radio environment that simulated real-world communications. We encouraged a variety of multi-armed bandit algorithms to communicate with each other in many different experiments. Our experiments allowed us to discover which of these intelligent algorithms were most durable, efficient, and cooperative over time. This inaugurates a new standard for radio communications.
DanceMob: Mobile Real-Time Motion Transfer	Kurt Keutzer [EECS]	Mobile devices, which are in heavy daily use, possess enough computing power to run applications based on machine learning, particularly neural networks. Team DanceMob has used a number of state-of-the-art techniques to do pose detection and motion transfer that can run in real-time on mobile phones. We have embedded these networks into an iOS application that lets you point your camera at a person and see someone else in the exact same pose. While we have focused on entertainment applications, this pipeline of shrinking neural networks is useful on a global scale to make any application involving neural networks significantly smaller, faster, cheaper, and more accessible to the population.
Rapidly Deployable Autonomous Mobile Robots for Disaster Rescue	Alice Agogino [ME]	The number of natural disasters affecting humanity has increased over the past 50 years. Climate change has increasingly led to stronger and more frequent disasters. These can be anything from earthquakes, to hurricanes, and also forest fires. Current solutions to monitor these disaster sites, such as drones, satellites, and existing ground units, lack in their current capabilities. This is where we come in. We have created an active tensegrity robot to better serve emergency personnel in exploring and monitoring disaster sites. More specifically, this new system will help inform first responders of hazardous ground conditions and aid in search and rescue operations.
Impact-Resistant Sensing Solution for Disaster Monitoring	Alice Agogino [ME]	Squishy Robotics is developing a cost competitive robotics solution for disaster management. The robots will remotely provide first respondents with situational awareness reducing human exposure to the hazard at hand. Squishy robots are drone deployed in swarms hundreds of meters above the danger area. Their unique tensegrity structure enables them to absorb the shock from the impact of these drops. Our tensegrity structure is made from 6 Carbon Fibre rods held in tension via 24 extension springs. Additionally, 12 springs connect the external skeleton of the robot to a central payload which can carry a variety of electronics such as cameras, chemical sensors, etc. These sensors provide first respondents with visuals and sensory information prior to physically entering a potentially hazardous situation. These electronics are also well-protected in the central payload enclosure, designed and manufactured in-house by our team. A complementary software package including an intuitive user-interface which displays the analyzed sensor data is also developed in-house.
Enhancing On-board Experiences in Autonomous Vehicles	Alice Agogino [ME], Euiyoung Kim [ME]	Autonomous vehicles are causing a paradigm shift in the mobility sphere that seeks to transform the very core of the status quo. Beyond just developing the algorithms to control the vehicle, there is also a unique opportunity to redefine the user experience for commutes in these self-driving vehicles. The Connected Sensors team has designed a vibration system that can be embedded into the car chairs of the can connect seamlessly to a smartphone. The system is aimed to notify users of any important events (sharp turns, alarms, messages) through vibrational haptic feedback. For our design, we have employed the Human Centered Design methodology to explore user needs and validate designed solutions. This enabled us to design solutions that are targeted at specific user needs which consequently translates to commercial viability.

Human-Centered Design (HCD) for Shared Autonomous Vehicles Mobility in a Dynamic Society	Alice Agogino [ME], Euiyoung Kim [ME]	In the United States, the increase in population of big cities have yielded in traffic growth, raising daily average commute time. The development of Autonomous Vehicles (AV) and ridesharing companies will play an active role in alleviating this problem. However, current ridesharing companies only offer static services that do not satisfy different user's needs. The scope of this research is to create a new rideshare solution that will be flexible and automatically reconfigurable, able to adapt to a diverse set of customers in the same physical space. The final solution consists of a "partition base" ridesharing AV which is able to provide different interior configurations according to the passengers that are sharing the ride. This solution was obtained as a result of utilizing a Human-Centered Design approach: through defining user pain points and testing possible solutions through experiential prototyping. By iterating different products from user-test feedback, the final service achieved by this research addresses the main concerns expressed by the users. This new ridesharing concept will be part of the solution to the future mobility problem.
Modeling the Energy Impact of Autonomous Vehicles	Matt Zebiak [GM], Wayne Delker [Fung]	The transportation sector was responsible for 28.5 percent of U.S. greenhouse gas emissions in 2016, but the automotive industry is trying to decrease its contribution to climate change. This project quantified the energy impact of utilizing electric autonomous vehicles (AV) over traditional human-driven vehicles through the development of two AV simulation models. The models input characteristics such as vehicle speed and motor torque of a human-driven trip on a Chevrolet Bolt and output optimized AV speed profiles. Two different models were designed and tested on data from a General Motors database: an interpolation model estimates the potential energy savings and model predictive control (MPC) understands the vehicle inputs required to generate optimized AV speed profiles.
Improve road safety using autonomous vehicles	Anil Aswani [IEOR]	The goal of our capstone project is to engineer a neural net model based on open-sourced self-driving car data to perform pedestrian detections along the path traveled by an autonomous car. Detecting pedestrians accurately and efficiently while the autonomous vehicle is traveling on the road will ensure the vehicle is able to react ahead of time and prevent any fatality. The model will be able to meet the objectives of detecting pedestrians in the images using Single Shot MultiBox Detector (SSD) and Regional Convolutional Neural Network (Faster RCNN), and choosing the algorithm with better performance metrics.
Neural Networks for Image Recognition in Transportation	Anil Aswani [IEOR]	Giants like Tesla, Uber and Waymo are developing some of the most advanced autonomous driving systems. Nevertheless, Self-Driving Vehicles (SDVs) are yet to be widely adopted by the general public. One chief reason is skepticism from the consumers, especially with regards to the fundamental gaps that exist in efficiency and security. Our team, AutoSense, is studying the technological and societal contexts that play into the success of SDVs. Moreover, we have identified pedestrian detection and road sign recognition as two essential features relevant to this purpose. The result of our study is a Deep Learning model that detects and identifies objects, as well as accurately predict safety hazards to improve driving control effectiveness.
Accurate Traffic Sign Recognition using Data Augmentation	Gabriel Gomes [ME], Lars Tornberg [Volvo], Sohini Roy Chowdhury [Volvo]	Imagine that you are driving on the highway against the sun, and sunlight shines brightly in your eyes. You did not see the speed limit sign and a policeman on the side of the road with a speed detector found you speeding. You did not do it on purpose but you have to pay the $200 fine. It could be worse: on average, about 3000 fatalities are caused by stop-sign accidents alone. In other words, missing a street sign problem can be much more serious. Our project is to teach AI in a robust way to help driver detect street signs better in all weather conditions which will help safe autonomous driving become a reality. With our project, we will be able to reduce traffic accidents and save your money, and more importantly, your life.
