Vol. 3, No.12, 2025

airis4D

We Dream, Design, Develop and Deploy the Future

Artificial Intelligence Research and Intelligent Systems

Artificial Intelligence & Machine Learning

Astronomy and Astrophysics

Biosciences

Computer Programming

Fiction

An airis4D Publication

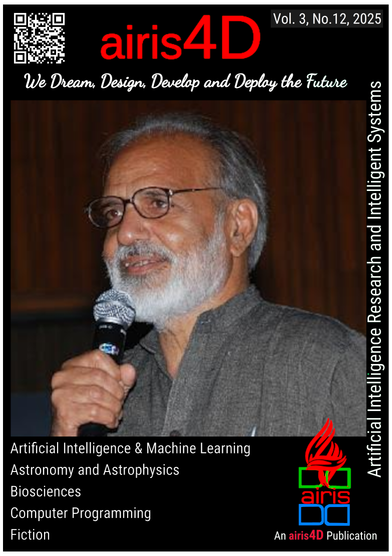

Cover page

Professor Naresh Dadhich was the second director of the Inter University Centre for Astronomy and Astrophysics (IUCAA, Pune). A renowned cosmologist who could see the geometric beauty of the General Theory of Relativity, in the way that very few others than Albert Einstein might have comprehended it. His demise on November 6th 2025, was a significant loss for humanity and the world of science and mathematics. We pay homage to his memories.

Journal Publisher Details

- **Publisher** : airis4D, Thelliyoor 689544, India
- **Website** : www.airis4d.com
- **Email** : nsp@airis4d.com
- **Phone** : +919497552476

Editorial

by Fr Dr Abraham Mulamoottil

AIRIS4D, VOL.3, NO.12, 2025

www.airis4d.com

This issue of AIRIS4D pays tribute to Professor Naresh Dadhich (1944–2025), a renowned Indian relativist and one of the founding pillars of the Inter-University Centre for Astronomy and Astrophysics (IUCAA), Pune. Professor Ajit Kembhavi remembers Professor Naresh Dadhich's significant contributions to gravity theory, black hole physics, Lovelock gravity, and brane cosmology. As a teacher and mentor, he inspired generations of students and collaborators, many of whom became leading scientists. He continued publishing until his final days. Dadhich played a central role in building IUCAA from its inception, promoting astronomy across India through extensive outreach and forming strong national and international collaborations—including with countries in Central Asia, Africa, Iran, and Turkey. His efforts helped integrate Indian institutions into major global projects such as SALT, the Thirty Meter Telescope, and the LIGO-India initiative. Known for his simplicity, rationalism, socialist ideals, and joyful personality, Dadhich maintained deep friendships across scientific and artistic communities. He is fondly remembered for his clarity in explaining complex concepts and his commitment to nurturing young talent. The issue also features a personal reflection by astronomer Joe Philip Ninan, who recounts childhood memories of Dadhich as an inspiring teacher who sparked his early interest in physics.

Arun Aniyan addresses "Few-Shot Learning (FSL)", one of the biggest limitations of traditional deep learning—its dependence on huge labelled datasets—by enabling models to learn new concepts from only a few examples. Standard deep learning struggles in real-world scenarios where data is scarce, costly, sensitive, or rare—such as medical imaging, scientific discovery, robotics, and personalised AI. FSL overcomes these challenges through meta-learning, metric-learning approaches like prototypical networks, adaptable methods like MAML, and data-generating techniques that expand limited datasets. By learning how to learn, FSL allows AI systems to generalise quickly from minimal data, making it essential for high-impact, data-constrained domains and a significant step toward more human-like, efficient, and generalizable artificial intelligence.

'The Time Arrow of Entropy" by Jinsu Ann Mathew explains how entropy—often described as disorder—gives direction to time by making many processes naturally irreversible. From broken glass to heat flow, systems tend to move from order to disorder, creating the "arrow of time." It shows how memory in natural and artificial systems preserves traces of the past, slowing change and shaping future behaviour. Language and society evolve in a similar forward-moving way, carrying history that prevents them from returning to earlier states. By observing how entropy fluctuates, trends, or suddenly jumps, we can understand the rhythms and transitions in physical, biological, and social systems. Ultimately, entropy becomes not just a measure of uncertainty but a lens for understanding how systems evolve.

The article "Unboxing a Transformer using Python - Part I" by Linn Abraham introduces how Vision Transformers (ViTs) adapt the transformer architecture—originally built for language translation—to image recognition tasks.

Unlike CNNs, ViTs excel at capturing long-range dependencies in large datasets but must relearn local pixel relationships due to fewer built-in spatial priors. Using a PyTorch implementation, the author walks through the structure of a ViT model for classifying solar active regions, highlighting key components such as patch embedding, positional embeddings, class tokens, transformer encoder blocks, and the MLP head. A detailed look inside the patch embedding layer shows how images are divided into 16×16 patches, flattened, normalized, and projected into a lower-dimensional embedding space to prepare them for transformer processing. This first part focuses on understanding the model architecture, setting the stage for deeper exploration of positional embeddings and transformer layers in future instalments.

Dusty Plasma by Abishek explains how dusty plasmas—ionized gases containing charged dust grains—form a uniquely complex, multi-scale system where grain charging, electromagnetic forces, and Debye screening shape particle behaviour. Dust grains acquire charge through electron/ion collection and photoemission, leading to forces such as electrostatic lift, ion drag, neutral drag, thermophoresis, and gravity that determine whether grains levitate, drift, or form ordered structures like plasma crystals. Strong coupling between grains produces rich collective behaviour, including phase transitions, dust acoustic waves, instabilities, and self-organized patterns. Because grain charge fluctuates with plasma conditions, microscopic charging dynamics strongly influence macroscopic transport and wave phenomena. Dusty plasmas play major roles in astrophysical environments—such as protoplanetary disks, cometary comae, and planetary rings—as well as in technology, where dust can both hinder semiconductor and fusion operations and serve as a valuable tool for diagnostics, nanoparticle synthesis, and the study of many-body physics.

Geetha Paul provides a comprehensive overview of diabetic retinopathy (DR), a major diabetes-related eye disease that damages retinal blood vessels and can lead to blindness. It explains how chronic high blood sugar causes microvascular injury, resulting in microaneurysms, haemorrhages, exudates, macular oedema, and in advanced stages, the growth of fragile new blood vessels (proliferative DR). Key associated conditions such as drusen, choroidal neovascularisation (CNV), and cystoid macular edema (CME) are also described. The article highlights symptoms, risk factors, diagnostic methods, and treatments, including anti-VEGF injections, laser therapy, steroids, and vitrectomy, emphasising that early detection and strict diabetic control can prevent the most severe vision loss. It also outlines how retinal image processing, machine learning, and deep learning—using steps like preprocessing, segmentation, feature extraction, and CNN-based automated grading—play a crucial role in detecting, classifying, and monitoring DR with high accuracy.

News Desk - Meomories Never Fade

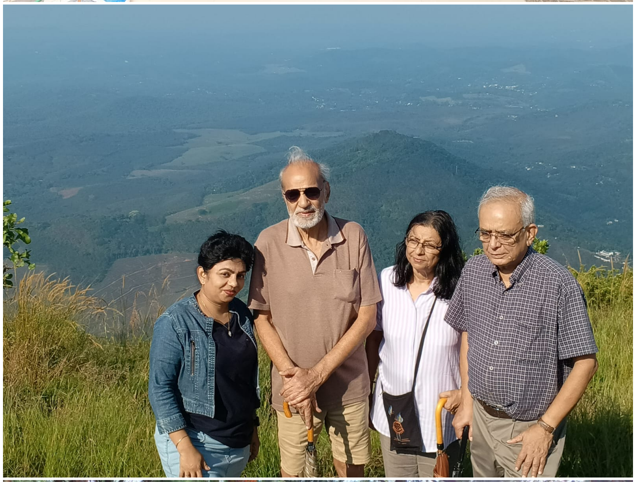

Professor Naresh Dadhich had spend three days as guest to airis4D in November 2024 with his wife Sadhana and Professor Kembhavi and his wife Asha. The memories of the days spend together will never fade away.

Contents

Professor Naresh Dadhich
September 1, 1944 – November 6, 2025

by Ajit Kembhavi

airis4D, Vol.3, No.12, 2025

www.airis4d.com

Professor Naresh Dadhich passed away, at the age of 81, on November 6, 2025 in Beijing while on an academic visit to China. His sudden demise is a tragic loss to the scientific community, and to the very large number of people from diverse fields and professions, in India and in many other countries, who were his friends, colleagues and collaborators. He worked mainly in general relativity and gravitation theory, and was one of the founders of the Inter-University Centre for Astronomy and Astrophysics (IUCAA), where he was Director during 2003-2009. He continued to be associated with the institute as an active scientist to the very end.

Dadhich was born and brought up in a village near Churu in Rajasthan, where his father was a priest. He left home for school at a young age, got his first degree in mathematics from BITS-Pilani and his M.Sc. from Vallabh Vidyanagar. He then reached the University of Poona (now SP Pune University, SPPU) in 1966, for research in mathematics. There he had the good fortune of becoming a Ph.D. student of the renowned mathematician and relativist Professor V. V. Narlikar, the father of Professor Jayant Narlikar. Naresh specialised in the very difficult area of general relativity, in which he became a great expert.

Soon after his Ph. D. in the early seventies, Naresh was appointed as a lecturer in the Department of Mathematics in SPPU, where his scientific journey truly began. He worked on classical and quantum aspects of gravity, with Lovelock Gravity being one of his favourites, brane world cosmologies, gravitational collapse, wormholes and the astrophysics of black holes. Over some years he gathered around him a number of highly talented young researchers, including Sanjeev Dhurandhar, B. S. Sathyaprakash and Patrick Dasgupta, who had obtained their Ph.D. from leading institutes in the country. There were also research students including Sanjay Wagh, Ravi Kulkarni, Sucheta Koshti and Varsha Daftardar. Some of these young people became internationally leading researchers in their areas. The work done during this period included the very interesting magnetic Penrose process for the extraction of energy from black holes. Naresh continued to publish until his last days, and had a paper accepted just a few days before he passed away. His essays for the Gravity Research Foundation annual competition got Honorable Mention several times, including in 2025, placing him amongst the oldest persons to be so honoured. He was President of the Indian Association of General Relativity and Gravitation many years ago.

