Abstracts - June 2025

Yasmin Jafari - Impact of Acquisition Time and Motion on Lesion Quantification in Prostate Cancer Using Dynamic 68Ga-PSMA-11 PET/MRI - Norrlands Universitetssjukhus (NUS)

This study investigated the effects of dynamic PET acquisition duration and patient motion on lesion quantification in 68Ga-PSMA-11 PET/MRI of 31 prostate cancer patients. Using List-Mode data, PET images were reconstructed into 2.5-minute (14 frames) and 5-minute (7 frames) intervals, covering a total scan duration of 35 minutes. Lesions were segmented using three thresholding strategies, frame-specific dynamic thresholding, constant SUVmax median-based thresholding, and fixed voxel count, implemented in imlook4d. Additionally, a centroid-based realignment method was applied in MATLAB to correct for inter-frame lesion displacement.

Quantitative metrics, including lesion volume, SUVmax, centroid position, were evaluated across frames. Motion correction improved quantification consistency in several cases, though the extent varied between patients. A statistically significant inverse correlation between lesion volume and SUVmax was observed, with smaller lesions exhibiting higher SUVmax values. Contrary to expectations, changes in bladder volume over time showed only weak, non-significant correlations with lesion displacement ($p > 0.05$).

The findings highlight the complex impact of physiological and patient motion on lesion quantification in dynamic pelvic PET imaging. Implementing effective motion correction and applying appropriate segmentation thresholds are essential for robust analysis, supporting the need for standardized protocols in dynamic PET imaging for prostate cancer assessment.

Peter Holt - Efficient Inventory Validation using Machine Vision – Combitech AB

Imagine a world where machines can perceive and understand the visual intricacies of our physical environment, much like humans do. No, this isn't a nod to the Terminator franchise, but rather the potential of Machine Vision (MV), a technology capable of transforming industrial Quality Control (QC). Historically, QC relied on manual inspections, which are time-consuming and prone to human error, leading to inefficiencies and increased costs. MV offers an alternative by automating these processes and ensuring precision.

While MV systems often utilize Artificial Intelligence (AI), they aren't exclusively dependent on it. Consider the barcode scanner, a type of MV system that operates without AI's superpowers. However, for more complex problems, AI becomes a useful ally. In advanced MV systems, AI is typically used for object detection, where cameras and algorithms interpret visual data. This allows the system to identify and differentiate objects, making it ideal for tasks like inventory validation in industrial settings.

These systems are generally expensive, but recent advancements in AI accelerators have made it feasible to develop cost-effective MV systems. This project involves designing and implementing an MV system to detect and differentiate various sizes of screws, nuts, and washers. To expedite detection, the AI HAT+, an add-on board for the Raspberry Pi 5, has been utilized. The AI HAT+ is equipped with an AI accelerator chip capable of performing 26 Terra Operations Per Second (TOPS). The key questions were: How are speed and performance affected? And do the speed gains outweigh the complexity of integrating the AI HAT+ in terms of development time and usability?

It was shown that the processor significantly enhances detection speed, making the MV system up to 300 times faster. However, deploying object detection models requires converting them to a specific format, which impacts performance, with the F1-score dropping from 0.9491 to 0.7972 for one of the models. The conversion process is complex and limits the variety of models that can be used. This was evident when a preferred model for the task was unavailable for conversion. Additionally, models crashed during conversion if trained beyond a certain number of epochs, and even a model expected to be convertible failed to convert successfully. While the AI HAT+ has limitations regarding model compatibility and conversion, its ongoing development and community support suggest a bright future, with the potential to become a standard component in machine vision systems.

Edvin Lind - Autonomous control of a rockbreaker - a physics-informed approach using reinforcement learning – Fysik institutionen, UMIT Research lab

A rockbreaker is a machine used in mining operations to break rocks that are too large to pass through a grizzly or to push mispositioned rocks into an ore pass. Currently, rockbreakers are commonly remote-controlled, but automating this operation is of interest since the operator can suffer from limited spatial awareness since visibility can be reduced. This can lead to accidents, decreasing the longevity of the machine. In this thesis, a high-level control policy has been developed in order to automate rockbreaker operations within a simulation environment. A hierarchical reinforcement learning approach was used to learn a high-level policy that controls when and where to apply a hammer or push action based on the current pile configuration, with the aim of maximizing the likelihood of removing a rock. The agent was trained on specialized maps highlighting viable action regions for each action and later transferred to a simulation environment for further evaluation. The policy, when evaluated on similar maps, showed that 92% of predicted positions contained a rock and 70% of the positions could lead to the theoretical removal of a rock. Transferring the policy to the simulation, the performance decreased to 68% of predicted positions containing a rock and 50% of the positions leading to the removal of a rock. The hierarchical approach showed promising results for rockbreaker automation despite the performance gap between maps and simulation. Further work is needed to reduce the gap and improve the transfer of the policy between environments.

Erik Lodén - Advanced photogrammetry and image analysis for 3D construction – Knightec Group

From 3D to 2D, and back again!

Remember the time when cinema implemented a strongly suggested dress code in the salons? I'm of course talking about those stylish black-rimmed plastic glasses with a slight gray tint (or red and blue for you who enjoyed it pre-2000) used for the oh so trendy 3D effect movies. This gimmick that has sadly not survived the test of time worked, not by magic, but simply by presenting two separate views for each eye, simulating the effect that makes us experience depth. But how was this video format captured? Well, simply by using 2 cameras next to each other, one for each eye!
But did this kind of imagery die alongside its Hollywood trend going dormant? No, far from it! It still thrives in the field of computer vision and is called Stereo Vision. In a nutshell, it gives the computer the vision that we humans have. By placing two cameras at a fixed distance from each other, pointed in the same direction, and letting a computer find a common point between the two views, the distance to the common point can be calculated!
Well, isn't that neat! We have taken images using ordinary cameras, so basically created 2D projections from a 3D object and then derived back 3D data from the 2D! Alright, now we have created a computer that can ""see depth"" with the same mechanics as we do. Have you ever thought about how it would feel if you had a third eye somewhere? Well, that's hard to test, but what we can do is add more ""eyes"" for the computer by adding more cameras! But how does this affect what the computer sees? Well, it simply sees more of the object! If we continue to add more and more cameras around a subject, we can create a setup that allows the computer to see the whole object. Remember that we can calculate the distance to a common point between two images, continue to do that for all images and now we have

3D data of a whole 3D object. So with the help of a bunch of 2D images, we have been able to create a 3D object, a method known as photogrammetry.
A neat feature of photogrammetry is that if you want to make a 3D model of say, a boulder (or something else that excels at being still), you only require one camera to take multiple images. Since the subject does not move, each image from a new view can represent a new camera! This means that you are able to create 3D models with a single camera and a computer, which is considerably cheaper than doing a conventional 3D scan with LiDAR equipment. However, the absolute accuracy of this method is uncertain. Therefore, this study was made where the 3D object created from such a photogrammetry process was compared to a test-geometry, a cube of 1[dm]. The results showed a mean error of 0.2[mm], a strong result showing potential for this process to be used as a viable option in accurate and affordable 3D reconstruction!