Remote visualization and control for an autonomous driving platform	Trevor Darrell [EECS], Gary Chen [EECS], Fisher Yu [EECS]	Autonomous vehicles have the potential to make transportation safer and more efficient. However, one challenge to their development is that research requires huge amounts of data to train and test the algorithms. Furthermore, the process of collecting this data is complicated and resources-demanding.To address this challenge, we have developed a system that enables researchers to collect data by remotely piloting a vehicle while viewing real-time images and information from the sensors on the car. Our solution encompasses the whole data collection process where sensor information is forwarded through the car’s onboard computer to the user via an online web server. It collects data such as the GPS position, live videos taken by the cameras, and the measurements of the LiDAR. This enables the user to see what the car is experiencing in a human-friendly way, and without the cost and and inconvenience of running the experiments in person. Our hope is that researchers such as those without enough funding to get an autonomous car prototype can benefit from our system. More efficient data collection will accelerate research in autonomous driving and bring us safer, faster transportation.

Accelerating the Creation of Larger Scale Image Sets	Trevor Darrell [EECS], Gary Chen [EECS], Fisher Yu [EECS]	The technology of autonomous driving stands to greatly benefit humanity. Figuring out how to teach the on-board computer to classify objects from pictures of the car's surroundings is one of the most important technical challenge in autonomous driving. Current machine learning algorithms are good at this but need more data to continue to improve. Publicly available image sets such as ImageNet cannot meet the demands of the industry, so we are working to create a large-scale database for the open-source community. Our project uses active learning to partially automate our data collection. With this magnification of human effort, this data-set will be unique in its size relative to other open-source projects and provide the community with an important base upon which they can customize their own systems to further the field of autonomous driving.
Sensor Fusion and Control for Autonomous Vehicles	Francesco Borrelli [ME]	The two most difficult aspects of autonomous driving are obstacle detection and control. We have developed a platform on the Robotic Operating System (ROS) that combines radar, lidar and camera data to categorize and analyze objects in real time; then feed the data into a Model Predictive Controller (MPC). We work closely with the Berkeley MPC lab to integrate our system with their research goals and use a G80 Hyundai car to test and validate results. The end result is a real time pedestrian and vehicle detection system and controller with interactive UI to guide the car autonomously around a city block. We hope that this platform and data can be used to enhance future controller snad work towards cars sharing this data to further optimize their path planning and safety features
Sensor Fusion and Perception for Autonomous Driving	Francesco Borrelli [ME]	Many current autonomous driving vehicles suffer from poor obstacle detection by single sensor as well as the incorrect navigation information brought by different map applications. Our team is dedicated to develop sensor fusion technology, which analyzes and integrates the data from LIDAR, radar and camera. This approach will improve the accuracy and reliability of object detection and localization for autonomous vehicles.
Highway Traffic Modeling to Improve Mobility	Gabriel Gomes [ME]	Inefficient traffic management of freeways based on outdated traffic models significantly contributes to traffic congestion. Utilizing traffic flow data available from the California Department of Transportation, we fed it into our Open Traffic Model to create a simulation model for a portion of the I-210 highway most prone to congestion. Upon using Reinforcement Learning to control the model, we have been able to devise better traffic control strategies to optimize traffic flow.
Artificial Intelligence to Improve City-scale Traffic Flow	Alexandre M. Bayen [EECS and CEE]	With the introduction of autonomous driving, we can begin to explore the impacts of this new technology and the existing problems that can be solved as a result. Our team has narrowed its focus to solving the issues of traffic and safety. Using reinforcement learning techniques, we have developed AI algorithms that reduce traffic congestion and increase safety by controlling traffic lights and autonomous vehicles. By developing a city-scale scenario, our team was able to train our network to optimize throughput that can be scaled to improve traffic flow and address safety in an urban setting. [1]
Mechanism for UAV to accomplish fast pick up	Koushil Sreenath [ME]	Currently, there's an algorithm that allows the quadrotor drones to pick up and drop off cargo non-stop. However, this algorithm could not be tested due to the lack of hardware support. Our team designed a mechanism that could allow the drone to pick up cargoe while consuming zero energy during the quadrotor operation. This could be a game changer for short-range delivery. Introducing a new method characterized by high efficiency and precision, our mechanism will also enable deliveries to occur in less populated areas and areas where roads are not accessible.
Improving Bus Performance by Prioritizing Buses at Intersections	Alexander Skabardonis [CEE]	Our team's objective is to improve the public transportation service in the city of San Francisco by giving buses priority over other modes of transportation at traffic signals, also known as Transit Signal Priority (TSP). Active TSP modifies signal timings to benefit buses. We have partnered with the San Francisco Municipal Transportation Agency (SFMTA) to simulate active TSP along Geary Street, one of San Francisco's busiest corridors. Successful implementation of active TSP would reduce delay for public transit buses without significantly delaying other modes of transportation. With simulations that effectively improve bus service, the SFMTA can move forward with real world implementation starting on Geary Street followed by other locations across the city. Successful implementation would benefit bus riders by reducing travel time and increasing bus reliability. By improving the bus service, the city is encouraging people to choose public transit over cars which results in a smaller carbon footprint.
Enhancing Urban Mobility with Connected and Automated Vehicles	Alexander Skabardonis [CEE] and David Kan [CEE]	Around 1.25 million people die in road crashes every year with 20-50 million being injured or disabled. 333 million tons of CO_2 is released into the atmosphere by cars annually, and people are spending up to 210 additional hours on the road due to congestion. It is time to revolutionize urban transportation to ease the traffic woes. Connected and Automated Vehicles (CAVs) provides an exciting opportunity to enhance urban mobility while reducing the environmental impact of transportation. Our team has developed a microsimulation model of CAVs behavior along a stretch of signalized intersections to validate and quantify the host of benefits that CAVs bring about in terms of improved safety, reduced travel time and delays and less vehicle emissions.

High performance nuclear rockets for deep space exploration	Massimiliano Fratoni [NE]	With current propulsion technology, it takes NASA's best rockets years to reach targets as close as the asteroid belt, rendering deep space missions infeasible. As the space industry grows there is increasing demand for propulsion systems that open up the depths of our solar system for endeavors like asteroid mining. Our concept engine leverages the high energy densities of nuclear fuels to produce a high performance nuclear rocket to deliver a potential solution to this problem. Using recent research, new engine designs, and rigorous simulations, we seek to establish the feasibility of such a rocket and bring the edges of space closer than they have ever been.