Naresh had very simple sounding explanations for profound and difficult to understand concepts like the universality of space and time, constancy of the velocity of light, the curvature of space-time and its manifestation as gravity, derivations of Einstein's field equations and so forth. He lectured on these matters to diverse audiences, ranging from professional scientists to students barely out of school, with the same words and elan. Over the last few months, I have heard him speak on these topics to early college students in Darjeeling and Masters and research students in Silchar. It was never clear to me how much of these lectures were actually understood by the young people listening to him, but they certainly enjoyed the experience of listening to a person who looked exactly like a scient should, and spoke from his heart, wholly believing everything that he said.

Around 1987, Naresh's life took a dramatic turn. Jayant Narlikar wanted to set up a inter-university centre for astronomy, which took form as IUCAA, and Naresh played a pivotal role in creating the place. He was the first person to be appointed on the roles of IUCAA, even before Jayant Narlikar's appointment as the Founding Director. Naresh worked tirelessly providing liaison between SPPU, UGC, the government of Maharashtra and several miniseries at the Centre. The very difficult task of setting up an institution with a complex structure was made so much easier because of his efforts. He helped in identifying and transferring a piece of land on the SPPU campus for the new institute.

The message that IUCAA had been created had to be taken far and wide in the country. Naresh and I travelled incessantly, sometimes accompanied by Jayant Narlikar and other colleagues, to the main cities, as well as to universities and colleges in the smaller cities and towns in far flung areas. We talked to the faculty and students there about the facilities that IUCAA offered. We convinced them that their own specialties in physics, mathematics and statistics could find wonderful applications in astronomy. Their expertise could be harnessed to solving the new problems emerging from the multitude of new telescopes and satellites, and the use of emerging information technology. Soon the tide turned, and more and more people stated visiting IUCAA regularly to collaborate with the people there. New groups of astronomers emerged in several universities and colleges, and now we have a thriving community contributing to the national astronomical effort. We were particularly successful in West Bengal, Assam, other North-Eastern states, Kashmir and Kerala, where there are a large number of talented young people looking for new opportunities. Naresh and I continued to travel to various centres until he left for China a few weeks ago. We did that together for 36 years, acquiring in the process many collaborators and friends. Our younger colleagues now continue the tradition, bringing to IUCAA increasing numbers of highly creative modern young people.

Naresh had deep and abiding friendships and collaborations with a number of relativists and astrophysicists in many countries. There was an unusual element to these collaborations. While he worked with leading scientists in Western countries, and visited them often, he also collaborated over decades with groups in various countries like Uzbekistan and other Central Asian republics, Iran, Turkey, Pakistan and South Africa. Faculty as well as students from these countries visited IUCAA regularly, and many Ph.D.s were produced through these interactions. His ties with South Africa began soon after apartheid ended and during his many visits to the country, he developed close ties with the intellectual elite there, including scientists from different fields, artists, constitutional court justices, senior science administrators, vice-chancellors and others, and he even had a meeting with Nelson Mandela soon after his release from prison.

Over the years, and especially when he was Director, Naresh helped in making large astronomical facilities available to IUCAA. He got the telescope at Girawali going, and he made IUCAA a partner in the Southern African Large Telescope. He also began the process of IUCAA and other institutions in the country becoming partners in the Thirty Meter Telescope project, and building a LIGO gravitational wave detector in India. Sanjeev Dhurandhar and Naresh tried to set a up a gravitational wave detector in the country about 25 years ago, but they were far ahead of their time. Their pioneering efforts were not in vain, however, since they laid the foundation for the approval of the LIGO-India project in 2016.

Naresh had intellectual convictions which went far beyond his scientific side. In spite of his early family background, he was a non-believer and rationalist. He was a committed socialist and firmly believed in the equality of all men and women, young and old, and rich and poor. He was an activist and in spite of his busy schedule, often participated in marches and demonstrations on a variety of causes, including the environment. He lived a simple life bordering on the spartan, but he liked his fun too. He greatly enjoyed going to parties and organising many parties himself. Given the current rightwards march all over the world, I used to tell him that he would soon be the last *jholawala* still standing. He would of course have found a kindred spirit in the newly elected Mayor of New York, but he passed away before I could tell him that.

Naresh had a vast circle of friends and admirers in his home city of Pune and elsewhere in the country. He had very close ties with people from the performing arts and theatre all over. For many years, great playwrights, artists and thespians visited and performed in IUCAA, providing a rich and multihued background to the excellent science being done there.

Naresh is survived by his wife Sadhana, who is a deeply committed social worker and activist, his daughter Juee who is an entrepreneur in Information Technology, and his son Nishith, who is a successful producer of movies and serials.

Naresh clearly was a person of many parts. His tall, handsome, voluble presence, his infectious laughter and gracious company will be missed by many for long.

About the Author

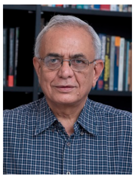

Professor Ajit Kembhavi is an emeritus Professor at Inter University Centre for Astronomy and Astrophysics and is also the Principal Investigator of the Pune Knowledge Cluster. He was the former director of Inter University Centre for Astronomy and Astrophysics (IUCAA), Pune, and the International Astronomical Union vice president. In collaboration with IUCAA, he pioneered astronomy outreach activities from the late 80s to promote astronomy research in Indian universities.

My Childhood Memories of Naresh Dadhich

by Joe Philip Ninan

airis4D, Vol.3, No.12, 2025

www.airis4d.com

My fond memories of Prof. Naresh Dadhich go back to the astronomy workshops that IUCAA used to conduct at Charal Mount, a hill campus near St. Thomas College, Kozhencherry, Kerala, in the late 90s. I was a middle school kid then, accompanying my father to these workshops. Dadhich was a towering figure to me—both figuratively and literally! All the kids adored him, and we would eagerly queue up to be picked up and swung around by him. Those were the best flying experiences ever!

He was also an extraordinary teacher. He would patiently explain the core concepts of the special theory of relativity and even the general theory of relativity at a level that any school kid could understand. No matter how tired he was after lecturing to the main workshop participants, he would happily listen to our naïve questions and explain things from first principles. His excellence as a teacher was evident not only in how he answered our queries, but also in how he encouraged us to think more deeply about each problem. Once, after explaining how the presence of any energy is sufficient to cause curvature in the spacetime metric and bend light, I remember asking him whether that meant two parallel light photons travelling billions of light years next to each other would be attracted to one another and end up bunching together. He smiled warmly, appreciated the question, and encouraged me to think about it more deeply and try to come up with an answer myself by the next day. I didn't have an answer then, but reflecting on that problem definitely sparked a deep curiosity to learn more physics throughout my school years.

Later in my life, I had the honour of meeting him multiple times at IUCAA during various stages of my career, and every interaction was insightful and inspiring. He was my childhood hero, who played a significant role in inspiring me to take up a career as a scientist.

Picture from the 1997 workshop at Charal Mount, I am standing third from the right in this photo, with my sister, mother and father between Prof. Ajit Kembhavi and Prof. Narsh Dadhich.

About the Author

Joe Philip Ninan, works on exoplanets, and astronomy instrumentation, He is currently a faculty in the Astronomy department at Tata Institute of Fundamental Research, Mumbai.

Part I

Artificial Intelligence and Machine Learning

Few Shot Learning

by Arun Aniyan

AIRIS4D, VOL.3, NO.12, 2025

www.airis4d.com

1.1 Introduction

One of the most significant and pervasive challenges in modern Artificial Intelligence, particularly within the field of deep learning, is the data scarcity problem. This issue refers to the fundamental limitation that, while incredibly powerful and capable of complex pattern recognition, traditional deep learning models are inherently data-hungry. They typically require an enormous corpus of labeled examples—often scaling into the thousands or even millions—to be trained effectively and to achieve robust, reliable generalization to unseen data.

This profound reliance on vast, meticulously curated datasets presents a major and often insurmountable bottleneck in a myriad of critical real-world scenarios. Consider domains such as medical imaging, where obtaining millions of examples of a rare disease is simply infeasible due to patient privacy, cost, or the sheer infrequency of the condition. Similarly, in specialized industrial applications, robotics, or natural language processing for low-resource languages, the cost and effort of manually labeling sufficient data are prohibitive. The data scarcity problem is not merely an inconvenience; it actively hinders the deployment of deep learning solutions in areas where data acquisition is inherently difficult, expensive, sensitive, or slow, thereby creating an urgent need for more data-efficient learning paradigms.

1.2 Practical Challenges of Standard Deep Learning

Standard deep learning models are fundamentally rooted in the principle of statistical learning. This paradigm dictates that a model's intricate network of parameters is fine-tuned through an iterative optimization process, specifically by minimizing a defined loss function across an expansive dataset. The efficacy of this process hinges on a critical trade-off: the volume of data must be immense enough to simultaneously prevent overfitting—the detrimental scenario where the model merely memorizes the idiosyncrasies and noise of the training data—and to guarantee robust generalization—the model's true measure of success, which is its ability to perform accurately and reliably on entirely unseen, out-of-sample data. The prerequisite for these massive datasets, however, introduces several significant and often prohibitive practical difficulties across various real-world applications:

- The Financial and Temporal Burden of Data Acquisition
  The process of acquiring raw data is often just the beginning. The subsequent step—labeling or annotating this data with accurate ground-truth information (e.g., classifying an image, transcribing audio, or delineating boundaries)—is an incredibly resource-intensive endeavor. It demands significant financial investment to hire and train human annotators and consumes substantial amounts of time, creating a bottleneck that delays the development and

deployment of new AI systems. This cost scales linearly, or often super-linearly, with the desired size and complexity of the dataset.

- Constraints Imposed by Privacy and Ethical Considerations
  In highly sensitive and regulated domains, such as healthcare (e.g., patient records, medical scans) and finance (e.g., transaction histories, personal credit data), the volume of accessible data is severely restricted. Strict government regulations (like GDPR, HIPAA) and corporate policies, driven by critical ethical concerns over individual privacy and data security, mandate data anonymization, aggregation, or outright restriction, making the assembly of massive, comprehensive datasets practically impossible.
- Challenges in Handling Novelty, Rarity, and Extreme Class Imbalance
  Deep learning struggles profoundly when faced with categories that are newly emerging or extremely rare. By their nature, these categories preclude the possibility of gathering the large, statistically representative datasets required by standard models. Examples include:
  1. Scientific Discovery: A new, never-before-seen viral strain or an anomaly in astronomical data.
  2. Ecology: A rarely observed or newly discovered animal or plant species.
  3. Security: A novel, sophisticated zero-day cyber threat or a rare type of equipment failure. The lack of historical data for these 'outlier' events makes the statistical learning approach ineffective, as the model has little to no data to learn from, making accurate prediction and classification nearly impossible.