Isak Eriksson - Drones for Precision Scanning of Large Objects and Route Planning in Enclosed Environments - Knightec Group

Unmanned aerial vehicles (UAVs) are increasingly used for a wide range of tasks, including inspections in enclosed industrial settings. In such scenarios, autonomous navigation is a valuable tool that reduces the operator's workload by enabling access to hard-to-reach areas and minimizing the need for manual labor. These environments typically lack reliable GPS signals and present physical and visual challenges, requiring advanced navigation methods for safe traversal of the environment. To address this, a lightweight, vision-based navigation system is proposed. The system relies on cameras to perceive depth and detect obstacles, while onboard processing allows the UAV to understand its surroundings in real time. It builds a 3D map of the environment and continuously updates its position within that map, enabling it to explore the environment without external guidance. To navigate, the UAV uses a strategy that focuses on unexplored areas within the field of view and chooses paths that maintain smooth and continuous motion. If no suitable path is directly visible, it then calculates a safe route to a new unexplored area using traditional path planning techniques. The system was tested in simulation and partially in real-world conditions. Real-world testing confirmed the accuracy of the mapping component, while the navigation functionality was evaluated in simulated environments. The results indicate that the system offers a valid approach for autonomous mapping and navigation in enclosed environments.

Gustav Jonsson - YOLO till visionsbaserad identifikation och lokalisering av jordbruksredskap - JOST AB

This project aims to use synthetic training data to classify and localize front loader implements using monocular vision. Blender is used to create a synthetic training dataset by modeling a virtual environment with a front loader and various implements. From this environment, a dataset is exported for training computer vision models.
The model architecture used is YOLOv11n (You Only Look Once). Two models are
trained: one for classifying implements and one for detecting the hooks and locking loops on the implements. The detection of the hooks and locking loops enables pose estimation of the implement relative to the camera using the Perspective-n-Point (PnP) method.
The models are first evaluated on animated data to assess their performance under ideal conditions with respect to environment, lighting, and camera mounting. The same models are then evaluated in real-world conditions to investigate the performance gap between synthetic and real data. For real-world evaluation, a Raspberry Pi 5 with a camera is used. A Raspberry PI 5 is a single-board computer that can be used in embedded systems to perform calculation
tasks. The detection models are assessed using confusion matrices, mean Average Precision (mAP), and F1 scores.

The results show strong performance in the synthetic environment. The implement classification model achieved a mAP of 0.993 and an F1 score of 1 at 85.3% of the confidence scale. The hook and loop detection model reached a mAP of 0.910 and an F1 score of 1 at 46.7% of the confidence scale. This shows that the models are capable of classifying implements and there hooks and locks with good accuracy in ideal conditions. However, performance dropped significantly in real-world conditions. The implement model scored 0.364 in mAP and an F1 score of 0.5 at 9.9% of the confidence scale, while the hook and loop model achieved a mAP of 0.0623 and an F1 score of 0.33 at 26.1% of the confidence scale. The acquired result in real-world conditions indicates that the models are capable of classifying implements, hooks and locks in real-world conditions, but with a significantly lower accuracy.
The pose estimation model showed better transferability. In both synthetic and real-world scenarios, the percentage error in distance estimation was around 2%, with a 95% confidence interval of −0.018 ± 0.072 m in the synthetic case and −0.027 ± 0.046 m in the real case. This indicates that the pose estimation model performs reliably in real-world applications.
In conclusion, the project demonstrates that models trained on synthetic data can achieve strong results under ideal conditions. While detection performance in real-world environments requires further improvement, the pose estimation component already shows promising potential for practical use.

Noak Wallgren - Pseudo-noise-modulated ground penetrating radar system using direct RF sampling - Guideline Geo AB

Ground penetrating radar (GPR) is a non-intrusive geophysical technique that employs a configuration of transmitting and receiving antennas to image subsurface structures. This is conventionally achieved by analysing the propagation and reflection of pulsed electromagnetic waves at dielectric interfaces. To mitigate drawbacks of conventional impulse radar, such as high power consumption, pseudo-noise (PN) modulated systems have been proposed as a more efficient, stable, and integrable alternative, albeit with greater complexity and processing demands. This thesis develops a prototypical PN-GPR system using mathematical sequences with impulse-like autocorrelation as excitation stimuli and evaluates its performance against commercial radar systems and simulated data. Designed for ultra-wideband operation centred around 100 MHz, the prototype demonstrated strong performance in subsurface feature detection. Buried reflectors at depths shallower than 8 meters were imaged with high resolution --- approaching 0.3 meters --- and multiple subsurface profiles with distinct boundaries were delineated with depths in agreement with expected values. This was particularly evident with longer stimuli, which demonstrated lower noise floors and performance comparable to other systems operating at higher peak voltages. The results were, however, highly sensitive to post-processing methodology, which --- combined with system constraints and a rudimentary data acquisition framework --- limited the conclusiveness of the evaluation. Similarly, electromagnetic compatibility analysis showed emissions largely comparable to those of a conventional system, although regulatory thresholds were occasionally exceeded. Nonetheless, enhancements in bitstream transmission and sampling rates were suggested to significantly improve system behaviour. If implemented, these findings suggest a practical viability of PN-based excitation for geotechnical surveying and utility mapping, where low-frequency, wideband GPR systems offer effective, field-deployable solutions.

Edvin Hortlund - Simulation of Heavy Vehicle Traversal on Deformable Terrain - Institutionen för Fysik, Umeå Universitet. UMIT research lab

This study investigates how terrain deformability affects the performance of heavy forestry vehicles using simulation-based analysis. A Komatsu 895 forwarder was simulated in AGX Dynamics across three test scenarios including obstacle traversal, turning, and uneven terrain. Each scenario was evaluated on both rigid and deformable terrains with varying soil deformability, vehicle speed, and load.Results show that terrain deformability is the primary factor influencing vehicle behavior. Higher compression indices lead to increased fuel consumption, deeper rut formation, and more unstable wheel-ground interactions.