Design and Optimization of Vertiport Locations for Urban Air Mobility	Jasenka Rakas [CEE]	Due to the ever-increasing population in urban environments, the transportation system needs to be redesigned to mitigate the congestion iand create an environment that is more city-centric and needs-driven. A promising, yet unexplored, alternative is to consider the air as a new means of transit. Our project tackles one of the biggest issues in Urban Air Mobility, which is to create a scalable optimization model that would give the user locations for vertiports -- platforms that allow aircrafts to take off and land verticall -- in the Bay Area. The model will then optimize the air traffic using these locations in addition to other constraints set forth.
Data Analytics for Better Decision-Making in Airport Infrastructure Investments	Jasenka Rakas [CEE]	In this project, the team is dedicated to find whether the vulnerability of critical aviation infrastructure system was improved by fortifications under severe weather events such as lightning strikes. The team started off by conducting a case analysis on Baltimore-Washington International Airport (BWI), where facility reliability and flight performances were examined. To obtain more convincing results, the team gradually expanded their expertise into the area of data science and network science. Using these concepts, the project team proposed investigating the relationships among all of the facilities located within the airports. The results of this project are building strong foundations for a prediction model in the future to help the FAA make smarter investment decisions.
Molten salt reactor production of Molybdenum-99	Massimiliano Fratoni [NE]	There are predicted future shortages in the global supply of Molybdenum-99 (Mo-99). Current Mo-99 production is reliant on fuel reprocessing of Highly Enriched Uranium (HEU) targets. By using Low Enriched Uranium in a molten salt reactor, my process will produce a stable Molybdenum Oxide solid that can be supplied to Tc-99m suppliers and end users. My project combines the Target supplier, Irradiation services supplier, and Mo-99 supplier services into one step. The process to extract, isolate, and stabilize Mo-99 has four steps: 1. Molybdenum Fluorides produced in the molten salt reactor are captured, along with other volatile fission products. 2. Fission product gases react with Chlorine triFluoride to upgrade any MoF4 or MoF5 to MoF6. 3. Molybdenum Fluoride is isolated by using a distillation column. 4. The isolated Molybdenum Fluoride is converted to a stable solid, Molybdenum Oxide, by hydrolysis with Hydrogen gas and steam.
Understanding Nuclear Interaction Uncertainties in Lead Systems	Massimiliano Fratoni [NE]	Many advanced nuclear reactor designs are using lead to transfer heat. However, there are large uncertainties in how neutrons interact with lead, which cause discrepancies between experiments and computational models. Performing a benchmark evaluation will investigate the causes for these discrepancies to minimize neutron interaction uncertainties, allowing for greater confidence in computational tools.
RaMHaM, A safer container to Handle Nuclear Spent Fuel Rodlets	Peter Hosemann [NE]	Los Alamos National Lab needs a container to transport radioactive spent fuel rodlets from Idaho National Lab to Los Alamos for scientific research. Without proper protection, direct contact with radioactive rodlets would cause fatal consequences to the worker operating on the rodlets. The project aims to design such a rodlet container that meets Nuclear Regulatory Commission (NRC) guidelines for transportation and research purposes. The improved container features increased radiation shielding, increasing the protection for radiological workers on the project, and remotely controlling the shutter mechanism from a safe environment. The final product makes transporting nuclear spent fuel rodlets from Idaho National Lab to Los Alamos National Lab more feasible. With rodlets become available at Los Alamos, nuclear scientists can conduct more neutron beamline tomographic research on those rodlets.
Electric Load Forecasting and Optimization for Distributed Energy Technologies	Seth Hoedl [Post Road Foundation]	Small electric utilities often lack the sensor infrastructure to monitor their load effectively. This lack of data hinders their ability to take advantage of new distributed energy resources that can improve grid resiliency and reduce peak demand. The Post Road Foundation team developed an electrical load forecasting model that uses limited data sets to accurately predict load across a year for rural electric utilities, which can eventually be used to optimize the deployment of distributed energy resources.

ParaSolar: Autonomous energy distribution network for efficient solar use	Gabriel Gomes [ME]	A microgrid is a small community of houses where each house has multiple energy source options that include trading energy between each other. In our microgrid demo, each house has a solar panel that is connected to the battery and every house has the ability to buy from another house. When a house needs more energy than its battery can provide, it will look for potential sellers and will choose the most affordable one. This is a step towards self-sufficient communities, reducing overall energy loss in transmission and saving money spent on energy usage.
ParaSolar: Solar-tracking umbrella for device charging in cafes	Gabriel Gomes [ME]	To keep our phones charged, we are forced to lug around bulky external batteries or frantically look for an outlet, which can be inconvenient or unreliable. Our project offers an additional environmentally-friendly option that can help take away the headache of keeping your precious tech charged while visiting a cafe. Our rotating solar tracking umbrella is capable of gaining about 20% efficiency over average solar panels, meant to be installed in public places like cafes to charge customer devices.
Vertical Axis Wind Turbines For Small Scale Efficiency	Alice Agogino [ME], Tom Flynn, Phil Marcus [ME], Bruce Webster	While wind power has become an increasing contributor to the global energy supply, traditional horizontal axis turbines are poorly suited to small scale applications due to lack of efficiency, danger to avian wildlife and disruptive noise and visual impacts. The Vertical Axis Wind Turbine (VAWT) Team has used computational as well as real world testing to analyze and improve the performance of vertical axis turbines. In addition to designing and testing a wind concentrator to increase airflow through the turbine the team has also studied and tested the effects of various multi-turbine arrangements. With these improvements the team aims to raise the per square foot output of vertical axis turbines to compete with other small scale renewable energy technologies.
Harness the Power of Machine Learning to Help People with Diabetes	Vikram Singh, Ricardo Abad, & Sarine Babikian [Glooko], Paul Grigas [IEOR]	Diabetes is the 7th cause of death in the U.S. According to the latest 2017 National Diabetes Statistics Report from Centers for Disease Control and Prevention (CDC), there are over 30.3 million individuals, nearly every one in ten people have diabetes in the U.S. Additionally, the average medical expenditures of diagnosed patients are about $7,900 per year. Healthcare technology companies are using big data to transform diabetes care. Glooko, a global diabetes data management company, synchronizes with over 200 different insulin pumps and provides broadly compatible web and mobile apps to monitor patients behaviors. Our team leverages Glooko's valuable data to build state-of-the-art machine learning models, aiming at understanding behavioral patterns to provide those patients with precise and customized recommendations. We are convinced that better diabetes management is an asset, with the improvement of patients' livings as well as reducing the economic cost of the society.
Data Driven Precision Healthcare Prediction	Anil Aswani [IEOR]	Healthcare prediction powered by machine learning is rising in recent years. Majority of researches have been focused on assisting in diagnosis and providing hospitals with predictions concerning insurance product costs, but few projects serve the patients directly. Our team has built machine learning models to serve as a low-cost and convenient pre-screening tool for people by predicting individual's probabilities of having a certain chronic disease, such as diabetes, mental health and thyroid disease and providing customized preventive measure recommendations.