In summary, the statistical appetite of standard deep learning forms a powerful barrier to entry for many critical applications, necessitating innovative approaches—such as Few-Shot Learning—that can learn effectively from minimal data.

1.3 Few-Shot Learning: A Paradigm Shift in Machine Learning

Few-Shot Learning (FSL) represents a critical evolution in the field of machine learning, emerging as a vital solution to the pervasive challenge of data scarcity. At its core, FSL is a specialized subfield dedicated to creating models capable of learning effectively from only a handful of labeled examples, often as few as one or five. This approach stands in stark contrast to traditional deep learning methodologies, which require massive, meticulously curated datasets (often thousands or millions of examples) to achieve high performance. FSL aims to mimic the remarkable efficiency of human learning, where a person can typically recognize and categorize a new concept—like a novel type of bird or a new gadget—after seeing just a single or a few instances.

The fundamental objective of FSL is not simply to achieve high accuracy with less data, but to instill a deeper, more adaptable form of intelligence in AI systems. This is often achieved through meta-learning, where the model is trained on a vast array of tasks to learn how to learn quickly, rather than just learning a single task. The knowledge gained from a diverse set of "training tasks" is then rapidly transferred to a new, data-poor "test task."

1.3.1 The Critical Need for FSL in Real-World Applications

The significance of Few-Shot Learning is most pronounced in real-world scenarios where the conditions for traditional big-data AI are difficult or impossible to meet. In these domains, data collection is inherently:

- **Expensive and Time-Consuming:** The manual labeling and verification of specialized data require expert human labor and significant financial investment.
- **Ethically Constrained or Private:** Certain data, such as individual health records or biometric profiles, are protected by privacy laws, severely limiting the size of available datasets.

- **Atypical or Rare:** By definition, some phenomena occur infrequently, making large-scale data collection fundamentally impossible.

FSL is therefore critical for enabling AI adoption across several high-impact sectors:

- **Medical Imaging and Diagnostics:** This is a particularly crucial application. FSL allows for the training of diagnostic models for rare diseases or novel strains of viruses where only a minute number of confirmed, expertly-labeled cases exist globally. It provides a path to deploy life-saving AI tools even when the data is extremely scarce.
- **Scientific Discovery and Conservation:** In fields like biology and astronomy, FSL can be used to rapidly identify and classify new species of plants, animals, or celestial objects based on initial, limited observations. This accelerates the pace of discovery and aids in conservation efforts for endangered or newly found organisms.
- **Robotics and Autonomous Systems:** For robots to be truly autonomous, they must be able to rapidly adapt to new, unseen objects, tools, or environmental conditions without needing hours of pre-training. FSL allows a robot to learn the function of a new grip or the layout of an unfamiliar room with minimal demonstration.
- **Personalized AI and User Experience:** FSL enables the development of user-specific models with exceptional efficiency. This is vital for personalization, where an AI assistant might learn a user's unique speech patterns, handwriting style, or specific product preferences with just a few examples, enhancing privacy and user experience.

By empowering models to glean meaningful insights from limited data, Few-Shot Learning represents a fundamental step toward achieving Artificial General Intelligence (AGI). It pushes AI closer to the adaptability, efficiency, and cognitive flexibility that define human learning.

1.4 Core Concepts in Few-Shot Learning

Few-Shot Learning typically involves two main phases: the Meta-Training Phase and the Few-Shot Task Phase.

1.4.1 Meta-Training Phase (Training the 'Learner')

In this phase, the model is trained on a large base dataset composed of many different classes. The goal is not to learn to classify these base classes perfectly, but to learn a robust and generalizable learning strategy or a good feature extractor.

- **Data Structure:** The base dataset is split into numerous "episodes" or "tasks." Each episode is a mini few-shot problem itself, designed to simulate the final few-shot task.
- **Objective:** The model learns a universal mechanism that can quickly adapt to a new task with minimal examples.

1.4.2 Few-Shot Task Phase (Applying the 'Learner')

This is the real test. The model is presented with a novel dataset containing classes it has never seen before during meta-training. This dataset is structured into a support set and a query set as shown in Table 1.1. This structure is often defined as an N-way K-shot

Table 1.1: Task Description

Set Name 1	Purpose 2
Support set $\mathcal{S}$	Contains the few labeled examples of the new
Query Set $\mathcal{Q}$	Contains unlabeled examples of the new classes f

classification task with:

- N-way: The number of novel classes.
- K-shot: The number of labeled examples per class (i.e., the "few shots").

An example would be 5 way 1 shot task, which means if $N = 5$ and $K = 1$, the model sees only one labeled example for each of the five new classes in the support set and must then correctly classify new examples in the query set.

1.5 Key Approaches to Few-Shot Learning

Several innovative techniques have been developed to tackle the FSL problem. The primary goal is either to find a better starting point for learning or to improve the classification process itself.

1.5.1 Metric-Learning Approaches

Metric-learning methods aim to learn an embedding space (a feature representation) where examples of the same class are clustered closely together, and examples of different classes are far apart.

- **Prototypical Networks:** This is a popular and intuitive approach. For a given few-shot task, the model calculates a prototype (the mean vector) for each class in the support set. New query examples are classified based on the distance (e.g., Euclidean distance) to the closest prototype.
- **Relation Networks:** Instead of just measuring distance, these networks learn a non-linear similarity metric. They take a pair of images (a support image and a query image) and output a score indicating how likely they belong to the same class.

1.5.2 Model-Agnostic Meta-Learning (MAML)

MAML is a highly influential technique that focuses on finding a set of initial model parameters that are very sensitive to change. This means that with only a few gradient descent steps and a few examples (K-shots), the model can quickly adapt and perform well on a new task.

The meta-training objective is to minimize the final loss after performing one or a few inner-loop gradient steps. It learns a universal initialization point for the learner.

1.5.3 Data Augmentation and Generative Models

Another strategy is to increase the amount of data by synthesizing new, realistic examples. While standard data augmentation (e.g., rotation, cropping) helps, FSL often uses more sophisticated generative models (like GANs or VAEs) to create synthetic examples of the few-shot classes, effectively turning the K-shot problem into a richer-shot problem.

1.6 The Few-Shot Learning Process Workflow

Few-Shot Learning involves a systematic, multi-step process. The steps below outline a typical metric-learning workflow:

1. **Meta-Training**
   Train a feature encoder network on episodes from the base dataset to learn a robust feature space.
2. **Task Formulation**
   Sample an N-way K-shot task, creating a Support Set S and a Query Set Q with novel classes.
3. **Feature Extraction**
   Use the trained encoder to map all images in S and Q into the embedding space.
4. **Prototype Calculation**
   For each class in S, calculate the class prototype (e.g., the average of the feature vectors).
5. **Prediction (Inference)**
   Classify each query example by finding the prototype it is closest to in the embedding space.

1.7 Conclusion

Few-Shot Learning (FSL) is not merely an incremental improvement over existing deep learning techniques; it represents a paradigm shift toward more efficient, adaptable, and human-like AI. We have seen that FSL directly addresses the most crippling limitation of standard deep learning—its insatiable demand for data—by enabling robust model performance from only a handful of examples. Through meta-learning strategies, metric learning, and sophisticated data synthesis, FSL allows models to acquire the crucial skill of learning how to learn.

The immediate and long-term potential of FSL is transformative, promising to unlock AI deployment in high-stakes, data-scarce domains previously considered intractable:

- **Accelerating Scientific Discovery:** FSL is poised to significantly accelerate research in areas like medical diagnostics for rare diseases, real-time identification of novel pathogens, and rapid classification in ecological and astronomical surveys, where data inherently exists in small batches.
- **Enabling Adaptive Robotics:** For autonomous systems and robotics, FSL provides the necessary cognitive flexibility for quick adaptation to new tools, environments, and tasks, moving robots closer to true independence and utility in unpredictable human settings.
- **Fostering Personalized AI:** FSL allows for the creation of hyper-personalized models that learn an individual's specific preferences, biometric data, or interaction styles with minimal data points, greatly enhancing user privacy and the quality of user-AI interaction.

In essence, FSL moves Artificial Intelligence beyond the brute-force statistics of "big data" toward a model of genuine intelligence characterized by efficiency and generalization. For the next generation of AI researchers, Few-Shot Learning represents a crucial, active frontier—a fundamental step in realizing the dream of Artificial General Intelligence, where machines can learn new concepts as rapidly and effectively as humans. The ongoing research into more sophisticated meta-learning architectures and generative models promises to further solidify FSL's role as the cornerstone of data-efficient machine learning.

1.8 References

- https://towardsdatascience.com/few-shot-learning-a-new-paradigm-in-ai-a5170f438a95
- https://arxiv.org/abs/1904.05046
- https://lilianweng.github.io/posts/2018-11-20-meta-learning/

About the Author

Dr.Arun Aniyan is leading the R&D for Artificial intelligence at DeepAlert Ltd,UK. He comes from an academic background and has experience in designing machine learning products for different domains. His major interest is knowledge representation and computer vision.

The Time Arrow of Entropy

by Jinsu Ann Mathew

airis4D, Vol.3, No.12, 2025

www.airis4d.com

Entropy is usually described as a measure of disorder or uncertainty. But real systems—whether physical, biological, or social—do not stay still. They move, change, and carry memories of their past. When we look at entropy through the lens of time, a deeper idea appears: change has a direction. Some things naturally happen in one way but not the reverse—heat spreads out, languages evolve, societies grow, and information slowly gets lost.

This one-way movement is known as the arrow of time, and entropy is what gives it direction. By understanding how entropy changes over time, we can see why systems age, why they resist going backwards, and how their past shapes what they become.