Vehicle speed had a stronger effect than load, with slower speeds causing greater terrain damage due to longer wheel-soil contact duration, consistent with the time relaxation mechanism implemented in the AGX terrain model. The study highlights the strengths of the AGX deformable terrain model in capturing vertical compaction effects but also notes its limitations, most notably, the absence of lateral soil deformation modeling. Future work could include a comparative analysis using a new AGX terrain model that incorporates shear deformation to more accurately represent complex wheel-soil interactions.

Christopher Johansson - AI For Kinetic Plasma Physics - Fysikinstutionen

The movement of charged particles in electromagnetic fields is an essential and complex component of plasma physics. The charged particles are the main drivers for a wide range of plasma physics phenomena and the ability to predict their motion gives great insight into these phenomena. Today, numerical methods like the Boris method is used in order to predict the particles trajectory, but these numerical methods scale poorly to simulations with millions of particles. This issue raises the question: can a machine learning model learn to perform particle trajectory prediction such that it outperforms the current numerical methods? In this Master's thesis, we develop a machine learning model and compare it to the existing numerical methods in terms of efficiency, accuracy and scalability. In order to evaluate the models performance in different scenarios, we test various conditions of the particle parameters and background fields. The RNN machine learning model that is developed will prove to be much slower than the existing numerical methods but also very accurate. We find that while the machine learning model is not a replacement for the existing numerical methods, it does prove to be a valid alternative. The numerical methods are subject to certain constraints that the machine learning model is not, giving it some advantages over the numerical methods. This Master's thesis thus concludes that a machine learning model can indeed, accurately though not efficiently, perform charged particle trajectory prediction.

Wilma Andersson - Predictive Modeling of Brightness in the Pulp Mill Bleaching Process Using Supervised Learning Methods - ABB

Pulp demand has increased in recent years due to rising environmental concerns and measures like plastic taxes aimed at reducing plastic use. These changes have made producing high-quality pulp essential to meet customer expectations and needs. At the same time, the industry generates large amounts of data during production, opening up opportunities to use AI for better monitoring and control of pulp quality.

In pulp mills, technicians adjust processes to maintain the desired quality. Since real-time measurements like brightness often don't match the more accurate, time-consuming, lab tests frequent sensor calibration is needed.

We explore methods to measure pulp brightness using machine learning models, focusing on making continuous real-time measurements to improve control of the production process. Accurately predicting pulp quality in real time is challenging in mills like SCA Ortviken due to shifting data and lack of steady-state conditions.

Our result supports earlier studies showing that machine learning models struggle in changing, non-steady environments. While they perform well with stable, known data, their performance drops when applied to new periods with shifting distributions. This drop in performance is reflected in cross-validation scores, which vary depending on whether the data is shuffled or not, highlighting the variability. Such variability highlights the bigger problem of making reliable predictions when conditions keep changing.

These results underline the need for adaptable models and highlight how digital technologies and AI can update traditional industries, improving efficiency and product quality.

Alexander Enare - Measuring species of lichen in forestry stands using detection algorithms - SLU, skogens biomaterial och teknologi (SBT)

Hanging lichen play a crucial role both as a winter food source for reindeer and as a sensitive bioindicator of environmental change, including rising nitrogen levels linked to climate change. Declines in lichen populations can serve as early warning signals of ecological imbalance.

However, previous methods for detecting hanging lichen species, such as Usnea spp. (US), Alectoria sarmentosa (AS), and Bryoria spp. (BS) rely on ocular tools like relascopes, which introduce human bias and provide limited spatial insight. In this thesis, we propose a machine learning based method for detecting and monitoring two key lichen groups, US and AS/BS, in forest environments using image analysis. While machine learning has previously been applied to detect ground lichen and berry shrubs, its application to hanging lichen represents new ground.

We created a custom dataset from high resolution forest imagery, as no suitable dataset existed, and trained a YOLOv8s object detection model. The model achieved an F1-score of 0.79 for US and 0.62 for AS/BS, the latter affected by low visual contrast in complex forest scenes. To support practical deployment, we devolped a user interface (UI) for streamlined inference, visualization, and data export. We standardized image collection using a custom-built tool that captured images in the four cardinal directions.

We applied our method in a pilot study at SLU’s Åheden research forest. Where the ns max zone exhibited the highest lichen presence, which is consistent with its goal of maximizing ecological value. Our approach delivers a practical, reproducible alternative to manual surveys and lays the groundwork for automated, long term lichen monitoring.

Julia Servin - Modeling fluid dynamics of Y-shaped surgical graft: For application towards selective antegrade cerebral perfusion - MT-FoU

Selective antegrade cerebral perfusion (SACP) using a Y-shaped surgical graft is one protective strategy used to maintain blood flow to the brain during aortic arch surgery. To enable patient-specific predictions of intraoperative blood flow during SACP, the flow dynamics of the Y-graft must be characterized. This thesis presents a combined experimental and numerical investigation of the Y-graft, with the objective of developing a validated numerical model that can be incorporated into vascular tree simulations through flow resistance-based boundary conditions. The model ability to predict pressure difference between the perfusion sites was evaluated using intraoperative data from 15 patients, achieving an average model error of \$0.07 \pm 3.36\$ mmHg. Numerically derived expressions of the graft resistance in the right branch was incorporated into patient-specific vascular tree simulations, yielding no significant change in errors compared to previous simulations without the graft incorporated. This work lays the groundwork for identifying remaining sources of error in patient-specific simulations, which may likely stem from assumptions in the vascular tree models or limitations in intraoperative measurements.

Moa Hjorth - Topology optimization of anisotropic superhydrophobic surfaces for transition delay - Design Optimization Group, UMIT Research Lab, Umeå University

In this master's thesis, the use of superhydrophobic surfaces (SHSs) to delay laminar-turbulent transition is studied. SHSs are inspired by the water-repelling properties of the leaves of the lotus plant, and have been shown to reduce drag and delay laminar to turbulent transition, both experimentally and numerically. In this work, direct numerical simulations are used to simulate channel flow with SHSs placed on the channel walls. The aim is to use topology optimization to design the macroscopic layout of anisotropic SHSs that delay K-type transition. The anisotropic SHSs consist of microscopic roughness elements in the form of ridges and are implemented through a Robin-type boundary condition. Two specific cases of anisotropic SHSs are studied, one where the ridges are aligned in the streamwise direction, and one where the ridges are aligned in the spanwise direction. Optimization of the SHSs is performed with gradient descent and adjoint methods, and a continuation strategy is applied to enable convergence to better designs. The two optimized surfaces succeeded in delaying K-type transition compared to both a no-slip surface and their homogeneous counterparts. For the case with streamwise aligned ridges, the transition was delayed 53% compared to the no-slip surface and 19% with regards to a homogeneous SHS with streamwise aligned ridges. For the case with spanwise aligned ridges, the transition was delayed 9.7% and 29%, respectively. The transition delay by anisotropic SHSs was attributed to an inhibited growth rate of the secondary instabilities.