Fit Your Genes	Elena Flowers & Kord Kober [UCSF School of Nursing], Anil Aswani [IEOR]	Obesity, a leading indicator for one's predisposition towards diabetes, is a disease whose etiology is complex due to a combination of biological, environmental, socio-economical, and behavioral risk factors. The complexity of interactions between biological and behavioral risk factors creates a challenge for clinicians in advising weight loss interventions which can be successful in treating their patients. We have developed machine learning models that account for the complexity of these interactions to identify the most predictive combination of behavioral and biological factors in weight loss. The result is a clinically useful tool that inputs a patient's step-count and gene expression data to predict whether or not a patient is likely to respond favorably to an intervention. Except for the candidates' demographic information like average step counts during trial, BMI index, gender, age etc., we also record their gene counts for thousands of different genes. To best describe the relationship between the large volume gene expression data and weight loss mechanism, we incorporate two feature engineering methods, the Principle Component Analysis and KEGG pathway, to represent candidates' gene expression in a handful way. We introduce 4 different classification models, Logistic Regression, Random Forest, Support Vector Machine and MediBoost, train them with demographic features and PCA and KEGG pathway gene features and test the corresponding model accuracy. Our final goal is to select an interpretable and accurate method which can distinguish influential factors in the weight loss process and develop it into a clinical useful tool.
Machine Learning to Predict Health Status	Anil Aswani [IEOR]	The World Health Organization redefined health as "complete status of physical, mental, and social well-being." Therefore, our team studied both chronic disease and mental health conditions. As they are preventable and cost U.S. trillions of dollars every year, we used machine learning to understand the relationship between health status and life behaviors. We built a robust (98.5% accuracy) health-status predictor through random forest model. Additionally, we created a user interface that will provide personalized life behavior recommendations within a few seconds.

Building a Question-Answering System to Simplify Access to Clinical Data	Gabriel Gomes [ME], Gundolf Schenk & Angela Rizk-Jackson [UCSF Bakar]	The Electronic Health Records (EHR) and Clinical Notes are not easily leveraged by clinicians and researchers due to technical burden. We sought to build a Questioning/Answering System where users can employ natural language to query both EHR, and clinical notes. This tool is built leveraging the UCSF Information Commons, Apache Spark, Amazon Web Services, Natural Language Processing (NLP), and other Machine Learning infrastructure to provide a website accessible service to medical experts at UCSF.
Silence excessive beeps to fight alarm fatigue in hospitals	Gabriel Gomes [ME], Xiao Hu & Michelle Pelter [UCSF School of Nursing], Alex Ackroyd	With 70% of the alarms identified as being false, nurses become increasingly insensitive to them and fail to distinguish real emergencies. Our team leveraged on machine learning and deep learning techniques to improve the classifiers of heart arrhythmias, focusing on PVCs. In addition, we enhanced the training process through the implementation of an active learning framework. As a result we were able to reduce the need for costly cardiologist annotations, without compromising the sensitivity of the alarms. The improved arrhythmia detection algorithm would solve problems associated with alarm fatigue and bring great relief to healthcare professionals, and now patients can experience better quality treatments.
IntelliCare - Machine learning to improve hospital heart rate monitors	Gabriel Gomes [ME], Xiao Hu & Michelle Pelter [UCSF School of Nursing], Alex Ackroyd	8 out of 10 bedside hospital alarms are false alarms. This costs the US healthcare system more than $2.5 billion annually and limits the effectiveness of nurses in real situations of urgency. Machine learning algorithms could help improve monitors but they require large annotated datasets. However, manual annotation of ECG records is a time-consuming process for cardiologists. In collaboration with UCSF, our team has developed a semi-automatic annotation tool focussing human resources on only the most informative ECG records. This will ultimately lead to the creation of a dataset 100 times larger than any other available resource. We developed a proof of concept by focusing on Premature Ventricular Contraction, the most frequent cardiac arrhythmia. Our team designed a variety of machine learning pipelines, including feature-based models and deep learning on time series data, to recognize irregular beats. These models were then incorporated into an active learning framework requiring minimal human input to classify unannotated data. The resulting annotated dataset will help improve existing hospital monitors and allow nurses to focus on patients most in need.
Beyond Hands: Prosthetics to Empower Differently-Abled Drummers	Alice Agogino [ME], Tomas Georgiou [ME]	Beyond Hands is a project that aims to enable and empower those who are differently-abled to be able to drum using comfortable, easy to use, and adaptable prosthetics. Existing drumming prosthetics in the market are either too expensive or are not accurate. Beyond Hands aims to create low-cost drumming prosthetics that operate accurately. To achieve this, the team has honed in on a crucial aspect of drumming that is paramount to a drummer's success: the difference between a double stroke and a single stroke. With a tighter grip on the drumstick, the drummer could play more rigid, controlled beats, which are known as single strokes. With a looser grip on the drumstick, the drummer could enable for the drumstick to play quicker, less controlled notes, known as the double stroke. The double stroke occurs when the drumstick strikes the drum pad, and since the user has a relaxed grip, the drumstick rebounds and strikes the drum pad in a quick succession, until the drummer tightens their grip to restrain the drumstick from striking the drum pad. To achieve this, the team has explored variable stiffness, which is a change in the material stiffness of a given material. Through the team's rapid prototyping process, a functional drumming prosthetic that accomplishes the single and double stroke differentiation was realized. We hope that its introduction can improve the user's quality of life.
A Pulse Detector for Monitoring Alcohol Consumption Based on Traditional Chinese Medical Diagnostics	Liwei Lin [ME], Junwen Zhong [ME], Yao Chu [ME]	Each year in the US, alcohol intoxication causes 2,000 deaths directly and 30,000 deaths indirectly in the U.S, leaving families and communities heartbroken. It is important for people to know their limits when drinking, but not everyone does. We are inspired by capability of pulse reading from Traditional Chinese Medicine and we have developed an alcohol consumption alert device based on users' pulse pattern analysis using machine learning and data mining. By wearing the pulse detector, the user's pulse pattern will be analyzed by an data trained algorithm and set a drinking threshold. After the user consumes enough alcohol to reach the threshold, the device will vibrate or beep to suggest that the user stop drinking. In the near feature, our team is looking forward to expanding the capability of the pulse detector to detect other medical diseases.

