

Introduction to AI, ML and Research 2024, **Inaugural Cohort Proceedings**





Proceedings of the BeyondAI: Introduction to AI and Research Programme

Authors:

Shaana Amarawickrama, Karuna Prakash, Osewuike Igue, Gulalayi Khankhel, Louis Vidal, Jack Woodhouse, Aditi Singh, Yonatan Yifat, Mena Zaied, Alazar Adane, Palak Kundu, Sumayah Adegbite, Adityan Srinivasan, Sadhana T. Kalidoss, Nayra Saadawy, Priyam Raul, Siddhart K. Gopal, Harini Muthurengan, Emeka Odiwe, Loulia Jaafari, Warenya K. Arachchillage, Ozodbek Eshmurodov, Arav Ajaykumar, Nafiul Haque, Shaurya Karmakar

Editors:

Filip Bár, Keisha Kwok, María Delgado Álvarez

December 2024

Preface

Welcome to the proceedings of the inaugural 2024 co-plored Machine Learning from both research and engi-Education from October 7th to December 6th, 2024.

This year, the programme brought together thirty exceptional young researchers, aged 14–22, from across the globe. Each participant was selected through a highly competitive process, standing out amongst 412 applicants from over 45 countries. To secure a place, candidates first their application. Only those who demonstrated strong independent learning skills and mastery of the foundational topics progressed to the *Course Stage*, where they further honed their knowledge before embarking on their research projects under the guidance of an academic mentor during the Research Stage.

The Preparation Stage, spanning six weeks, formed the cornerstone of the selection process. Depending on their prior educational background, applicants engaged in self-directed study to build or reinforce their understanding of Linear Algebra, Calculus, Multivariable Calculus, Python programming, and LATEX typesetting, while also exploring fundamental AI concepts through curated resources. Only those who demonstrated sufficient proficiency in the basics across all required disciplines earned admission into the programme.

During the four-week *Course Stage*, participants ex-

hort of the BeyondAI: Introduction to AI and Research neering perspectives. The research-focused component (BeyondAI) programme, organised by ThinkingBeyond emphasised conceptual understanding and mathematical modeling, introducing core ideas in a way accessible to newcomers. Meanwhile, the engineering component focused on the numerical implementation of these models, bridging the gap between theory and practical application.

Beyond building subject-specific knowledge and skills, had to complete the rigorous *Preparation Stage* as part of the programme introduced participants to the world of academic research. A series of workshops equipped them with essential research skills, including scheduling and time management, effective learning strategies, personal knowledge management, project management, scientific writing, and research presentation. Participants worked individually and collaboratively, undertaking tasks such as writing mock research papers, designing research posters, and coding their own multilayer perceptrons (MLPs) from scratch.

> Those who successfully completed the Course Stage advanced to the five-week-long *Research Stage*. Working in teams of two, participants delved into Machine Learning projects under the mentorship of experienced academics. Alongside their research, they attended expertled talks on advanced topics and real-world applications, further broadening their understanding of the field.

Machine Learning and Artificial Intelligence are

form industries and redefine problem-solving across domains. BeyondAI serves as a vital platform for nurturing young talent, providing them with a strong foundation in AI, exposure to academic research, and essential skills for intellectual growth. Our mission is to empower participants by fostering their expertise and curiosity, equipping them with the knowledge and experience necessary to excel as future subject experts, researchers and innovators.

We extend our deepest gratitude to the ThinkingBeyond team and the many volunteers whose dedication made this programme possible. Special thanks go to our academic mentors: Dr. Devendra Singh Dhami, Dr. Helena Bahrami, Dr. Filip Bár, Dr. Matej Cief, Adeveni Damilare Adeoye, MSc, Emilie Gregoire, MSc, Barbora Barancikova, MSc, and Matthew Pugh, BE. Their guidance and support were invaluable in ensuring the success of our participants. We also extend our appreciation to the participants themselves, whose dedication and hard work form the foundation of these proceedings.

The research projects presented in this volume reflect a diverse range of topics and applications, from

rapidly evolving disciplines with the potential to trans- fundamental studies on Deep Learning optimisers and the phenomenon of Double Descent to explorations of advanced architectures such as Geometric Clifford Algebra Networks. These projects represent a significant leap beyond the Course Stage, posing considerable challenges—particularly for participants who were new to Machine Learning at the start of the programme.

> We commend all participants for their perseverance and achievements in overcoming these challenges. As you explore these proceedings, we invite you to appreciate the depth of their intellectual efforts and contributions. We are proud to showcase their work and hope it serves as inspiration for aspiring researchers in STEM.¹ We are confident that these young minds will continue making meaningful contributions to the field and look forward to witnessing their future successes.

Chairs of BeyondAI

Dr. Filip Bár and Keisha Kwok

December 2024

¹The code for each project can be found on https://github.com/ThinkingBeyond/BeyondAI-2024.