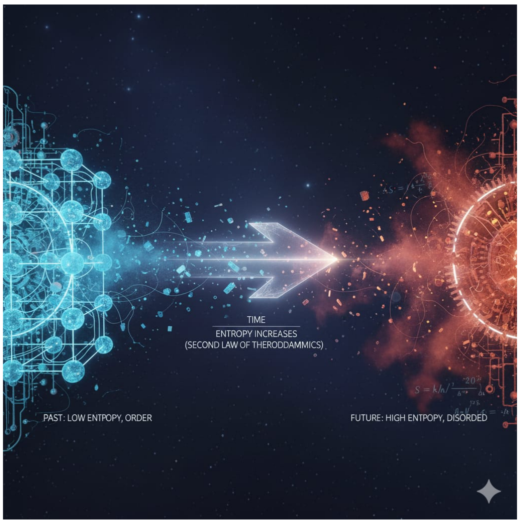

(image courtesy:AI-generated image)

Figure 1: The thermodynamic arrow of time

2.1 Irreversibility: Why Time Has a Direction

Irreversibility means that some things in life happen in only one direction. Even though the basic rules of physics can run forward or backward, the world we see doesn't behave that way. Time feels like it moves forward because certain changes naturally go from order to disorder and almost never the other way around. This is why we remember yesterday but not tomorrow, why things age instead of becoming new again, and why many processes cannot simply "undo" themselves.

A simple example is a glass falling and breaking. When the glass is whole, all its pieces are arranged in a very neat, organized way. Once it breaks, the pieces scatter in many different positions. It's easy for a whole glass to break because there are many ways for the pieces to land, but it's almost impossible for the scattered pieces to jump back and form the original glass again. So in everyday life, we only see the breaking, not the un-breaking.

Another familiar example is heat. If you place a hot spoon in cold water, the heat spreads into the water until both become warm. But the warmth never suddenly leaves the water to make the spoon hot again. The mixing of heat happens naturally, but the reverse doesn't. These everyday experiences show us why time feels like it goes in only one direction: some changes happen easily, while their opposites are so unlikely that we never see them. That is what gives time its direction.

2.2 Memory in Natural and Artificial Systems

Memory in a system means that its past affects how it behaves today. Some systems hold on to their previous states for a long time, while others forget quickly. This "memory" doesn't always mean conscious remembering—it can simply be a pattern, a structure, or a change that stays even after the original cause is gone. Because of this, systems do not start fresh at every moment. Their history continues to shape what they do next.

In nature, many things show this kind of memory. A magnet keeps its direction even after the external force that aligned it is removed. Trees remember past seasons by forming growth rings that reflect dry or rainy years. Even the human body carries memory—your immune system "remembers" past infections and reacts faster when it encounters the same germ again. These are all examples where the past leaves a mark that influences the future. Artificial systems, like machines and computers, also have memory. A simple thermostat "remembers" the previous temperature setting and adjusts the room based on that. A trained neural network remembers patterns from the data it learned and uses them to make predictions. Even phone keyboards have memory: they learn the words you type often and suggest them more quickly. In all these cases, the system's past experience guides how it behaves now.

Memory slows down how quickly a system changes. Instead of jumping immediately into randomness or disorder, the system holds on to what it has learned or experienced. This is why memory plays an important role in how systems evolve over time—it keeps the past alive and shapes the direction of future behaviour.

2.3 History in Language and Society

History plays a powerful role in shaping both language and society. Neither of them starts from scratch at any moment; instead, they grow on top of what came before. This means that the way people speak, behave, and organize themselves today is strongly influenced by choices, habits, and patterns that developed over many years. Because of this deep influence of the past, changes in language or society usually happen slowly and rarely return to earlier forms.

In language, history leaves clear traces. Words change their meanings over time, but they do not suddenly become what they were centuries ago. Grammar also shifts gradually, building on earlier usage rather than reinventing itself. For example, English no longer uses words like "thou" and "thee," and it will not suddenly bring them back into everyday speech. New words and styles appear, but they grow out of older ones. The language we speak today is a layered record of how people spoke in the past. Society shows the same pattern. Customs, laws, and social norms evolve over time, shaped by many years of shared experience. Traditions—such as festivals, greetings, or family roles—continue because they have been passed down through generations. Even when societies change, they do so gradually, carrying forward parts of their history. For instance, cities develop new technologies and lifestyles, but old neighbourhoods, cultural habits, and social expectations often remain. The past acts like an anchor, influencing how people behave and how they make decisions.

Because language and society carry their history with them, they rarely move backward. Change happens in a forward direction, building on what already exists. This is why history becomes an active force: it keeps certain patterns stable, guides how new ones form, and shapes the way communities grow and communicate

2.4 Entropy Over Time: Patterns, Trends, and Transitions

As we follow entropy through time, we begin to notice the patterns in how systems change. Entropy does not stay fixed; it shifts in response to the small and large events happening inside the system. Sometimes these changes are gentle and gradual, and sometimes they are sudden. Watching these changes helps us understand the rhythm of a system's life—how it grows, adapts, or stabilizes.

In many systems, entropy moves in small daily fluctuations. These little rises and falls come from ordinary variations: a slight change in how people speak, a tiny shift in temperature, or a brief change in social behaviour. These small movements may seem trivial, but they show that the system is active and constantly responding to its environment. Over longer periods, clearer trends appear. Some systems slowly become more uncertain or diverse—for example, a language steadily adding new words or a community becoming more interconnected over years. Other systems might become more organized, as when a group settles into new routines or when a learning system gradually becomes more confident and consistent. These long-term trends show the direction a system is heading.

At times, however, entropy does not change slowly at all. Instead, it jumps. These jumps often signal major turning points—a sudden shift in public mood, a rapid change in technology, or a natural system reaching a threshold and reorganizing itself. These transitions are important because they mark moments when the usual patterns break and something new begins. By looking at entropy over time, we gain a clearer picture of how systems behave. It becomes easier to see when they are stable, when they are slowly drifting, and when they are on the edge of change. Entropy, viewed this way, becomes not just a measure of uncertainty but a tool for understanding how systems move and evolve.

2.5 Conclusion

Looking at entropy through the lens of time helps us understand why systems move the way they do. Whether we study physical processes, languages, or societies, we find the same pattern: the past shapes the present, and changes tend to move in a forward direction. Irreversibility, memory, and history each show us how systems carry traces of what they were, and how these traces guide what they become. By watching how entropy shifts—slowly, quickly, or suddenly—we can see not just the current state of a system but the path it is taking. Entropy, in this sense, becomes more than a measure of uncertainty. It becomes a way to read the story of a system's evolution, showing how small actions build into larger patterns and how the arrow of time leaves its mark on everything that changes.

References

- Entropy - Physics
- Entropy Is Universal Rule of Language
- Hidden Information, Energy Dispersion and Disorder: Does Entropy Really Measure Disorder?
- Where Does Our Arrow Of Time Come From?
- Can Time really heal anything?

About the Author

Jinsu Ann Mathew is a research scholar in Natural Language Processing and Chemical Informatics. Her interests include applying basic scientific research on computational linguistics, practical applications of human language technology, and interdisciplinary work in computational physics.

Unboxing a Transformer using Python - Part II

by Linn Abraham

AIRIS4D, Vol.3, No.12, 2025

www.airis4d.com

Code: Position and Class Embedding

```
1  class ViT(nn.Module):
2      def __init__(self, *, image_size, patch_size, num_classes, dim, depth, heads, mlp_dim, pool = 'cls',
   ↪   channels = 3, dim_head = 64, dropout = 0., emb_dropout = 0.):
3          super().__init__()
4          num_patches = (image_height // patch_height) * (image_width // patch_width)
5          ...
6          self.pos_embedding = nn.Parameter(torch.randn(1, num_patches + 1, dim))
7          self.cls_token = nn.Parameter(torch.randn(1, 1, dim))
8          ...
9
10     def forward(self, img):
11         x = self.to_patch_embedding(img)
12         b, n, _ = x.shape
13         cls_tokens = repeat(self.cls_token, '1 1 d -> b 1 d', b = b)
14         x = torch.cat((cls_tokens, x), dim=1)
15         x += self.pos_embedding[:, :(n + 1)]
16
```

Figure 1: Position and Class embedding in ViT

3.1 Introduction

In the first part of this article, we looked into the patch embedding layer part of the Vision Transformer. In this article, we continue from there and look into how the position and class information are learnt and then onto how the transformer is implemented.

3.2 Position & Class Embedding

By creating patches and flattening them out, we have destroyed the position information of the patches in the original 2D grid. Also, since these patches are processed in parallel by the attention mechanism, even the 1D ordering of the patches is meaningless. But the spatial position is a physical information that we want the model to learn. Additionally, since this is a classification model, we also want the class label to be learnt by the model. Thus we want to to have placeholders for these two important pieces of information in the embedding and then let the model learn it.

To see how how this is implemented, note that we first embed the class token. The class token is defined and initialized as a learnable parameter of shape $[1, 1, dim]$ (see, line 7 in the code). The `repeat` function takes this class token and repeats it along a dimension as specified by the pattern. The new variable b which defines the number of repetitions to be made along the zeroth axis is passed as an input to the function. These are then concatenated to the patch embeddings along its 1st dimension which is the sequence dimension with size equal to number of patches in the image. Thus the new shape becomes $[b, n + 1, dim]$ with the first token in each sequence being the class token followed by the patch tokens. The positional embedding that we defined earlier is now added to the output of the class-token embedded tensor. Note that this is an element-wise addition and not a concatenation as with the case of the class token. Also note that there is broadcasting happening in the zeroth dimension (the 1 gets broadcasted to the batch size). Other than this, there is no change in the output shape as a result of this addition. Also we require n+1 in the sequence dimension (which might be to allow for cases where the number of patches are variable or something like that).

3.3 The Transformer

Let's see what has happened so far. We started with a batch of images, which have been made into a batch of sequences with each sequence containing the image patches, that have been flattened along the height, Source Code: Defining the Transformer

```
class Transformer(nn.Module):
    def __init__(self, dim, depth, heads, dim_head, mlp_dim, dropout = 0.):
        super().__init__()
        self.norm = nn.LayerNorm(dim)
        self.layers = nn.ModuleList([])
        for _ in range(depth):
            self.layers.append(nn.ModuleList([
                Attention(dim, heads = heads, dim_head = dim_head, dropout = dropout),
                FeedForward(dim, mlp_dim, dropout = dropout)
            ]))

    def forward(self, x):
        for attn, ff in self.layers:
            x = attn(x) + x
            x = ff(x) + x

        return self.norm(x)
```

Figure 2: Definition of the transformer in ViT

width and channel dimensions, and with an extra class token added to it. What remains to be explained is the transformer block, which we partly saw in the beginning of this article.