Medical Radioisotope Production through Novel Neutron Generator Technology	Karl Van Bibber [NE], Lee Bernstein [NE], Jon Morrell [NE]	Targeted alpha therapy has immense, demonstrated potential for treating late-stage cancers with remarkable precision at the cellular level. Actinium-225 is a promising radioisotope for targeted alpha therapy as it is easily conjugated with existing carrier proteins that uniquely identify cancerous cells in the body. Delivered to the cancer site by the carrier proteins, the actinium alpha decays and releases high energy, low range particles effective for killing cancer cells while minimizing effects to healthy tissues. Current use of this therapy is inhibited by an inability to produce Actinium-225 at a scale suitable for medical treatment. Additionally, several proposed production methods contain inherent, chemically inseparable byproducts with high potential for long-term side-effects. Our project aims to generate isotopically pure Actinium-225 with the redesign of a high flux neutron generator. The neutron generator utilizes fusion to produce high energy neutrons that interact with natural radium and remove a neutron from the Radium-226 nucleus (226Ra(n,2n)). Radium-225 subsequently decays with a mean lifetime of 20 days into the desired Actinium-225. A reliable supply of Actinium-225 will not only enhance the treatments of several diseases such as prostate cancer and leukemia, but will also assist research into treating viral infections and parasitic diseases.
Finding anti-aging solution based on young mice blood	Irina Conboy [BioE], Chao Liu [BioE]	Longer lifespan is often taken to be one of humanity's greatest wishes. With the rapid development of disease diagnostic methods and therapies, the human lifespan has increased dramatically over the last century. However, this increase presents new challenges: as people grow older, the number of gene mutations -- leading to cancer and neurodegenerative disease -- increase. Similarly, the wound healing system becomes less efficient due to a lack of healthy stem cells. We used parabiosis, a technique that connects blood vessels of young and old mice, to study whether some functional growth factors, characteristic of young mice, can improve wound healing in older mice. Our project aims to identify and isolate these factors for preclinical and clinical studies to slow down aging.
Pipette Aid - Affordable Virtual Assistant for Laboratory Experiment	J. Christopher Anderson [BioE]	To achieve valid conclusions, scientific researchers repeat experiments over and over to verify reproducibility of results. However, repetition of tedious experiment procedures is exhausting and stressful, increasing human error and threatening the quality of research and publications. Our solution, Pipette Aid, is a virtual assistant in the form of a LCD display, implemented with a semi-automated Java algorithm. It incorporates an event-driven GUI that walks users through complicated protocols, step by step. To begin, necessary materials are placed at designated areas on the display. At each step, target reagent tubes and appropriate wells on microplates are lit up with LCD backlight, eliminating user's burden to constantly read labels and count rows and columns. Experiment instructions and progress are also shown on the display, acting as a real time placeholder, lowering human errors in case the user is tired or distracted. Once completing the current step, simply pound on a large button to let Pipette Aid know you are ready to advance to the next. Pipette Aid is simple to use and cost friendly, perfect for laboratories on the budget. Together, we can reduce human error and add security to scientific conclusions, one step at a time.
Mechanically-Stable Graft Design for Hip Replacements	Grace O'Connell [ME]	Approximately 150,000 total hip replacements take place each year in the US for patients under the age of 65. These procedures are necessary due to damage to the hip joint caused by osteoarthritis, hip trauma, fracture, or other conditions. After 20 years with the implant, most of these patients will have undergone revision or replacement surgeries, which have a high risk of complications. Therefore, there is a need for a longer lasting hip replacement for younger patients, to reduce the need for these additional surgeries. Our project aims to find a mechanically-stable design of a bone and cartilage allograft to be used instead of synthetic hip implants. This alternative is beneficial for younger patients because it is expected to last for life after complete integration. We 3D modeled press-fit graft designs, and used finite element analysis to determine an optimal geometry. The success of this project will result in a hip graft which can support limited activities during the healing process. After the recovery period, patients will be able to return to a normal, active lifestyle.
Knee-X	Homayoon Kazerooni [ME]	Exoskeletons have been used in the industrial and medical space for years, but we are bringing exoskeletons to the recreational market with Knee-X. Knee-X can support active senior lifestyles by reducing muscle fatigue, thus, they can feel confident living a better life. Our goal is to use light but tough materials with compact mechanical design to create an exoskeleton system powered by a battery and supported with electronic controls.
Velocity: Making stationary bikes fun with virtual reality	Coleman Fung [Blue Goji], Gabriel Gomes [ME]	Imagine biking down a mountain path, fighting to keep your bike stable over the rough terrain. Compare those thrills to the monotony of using stationary bikes at gyms. Team Velocity has brought these two experiences together to create a fun and immersive exercise experience. The Velocity Bike integrates a traditional exercise bike with virtual reality gaming, all while maintaining a safe environment. Our team has developed haptic systems to provide tangible feedback to the player, allowing them to feel every bump in the road. With pedaling and steering synchronized with gameplay, riders will be totally absorbed into the virtual environment. The Velocity bike is here to make exercise fun again.

Improving Mobility for Elderly Users of Blue Goji's Infinity Treadmill	Coleman Fung [Blue Goji], Gabriel Gomes [ME]	"A growing number of older adults fear falling and, as a result, limit their activities and social engagements." - National Council on Aging One in four Americans aged 65+ falls each year. Falls are a leading cause of fatal injury among elderly individuals. A fall and subsequent injury reduces elderly individuals' confidence in walking and increases their fear of falling. This can severely limit their mobility and quality of life. Team Infinity aims to build confidence in elderly individuals' abilities to maintain an active lifestyle. By using sensors on Blue Goji's Infinity Treadmill, we build data-driven models which characterizes gait and posture that are associated with mobility. By integrating our models with the treadmill and tracking these metrics over time, we can measure the rehabilitative intervention effects, helping elderly users regain control in their everyday lives.
Workoutimal: Design and Construction of a Fatigue-Responsive Exercise Bike	Gabriel Gomes [ME]	Exercise machines help people shape and condition their bodies. But more than 60 percent of people find it hard to exercise efficiently. This smart bike gives people easy access to personalized workouts in their homes, at a lower cost and in less time than traditional exercise bikes. The goal of the WorkOutimal project is to pave the way for a groundbreaking class of intelligent self-optimizing exercise machines. The machines will automatically adapt their exercise workout routines to the health and fitness history of the user, as well as to their current level of fatigue. During the exercise, the machine collects the physical data of the user in real time, such as heart rate, speed and foot pressure, which an algorithm uses to determine the condition and fatigue level of the exerciser. The intelligent algorithm inside the controller processes the data and generates suggestions for the user to reach their target exercise goals. In conclusion, WorkOutimal provide users with a solution to exercise in a safe, effective, and cost efficient way.
BackTracker: A Smart Back Brace for Post-Surgical Tracking and Assessment	Ruzena Bajcsy [EECS], Robert Matthew [EECS]	About 19% of patients of spinal fusion surgery must undergo reoperation. As there is a lack of understanding of how patient lifestyles affect their spinal fusion, better tools are needed for doctors to track and assess their post-surgical patients. BackTracker combines Inertial Measurement Unit sensors (IMUs) with a traditional back brace to generate useful features like step count, action classification, and changes in posture. This system is low-power, highly portable, automated, and can classify and quantify patients' movements with >90% accuracy. By providing access to more high-quality data, BackTracker will enable doctors to learn about how patients are recovering, personalize post-surgical therapy and improve our understanding of the causes of reoperation.