Contents

S. Amarawickrama, K. Prakash, An Elementary Proof of the Universal Approximation Theorem for	
Multilayer Perceptrons	1
O. Igue, G. Khankhel, Weight Initialization for MLPs	2
L. Vidal, J. Woodhouse, Comparing Transformers to LSTMs with Attention	3
A. Singh, Y. Yifat, A Comparative Analysis of Optimizers for Classification with CNN	4
P. Kundu, S. Adegbite, Non-Linear Classifiers	5
M. Zaied, A. Adane, Cooperative Control of Multi-Agent Systems: An Optimal and Robust Perspective	6
A. Srinivasan, S. T. Kalidoss, Graph Neural Networks for Anomaly Detection in Dynamic Graphs.	7
N. Saadawy, P. Raul, Linear Search vs Grover's Algorithm	8
S. K. Gopal, H. Muthurengan, MNIST Digit Classification: Comparing Classical and Quantum Ap-	
proaches to Hyperparameter Tuning	9
E. Odiwe, Kolmogorov Arnold Networks vs Multi-Layer Perceptrons	10
L. Jaafari, W. Kasthuri Arachchillage, Geometric Clifford Algebra Networks: Bridging Geometry and	
AI	11
O. Eshmurodov, A. Ajaykumar, Overfitting vs Double Descent	12
N. Haque, Early Detection of Diabetic Retinopathy Using Machine Learning	13
S. Karmakar, Double Descent vs Overfitting in Deep Learning	14



An Elementary Proof of the Universal Approximation Theorem for Multilayer Perceptrons

Mentors: Mr. Matthew Pugh, Dr. Filip Bar

Shaana and Karuna





4**0**1



0-

0

1. ABSTRACT

-0

Students: Osewuike Igue, Gulalaiy Khankhel

This project investigates how weight initialization techniques affect the training dynamics and performance of MLPs, focusing on convergence, aradient stability, and accuracy. We also explore their interaction with activation functions, optimizers, and model depth to determine optimal strategies.

2. METHODOLOGY

We ran 3 main experiments to analyze the performance of Weight Initialization Methods including Zero Initialization, Random (Gaussian) Initialization, Xavier Initialization, and He Initialization ([2], [3]).

The effect of different activations (ReLU, Tanh, Sigmoid) on the performance of different weight	The effect of different optimizers (SGD, Adam) on the performance of different weight initialization prothede	The effect of MLP depth on weight initialization performance in a shallow and deep
initialization methods.	initialization methods.	architecture.

The MNIST Handwritten Digit Dataset was normalized and split into training and testing sets. Models were trained with a learning rate of 0.01, batch size of 150, and Cross Entropy Loss. For Activation and Optimizer Experiments, we used a 2-hidden-layer MLP, and for the MLP Depth Experiment, we compared 1 vs 6 Hidden layers. The Main Performance metrics included accuracy and convergence speed, with results visualized through graphs for test/training loss over epochs, test accuracy over epochs, weight distribution, and activation distribution.





Optimizers Experiment



Mentor: Barbora Barancikova

4. CONCLUSION

0-

0-

Check out

our code!

SEC.

Our research shows that Xavier and He Initialization perform well across various setups: Xavier suits tanh and sigmoid, while He is ideal for ReLU. Random Initialization works for shallow models but risks **aradient** issues in deeper ones, while Zero Initialization fails entirely as it prevents learning. He initialization excels by maintaining gradient flow and avoiding inactive neurons hence ensuring faster convergence with ReLU. Xavier balances gradient propagation in deeper networks leading to better stability and test accuracy. In conclusion, Xavier and He are the best choices, Random is viable for simple models, and **Zero** should be avoided ([1], [2], [3]).

5. FUTURE RESEARCH

Future research on this topic could explore additional weight initialization techniques, such as LeCun or Orthogonal. Moreover, our findings could be extended to more **complex** tasks and diverse datasets to test their **generalisability** and provide more insight into how different initializations interact with various architectures. This could help refine best practices across a wider range of applications ([1], [3]).

References:

• [1] Haykin, S. S. (2009). Neural networks and learning machines. Pearson Education.

- [2] Glorot, X., & Bengio, Y. (2010). Understanding the difficulty of training deep feedforward neural networks
 - [3] He, K., Zhang, X., Ren, S., & Sun, J. (2015). Delving deep into rectifiers: Surpassing human-level performance on imagenet classification.
 - Fig. 1 from https://www.javatpoint.com/multilayer-perceptron-in-tensorflow.

How do Transformers compare to LSTMs with attention in performance and efficiency for text-based sentiment analysis?

Authors: Louis Vidal, Jack Woodhouse Mentor: Matei Cief



Bahdanau, D., et al. (2015). Neural Machine Translation by Jointly (Reference: https://arxiv.org/abs/1409.0473)



A Comparative Analysis of Optimizers for Classification with CNN

BeyondAI 2024: Project 9B

A. Singh, Y. Yifat; Mentor- Mr. Adeyemi D. Adeoye

I. INTRODUCTION

Optimization is a crucial aspect of Deep Neural Networks (DNNs). During training, an optimizer adjusts the weights and biases of the DNN to minimize a loss function.

minimize
$$f(\theta) = \sum_{i} \ell(\hat{y}(\theta; x^i), y^i)$$

where $\boldsymbol{\theta}$ is the vectorized model parameters.