After the class and position embedding steps, a dropout layer is added before being fed into the transformer block. Within the transformer block we first see a normalization layer. After this comes the `to_qkv` function which is does a linear projection to a dimension which is three times the inner dim. This inner dimension is the product of the dimension per attention head times the number of attention heads. So this is where the Q, K and V matrices are learnt. After this it is again split into the three different matrices. And then reshaped to split into each individual head so that the dot product of the q and k vectors can be taken. The dot product is implemented as a matrix multiplication with one vector transposed. The transpose is taken by swapping the last two dimensions. At this stage the output is also scaled for stability. This is followed by another dropout layer. After this a weighted addition of the value vectors with the attention scores are taken using regular matrix multiplication. The final rearrangement converts the single head into the multi-head output. If more than one head is used add another projection from the inner dimension to the input dimension otherwise just use an identity operation.

3.4 The Final layers

After the transformer block we have a pooling layer. Two valid options are mean or class. If mean, all the tokens are averaged else the class token is taken as the representation for the image. This is followed by the MLP head which is a linear layer with size equal to the number of classes. In between we have a `to latent` layer which is just an identity operation in this case. This is done so that it can be swapped with something else if required without changing the forward function.

Source Code: Defining Attention

```
class Attention(nn.Module):
    def __init__(self, dim, heads = 8, dim_head = 64, dropout = 0.):
        super().__init__()
        inner_dim = dim_head *  heads
        project_out = not (heads == 1 and dim_head == dim)

        self.heads = heads
        self.scale = dim_head ** -0.5

        self.norm = nn.LayerNorm(dim)

        self.attend = nn.Softmax(dim = -1)
        self.dropout = nn.Dropout(dropout)

        self.to_qkv = nn.Linear(dim, inner_dim * 3, bias = False)

        self.to_out = nn.Sequential(
            nn.Linear(inner_dim, dim),
            nn.Dropout(dropout)
        ) if project_out else nn.Identity()

    def forward(self, x):
        x = self.norm(x)

        qkv = self.to_qkv(x).chunk(3, dim = -1)
        q, k, v = map(lambda t: rearrange(t, 'b n (h d) -> b h n d', h = self.heads), qkv)

        dots = torch.matmul(q, k.transpose(-1, -2)) * self.scale

        attn = self.attend(dots)
        attn = self.dropout(attn)

        out = torch.matmul(attn, v)
        out = rearrange(out, 'b h n d -> b n (h d)')
        return self.to_out(out)
```

Figure 3: Definition of Attention in ViT

References

[Goodfellow et al.(2016)Goodfellow, Bengio, and Courville] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. Adaptive Computation and Machine Learning. The MIT press, Cambridge, Mass, 2016. ISBN 978-0-262-03561-3.

[Dosovitskiy et al.(2021)] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. AN IMAGE IS WORTH 16X16 WORDS: TRANSFORMERS FOR IMAGE RECOGNITION AT SCALE. 2021.

[Chollet(2025)] Francois Chollet. *Deep Learning With Python, Third Edition*. MANNING PUBLICATIONS, S.l., 2025. ISBN 978-1-63343-658-9.

[Alammar()] Jay Alammar. The Illustrated Transformer. https://jalammar.github.io/illustrated-transformer/.

About the Author

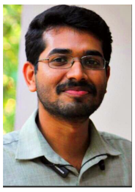

Linn Abraham is a researcher in Physics, specializing in A.I. applications to astronomy. He is currently involved in the development of CNN based Computer Vision tools for prediction of solar flares from images of the Sun, morphological classifications of galaxies from optical images surveys and radio galaxy source extraction from radio observations.

Part II

Astronomy and Astrophysics

Dusty Plasma

by Abishek P S

AIRIS4D, VOL.3, NO.12, 2025

www.airis4d.com

1.1 Introduction

Plasma is a distinct state of matter that arises when a gas becomes highly ionized, producing a mixture of free electrons and ions. Although filled with charged particles, plasmas are generally quasi-neutral, meaning that positive and negative charges balance out on large scales. A key feature of plasma is Debye shielding, where electric fields are screened over a characteristic distance known as the Debye length [1]. This prevents long-range electrostatic interactions from dominating its behaviour. Remarkably, plasma makes up about 99.9% of the visible universe, encompassing stars, interstellar gas, and the solar wind. Plasma can be classified into several types based on their temperature, ionization, density, collisions, and magnetic properties. The main categories include cold plasma, warm plasma, hot plasma, ultracold plasma, fully ionized plasma, partially ionized plasma, collisional plasma, non-collisional plasma, neutral plasma, non-neutral plasma, high/low density plasma, magnetic plasma, and dusty plasma [2,3]

Dusty plasma, also called complex plasma, is a type of plasma that contains tiny solid particles ranging from nanometres to micrometres in size that become electrically charged and interact with the surrounding ions and electrons. These dust grains can originate naturally, such as in interstellar clouds, planetary rings, and Earth's mesosphere, or be artificially introduced in laboratory experiments. Once charged, the dust particles significantly alter the plasma's behaviour by disturbing the balance of charged particles, leading to new phenomena not seen in ordinary plasmas. For example, dusty plasmas can form plasma crystals, where dust grains arrange themselves into ordered lattice structures due to strong electrostatic coupling. Depending on conditions, they may behave like a gas, liquid, or solid, making them a fascinating system for studying phase transitions and collective effects. Dusty plasmas are important in astrophysics because they help explain processes in star formation, comet tails, and cosmic dust clouds, while in technology they appear in semiconductor manufacturing, fusion devices, and even spacecraft propulsion [2]. Their study bridges plasma physics, materials science, and astrophysics, offering insights into how matter behaves under extreme and complex conditions.

1.2 Formation of Dusty Plasma

Dusty plasma formation begins with the introduction or creation of solid particles (dust grains) in a gas that can be ionized; these particles may be naturally present (cosmic dust, cometary ejecta, planetary rings, mesospheric aerosols) or produced in situ by mechanical, chemical, or sputtering processes in laboratory and industrial discharges. When the background gas is ionized by thermal energy, electric discharges, photoionization, or energetic particle flux the resulting free electrons and ions collide with and are collected by the dust grains, so that grains rapidly acquire net charge (typically negative, because electrons are far more mobile than ions). Additional charging channels such as photoemission (from ultraviolet illumination) and secondary electron emission (from energetic ion or electron impacts) can modify the sign and magnitude of the grain charge, producing spatial and temporal charge variations. As grains charge, they perturb the local plasma: mobile electrons and ions rearrange to screen the grain potential over a Debye length, so interactions between grains are screened Coulomb (Yukawa) interactions rather than pure Coulomb forces [4]. In many laboratory and space settings, the balance of forces on a grain-electrostatic lift from sheath or ambipolar fields, gravity, ion drag, neutral gas drag, and thermophoretic forces determines whether grains levitate, drift, or settle; in plasma sheaths above electrodes, for example, electric fields can levitate negatively charged grains against gravity, enabling stable suspensions and the formation of ordered structures. As the dust density and coupling strength increase, strong electrostatic coupling between grains can drive transitions from gaseous to liquid-like and even crystalline arrangements (so called plasma crystals), while collective modes unique to dusty plasmas such as dust acoustic waves emerge because the massive, charged grains introduce new low-frequency dynamics. The charging process itself is often time-dependent: fluctuating plasma conditions, variable illumination, and grain motion cause charge to vary on timescales comparable to or longer than grain dynamical times, coupling microscale charging kinetics to macroscopic transport and instabilities. In many natural environments (protoplanetary disks, cometary comae, interstellar clouds) coagulation, charging, and plasma drag together influence grain growth and dynamics, while in technological plasmas (semiconductor processing, fusion edge plasmas) dust formation and charging can degrade performance or seed new instabilities. The net result is a self-consistent, multi-scale formation process in which ionization, grain charging, screening, and force balance produce a complex medium whose structure and waves differ qualitatively from ordinary two-component plasmas.

1.3 Forces acting on dust particles

Dust particles immersed in an ionized gas rapidly acquire charge through electron and ion collection, photoemission, and secondary electron emission, so that electromagnetic forces often dominate their dynamics; the electrostatic force $F_E=Q_dE$ produced by local electric fields can levitate negatively charged grains in sheaths or drive them along field gradients, while gravity $F_g=m_dg$ pulls them downward and sets a baseline for levitation height. Streaming ions transfer momentum to grains producing ion drag (with collection and orbital components) that can push grains in the ion flow direction and modify equilibrium positions; collisions with neutral gas molecules produce neutral (viscous) drag that damps grain motion and controls relaxation timescales. Temperature gradients in the neutral gas give rise to a thermophoretic force that drives grains from hot to cold regions, and radiation pressure or photoelectric effects can both exert direct momentum transfer and change grain charge, altering electrostatic responses. Charged grains interact with one another through Debye-screened Coulomb (Yukawa) potentials, so interparticle electrostatic forces can be strongly repulsive and, at high coupling, produce ordered “plasma crystal” lattices or liquid-like behaviour; the effective coupling is governed by grain charge, spacing, and the Debye screening length [5,6]. Crucially, charging is often time dependent fluctuations in plasma density, temperature, illumination, or grain motion cause the grain charge to vary on timescales comparable to dynamical times, coupling microscopic charging kinetics to macroscopic transport, instabilities, and wave phenomena such as dust acoustic waves. In laboratory setups, the balance of electrostatic lift, ion drag, neutral drag, thermophoresis, and gravity determines whether grains levitate in sheath regions or settle, and microgravity experiments reveal three-dimensional structures otherwise distorted by gravity. In natural environments (cometary comae, planetary rings, interstellar clouds, mesosphere) the same forces, together with coagulation and plasma chemistry, control grain growth, transport, and radiative effects. Understanding dust dynamics therefore requires treating electromagnetic, collisional, and external forces simultaneously, using kinetic, fluid, and molecular-dynamics approaches to capture the multi-scale, strongly coupled behaviour unique to dusty plasmas.