Filip Nordvall - A CFD Analysis of Thrust Losses in a Solid Rocket Motor Nozzle - Totalförsvarets forskningsinstitut

Rocket propulsion is an ancient discipline built on centuries of innovation. Yet, the internal physics of rocket motors remains incompletely understood, with modern development still relying heavily on idealized models and empirical correlations. Recent advances in computational power have enabled a new approach that offers the potential to uncover the complex flow phenomena inside rocket motors: Computational Fluid Dynamics (CFD). This work presents a comprehensive numerical framework for predicting the performance of a converging-diverging de Laval nozzle operating with a composite propellant. The simulation captures compressible, reactive flow and the effects of suspended liquid aluminium oxide particles, employing Large Eddy Simulation (LES) with an implicit sub-grid formulation (ILES) to resolve the unsteady three-dimensional dynamics. Chemical kinetics are governed by a 97-reaction kinetic mechanism model and integrated using a third-order Rosenbrock ODE solver. Dispersed alumina particles in the nozzle flow are tracked with a Lagrangian Particle Tracking (LPT) module that exchanges momentum and heat with the gas phase. The method is validated against experimental data from a static-fire test of a Ballistic Test Motor (BTM) conducted at the Swedish Defence Research Agency (FOI). Simulated nozzle efficiency reaches 97.1%, matching the experimental result of 96.0%. Beyond serving as an accurate physics-based alternative to traditional one-dimensional ballistics codes, the framework provides spatially and temporally resolved flow fields that enhance physical insight. Though applied here to a lab-scale BTM, the methodology is extendable to tactical-scale motors, supporting future propulsion design and test-data interpretation.

Niklas Edlund - Predicting parameters of Bioreactor Dynamics using inverse Fourier Neural Operators - Sartorius Stedim Data Analytics

In this thesis we propose a data-driven framework for quick estimation of parameters of bioreactor dynamics with a one-dimensional variational autoencoder (VAE) to impute both missing variables and datapoints. These reconstructed trajectories then fed to a type of Fourier Neural Operator [8] (FNO) designed specifically to solve the inverse problem of ODEs. We focus on immortalized cell cultures in fed-batch reactors whose dynamics are assumed to be governed by a certain set of ordinary differential equations describing the density of viable, dead, lysed cells, byproduct accumulation and substrate consumption. To solve the inverse problem, inferring parameters from the bioreactor dynamics, we generate a large amount of synthetic data and train a model to first reconstruct full trajectories and then train the FNO to predict the underlying parameters of the system. This novel architecture is compared to a more naive approach using only a VAE and a multi-layer perceptron to infer the variables from the latent space. A curriculum learning scheme is used to mimic sparse time measurement data. At inference, our model delivers parameter estimates in roughly 10 ms on a CPU and generates stochastic trajectory ensembles. Validation on fed-batch data demonstrates close agreement between trajectories that arise from the predicted parameters and the true trajectories. We discuss how our approach can greatly reduce the time required to estimate parameters from bioreactor data and outline future improvements like filtering dynamically irrelevant parameter regions and improving the VAE part of the network.

Johan Borgelin - Simulation of particle accumulation during thermal runaway - Scania

A recent surge in fire accidents as a result of thermal runaway (TR) occurring in battery packs has led to companies studying the event using simulations to improve the safety of their electric vehicles. TR is the evacuation of hot particles and gases from a cell as a result of electrical, mechanical or thermal overstressing. Due to computational costs, these simulations often neglect particles. Studies have shown, however, that accumulation of particles during TR can ignite the surrounding gas, melt walls and short circuit the battery pack, escalating the event even further. The aim of this thesis is therefore to find a fitting simulation method capable of predicting particle accumulation in a battery pack during TR. The goal is to use STAR-CCM+ to deliver a fast and accurate model where simulation settings such as physics, mesh and boundary conditions are included in the deliverable. Lagrangian multiphase was used together with a sticking criteria consisting of the particle's mass, velocity, temperature and incidence angle as parameters. To validate the sticking criteria, results from an experimental TR test were used. In the test, x grams of gases and particles were ejected from a cell into a test channel, where y grams of particles accumulated, distributed over four different regions. To speed up the simulation, removal of stuck particles was incorporated in the model, resulting in a decrease in runtime of 45%. A method of lowering the porosity in cells of deleted particles was investigated in an attempt to improve accuracy, but paused after news of an upcoming STAR-CCM+ release making the method obsolete. The final sticking criteria simulation showed a total accumulation of z grams of particles, with deviations of a, b, c and d grams in the four regions compared to the experimental test. The results are promising relative to the total of x grams being ejected from the cell. However, questions regarding the generality of the sticking criteria remain, where improvements such as replacing the temperature parameter with a conjugate heat transfer parameter are discussed.

Malin Gidlund - Statistical analysis of terrestrial bow shock ramp generator behaviour using bow shock database created from Magnetospheric Multiscale mission measurements - Umeå Universitet, Institutionen för Fysik

The terrestrial bow shock is a collisionless shock formed in the interaction of the solar wind with the Earth's magnetosphere. At the bow shock, the solar wind is rapidly compressed and heated, and the magnetic field is piled-up. This forms the bow shock ramp. Within the fluid model description of a plasma, we expect the ramp to act as a generator, converting kinetic energy into thermal and electromagnetic energy. This generator is then expected to connect to load regions, either directly upstream and downstream of the ramp or at more distant locations, where the energy can be transferred back to the particles through acceleration. In reality, this textbook behaviour, of a generator ramp, is not guaranteed.

The goal of our statistical analysis is to study the power density ($E \cdot J$) in the terrestrial bow shock ramp and surrounding plasma regions, in order to explain the distribution of textbook and non-textbook behaviours and the variation in ramp generator capacity for textbook ramps. We therefore study the power density's dependence on upstream and downstream disturbances, the angle between the bow shock normal and the interplanetary magnetic field θ_{Bn}, the upstream ion beta β and the Alfvén Mach number M_A.