An optimization algorithm may perform better under one setting than another. Therefore, a comparative study of different algorithms is important for identifying an optimal setup for different algorithms for a particular task and model architecture. In this study, we aim to evaluate the performance of the six popular optimizers (SGD, Adam, AdaGrad, AdaDelta, RMSprop and Nadam) on a classification task using a Convolutional Neural Network (CNN) architecture. These six optimizers have been chosen because they are widely used in Deep Neural Networks. Each optimizer has its characteristic strengths and limitations, and oftentimes, one addresses the limitations of another. The choice of an optimizer directly influences the generalization properties and robustness of the final trained model, and hence it is important to choose an optimizer which is not only computationally efficient but also enhances the model performance.

II. PERFORMANCE COMPARISON & METRICS

The performance was assessed using four key evaluation metrics: accuracy, precision, recall and F1-score. Accuracy measured the overall correctness of predictions. Precision evaluated the proportion of true positives among predicted positives. Recall determined the proportion of true positives correctly identified. The F1-score balanced precision and recall. These metrics collectively provided a comprehensive evaluation of each optimizer's effectiveness in classification tasks. Experiments were conducted to analyze the effect of learning rate (for SGD) and batch size on model performance.



IV. MODEL ARCHITECTURE



V. RESULTS

As shown in figure 4.2, AdaGrad achieved the highest accuracy of 79.09% and F1-score of 79.33%. The highest precision was achieved by Adam which was 84.23%. Nadam excelled in recall with 81.62%. It was found that SGD achieved highest accuracy with 0.01 learning rate. In figure 4.3, it was observed that AdaGrad demonstrated high accuracy across various batch sizes, with the best performance (78.15%) at the batch size of 32. RMSprop and Nadam performed better at higher batch sizes (512), with accuracies of 79.58% and 78.6%, respectively. Adam peaked with a batch size of 128, achieving an accuracy of 78.0% which showed its adaptability to different batch sizes. AdaDelta maintained stable but relatively lower performance across all batch sizes.



Optimizers

Figure 4.2: Comparison of Optimizers based on the performance metrics



Figure 4.3: Comparison of Optimizers based on the batch sizes







Optimizer	Description	Update Formula				
Adam	Combines momentum and adaptive learning rates for fast and robust convergence which is ideal for NLP, computer vision and large datasets. However, it tends to overfit smaller datasets and has higher computational costs.	$\begin{split} \theta_t &= \theta_{t-1} - \frac{\eta}{\sqrt{\hat{v}_t} + \epsilon} \hat{m}_t, \hat{m}_t = \frac{m_t}{1 - \beta_1^1}, \hat{v}_t = \frac{v_t}{1 - \beta_2^1} \\ m_t &= \beta_1 m_{t-1} + (1 - \beta_1) g_t, v_t = \beta_2 v_{t-1} + (1 - \beta_2) g_t^2 \end{split}$				
Nadam	Adapts learning rates based on recent gradient magnitudes that makes it effective for non-stationary data, especially in RNNs and reinforcement learning. It is sensitive to hyperparameters.	$ \begin{aligned} \theta_t &= \theta_{t-1} - \frac{\eta}{\sqrt{\hat{v}_t} + \epsilon} \left(\beta_1 \hat{m}_t + \frac{(1 - \beta_1)g_t}{1 - \beta_1^t} \right) \\ \hat{m}_t, \hat{v}_t, m_t, v_t \text{ as in Adam} \end{aligned} $				
RMSprop	Introduces Nesterov momentum for faster convergence, commonly applied in computer vision and speech recognition. Despite its speed, it requires careful hyperparameter tuning and is computationally intensive.	$\theta_t = \theta_{t-1} - \frac{\eta}{\sqrt{v_t} + \epsilon} g_t, v_t = \beta v_{t-1} + (1 - \beta)$				
AdaDelta	Addresses the decay issue in AdaGrad by normalizing updates with a moving average of gradients, making it suitable for RNNs and LSTMs, albeit with complexity and instability for non-convex data.	$\begin{split} \Delta \theta_t &= -\frac{\sqrt{s_{t-1}+\epsilon}}{\sqrt{v_t+\epsilon}}g_t, v_t = \gamma v_{t-1} + (1-\gamma)g_t^2, \\ &s_t = \gamma s_{t-1} + (1-\gamma)(\Delta \theta_t)^2 \end{split}$				
SGD	Uses mini-batches for gradient descent, offering simplicity and efficiency. It converges to the global minimum for convex functions but may oscillate on non-convex ones. Commonly used in linear regression and deep neural networks.	$ heta_t = heta_{t-1} - \eta g_t$				
AdaGrad	Adjusts the learning rate based on the historical sum of squared gradients, making it effective for sparse data like NLP but prone to slowing down over time.	$\theta_t = \theta_{t-1} - \frac{\eta}{\sqrt{G_t} + \epsilon} g_t, G_t = G_{t-1} + g$				