Machine Learning for Seizure Treatment Implants	Rikky Muller [EECS]	Neural disorders such as epilepsy affect millions of Americans and even more worldwide, both physically and socially. Our project aims to develop a seizure detection model that can be integrated into an implantable device, inhibit seizures when one is occurring, and remain dormant the rest of the time when no seizure is happening for all patients. This is done using machine learning algorithms that are trained using feature calculations performed on raw EEG data that can help classify sections of the same data as a seizure or non-seizure. After selecting an algorithm along with its optimum parameters, based on its accuracy and false detections results, we can potentially create a hardware architecture for implantable devices that can utilize this capability and thus enable timely detection and inhibition of seizures.
Exploring post-stroke neural recovery via exogenous electromagnetic field therapy	Syed Hossainy [BioE]	Ischemic stroke, characterized by the restriction of blood flow to the brain due to mechanical blockage via blood clot, leads to significant mortality in both developed and emerging nations. Current therapies for ischemic stroke are flawed due to the short time window they can be administered and loss of neural function after oxygen is reintroduced. Our team aims to develop an adjunctive therapy using low-frequency electromagnetic field (LFEMF) stimulation to support neural rehabilitation, thereby alleviating patients' disability profile. It is our hope to establish a foundation for future exploration of LFEMF as a new era of stroke treatment.
Vision Correcting Displays	Brian Y. Barsky [EECS]	Sight is fundamental to the human experience. Most people will need to correct their vision at certain points in their lives. Modern vision science has solutions to vision problems such as short and far sightedness, including eyeglasses, contact lens and eye surgery. However, all these technologies have inevitable shortcomings: inconvenience, discomfort, possible side effects, and the inability to correct high-order aberrations such as coma. Moreover, about 90 per cent of people suffering from visual impairments live in developing countries, lacking access to reliable eye care or any vision correction solution. Usually, vision care providers use tools to adjust the patient's vision of the world. Our team asked the question: what if we tried to adjust the world to the viewer, rather than the viewer to the world? To enable patients to regain sharp vision, we created a screen mask which bends the light with the help of a computational model, to the exact opposite needed to correct for common vision problems, Our project aims to free people from glasses, contact lenses, and other costly vision correction tools. Instead, users can just add a mask on their screen, and they won't need to wear anything to see the screen perfectly. According to our user research, people consistently report clearer vision when using Vision Correcting Displays. In addition, this project may also help people suffering from rare eye diseases to regain clear vision without wearing special devices or having eye surgery.

Retinal Image-Based Eye Tracker for Neurological Diagnoses [CLT]	Christy Sheehy [Light Technologies] and Austin Roorda [UC Optometry]	The burden of multiple sclerosis (MS) costs the US healthcare system upwards of $28B annually. Delayed MS diagnoses and unpredictable disease progression create poor patient outcomes and inefficiencies across the healthcare system. There is a lack of objective, cost-effective, and non-invasive tools capable of monitoring disease progression and tracking treatment efficacy. Our industry partner, C. Light, is building an instrument that extracts eye motion over 120x more accurately than current devices on the market and uses disease prediction algorithms to efficiently, objectively, and non-invasively aid in the prognosis and monitoring of MS. C. Light will provide value throughout the care chain of MS: better outcomes and peace-of-mind for patients, treatment feedback and revenue for physicians, reduced overall healthcare costs for payers, and fast, objective outcomes measurements for therapeutic developers.
Machine learning to discover frauds on the bitcoin blockchain	Kevin Feng [Beam], Anil Aswani [IEOR]	Cryptocurrencies, which are built on blockchain technology, solve the problem of double spending and digital authentication, giving legitimacy in its claims to be the future of transactions. However, since its introduction in 2009, Bitcoin has majorly grown to be associated with illegal activities over the deep web, including money laundering, human trafficking, weapons and drug sale, and terrorist financing. While there exist robust measures for detecting fraudulent transaction in the fiat money realm, there is no compelling equivalent for cryptocurrencies. This is where our project comes in: a technology that could filter or better flag transactions that might be involved in money laundering or other illicit activities. This technology could prevent banks from incurring billions of dollars worth of fines and, in the long run, might even reduce illicit activities that can be carried out with ease today, and finally help bitcoin be adopted as a legitimate currency.
T-REX: A Security-focused Cryptocurrency Exchange	Dawn Song [EECS], Jian Liu [EECS]	Existing cryptocurrency exchanges are suffering from hacking events, which indicates insufficient security and weak anonymity. Our team has worked on a decentralized, real-time cryptocurrency exchange that combines Trusted Execution Environment with off-chain channels to prevent attacks and adopts a highly efficient anonymity algorithm to obscure transaction identities. The exchange service provides strong security and anonymity with high performance.
Exploiting the Power of Neural Networks for Financial Technology	Xin Guo [IEOR]	Machine learning is one of the most disruptive fields in the current technology space. Machine learning techniques are increasingly being used to predict the behaviour of financial markets. Methods ranging from traditional statistics to modern deep learning methods, all corners of machine learning are being tested, with one aim: to predict future financial behaviour. But a problem which is common to all machine learning techniques is the bottomless appetite for data. The performance of machine learning models increases with the amount of data it's trained on. But the rate at which the data are available to us is not sufficient for many prediction algorithms. We plan to solve this problem by deliberately synthesising fake data, which are similar to the real data, so that we have more data to train machine learning models on. We make sure that the generated data capture all the essential characteristics of the real data, by employing quantitative measures, like comparing the likelihood and distribution between generated data and the real ones. Our goal is to strike an optimal tradeoff between the dissimilarity and explainability of the generated data, as compared to the real ones.
Generating Simulated Financial Data	Xin Guo [IEOR]	A credit default swap (CDS) is a financial derivative or contract that allows an investor to "swap" or offset his or her credit risk with that of another investor. As a significant measurement for the credit risk and a signal of the credit risk for the market, Credit Default Swap data set (CDS) accounts for 60% of the money market. However, historical CDS data set is insufficient for professional analysts for detailed analysis and risk early warning. The first step is to mine practical meanings and relationships among all these financial orders. We would apply some statistically analysis and basic Machine Learning models for simple prediction and backtest. Moreover, our major purpose is to apply the latest neural network model, GANs (generative adversarial neural networks) into current CDS data to generate more emulational data on time series, which could be used for further risk management all over the world. On the other hand, our project could be used to predict CDS price and generate trading strategy for arbitrage as expected. Furthermore, our project could automatically generate annalistic reports about CDS and risk warning periodically
Investment Decisions Made Easy: Corporate Bond Transactional Cost Estimation	Xin Guo [IEOR]	Unlike equities investment such as stocks, corporate bonds are usually traded in over-the-counter market, the lack of transparency in terms of cost and pricing make it necessary to assess the quality of the trades. This project aim to improve the availability of crucial trading information to not only institutional but especially individual investors through extracting information from the enhanced TRACE database and utilizes varias machine learning algorithms to model and perform transaction cost analysis.