In order to make this statistical analysis we require a reliable dataset of ramp intervals, with associated upstream and downstream intervals for parameter calculations. We therefore create an in-house database of Magnetospheric Multiscale mission bow shock ramps, using approximate ramp timestamps from the SHARP Shock database. For our analysis, we use a ramp interval dataset of 217 bow shock crossings from our database, and θ_{Bn}, β and M_A imported from the SHARP Shock database.

Our results show that only ~ 60% of the ramps exhibit a textbook generator be- haviour. We hypothesize that the non-textbook behaviour is due to disturbances upstream and downstream of the ramp, since the disturbances can cause a solar wind kinetic energy to be converted before it reaches the ramp. In our study we see a decrease in generator capacity at the ramp with increased disturbances in the surrounding plasma regions, as we expect. From our study, we also note a decrease in ramp generator capacity for increasing β and MA. We suggest that the power density's dependence on β and MA, is mainly due to a dependence on the upstream magnetic field magnitude B. Since B downstream is on aver- age three times as large as B upstream, and the current density J is calculated from the B variation over the ramp, we hypothesize that the magnitude of J at the ramp (and therefore also the power density) depends on B upstream. This hypothesis is supported by our results.

Sebastian Lundström - Quantum Relativistic Theory: Corrections to Vlasov Theory in the Weakly Coupled $e^- – e^+$ Limit - Institutionen för Fysik, Umeå Universitet

The Vlasov equation governs classical collisionless plasmas but fails to capture quantum effects that can significantly alter plasma dynamics. Known effects arise at high densities n_0 (where wavefunctions overlap), low temperatures (T ≤ T_F, where degeneracy suppresses thermal excitation), and under strong fields near the Schwinger limit (E_cr, where pair production becomes significant). However, the combined regime of weak fields ($E \ll E$_cr) and low densities–where long plasma periods allow even weak fields to perturb the Dirac vacuum–remains largely unexplored. To address this gap, we study this regime in this work. By applying the Dirac–Heisenberg–Wigner (DHW) formalism to a onedimensional electrostatic field, we derive new equations in the Weakly Coupled Electron-Positron (WCEP) limit—valid for weak ($E \ll E$_cr), slowly varying ($\partial/\partial t \ll$ 1/τ_C) fields. Building on this, by assuming homogeneity ($\hbar\nabla$_p∇_r $\ll$ 1), we derive an evolution equation involving cubic nonlinearities. Although the leading nonlinearities cancel, the remaining terms reveal signatures of vacuum polarization—a known quantum effect—along with additional contributions that may indicate a previously unrecognised form of vacuum–matter interaction. The latter depends on the initial distribution of real particles, such that in the absence of matter it vanishes completely. Notably, despite recent advancements in laser technology, current field strengths remain well below the Schwinger limit (E_cr), making these weak-field quantum corrections not only theoretically interesting but also experimentally relevant. To conclude, although the quantum effects in this regime are small, Vlasov theory does not capture them. Thus, for a more precise description of weak-field, low-density plasmas, these quantum corrections should be included. Although our case is highly simplified, the emergence of quantum effects in this regime helps us identify where quantum theories should be chosen over classical Vlasov theory, and calls for further investigation of the regime. Moreover, a deeper understanding of the Dirac vacuum could benefit not only quantum electrodynamics (QED) but also other fields, including the design of new Dirac materials—graphene being a noteworthy example.

Henrik Berggren - Optimization of a hydraulic block with respect to additive manufacturing - BAE Systems Hägglunnds AB

Additive manufacturing (AM) is a method of production that has been under steady development during the last decades. It has shown to be a technology with great potential since it has many benefits such as short lead times, the ability to utilize various materials and a completely different set of design constrictions. Many different industries today use the technology to various degrees. One example is 3D-printing plastic models of components which can be used as prototypes. Still there is a desire to use AM more directly by making parts that are made of metal. A part created with AM that is made from metal can be produced faster and with a completely different design than what traditionally would be possible. This thesis aimed for BAE Systems Hägglunds AB (BAE) to explore the possibility of creating a component that is featured in one of their vehicles with the AM technic of powder bed fusion (PBF). PBF have seen much development and is suitable for smaller objects with higher detail. One component that fulfilled these requirements was a hydraulic manifold which was used in the vehicle and would be a great example of what would be possible to achieve with AM. BAE was interested in what benefits, possibilities and changes to the workflow AM would bring with it. To answer this research was done and the hydraulic manifold was redesigned adapting to the constraints of AM but also implementing aspects so to reduce the pressure drop and reduce the weight by topology optimization (TO). During the redesign, ongoing conversations was had with machinist and a supplier to print this design so to also eventually see how the postprocessing would be done. The redesign resulted in a new hydraulic manifold that would be swappable with the already existing block. It was designed with a new material in mind going from aluminum to steel and with the TO it was still reduced a little in weight. The design had rerouted channels for the hydraulic fluid to experience less pressure drop and with a new channel profile to be more printable. The thesis ends with recommendations of what final steps should be done when the part is printed. The printing and machining of the design will be done towards the end of the summer by other people. This since there were no earlier timeslots to have the part printed.

Henrik Rönnqvist - Induced currents and joule heating in reactor hall floor - Hitachi energy

This study investigates Joule heating in the reinforcement grid, found in the concrete floor, of reactor halls within high voltage direct current (HVDC) converter stations. The analysis is conducted using two finite element models developed in COMSOL Multiphysics; a detailed realistic model, and a simple model that approximates the grid as a homogenized block with equivalent electromagnetic properties. The purpose is to evaluate the magnetic field generated by the reactors and induced currents in the floor reinforcement, and to quantify the associated heat losses. Key modeling assumptions and design decisions are validated through convergence studies, benchmarking of conduction model for the rebar (Amp`ere's law, impedance boundary condition and transition boundary condition), and sensitivity analysis with respect to grid geometry and material parameters. Trends in heat generation are found to be primarily linear with respect to surface current density, enabling both models to be summarized by general predictive equations. A comparison between the two models reveals a consistent offset in predicted heat losses, with the simple model underestimating losses by approximately 16% across the tested range. This offset is shown to be systematic and can therefore be compensated for in practical applications. The resulting framework provides a scalable and adaptable toolset for assessing heat generation in reactor hall floors, supporting future design strategies.