Authors: Palak Kundu, Sumayah Adegbite Non-Linear Classifiers

Mentor: Miss Barbora Barancikova

INTRODUCTION

Linear Classifiers \rightarrow Struggle with Non-linear Data

Non-linear Classifiers \rightarrow Evaluate Accuracy \rightarrow Evaluate Complexity \rightarrow Evaluate Interpretability

Compare on Different Datasets → Guide Selection of Best Method for Complex Data

METHODOLOGY

We evaluated classifier performance on three diverse datasets, each with varying complexity and feature types. The datasets were split into 70% training and 30% testing sets, ensuring that training data was used for model training, while testing data evaluated the model's generalization.

CLASSIFIERS EVALUATED



F1-Score | Training Time





The bar chart and line graph represent the accuracy scores and training time of the datasets shown in the image beside them, respectively. Each of the bars and points in order represents (LR)Logistic Regression (Linear), SVM (RBF kernel), (DT)Decision Tree, (RF) Random Forest, (GB) Gradient Boosting, and (NB) Naive Bayes.

accuracy and robustness while reducing overfitting, though they are more computationally expensive than simpler classifiers





0.92 0.98 0.87 0.96 0.96

- Cortes, C., & Vapnik, V. (1995). Support-vector networks. Machine learning, 20(3), 273-297
- Breiman, L. (2001). Random forests. Machine learning, 45(1), 5-32.

REFERENCES

 Murphy, K. P. (2012) Machine learning a probabilistic perspective.

MIT pres



A programme led by **ThinkingBeyond**

CONCLUSION

• Logistic Regression and Naive

Random Forest and Gradient

Boosting are effective for noisy or

imbalanced datasets, balancing

FUTURE WORK

• Trying Out Other Types

of Classifiers such as

• Evaluation on More

Diverse Datasets

neural networks

Bayes are faster but less

accurate.

Cooperative Control of Multi-Agent Systems: An Optimal and Robust Perspective

BeyondAI, Thinking Beyond 8B

Introdution

The coordination of multiple autonomous agents, such as unmanned aerial vehicles (UAVs), in dynamic and complex environments, has become a critical challenge in modern robotics. UAV swarms are increasingly being utilized for a variety of applications, including search and rescue operations, environmental monitoring, and infrastructure inspection. However, achieving effective collaboration among multiple UAVs while maintaining efficiency, avoiding collisions, and adapting to unpredictable environmental conditions remains a complex task

Inspired by coordinated behaviors observed in nature, such as flocking in birds and schooling in fish, the study of multi-agent systems (MAS) has grown in prominence. These natural systems exhibit remarkable levels of cooperation, where individual agents follow simple rules yet achieve globally optimal behaviors. The study of multi-agent coordination, particularly in mobile autonomous systems, has emerged as a key area of research in fields such as control theory, robotics, and optimization.

The problem statemant

ņ

In Multi-agent reinforcement learning, agents exist in different scenarios such as working collaboratively to tackle a given problem, competing with each other to get an optimal reward, or doing both works together or competing for each other like in football. In these processes of learning we have various problems in order to train how they work effectively and efficiently

When we work in real-time continuous environments such as self-driving car require immediate responses to change of environment. This require a huge amount of sensor data analysis to execute the best action in a given state

As the number of agents in a system increases, managing their inte grows exponentially. This results in communication network overload and loss of performance leads to suboptimal outcomes.





The coordination of drones in dynamic and cluttered environments presents several challenges, especially when tasks require collaboration among multiple agents. Traditional single-agent systems fail to meet the complex demands of real-time navigation, obstacle avoidance, and collaborative task execution Communication



TD Loss= $\mathcal{L}_{in}(\theta)$

0 (6 a. 0)

Methodology

Cluster Space Control Framework

The method relies on the definition of a vector of CS variables that describes the position, orientation, and shape of a virtual kinematic mechanism representing the formation, allowing for full specification and control of the system. For a group of N robots with m degrees of freedom (DOF) each, the framework defines a set of kinematic transforms c = KIN(r) -where r is the vector of (N × m) robot space variables, (i.e., position and orientation of the robots), and c is a vector of (N × m) cluster space variables- as well as a set of inverse kinematic transforms r = INVKIN(c). Additionally, relations between velocities in both space.