Predicting Repeat Borrowing using Machine Learning	Industry advisors: Ben Croes, David Paulsen, Jason Kremer [Prosper], Lee Fleming [IEOR], Alex Beliaev [Fung]	Partnering with Prosper Marketplace -- a leading peer-to-peer lending company -- we developed machine learning models to gain insights into what type of borrowers are more likely to return to Prosper in the future. To identify these borrowers, we sought to identify which customer engagement behavior features suggest a higher likelihood of returning. Our models will help to identify what customers should be targeted in Prosper's future marketing campaigns.
Improving Customer Experience using Machine Learning	Damien Thoulouse [Asurion], Alex Beliaev [Fung]	A major drawback in tech chat support is that customers who ask difficult questions do not receive responses fast enough from the tech experts. A system that incentivizes difficult questions can significantly reduce the turnaround time on these questions and drive up customer satisfaction. Our team is disrupting this problem space by using advanced Natural Language Processing techniques and Neural Networks. We use state-of-the-art sentence embedding methods developed by Google to represent chat transcripts and train Neural Networks on these embeddings to gauge question difficulty. The final result is an API that can smartly identify difficulty of customer questions in real-time to develop specialized pricing schemes.
Analyzing and Simulating Corporate Technology Strategy in the Aerospace & Defense Industry through a Serious ("War") Game	Troy Shideler & Chistopher Eastman [Lockheed Martin], Wayne Delker [Fung]	The Aerospace and Defense (A&D) industry is a complex, competitive, and risky environment. As such, A&D decision makers need to be able to assess industry trends and value creation patterns in order to construct strategies and resource allocation plans that build up their company's portfolio while maximizing returns and diminishing risks. The goal of this project was to build, run, and analyze a re-playable Serious ("War") Game for the industry. Our game provides decision makers with a safe and simulated environment to test different strategies and resource allocation plans. This environment allows them to understand the implications of challenges related to marketplace, industry, and government factors while implementing corporate investment strategies around Emerging Technologies (with a particular focus on autonomous technology). Players can test methods of developing talent, technology capabilities, mergers and acquisitions, and influence in order to bid for and fulfill revenue-generating contracts. The data captured from the simulation provides insights on value creation patterns and decision-maker thought processes/ biases. To incorporate realism into the simulation, game mechanics were backed by industry research done by the team. Guidance from industry sponsors and insights from data analysis both contributed to realism and data-driven conclusions.
Leveraging Machine Learning to Detect Fake Tweets	Gireeja Ranade [EECS]	With the internet becoming ubiquitous, an increasing number of people rely on social media to stay updated with the latest news. However, that also makes social media a breeding ground for people wanting to push fake content and political propaganda. Users can be easily misled by those fake news that can eventually influence the perception of the society. Our team aims to find an efficient way to to identify such fake news spreading on the social media based on the content itself (and not user characteristics) to help improve the experience of the users.
Scalable Deep Learning for Distributed Systems	John Canny [EECS]	As machine learning (ML) models become increasingly sophisticated, greater computational resources are needed. Since the introduction of hidden units, inspired from the human nervous system, artificial neural networks (a class of ML) have doubled in size roughly every 2.4 years resulting in networks as large as 16 million neurons. Even powerful supercomputers struggle to process networks this big, leading to model training times of several weeks. The potential of ML is thereby limited by the time required to train models. Simultaneously, cloud computing services have caused a sharp decline in the cost of computing infrastructure needed to be maintained, making it easier to deploy larger models. The field of distributed ML leverages distributed resources usually through cloud computing. Algorithms are sped up by running them in parallel on multiple machines and aggregating the results at a later step. The BIDMach framework, developed by Prof. John Canny's lab, is state-of-the-art for most machine learning tasks, including distributed ML. BIDMach uses Elastic Averaging Stochastic Gradient Descent and modified All-Reduce algorithm to efficiently distribute, run and aggregate neural networks. We have been able to demonstrate ~50% decrease in training time for 60% training accuracy for the popular image recognition network AlexNet when using BIDMach on a cluster of 4 machines compared to a single machine.

Second-Order Methods for Deep Neural Network Training	Michael Mahoney [EECS], Amir Gholami [EECS]	Modern development of AI technology is limited by the time and expense required to train deep neural networks. In recent years, optimization researchers have proposed that second-order training methods may help to alleviate the burden of training. Second-order methods are based on higher-order derivatives and may help to significantly accelerate the process. We have conducted a comprehensive investigation of various state-of-the-art methods in this field. Our work has resulted in software, experimental results, analysis, and practical insight for deep learning practitioners interested in employing second-order methods.
Aika: An End-to-end Automated Machine Learning (AutoML) Framework	Dawn Song [EECS], Mitar Milutinovic [EECS]	AutoML provides an automatically optimized machine learning solution for real-world data inference problems without human-involved process. Our capstone project, Aika system, is established on the idea of AutoML and involves around 200 people for development and maintenance. Our team has worked on improving the generality and completeness of Aika by contributing more primitives, the smallest functional units of a machine learning pipeline, to the system. At the same time, we have practically applied Aika to compete for real-world data science problems and improved the prediction performance after we added new primitives.
Aika: An Automated Machine Learning Framework	Dawn Song [EECS], Mitar Milutinovic [EECS]	The developing information technology generates 2.5 quintillion bytes of data. In 2016, it took a data scientist 5 months on average to complete a data analytics circle. A complete data analytics circle, including data exploring, cleaning, model fitting, and outputting results, is called a data science pipeline. Our solution is an automatic data science pipeline search system, called Aika that automatically solves data science problems and generates the desired output, like a real data scientist. Aika stores thousands of past dataset solutions including biology, sociology, math, image, and other popular areas and applies the D3M data model to search for the best answers to input problems. The auto machine learning system Aika could solve problems 90x faster than current means, which could allow non-tech background people to take advantage of data science and reduce a great amount of manpower.
Machine learning pipelines for automating the data analysis process	Dawn Song [EECS], Mitar Milutinovic [EECS]	Imagine you are a small business owner, it can be costly to hire a data scientist. Our team has developed an automatic tool for you to do data mining on your own. By defining a pipeline language to and build an engine to devise and test these pipelines, we can automate the whole process of data mining.