Gustav Björkman - Autonomous Drone Scans Of Large Objects Using Vision-Based Object Detection - Knightec Group AB

Unmanned aerial vehicles, commonly referred to as drones, is an area of research, technology and use that is growing rapidly in commercial and industrial applications such as remote inspection and sensing. In conjunction with new algorithms and technology related to autonomy emerging, new applications using autonomous drones are constantly being developed and sought-after.

This thesis explores equipping a drone with a camera and a vision-based object detection model in YOLOv5n to perform autonomous navigation and scans of large objects. The method uses YOLOv5n inference to detect the object and navigates using PX4-autopilot and ROS 2 nodes based on the detection data and pre-set user parameters. To evaluate the implemented method, drone simulations are run in Gazebo, where a custom-made object has been implemented and used as training data using the YOLOv5n model to simulate a large object. The simulations highlight how the drone identifies the object and travels towards it before starting a scanning sequence. Furthermore, the scanning sequence consists of an orbital motion with increasing height around the center of the object while capturing images. A slight discrepancy was observed between the orbital motion and the center of object and was attributed to sensor calibration and distance estimation errors. The method proposed is seen as a promising step towards creating a complete solution capable of performing autonomous scans of large objects using drones. However, the method will need to be physically implemented and tested to ensure validity under real-life conditions. Further possible improvements and areas of exploration include equipping additional sensors or training the model to consider object avoidance.

Emanuel Stenmark - Magnetic Resonance Imaging and Computational Fluid Dynamics for non-invasive assessment of venous stenoses in Idiopathic Intracranial Hypertension - Medicinsk teknik, forskning och utveckling (MT-FoU), Norrlands Universitetssjukhus

Idiopathic intracranial hypertension (IIH) is a condition characterized by elevated intracranial pressure (ICP) in the absence of an identifiable cause. Venous sinus stenoses have emerged as a possible hemodynamic factor associated with increased ICP in a subset of patients. Clinical assessment often relies on invasive manometry to assess pressure differences and guide the decision to perform venous sinus stenting. However, intravascular catheters introduce procedural risks and may themselves perturb venous pressure dynamics.
This study investigates the possibility of using a non-invasive method combining four-dimensional phase contrast magnetic resonance imaging (4D PCMRI) and computational fluid dynamics (CFD) to estimate pressure differences along the dural venous sinuses in IIH patients. Ten subjects, five diagnosed with IIH and five healthy controls, underwent 4D PCMRI at Norrlands Universitetssjukhus. The venous system was segmented manually using Synopsys ScanIP™, from which CFD-compatible meshes were generated. Flow measurements were semi-automatically obtained from 4D PCMRI velocity data, and numerical simulations were conducted in COMSOL Multiphysics™ to compute pressure differences over the dural sinuses.
The results show that the highest simulated pressure differences across the venous sinuses were significantly elevated in IIH subjects compared to controls ($p = 0.0013$, Mann–Whitney U test). All IIH subjects also exceeded the empirically derived threshold of 1.59 mmHg. A strong positive correlation was observed between maximum pressure gradients and clinically measured ICP (Pearson's $r = 0.869$, $p = 0.0011$). Sensitivity analysis further confirmed that even minor changes to vessel geometry (e.g., a 1 mm reduction in stenosis diameter) substantially altered simulated pressure gradients, highlighting the method's responsiveness to anatomical variance.
These findings support the use of simulated venous pressure differences as a promising non-invasive marker for identifying pathological hemodynamics in IIH. In addition to its diagnostic potential, the study also contributes to the understanding of IIH pathophysiology by illustrating how elevated venous pressure can disrupt cerebrospinal fluid dynamics and lead to sustained intracranial hypertension. Clinical implementation remains hindered by the segmentation's reliance on user expertise and the complexity of CFD modeling. Future efforts should therefore prioritize automation and robust validation against conventional measurement methods.

Henrik Wahlström- Particle-in-cell plasma model optimization for GPUs - Institutionen för fysik

Amitis is a 3D GPU-based hybrid PIC model for plasma physics. The solver moves particles, solves field equations and maps quantities between free-flowing particles and a discrete lattice of nodes. The mapping steps of PIC models are bottlenecks when implemented on the GPU, and in thesis thesis we explore methods of alleviating one of those bottlenecks, specifically mapping particles charges to charge densities.

We try two different methods, both implementing reduction algorithms in the small but fast shared memory available on NVIDIA GPUs. The first method requires a minimal amount of shared memory but depends on sorting the particles each iteration. The second method requires a larger amount of shared memory but does not require sorting as often.

The first method maps the charge density effectively, but the preprocessing step of sorting particles each iteration makes the method in total slower than Amitis' implementation. The second method significantly reduces the performance penalty from sorting but due to the large amount of shared memory needed, only 60\% of the GPUs resources can be utilized, making the performance worse or on par with the Amitis' implementation.

Shared memory is an effective tool, but to implement an effective reduction algorithm while not depending heavily on slow sorting requires more shared memory than is available. There is also no clear trend for the growth of the shared memory with newer releases of GPUs. In conclusion, the mapping steps of 3D PIC methods does not benefit from reduction algorithms in shared memory due to hardware limitations.

Theodor Jonsson - Addressing Sparsity in Multi-Market Retail: A Vision-Language Model Enriched Transformer for Robust Personalization - Sift Lab AB

Personalized recommendation systems are crucial in e-commerce for driving sales and enhancing customer experience.
While they are vital for suggesting relevant products from vast catalogs, traditional linear recommendation models often struggle with complex user behavior, products with limited interaction histories, and filter bubbles that limit product discovery.
Transformer-based architectures offer a promising alternative and have demonstrated strong performance on session-based data.
However, their effectiveness on sparse transaction data across diverse markets remains largely untested.
To address this gap, we developed a personalized recommendation engine using the transformers4rec library enriched with semantic meta-embeddings derived from vision-language models.
We integrated this system into the production pipeline at Sift Lab, a data analytics platform specializing in predictive analysis across various e-commerce markets.
Benchmarking this model on real-world datasets against Sift Lab's existing linear embedding aggregation model, using Normalized Discounted Cumulative Gain (NDCG@10) as the standard ranking metric, achieved compelling results: an average NDCG@10 improvement of 94.9\% and a significantly increased ability to handle product sparsity.
These outcomes demonstrate that fusing semantic enrichment with behavioral modeling can mitigate effects of data sparsity and unlock more personalized, diverse recommendations, especially for novel products, offering a robust tool for the next generation of retail AI.