Geometric Coordination & Optimal Control

Geometric Coordination

Geometric principles such as similarity transformations and manifold optimization provide a robust foundation for multi-agent coordination. These methods align for maturagent coordinators, mese metuosa and individual agents' trajectories and maintain formations through precise transformations in Euclidean and Riemannian spaces. For multi-agent systems (MAS), geometric coordination ensures that each agent adjust: its position in relation to others, maintaining formation and avoiding collisions during navigation. Key Formula: $T = \{R, t\}, R \in SO(n), t \in \mathbb{R}^{n}$

· T: Trans

Integration of RL Framework in UAV Navigation

This diagram illustrates the overall framework for integrating Reinforcement Learning (RL) into UAV navigation, highlighting the interaction between the environment, the agent, and the reward function. It provides a structured approach to ensure that the UAV effectively learns and performs tasks such as collision avoidance, localization, and waypoint navigation.



Optimal Control

Optimal control minimizes a defined cost function ensuring efficient resource use and high-performance task execution. In MAS, control laws are derived to balance operary consumption are derived to balance energy consumption, formation maintenance, and collision avoidance The theory ensures that each drone's trajectory i optimal with respect to the task requirer

Simulation & results

This study explores the integration of reinforcement

 $L^{PG}(\theta) = E_t [\log \pi_\theta(a_t | s_t) * A_t]$

Mr Matej Cief

Menna Zaied and Alazar Adane

PPO is related to policy optimization where strategy determines the agent's actions in a given state through training, the policy update gradually and this ensures stability in the policy function. One of the main advantages of using PPO is its effectiveness in continuous action space such as autonom

vehicles. PPO uses a stochastic policy, which supports exploration, while the ensures the exploitation of

PPO Algorithim

learned knowledge which doesn't lead to large amount of change

ReLu Rel

learning (RL) and geometric methods to address these challenges, focusing on enabling real-time obstacle avoidance and efficient task scheduling in UAV swarms By leveraging the strengths of MAS and advanced mathematical techniques, we propose a framework that optimizes resource utilization and ensures less coordination in dynamic enviro



Our results

Drone Navigation Results Using PPOT n (PPO). Metrics such as FPS, tir ments in navigation efficiency and stability in a dynamic 3D



Reference

2hang, W., & Miele, A. (2020). Geometric Methods for Multi-Agent : Springer. Shishika, R., & Chei, H. (2017). On the Optimal Centrol of Multi-Age Communication. arXiv. • Chen, X., & Zhang, Y. (2023). An Overview of Coordination Techniques for Multi-Agent 5 Satton, R. S., & Barto, A. G. (2018). Reinforcement learning: An introduction. NIT Pro-Long, P., Fan, T., Lian, X., Liu, W., Zhang, H., & Pan, J. (2028). Towards optimally decord



Abstract

We implement a Spatio-Temporal Graph Neural Network (SpatioTemporalGNN) to understand spatio-temporal relationships between nodes in a dynamic graph. To realise DGNNs, we utilise satellite trajectories and the prediction of potential collisions as edge features. By utilising Two-Line Element (TLE) data, we model satellites as nodes in a dynamic graph and encode their relative spatial interactions as edge features. Our approach demonstrates the capability of graphbased learning to capture spatiotemporal dependencies, enabling the prediction of anomalies or in our application, impending satellite collisions.

Dynamic Graph Neural Networks

Graph Neural Networks (GNNs) are a powerful tool for learning from graph-structured data, where nodes represent entities and edges capture relationships between them. They are widely used across various domains, such as social networks, molecular biology, and recommendation systems, due to their ability to effectively model complex relationships and interactions. This is particularly important in real-world scenarios where graphs are not static but change as interactions and relationships evolve.

DGNNs represent entities as nodes with feature embeddings and encode evolving relationships through dynamically updated edges. These dynamic edges can reflect changes in connections, interactions, or weights over time, allowing the model to adapt to varying graph structures. We use GCNConv layers, which aggregate information from neighboring nodes and edges, to effectively learn and propagate both spatial and temporal patterns across the graph. By leveraging this approach, DGNNs can capture the interplay between the structural and temporal aspects of the data, making them capable of extracting meaningful insights from time-evolving datasets and adapting to the dynamics of complex systems.

GNNs for Anomaly

Detection in Satellite

Future Work *L*

• Extend the model to integrate external factors such as space weather and satellite

• Improve edge representations using

• Apply the framework to broader datasets,

more rigorous training

including simulated collision events, for

attention mechanisms for better and

 \bullet

Trajectories

_> Methodology

We approach this problem by working with a spatio-temporal graph neural network model that utilizes Graph Convolution via the GCNConv function in the PyTorch-Geometric library. This model is designed to capture both spatial and temporal dependencies in dynamic graph data. The first GCN layer transforms the input node features into an intermediate representation by aggregating information from neighboring nodes, effectively encoding the local graph structure.

The second GCN layer further refines this representation, capturing deeper spatial dependencies and enhancing the model's ability to discern intricate patterns. By normalizing the feature distribution, we maintain training stability, improve convergence, and ensure consistent performance across varying graph structures. This approach allows us to effectively handle the complexities of evolving graph data and extract meaningful patterns.



Model Architecture [4]

We calculate our loss using Mean Squared Error (MSE), a widely used metric for regression tasks that measures the average squared difference between predicted and actual values.