Developing Interactive Visualizations for Natural Language Processing Research	John Canny [EECS], Philippe Laban [EECS]	Advances in natural language processing have led to revolutionary consumer products like translation apps and text-to-speech software, but the "blackbox" of deep learning means that researchers may not understand how their models achieve great results. Our team has built an interactive visualization tool to show how the Transformer model, a state-of-the-art text summarization model, learns. Our results show that the Transformer model is able to learn different parts of semantic, high-level sentence structure like parts of speech and grammatical structure, allowing it to produce convincing summaries of complex text articles.
Examining the Water Quality Disparity in California	David Culler [EECS]	Half a million people in California are currently threatened by contaminated water, with the risk of exposure distributed unevenly among communities with different socio-economic status. Our team examined the correlation between the risk of exposure to contaminated water and the socio-economic status of communities, and our initial analysis suggests that communities with disadvantaged socio-economic status tend to have more water quality issues. Furthermore, we study Kern County as an example of such a community with both disadvantaged socio-economic status and poor water quality. Utilizing the annual water quality data published by the government, we illustrate the story of water quality disparity in California through data story-telling.
3D Volumetric Static Layout Reconstruction for Augmented Reality	Allen Yang [EECS]	Augmented Reality (AR) is evoking more excitement in recent years, powering many researches and industrial applications across social interaction, health care, media and education. There arises a demand for better development environment for 3D content creation. OpenARK plans to further the mainstream adoption of AR technologies by making AR development more accessible to software engineers. We are building an open source SDK to provide a free, flexible, and hardware agnostic tool to facilitate any developer in recreating real-world scenes as 3D images, meshing with functionalities such as point cloud sampling, mesh generation, and mesh simplification. Our product will allow more developers to express their ideas through AR/VR, and therefore give consumers more variety and volume of AR/VR content.
Towards Precise 3D reconstruction using sensor fusion	Allen Yang [EECS], Oladapo Afolabi [EECS]	Precise 3D reconstruction is hard without high-end sensors. The rise in convolutional neural network and other machine learning algorithms have changed the way people do 3D reconstruction. However, those algorithms need large amount of data to be effective. Our project serves to solve the data bottleneck by proving the concept of a sensor fusion solution. We combine cameras with different costs and accuracies together to capture data for the same object simultaneously. This will satisfy the need for different kinds of data collection goals and researches. We achieved our goal by various iterations of hardware and software designs. Some of the outputs are camera mounting, camera syncing circuits and code and data format designs.

Immersive Semi-Autonomous Aerial Command System	Allen Yang [EECS]	As technologies around of Unmanned Autonomous Vehicle (UAV) have become more prevalent, people's imaginations have shifted from ground to space. Controlling UAVs, however, is usually difficult. This difficulty has to do with the fact that traditional control interfaces are not suited for dealing with a fleet of drones performing complicated tasks. In this project, we introduce a new Virtual-Reality-based platform and interface for controlling and interacting with drones to handle 3D tasks more intuitively and efficiently. Our system would allow users to wear a VR headset and control the drones using VR controllers in an immersive environment.
Marine Ecology Virtual Reality Learning Experience	Kathryn Quigley [Lawrence Hall of Science], Allen Yang [EECS]	Pheena is the fin whale sculpture at the entrance plaza of the Lawrence Hall of Science (LHS). All the museum visitors love her. Working with the museum staff, our team worked to improve a Virtual Reality (VR) experience, which allows museum visitors, mainly children (aged 7-13), to learn about the marine ecosystem that Phenna lived in. Our team's challenge was to improve an existing VR learning experience at LHS that allows players to explore the underwater environment. We aimed to make the experience more entertaining for children, and improve the locomotion system to work with players of all ages and abilities. We added fun features and tasks to the experience, including counting marine animals, catching them, or measuring the state of the sea environment, which all mimic the real work of marine biologists. We also investigated the performance and preferences of different age groups for controlling the locomotion in this VR experience. The current result indicates that children preferred the implementation of assisted swimming motion to move around in the virtual environment. VR learning experience is a whole new frontier for education, and the experiences like ours are opening up new spaces for museum visitor experience. The research results could also be a reference for locomotion development in other VR experiences.
Tools for Strategy Consulting: USPTO	Lee Fleming [IEOR] & Andrew Toole, Nicholas Pairolero, and James Forman [USPTO]	The established routing process to a patent's approval from its issuance, at the USPTO (United States Patent and Trademark Office), is cumbersome and time consuming. What makes the process arduous is the involvement of many supervisors between its initial and final stage. We intend to use our technology to make this routing process more efficient by providing a Virtual Reality platform to allow a seamless transition in the life of a patent from its issuance to its approval for the examiners at the patent office. Our virtual reality platform provides an accurate visual representation of a patent under consideration surrounded by patents addressing similar technologies and allows for exhaustive analysis of documents using Natural Language Processing. We intend to assist the examiners by suggesting patents pertaining to topics they specialise in by considering information and data gathered through years by the USPTO.
Tools for Strategy Consulting - Siemens	Lee Fleming [IEOR], Florian Michahelles [Siemens]	Patents use technical and legal jargon that is extremely hard to understand, and it is even harder to understand the relationship between patents. As technology evolves, patent strategies are becoming more and more important. Many companies hold large patent portfolios but do not have the expertise to effectively analyze the relationship between their own patent and their close competitors'. For example, our client, Siemens, used to collaborate with Metaio in the field of Augmented Reality. However, Metaio was later acquired by Apple, which means there could be potential risks of infringement if Siemens or Metaio continue to use their shared technology. This project aims to develop a comprehensive solution, or a 'tool', that can process a large amount of patent date and visualize the relationship between data in a straight forward and interactive manner. This goal is achieved by first using a Natural Language Processing algorithm specifically tailored to patent data to make the patents understandable to the computer. Various unsupervised machine learning algorithms are then used to determine the relationship between patent. Finally, said relationships are visualized in a Virtual Reality environment, which would help the user intuitively identify patents of business interests, such as potential infringement.
Virtual Reality Patent Analytics for Technology Assessment	Lee Fleming [IEOR]	Google's patent based acquisitions range from a successful 500 million dollars acquisition of DeepMind, to a failed 12.5 Billion dollars Motorola acquisition. Google lacks access to a patent analysis and visualization tool that aids the due diligence process for patent based acquisitions. We are changing the future of Patent analytics, with a state-of-the-art platform in virtual reality (VR) for fully immersed exploratory data analysis. Our team used Term Frequency–Inverse Document Frequency (TFIDF) to represent a similarity score between a preselected set of Patents, K-means to cluster the feature matrix and used Principal Component Analysis (PCA) to reduce the dimension from 33,000 initial features to a 3D vector and visualize the patent portfolio in a VR environment.