Philiph Lundberg - Reality to Simulation: A Scene Understanding Approach to 3D Log Pile Scene Reconstruction - Institutionen för fysik

This thesis presents a pipeline for physically accurate reconstruction of log pile scenes from RGB-D data, connecting real-world perception and physics-based simulation. The proposed method integrates SAM-6D, a zero-shot 6D pose estimation framework, with AGX Dynamics, a high-fidelity physics engine. Starting from RGB-D images and a CAD model reference of a log, SAM-6D identifies and performs 6D pose estimation for each individual log. The resulting poses and segmentation masks are used to infer the terrain beneath occluded regions through interpolation, generating an initial heightfield guess. Direct simulation of the predicted scene, however, does not often result in stable configurations. To address this, a heightfield optimization process is introduced. The terrain under each log is perturbed locally, and candidate configurations are evaluated in simulation using a loss function that penalizes deviation from the predicted poses, accumulated linear and angular velocity after spawn, and terrain distortion. The system is evaluated on synthetic log pile scenes under varying conditions in three different tests: AGX generated log pile scenes, repeated optimization on a poorly performing configuration, and added environmental complexities using Blender. Results show that the optimized simulations achieve median position errors of 18 mm, which is 7% error relative to the chosen log diameter, and angular deviations below 1° after letting the logs settle for 78 AGX-generated scenes, with a resulting performance increase of 59%. The heightfield optimization also demonstrates consistency across 57 repeated runs on a log pile configuration with poor initial stability, resulting in a 94% improvement. The pipeline successfully segmented all three logs in 5 out of 10 Blender generated scenes. For these, the optimized simulations achieved a median position error of 23% relative to the log diameter and angular deviations of 1.6°."

Joel Vik - CFD investigation of spillway flows and structural stressors - Evaluating operational conditions and aerator integration - Vattenfall R&D

Hydropower remains a cornerstone of Sweden's renewable energy system, with spillways playing a vital role in ensuring dam safety by safely conveying excess water downstream. However, long term exposure to extreme hydraulic conditions can lead to structural deterioration in spillway components.

This study presents a numerical investigation of hydraulic conditions within a spillway exhibiting signs of structural deterioration. Using computational fluid dynamics (CFD) in ANSYS Fluent, flow variables were analyzed under various operational and geometric scenarios. The objective was to identify mechanisms contributing to material erosion and to evaluate potential mitigation strategies.

A two-dimensional model was initially developed and validated against discharge measurements. Mesh sensitivity studies were conducted to ensure numerical reliability, followed by simulations assessing the effects of varying tailwater levels and gate configurations. Results indicated that reduced tailwater levels and partial gate closures amplified damaging hydraulic conditions in varying degrees. In contrast, higher downstream water levels generally improved flow stability and reduced risk factors.

To address the identified vulnerabilities, a three-dimensional model incorporating an aerator structure was developed. The addition of the aerator successfully eliminated critical low-pressure zones near the spillway crest and both reduced- and redirected peak hydraulic loads to the nappe impact region, significantly lowering cavitation potential and associated risks. Two aerator configurations with varying dimensions were tested and produced nearly identical results, highlighting the the potential for design flexibility and the value of exploring further variations for optimization.

The results suggest that targeted geometric modifications can effectively improve spillway resilience, offering a foundation for future design improvements and operational guidelines aimed at minimizing long term structural damage.

Johan Jonsson - Energi- och effektiviseringsanalys av mobilhydraulik i frontlastare genom modellering och simulering - JOST Umeå A

På grund av klimatomställningen står jordbrukssektorn inför stora utmaningar. Energieffektiviteten hos traktorer blir allt viktigare i och med elektrifieringen, och frontlastare har en central roll i traktorers användning. Genom modellering och simulering har det här examensarbetet undersökt möjligheterna att förbättra energieffektiviteten hos frontlastare. Först genom att kartlägga energiåtgången för dagens lastkännande system med hjälp av loggade körcykler från en verklig traktor och frontlastare. Resultatet visade att mest energi går åt till strypning i ventiler och slangar, ca 50 \% av den mekaniska energin in till pumpen. Den hydrauliska effektiviteten var mellan 32-46 \%. Sedan utvärderades möjligheten att förbättra dagens system genom mjukvaruoptimerad styrning och elektronisk lastkänning, e-LS. En styralgoritm togs fram och resultatet jämfördes med dagens system. Resultatet visar att det finns potential att minska energiförbrukningen med upp till 43 \% för att utföra samma arbete, på grund av minskade strypförluster. Möjligheten att regenerera energi analyserades även. Ca 20 \% av den ingående mekaniska energin är möjlig att regenerera, vilket har potential att minska energiförbrukningen med runt 8 \%. För att validera resultaten i verkligheten bör verkliga tester av e-LS genomföras och fler körcykler bör loggas från verklig användning för att kartlägga energiförbrukningen närmare. Resultatet visar dock tydligt att dagens system har stora energiförluster genom strypning som det finns potential att minska.

Teresia Granberg - Computation of cerebral arterial pressure drop using 4D flow MRI - Medicinsk teknik, FOU

Accurate estimations of cerebral pressures are essential in clinical contexts where cerebral perfusion might be disrupted, such as during vascular surgeries involving temporary interruption of blood flow. This is particularly important since direct measurements require invasive procedures, which always entails risks. A previously established method combines computational fluid dynamics (CFD) with 4D Flow magnetic resonance imaging (4D Flow MRI) and computed tomography angiography (CTA) to estimate cerebrovascular pressure drops and resistance. However, this approach is limited by its reliance on ionizing radiation and time-consuming manual work. This study investigates alternative models for estimating relative pressure drops across different cerebral arteries using only 4D Flow MRI data. By evaluating the performance of the reduced Bernoulli (RB) model, the Stokes Estimator (STE), and the virtual work-energy relative pressure (vWERP) method, the aim is to assess whether these techniques can serve as viable alternatives to CFD-based simulations. The results show that both STE and vWERP offer promising pressure estimates in the larger cerebral arteries, while the RB model underestimates the pressure drops but remains the simplest to implement. However, none of the models could consistently capture the expected pressure drops in the smaller anterior cerebral artery, indicating a need for further investigation. Differences in model accuracy were likely influenced by mask segmentation variability and artery-specific flow characteristics. The findings suggest that pressure estimation using only 4D Flow MRI is promising and could be applicable to cerebrovascular assessment in clinical research. Additionally, broader studies are required to identify robust and standardized strategies for geometry segmentation that can improve precision and reproducibility across different vessel sizes and flow conditions.