Our testing metrics include Mean Average Error (MAE) for evaluating prediction accuracy, along with F1-score, Recall, and Precision, which provide a comprehensive assessment of model performnce.

Created By : Adityan Srinivasan & Sadhana Thirumangai

Group 7B : Graph Neural Networks for Anomaly Detection in Dynamic Graphs M

-1

References <

Representation Learning on Large Graphs.

Tracking

1.Kipf, T. N., & Welling, M. (2017). Semi-Supervised

2. Spacetrack.org, Two-Line Element (TLE) Data for Satellite

3. Hamilton, W. L., Ying, Z., & Leskovec, J. (2017). Inductive

4.Kim, Minseong & Lee, Jaeseung & Kim, Jibum. (2023).

GMR-Net: GCN-based mesh refinement framework for elliptic PDE problems. Engineering with Computers.

Classification with Graph Convolutional Networks.

Mentor : Dr Helena Bahrami



We utilized the NORAD dataset's TLE data of active

satellites, which provides detailed two-line element (TLE)

information for orbital modeling. Working with the SGP4

propagation algorithm, we converted TLE data into

Cartesian coordinates to represent satellite positions in

three-dimensional space. Using this ground truth, we

generated a temporal sequence of graphs, where edges

represent relative distances between satellites, capturing

Each snapshot of the dynamic graph was taken daily,

enabling the model to process evolving spatial and

temporal patterns effectively. We normalized features to

ensure stable training and split our data into testing and

their spatial relationships over time.

training sets for robust evaluation.

Data

We observe that our model is able to accurately capture the spatiotemporal dependencies between satellites overhead and effectively predict collisions between satellites, which we have defined as anomalies. Impressively, our model achieves a 99% accuracy in identifying these anomalies, showcasing its reliability and precision in high-stakes scenarios.

This novel application of DGNNs provides an opportunity to further develop by incorporating additional features, potentially improving predictive performance and expanding the model's applicability.

Code

A programme led by ThinkingBeyond

7



A programme led by **ThinkingBeyond**

10









Traditional models for robot dynamics are based on complicated kinematic equations, geometric Clifford algebra networks (GCANs) are an efficient alternative. Our research focuses on developing GCAN models for use in for predicting motion and forces in robots. Existing work on GCANs utilize symmetry group transformations through aeometric alaebras with the use of group action layers that



a terround

GCANS

BRIDGING

GEOMETRY

& AI

3 Obser

geometric algebras with the use of group action layers that combine multivector transformations with specified group actions (Ruhe, 2023). (Low, n.d)



Fig 2: Planes (Ruhe, 2023)

Initially, we aimed to incorporate GCANs in robotic arms. However, GCANs could be more useful in outdoor robotic systems. Because of time restrictions for the research stage, we focused on implementing a Geometric Clifford Algebra - Multilayer Perceptron (GCA-MLP) and testing it for various tasks like projections and object recognition. In our GCA-MLP, the multivector(M) and the standard deviation(std(M)) of its components are used to normalize the multivector across the batch dimension.

The activation function MSiLU along with the loss function Mean Squared Error(MSE) is used to predict trajectories. MSE helps find the error between the target position and predicted position for each sample in the batch.

For the robot dynamic application, we used cross entropy loss along with the NAO dataset on robotic arms. This was used to observe how accurate our MLP was at classifying objects based on their position or rotation in space



Even though GCANs are able to handle 3D problems, they're not great with 2D points. The loss values remained consistently high, this may be caused due to the limited dimensionalities of the 2D Euclidean space, and lack of geometric depth.



While our MLP does not work well in 2D spaces, combining non geometric features proved to improve the accuracy and reduce the loss significantly. Though this is interesting, because of the computational overhead, we decided to not carry on with this approach.

References: Löw, T. Calinon, S. (n.d.). Geometric algebra for optimal control in robotics using gafro./diap Research Institute, Martigny, Switzerland.

scan for our code!

Ruhe, D., Gupta, J. K., de Keninck, S., Welling, M., & Brandstetter, J. (2023). Geometric Clifford Algebra Networks. arXiv. https://arxiv.org/abs/2302. 06594

The GCA-MLP proved to be highly accurate with extracting, combining and classifying geometric features like rotation and position in space. We plan to develop this to be used in aerial robots for pose classification and anomaly detection in the future.



D

(Generated by Matplotlib)

Our MLP is moderately accurate at trajectory prediction, however, the predicted trajectory goes off track after a few time steps. We tried fine-tuning the hyperparameters and adding an attention layer to our code. Though we were able to improve the accuracy and minimize the loss, it wasn't highly accurate

otions



Mentors: Dr.Filip Bar, Mr. Matthew Pugh

Loulia J. & Warenya K.