1.4 Collective behaviour of dusty plasma

Dusty plasma collective behaviour arises when many charged grains interact through screened Coulomb (Yukawa) potentials[7], so that the ensemble can behave like a strongly coupled many-body system whose macroscopic properties depend on grain charge, interparticle spacing, and the Debye screening length; under strong coupling the dust component can undergo phase transitions from gas-like to liquid-like and even to ordered solid-like states known as plasma crystals, where grains form lattice structures and support phonon-like excitations. Because dust grains are heavy, they introduce low-frequency collective modes such as the dust acoustic wave, whose dispersion and damping are governed by dust mass, charge, neutral drag, and ion dynamics; these modes couple to ion-acoustic and sheath oscillations, producing modified wave spectra and multi-scale interactions. Charging dynamics and charge fluctuations play a central role: time-dependent collection of electrons and ions, photoemission, and secondary emission cause grain charge to vary, which feeds back on interparticle forces and can seed instabilities, mode coupling, and anomalous transport. External fields and flows electric fields in sheaths, ion streaming, neutral gas flows, and temperature gradients drive directed transport through ion drag, thermophoresis, and electrostatic forces, producing layered levitated structures, voids, and flow-induced ordering; ion streaming in particular can destabilize ordered states and excite dust-density waves. Collisional damping with neutrals and frictional forces set relaxation times and determine whether collective excitations are underdamped or overdamped, while strong coupling enhances correlation effects such as caging, long-range order, and non-Newtonian rheology in dusty plasma liquids. Nonlinear phenomena are common: solitary structures, shocks, vortices, and self-organized patterns emerge from the interplay of long-range screened interactions, dissipation, and external forcing. Laboratory experiments exploit the visualisability of individual grains to study many-body physics directly tracking particle trajectories reveals melting, crystallization, defect dynamics, and transport at the single-particle level while microgravity experiments remove gravitational distortion and expose true three-dimensional collective phases. In astrophysical and technological contexts, collective dusty plasma behaviour influences coagulation and growth of grains in protoplanetary disks, the formation of spokes and structures in planetary rings, and contamination and instability issues in plasma processing and fusion edge plasmas. Overall, the collective dynamics of dusty plasmas combine strong coupling, screened interactions, time-dependent charging, and multi-scale coupling to produce waves, phase transitions, instabilities, and self-organization that make them a uniquely accessible platform for studying complex plasma phenomena

1.5 Dust Acoustic Waves

Dust acoustic waves arise in a dusty plasma when the massive charged dust grains act as the primary inertial component and the lighter plasma species (electrons and ions) supply the restoring pressure, producing a new low-frequency branch of collective oscillation distinct from ordinary ion-acoustic modes [8]. Key properties include a phase velocity that is typically much smaller than ion thermal speeds, a dispersion relation that depends explicitly on dust charge, dust mass, dust density, and the Debye screening length, and a frequency range often in the Hz to kHz band in laboratory conditions. Because dust grains are heavy, the characteristic frequency is low and the waves are easily observable by direct imaging of particle motion, which has made dust acoustic waves a powerful diagnostic in experiments. Damping and modification of the wave arise from neutral gas friction (viscous drag), ion drag, and especially dust charge fluctuations: time-dependent charging processes (electron/ion collection, photoemission, secondary emission) change the grain charge on timescales comparable to the wave period and can either damp or destabilize the mode depending on plasma conditions [9]. In flowing plasmas, ion streaming relative to the dust can drive instabilities that amplify dust acoustic waves or produce shock-like structures; nonlinear evolution yields solitary waves, shocks, and vortices in strongly driven regimes. Theoretical descriptions typically start from fluid or kinetic models that include dust continuity and momentum equations coupled to Poisson's equation with Debye screening, and refinements incorporate collisional absorption, charge variation, strong coupling (Yukawa interactions), and finite grain size effects. Experimentally, dust acoustic waves were predicted and then observed in laboratory discharges where controlled dust injections and sheath electric fields produce levitated dust layers; modern experiments have visualized wave propagation, measured dispersion relations, and demonstrated nonlinear solitary structures and cylindrical solitons. In space and astrophysical settings, dust acoustic–like modes can influence dust transport, coagulation, and structure formation in cometary comae, planetary rings, and dusty regions of the interstellar medium. Overall, dust acoustic waves are a hallmark collective phenomenon of dusty plasmas [10], linking microscopic charging physics to macroscopic wave dynamics and enabling direct study of many-body and nonlinear processes

1.6 Conclusion

Dusty plasmas find broad applications across both natural-science and engineering domains: in astrophysics they help explain processes in protoplanetary disks, cometary comae, planetary rings, and interstellar clouds where charged grains affect coagulation, radiative transfer, and structure formation; in industry, dust formation and charging are central concerns in semiconductor manufacturing and plasma processing because contaminant particles degrade device yield, while in fusion devices dust from wall erosion influences plasma purity and safety and must be managed; in laboratory research dusty plasmas serve as an accessible tabletop platform for studying strongly coupled many-body physics, enabling direct visualization of phenomena such as plasma crystals, melting, defect dynamics, and wave propagation (e.g., dust acoustic waves) that inform both basic science and applied diagnostics. Key technological uses include using controlled dusty plasmas for nanoparticle synthesis and surface modification, employing dust as a diagnostic tracer to map electric fields and flows, and exploiting ordered dust structures to study phase transitions and transport at the single-particle level. Operational risks notably particle contamination, discharge nonuniformity, and dust-driven instabilities in processing and fusion environments mean that practical deployment requires active monitoring (optical imaging, laser scattering, probes), cleanliness protocols, and mitigation strategies such as controlled gas flows, electrostatic confinement, and periodic cleaning to limit performance loss and safety hazards.

Dusty (complex) plasmas are both a challenge and an opportunity: they complicate many plasma technologies by introducing additional charge, mass, and transport channels, yet they also provide a uniquely visible and tuneable system for exploring collective phenomena and for enabling novel material and diagnostic techniques. Understanding and controlling dust charging, screening, and force balances is essential to harness benefits while minimizing harms; progress depends on integrated approaches that combine experiments (including microgravity studies), kinetic and fluid modelling, and molecular-dynamics simulations to capture multi-scale coupling, charge dynamics, and strong-coupling effects. The field therefore sits at the intersection of astrophysics, materials science, and plasma engineering, offering both fundamental insights and concrete applications when dust is measured, modelled, and managed effectively.

References

1. Francis, F, Chen., (2015). "Introduction to Plasma Physics and Controlled Fusion (3rd ed.)". Springer Cham.
2. Mendis, D.A. (1979). "Dust in cosmic plasma environments." Astrophys Space Sci 65, 5–12
3. Shukla, P. K., & Mamun, A. A. (2015). "Introduction to dusty plasma physics."- CRC press.

4. Ignatov, A. M. (2005). "Basics of dusty plasma." Plasma physics reports, 31(1), 46-56.
5. Shukla, P. K., & Eliasson, B. (2009). "Colloquium: Fundamentals of dust-plasma interactions." Reviews of Modern Physics, 81(1), 25-44.
6. Lampe, M., Joyce, G., Ganguli, G., & Gavrishchaka, V. (2000). "Interactions between dust grains in a dusty plasma." Physics of Plasmas, 7(10), 3851-3861.
7. Thomas, H., Morfill, G. E., Demmel, V., Goree, J., Feuerbacher, B., & Möhlmann, D. (1994). "Plasma crystal: Coulomb crystallization in a dusty plasma." Physical Review Letters, 73(5)
8. Rao, N. N., Shukla, P. K., & Yu, M. Y. (1990). "Dust-acoustic waves in dusty plasmas." Planetary and space science, 38(4), 543-546.
9. Rosenberg, M., & Kalman, G. (1997). "Dust acoustic waves in strongly coupled dusty plasmas." Physical Review E, 56(6), 7166.
10. Merlino, R. L. (2014). "25 years of dust acoustic waves." Journal of Plasma Physics, 80(6), 773-786.

About the Author

Abishek P S is a Research Scholar in the Department of Physics, Bharata Mata College (Autonomous) Thrikkakara, kochi. He pursues research in the field of Theoretical Plasma physics. His works mainly focus on the Nonlinear Wave Phenomenons in Space and Astrophysical Plasmas.

Part III

Biosciences

Diabetic Retinopathy- Medical Image Processing

by Geetha Paul

airis4D, Vol.3, No.12, 2025

www.airis4d.com

1.1 Introduction

Diabetic retinopathy (DR) is a progressive eye disease and a leading cause of vision impairment and blindness among individuals with **diabetes mellitus**. This microvascular complication primarily affects the retina, the light-sensitive tissue lining the back of the eye, by damaging its small blood vessels due to prolonged exposure to high blood glucose levels. DR can develop in both type 1 and type 2 diabetes and typically advances over the years, often without initial symptoms. As diabetes prevalence grows globally, diabetic retinopathy has become a major public health concern, particularly impacting adults of working age. The primary risk factors include the duration of diabetes, poor glycemic control, hypertension, and dyslipidemia. Early identification and management, through regular eye examinations and optimised diabetic care, are essential to reduce the burden of visual disability caused by this disease.

The pathogenesis of diabetic retinopathy involves continuous injury to the retinal microvasculature. Chronic high blood sugar levels induce metabolic and biochemical changes, resulting in thickening of the capillary basement membrane, loss of pericytes, and the formation of microaneurysms. As the disease progresses, vascular permeability increases, resulting in fluid leakage, retinal swelling (macular oedema), and eventually ischemia due to capillary non-perfusion. In advanced stages, marked by proliferative diabetic retinopathy, abnormal new blood vessels proliferate on the retinal surface in response to ischemia. These vessels are fragile, prone to bleeding, and can precipitate complications such as vitreous haemorrhage or retinal detachment, ultimately risking permanent vision loss if untreated.

Clinically, DR is classified into non-proliferative (NPDR) and proliferative diabetic retinopathy (PDR) stages. NPDR is characterised by microaneurysms, haemorrhages, hard exudates, and sometimes macular oedema, while PDR is distinguished by neovascularisation. The onset and progression of DR may be insidious; therefore, annual comprehensive dilated eye examinations are recommended for all diabetic patients. Treatment options include optimising systemic risk factors, intravitreal pharmacotherapy (e.g., anti-VEGF agents), laser photocoagulation, and surgical intervention for severe cases. Despite advances in therapy, prevention remains the cornerstone of reducing the impact of diabetic retinopathy by achieving strict glycemic, blood pressure, and lipid control.