Rithu Krishna Kumar - Detection of Link Failures in High-Density Mesh Networks with Machine Learning - Hitachi Energy

As a computer network grows in size and the number of connections increases, the likelihood of link failures increases, potentially leading to communication inefficiencies and data loss. This thesis aims to investigate whether machine learning can be used to detect link failures in wired mesh computer networks, based on packet latencies and order of arrival as input. It also explores whether the number and placement of routers can help machine learning models detect link failures more effectively. A simulation environment was built to simulate the traffic data, and Random Forest and XGBoost were the machine learning algorithms implemented. The results indicated that not all link failures could be detected and that the same pattern of detectable and undetectable links persisted as the network grew. Router placement significantly influenced detection: placing two routers diagonally doubled the number of detectable links compared to using only one, while adding more than two routers resulted in only marginal improvement. Both machine learning models detected the majority of link failures successfully, but consistently struggled with a set of undetectable links. These were sometimes misclassified as no link failures or confused with other undetectable link failures. This suggests that input features such as packet latency and arrival order, derived from shortest path routing, do not exhibit sufficient variation to enable the detection of all link failures.

Alfred Edenhagen - Algorithm Investigations for 5G Positioning in Indoor Scenarios - Ericsson

The demand for high-precision positioning has been growing in the last few years, due to the increased demand for location-based services. While the Global Navigation Satellite System (GNSS) can offer precision on the scale of meters when estimating positions outdoors, modern industries require a high precision in indoor deployments. Cellular communiation is considered to be a viable candidate for positioning, via the 5G network. In this thesis, we have investigated different methods for how the current accuracy of 5G positioning can be improved. We have analyzed measurement data from a setup using one of Ericssons indoor cellular systemsto track a moving User Equipment (UE). We have investigated if the rise time and the power ratio of the measured signals can be used to determine weights in the Weighted Least Squares (WLS) method, which is used in the estimation of the UE position. It has also been investigated how the positioning error decreases as the number of measurements increase, as well as two special cases of the deployment of the Radio Dots. Additionally, the area in where the UE is expected to be, was limited in order to limit he effects of Geometric Dilution of Precision (GDOP). The results show that neither the rise time nor the power ratio, had any positive effect upon the performance. The performance accuracy has been shown to increase as the number of measurements increases.

Theo Hägglund - Investigating the influence of key components on the performance of AEM water electrolysers: Development of a standardized testing protocol - Fysikinstutionen

The demand for green and sustainable energy sources is steadily increasing in rate with the increasing energy usage. A promising solution for this problem is hydrogen production by water electrolysis. Anion exchange membrane water electrolysis (AEMWE) is a form of zero gap electrolyzer that uses low cost and earth abundant materials to produce hydrogen. The thesis investigates the development of a standardized testing protocol for an AEMWE cell. This incudes investigating the operating conditions as well as the influence of key components affecting the cell performance. This protocol will be used in studying materials used as electrocatalysts as well as evaluating the performance of the membrane electrode assembly (MEA). By testing the cell under varying operating conditions such as temperature, flow rate, pressure and KOH concentration, optimal operating parameters were found. In addition to this, protocol steps were incorporated to improve the activity of the MEA, such as conditioning and stressing. The resulting operating conditions landed in a flow rate of 1 ml/min, a temperature between 60°C to 80°C, a bolt torque of 3.5 Nm and an electrolyte concentration of 1M KOH. It is used together with a testing protocol that includes membrane conditioning and stressing, impedance and stability tests along with polarization curves in order to evaluate the performance of the cell.

Victor Lindholm - Evaluating LINe in Semantic Segmentation OOD Detection and Improving Performance with Pairwise Activation Statistics - Datainstutitionen

In artificial intelligence, Out-of-distribution (OOD) detection refers to the task of identifying objects that are not part of the training data. Failure to detect such anomalies remain a major safety concern for vision systems that operate in open environments, such as autonomous cars. While many OOD detectors have been designed for whole-image classification, far fewer look at the pixel level that semantic-segmentation networks require. This thesis studies whether a recent and promising LINe detector—originally proposed for image classification—can be transferred to segmentation and how its performance can be improved.

An encoder-decoder convolutional neural network, DeepLab v3+-R101, pre-trained on Cityscapes, is used as the segmentation model. The LINe method is adapted by creating class-specific activation and weight masks for every pixel, followed by log-sum-exp scoring. I then introduce a complementary pairwise co-activation score that flags pixels whose channel pairs are rarely active together in in-distribution data. Temperature scaling and light Gaussian blurring are applied as post-processing steps.

Experiments on the Fishyscapes Lost and Found dataset show that the plain LINe method transfer already outperforms standard post-hoc baseline methods (soft-max confidence, max-logit, Euclidean feature distance) althougth not as much as in the original classification implementation. On its own LINe got a false positive rate(FPR95) of 18.7 at 95\%true positive rate and a 24.6 \% average precision(AP) on the Lost and Found validation set.

When introducing the co-activation score and extra fine-tuning, performance increases significantly to FPR95 of 5.7\%, and AP to 53.6\%, showing the great potential of this new method.

Arvid Sehlberg - Energy Storage Solution for Variable Frequency Driven Shipboard Cranes - MacGregor Sweden AB

Efficient handling of goods is important to the smooth functioning of global supply chains. Shipboard cranes play an important role in loading and unloading cargo, and their reliability and ability to operate make them crucial tools. Electric shipboard cranes generate a lot of electricity when lowering the load. Today, most of this energy is wasted. In this thesis, a short-term energy storage solution for variable frequency driven shipboard cranes is investigated, with the goal of improving energy efficiency and reducing the environmental impact of shipping operations.
An accurate simulation of the power consumption of a crane was generated and two peak shaving algorithms were tested on the simulated data. Both algorithms accomplish the goal of removing negative power peaks. One algorithm used a forecast of future power demand, and it worked better than the other algorithm if the forecast was good, but it was poor at reducing the negative power peaks if the forecast was less accurate. A simulation of three cranes that share the same generator was also performed. With three cranes the capacity of each crane's energy storage could be reduced because they could balance the energy use from each other.
An investigation of whether lithium batteries or supercapacitors are the best type of storage for this application was also conducted. While supercapacitors are costlier per kilowatt-hour, their performance under frequent load cycles and superior safety make them the better choice.
Though further work is needed to go from simulation to real-world deployment, this thesis lays a solid foundation for developing an energy storage system for variable frequency driven shipboard cranes.