In this study, we investigate the interplay between double descent and regularization across shallow neural networks, decision trees, and ResNet architectures, using both synthetic and real-world-inspired datasets. Our experiments span multiple training regimes, including epoch-wise and model complexity-driven setups, revealing architecture-specific manifestations of double descent. Notably, dropout proves effective in mitigating overfitting in high-complexity regions, while weight decay demonstrates consistent regularization benefits across simpler models. Additionally, unexpected anomalies, such as test loss outperforming train loss, were observed in specific configurations, shedding light on the nuanced dynamics of regularization.

Findings

Polynomial regression was used as it provides a well-

controlled and mathematically grounded setting for

observing double descent. The model complexity can be

easily modulated by adjusting the polynomial degree,

allowing us to capture the behaviour at and beyond the

interpolation threshold. Polynomial regression is a classic

example where the double descent phenomenon is

theoretically understood, providing a baseline for comparison

with more complex models. Its simplicity made it ideal for

verifying foundational behaviours of double descent before

extending the experiments to neural networks and decision

trees.

Overfitting vs Double Descent

Introduction

Machine learning has evolved, leading to the development of overparameterized models that fit training data exactly while achieving strong generalization performance. This paradoxical behaviour challenges traditional notions of the bias-variance tradeoff, which has guided our understanding of model complexity and generalization dynamics.

Double descent refers to a behaviour in model performance as complexity increases. Traditional machine learning theory suggests a U-shaped curve where test error decreases up to an optimal complexity level before rising due to overfitting. However with double descent, test error initially decreases, rises, and then decreases aaain.

This study addresses the following research question: How do regularization techniques (such as dropout and weight decay) influence the double descent phenomenon and	U = 6 foreiter 9 Modern Exercised 9 Modern E
mitigating overfitting across different deep learning model architectures?	32 1 10 20 20 Made Stor (90:09)

Methodology

Synthetic datasets were generated to simulate regression and binary classification tasks. The primary data generation process followed

$[y = \sin(2\pi x) + \epsilon, \quad \epsilon \sim \mathcal{N}(0, \sigma^2)]$

where ϵ represents Gaussian noise. The classification tasks made a threshold on the regression outputs to create binary labels. Shallow networks were chosen to analyse dropout effects because they are simpler compared to deeper architectures. The smaller number of layers reduces the variables, making it easier to understand the role of dropout. Shallow networks are computationally efficient to train, which allowed us to explore a wide range of dropout rates and configurations across different training regimes. This choice aligns with the study's aim to explore the interaction between double descent and regularization.

Conclusion

0.2), but higher rates introduced instability and underfitting.

Weight decay smoothed loss curves and improved

generalization, though excessive regularization suppressed

double descent, leading to underfitting.

Combined regularization strategies provided the most

consistent performance, balancing overfitting and

underfitting across diverse configurations.











References

[1] Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machinelearning practice and the classical bias-variance trade-off. Proceedings of the National Academy of Sciences, 116(32):15849-15854, 2019.

[2] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. Deep Learning. MIT Press, 2016. Website.

[3] Jared Wilber and Brent Werness. Double descent: A visual introduction 2021

Future Work

Expand experiments to larger and more diverse datasets to validate findings and uncover additional patterns in double descent behavior.

Investigate additional regularization techniques, such as batch normalization, early stopping, and data augmentation, to assess their interaction with double descent

Extend the study to modern architectures, such as transformers, to evaluate the broader applicability of findings.

Explore adversarial training and other robust optimization methods as potential mitigators of overfitting in double descent scenarios

Future Work Although ResNet experiments did not yield clear double descent trends future studies could explore this phenomenon with larger computational resources, alternative datasets, or simplified ResNet variants. Investigating how modern architectures like ResNet interact with regularization techniques remains an open auestion with significant practical implications

- Double descent was observed in both model complexity and epoch-wise training, with test loss showing characteristic peaks at the interpolation threshold.
- Dropout effectively reduced overfitting at moderate rates (0.1-





EARLY DETECTION OF DIABETIC **RETINOPATHY USING MACHINE LEARNING**

Diabetic retinopathy (DR) is a leading cause of blindness, accounting for over 1 million blindness cases and 3.28 million severe vision impairments globally. Early detection can significantly mitigate its severity and impact. This research conducts a comparative analysis of machine learning models, specifically Convolutional Neural Networks (CNNs) and Vision Transformers (ViTs), to evaluate their effectiveness in detecting DR and similar medical conditions in low-resource settings or clinical practice.

Stages of Diabetic Retinopathy



Fig 01: DR Stages

HOW DETECTION WORKS

Machine learning models like CNNs and ViTs analyze retinal images to detect key indicators of DR, such as microaneurysms, hemorrhages, and exudates: • Stage 1 (Mild): Characterized by 1 or 2

- microaneurvsms.
- Stage 2 (Moderate): More than 20 microaneurysms and occasional hemorrhages.
- Stage 3 (Severe): Numerous hemorrhages in all quadrants, along with venous beading.
- Stage 4 (Proliferative DR): Growth of new, fragile blood vessels (neovascularization) prone to leakage.