2. Non-proliferative diabetic retinopathy

Non-proliferative diabetic retinopathy (NPDR) is the early stage of the disease in which symptoms will be mild or nonexistent. In NPDR, the blood vessels in the retina are weakened. Tiny bulges in the blood vessels, called microaneurysms, may leak fluid into the retina. This leakage may lead to swelling of the macula.

3. Proliferative diabetic retinopathy

Proliferative diabetic retinopathy (PDR) is the more advanced form of the disease. At this stage,

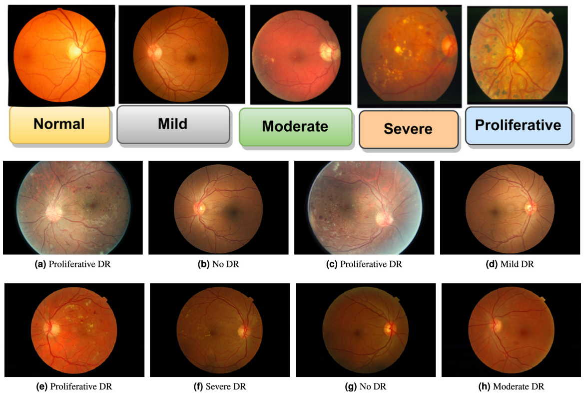

Figure 1: High quality to low quality of DR Retinal images,

Image courtesy:https://www.nature.com/articles/s41598-021-93632-8?fromPaywallRec=false

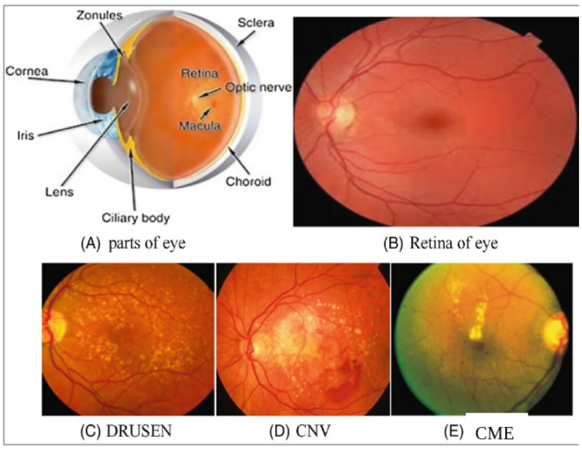

Figure 2: (A) Illustration of the parts of the eye, (B) Retina of the eye, (C) Drusen are yellow deposits under the retina, (D) CNV (Choroidal Neovascularisation) is abnormal blood vessel growth beneath the retina, causing leakage; and CME (Cystoid Macular Oedema): is fluid accumulation causing retinal swelling.

Image courtesy: https://www.nature.com/articles/s41467-021-23458-5

circulation problems deprive the retina of its oxygen supply. As a result, new, fragile blood vessels can begin to grow in the retina and into the vitreous, the gel-like fluid that fills the back of the eye. The new blood vessels may leak blood into the vitreous, causing vision to become cloudy

. Diabetic retinopathy is a leading cause of blindness resulting from damage to retinal blood vessels caused by diabetes. Image processing plays a central role in detecting, grading, and monitoring diabetic retinopathy using retinal fundus photographs.

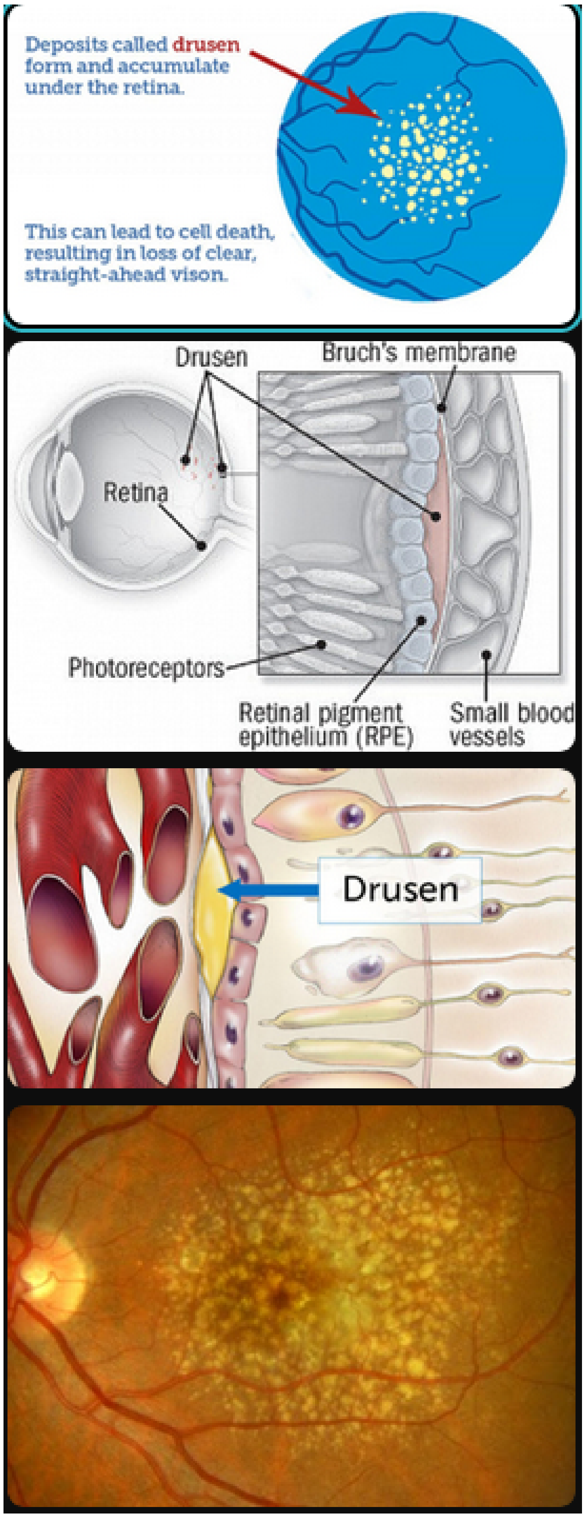

Figure 3: Illustration of Drusen and its formation beneath retina.

Image courtesy: https://www.brightfocus.org/resource/why-is-my-doctor-always-talking-about-drusen/

1.2 Drusen

These are small yellowish deposits of cellular debris, proteins, and lipids that accumulate between the retina and the underlying layer called Bruch's membrane. Drusen are a hallmark feature, especially in age-related macular degeneration, but can also be seen in other retinal conditions. They appear as small spots beneath the retina and indicate waste accumulation resulting from dysfunction of the retinal pigment epithelium.

1.3 CNV (Choroidal Neovascularization):

This is the abnormal growth of new blood vessels from the choroid layer beneath the retina into the subretinal space. CNV is often associated with advanced forms of macular degeneration and can cause leakage and bleeding, leading to vision loss. It is a serious complication that can be seen in diabetic retinopathy as well.

1.4 CME (Cystoid Macular Edema):

This refers to swelling in the macula (central retina) caused by fluid accumulation in cyst-like spaces. CME is a common consequence of diabetic retinopathy that results from vascular leakage and inflammation. It leads to central vision impairment and is a major cause of vision loss in DR.

1.5 Exudates

Exudates in the context of diabetic retinopathy are deposits composed primarily of lipids (fats) and proteinaceous material that leak out from damaged blood vessels in the retina. These deposits accumulate in the outer layers of the retina and appear as yellow or white patches on retinal images.

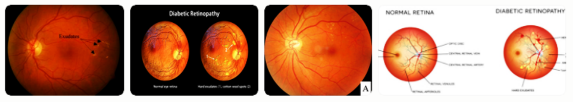

Figure 4: Exudates in Dibetic Retinopathic eye.

Image courtsey: https://pmc.ncbi.nlm.nih.gov/articles/PMC4568614/

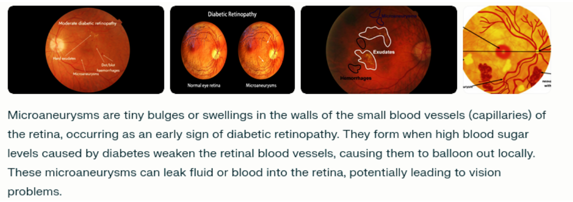

Figure 5: Microaneurysms in Diabetic Retinopathic eye

Image courtsey: https://onlinelibrary.wiley.com/doi/10.1155/2023/1305583

They are caused by the breakdown of the blood-retina barrier due to damaged retinal capillaries, allowing serum proteins and lipids to escape and settle in the retinal tissue. Exudates often indicate leakage and swelling in the retina, and their presence, especially near the macula (central retina), can lead to significant visual impairment.

1.6 Microaneurysms

Microaneurysms are tiny bulges or swellings in the walls of the small blood vessels (capillaries) of the retina, occurring as an early sign of diabetic retinopathy. They form when high blood sugar levels caused by diabetes weaken the retinal blood vessels, causing them to balloon out locally. These microaneurysms can leak fluid or blood into the retina, potentially leading to vision problems.

1.7 Haemorrhages

Haemorrhages in diabetic retinopathy refer to bleeding that occurs when the fragile blood vessels in the retina break and leak blood into the retinal tissue. They can appear in various forms depending on their location and size:

Dot and blot haemorrhages: Small, round haemorrhages found in the deeper layers of the retina, typically resulting from the rupture of microaneurysms.

Flame-shaped hemorrhages: Occur in the

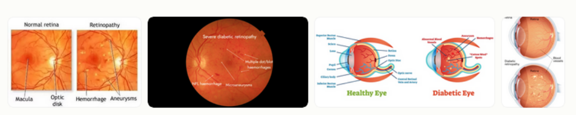

Figure 6: Haemorrhages in Diabetic Retinopathic eye.

Image courtsey: https://www.goodeyes.com/diabetic-retinopathy/

superficial layers of the retina, resembling flame shapes along the retinal nerve fiber layer

Haemorrhages are a sign of worsening diabetic retinal damage and are associated with increased vascular permeability and vessel wall weakening. In severe diabetic retinopathy, especially proliferative diabetic retinopathy, new fragile blood vessels may form and bleed into the vitreous humour, causing vitreous haemorrhage, which can lead to vision loss. The presence and extent of haemorrhages are important clinical indicators used in diagnosing and staging diabetic retinopathy.