These features are analyzed to classify disease. severity, aiding early intervention and treatment.



95%) all across but validation accuracy drops significantly as the dataset size decreases. The ViT model, however, exhibits stable validation accuracy (~63-94%) even with smaller datasets.

MultiClass Comparision (DR Stages)

Fig 06: CNN vs ViT Comparison Both CNN and ViT models display similar trends for F1-score, Sens Specificity (in Fig 06). However, the CNN's Training Accuracy is

increasing while validation accuracy plateaus at 70-75%, while the ViT

model oscillates around 73% for both training and validation, indicating more consistent but limited performance.

Analysis

severity:

Overall, the ViT model shows more consistent performance with training and validation accuracies around 70-73%, indicating better generalization and less overfitting. However, the confusion matrix reveals misclassification between "Moderate" and "No DR" categories. In contrast, the CNN model has higher training accuracy (~90%) but determine U(20, USP) uncensition powerfittion stagnant validation accuracy (70-75%), suggesting ove



Conclusion

In the context of detecting medical conditions like diabetic retinopathy, Vision Transformers (ViTs) demonstrate better generalization, making them potentially more robust for diverse, unseen cases. However, this advantage heavily relies on pretraining with large, related datasets. Convolutional Neural Networks (CNNs), despite a tendency to overfit, excel at extracting local features and distinguishing specific stages. For rare or underrepresented cases, CNNs can outperform ViTs if augmented datasets and regularization techniques are carefully applied. Both models demonstrate complementary strengths, and their selection should depend on the dataset size, available resources, and application requirements.

References

- Vishal Awasthi a, Namita Awasthi b, Hemant Kumar c, Shubhendra Singh c, Prabal Pratap Singh d. (2021). Optimized vision transformer Gradiabetic retinopathy detection using Harris Hawk optimization.
 Gulshan, V., Peng, L., Coram, M., Stumpe, M. C.,
- Vur, D., Narayanaswamy, A., & Webster, D. R. (2016). Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. Jama, 316(22), 2402-2410.
- Ting, D. S. W., Cheung, C. Y. L., Lim, G., Tan, G. S.
 W., Quang, N. D., Gan, A., & Wong, T. Y. (2017).
 Development and validation of a deep learning system for diabetic retinopathy and related eye diseases using retinal images from multiethnic populations with diabetes. Jama, 318(22), 2211popula 2223.
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In Advances in neural information processing systems (pp. 1097-1105).

Author MD Nafiul Hague Mentor



DOUBLE DESCENT VS OVERFITTING IN DEEP LEARNING

School Sc

Scan here to view the code

I. Introduction

Overfitting, a classical issue in deep learning, occurs when models perform well on training data but poorly on unseen data. Double descent challenges this notion, revealing that model performance can improve beyond the interpolation threshold as capacity increases, offering new insights into the dynamics of overparameterized models.

II. Methodology

- Model Design: A flexible feedforward neural network (FlexibleNN) was developed to study the effects of model complexity on performance.
- Data Preparation: The MNIST dataset was used for training and evaluation, consisting of grayscale images of handwritten digits.
- Training and Evaluation Pipeline: The model was trained using cross-entropy loss and SGD optimizer, with training and testing across epochs to analyze performance trends like overfitting and double descent.
- Visualization: Training and test losses were plotted across epochs for each model configuration. This allowed the identification of overfitting regions and the double descent behavior as model complexity increased.

III. Findings

• Effect of Model Complexity:

- Increasing the number and size of hidden layers initially improved performance, reducing both training and test loss.
- Beyond a certain complexity, overfitting occurred, evidenced by a significant divergence between training and test loss.
- Double Descent Phenomenon: For highly complex models, test loss showed an initial increase (overfitting) followed by a decrease as model capacity grew further, demonstrating double descent behavior.

• Performance Trends:

- Simpler models struggled to capture patterns effectively, resulting in higher training and test loss.
 Intermediate-complexity models balanced generalization and overfitting, achieving optimal performance.
- Loss Visualization: Training and test loss curves provided clear visual evidence of overfitting and the transition to the double descent phase as model complexity increased.

IV. Conclusion

This study demonstrated the double descent phenomenon, where increasing model complexity initially caused overfitting but later improved generalization. These findings highlight the complex relationship between model capacity and performance.

V. Double descent graph



VI. References

1.Wilber J., & Werness B. (December 2021). Double Descent. MLU-Explain[Online]

2. Belkin, M., Hsu, D., Ma, S., & Mandal, S. (2019). Reconciling modern machine-learning practice and the classical bias-variance trade-off. Proceedings of the National Academy of Sciences, 116(32), 15849-15854.

3. Deep Double Descent: Where Bigger Models and More Data Hurt <u>Preetum Nakkiran, Gal Kaplun, Yamini Bansal, Tristan Yang, Boaz</u> <u>Barak, Ilya Sutskever</u>

Author: Shaurya Karmakar Mentor: Adeyemi D. Adeoye