1.8 Causes and Progression

High blood sugar from diabetes damages retinal blood vessels, causing increased permeability and blood vessel loss, which leads to retinal ischemia and swelling. The retina reacts by growing abnormal new vessels that are fragile and prone to bleeding. Vision loss occurs particularly when macular edema develops or in the advanced proliferative stage with retinal detachment risks.

1.9 Symptoms

Common symptoms include blurry or cloudy vision, floaters, dark spots or areas in the vision, and eventual vision loss. Early stages may be asymptomatic, making regular screening important for people with diabetes.

1.10 Diagnosis

Diagnosis involves dilated eye exams using ophthalmoscopy or retinal imaging to detect microaneurysms, haemorrhages, exudates, and new vessel growth. Regular screening is crucial since early diabetic retinopathy can be managed before significant vision loss.

1.11 Treatment and Management

Treatment aims to manage diabetes (maintaining blood sugar, blood pressure, and cholesterol control) and directly address eye damage.

Treatments include Anti-VEGF injections to reduce abnormal blood vessel growth and macular swelling, Steroid injections to reduce inflammation and swelling, Laser photocoagulation to seal leaking vessels and prevent bleeding, and Vitrectomy surgery for severe cases with vitreous haemorrhage or retinal detachment.

Strict glycemic control and routine eye exams are essential to prevent progression. Early detection and intervention can prevent 90% of severe vision loss cases. Overall, diabetic retinopathy requires a combination of systemic disease management and targeted ocular treatments for best outcomes

1.12 Role of Image Processing in Diabetic Retinopathy

Retinal imaging techniques, such as fundus photography and optical coherence tomography (OCT), provide high-resolution images of the retina.
Image processing enhances these images through noise reduction, contrast adjustments, and colour channel extraction to highlight key retinal structures.
Segmentation algorithms isolate anatomical features like blood vessels, the optic disc and detect lesions (microaneurysms, haemorrhages, exudates), which are signs of DR.

1.13 Automated Detection and Classification

Machine learning and deep learning models, especially convolutional neural networks, analyse processed images to automatically detect and classify

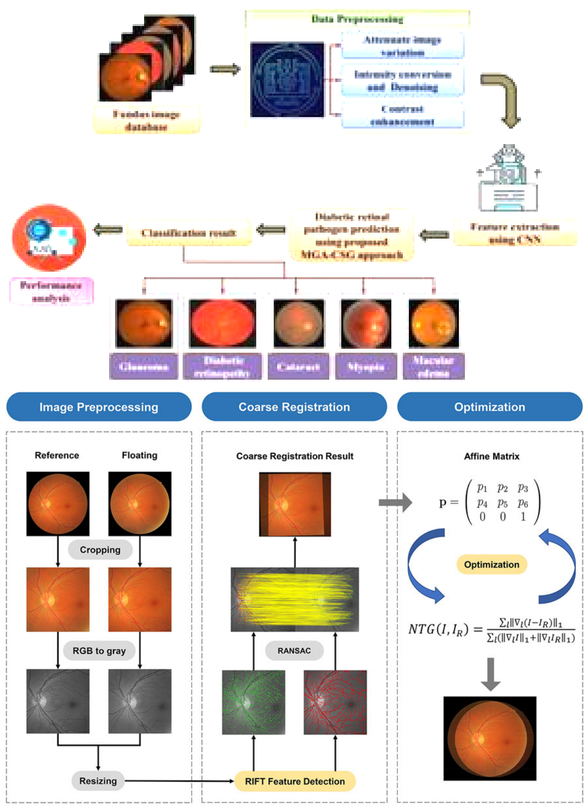

Figure 7: Illustration of the steps in DR -Image processing

Image courtsey: https://www.nature.com/articles/nrdp201612

DR stages.

These systems recognise subtle changes in retinal features associated with early and advanced DR that might be missed in manual examination.

Automated grading and early detection improve screening efficiency, allowing timely intervention to prevent severe vision loss.

1.14 Steps in Diabetic Retinopathy Image Processing

Preprocessing: Involves green channel extraction for higher contrast, histogram equalisation for brightness normalisation, resizing images, and denoising to enhance image clarity.

Feature Extraction: Quantifies biological features such as exudates, microaneurysms, haemorrhages, blood vessel area, and bifurcation points. These features are critical markers of retinopathy and its severity.

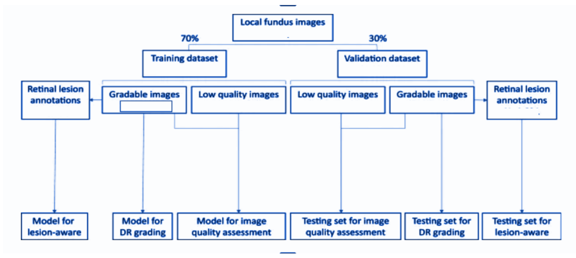

Figure 8: The local fundus image dataset was randomly divided into training and validation datasets. Seventy per cent of the total images in the dataset were used for training the image quality assessment sub-network. The lesion detection sub-network was trained using gradable images with annotations of retinal lesions. Then, the total gradable images in the training set were used to train the DR grading sub-network. All images in the local validation dataset were used to test the image quality sub-network. Finally, the gradable images labelled with retinal lesions were used to test the lesion detection sub-network. DR, diabetic retinopathy.

Image courtesy:https://www.nature.com/articles/s41467-021-23458-5

Segmentation: Localises regions of interest (optic disc, fovea, lesions) using thresholding and edge detection algorithms.

Classification: Machine learning and deep learning models (like CNNs, SVMs, and RSG-Net) automatically detect and classify stages of diabetic retinopathy by analysing extracted features.

Grading: Automated systems grade images into retinopathy stages (none, mild, moderate, severe, proliferative) to guide clinical decision-making.

1.15 DR grading pipeline end-to-end

Data Collection and Preparation

Acquire a large annotated dataset of retinal fundus images graded by DR severity levels (e.g., from public datasets like Kaggle's EyePACS).

Preprocess images by resizing (commonly to 224x224 pixels), normalisation, noise reduction, and enhancing contrast.

Utilise data augmentation techniques (such as rotation, flipping, and zooming) to enhance training data diversity and mitigate overfitting.

Model Architecture and Training

Choose a robust convolutional neural network

(CNN) architecture such as ResNet-50 or a specialized model like RSG-Net for feature extraction.
Apply transfer learning by fine-tuning a pretrained backbone on DR data to leverage prior knowledge.
Train the model for multiclass classification corresponding to DR grading stages (no DR, mild, moderate, severe, proliferative).
Use suitable loss functions (categorical cross-entropy for multiclass) and optimisers like Adam or SGD with tuned learning rates.
Implement regularisation techniques like dropout and batch normalisation to prevent overfitting.
Monitor validation metrics to apply early stopping or learning rate adjustments during training.

Model Evaluation and Validation

Evaluate performance using accuracy, AUC score, sensitivity, specificity, and confusion matrices.
Validate the model on separate test data to ensure generalisation.
Utilise robust metrics to evaluate misclassifications between adjacent DR grades and mitigate class imbalance.

Pipeline Automation and Deployment

Build a workflow to handle data input, preprocessing, prediction, and grading output. Optionally integrate a user interface for uploading images and displaying results.
Containerise the model using Docker and deploy on cloud platforms for scalability.
Implement continuous monitoring and retraining based on new data.
This pipeline encapsulates automated DR detection with high accuracy and practical usability, enabling timely diagnosis and referral for diabetic retinopathy patients.

References

https://www.nature.com/articles/s41467-021-23458-5

https://eyewiki.org/Age-Related_Macular_Degeneration

https://www.sciencedirect.com/org/science/article/pii/S1546221824007951

https://pmc.ncbi.nlm.nih.gov/articles/PMC10411652/

https://www.health.harvard.edu/newsletter_article

https://biomedpharmajournal.org/vol10no2/diabetic-retinal-fundus-images-preprocessing-and-feature-extraction-for-early-detection-of-diabetic-retinopathy/

About the Author

Geetha Paul is one of the directors of airis4D. She leads the Biosciences Division. Her research interests extends from Cell & Molecular Biology to Environmental Sciences, Odonatology, and Aquatic Biology.

About airis4D

Artificial Intelligence Research and Intelligent Systems (airis4D) is an AI and Bio-sciences Research Centre. The Centre aims to create new knowledge in the field of Space Science, Astronomy, Robotics, Agri Science, Industry, and Biodiversity to bring Progress and Plenitude to the People and the Planet.

Vision

Humanity is in the 4th Industrial Revolution era, which operates on a cyber-physical production system. Cutting-edge research and development in science and technology to create new knowledge and skills become the key to the new world economy. Most of the resources for this goal can be harnessed by integrating biological systems with intelligent computing systems offered by AI. The future survival of humans, animals, and the ecosystem depends on how efficiently the realities and resources are responsibly used for abundance and wellness. Artificial intelligence Research and Intelligent Systems pursue this vision and look for the best actions that ensure an abundant environment and ecosystem for the planet and the people.

Mission Statement

The 4D in airis4D represents the mission to Dream, Design, Develop, and Deploy Knowledge with the fire of commitment and dedication towards humanity and the ecosystem.

Dream

To promote the unlimited human potential to dream the impossible.

Design

To nurture the human capacity to articulate a dream and logically realise it.

Develop

To assist the talents to materialise a design into a product, a service, a knowledge that benefits the community and the planet.

Deploy

To realise and educate humanity that a knowledge that is not deployed makes no difference by its absence.

Campus

Situated in a lush green village campus in Thelliyoor, Kerala, India, airis4D was established under the auspicious of SEED Foundation (Susthiratha, Environment, Education Development Foundation) a not-for-profit company for promoting Education, Research. Engineering, Biology, Development, etc.

The whole campus is powered by Solar power and has a rain harvesting facility to provide sufficient water supply for up to three months of drought. The computing facility in the campus is accessible from anywhere through a dedicated optical fibre internet connectivity 24×7.

There is a freshwater stream that originates from the nearby hills and flows through the middle of the campus. The campus is a noted habitat for the biodiversity of tropical Fauna and Flora. airis4D carry out periodic and systematic water quality and species diversity surveys in the region to ensure its richness. It is our pride that the site has consistently been environment-friendly and rich in biodiversity. airis4D is also growing fruit plants that can feed birds and provide water bodies to survive the drought.