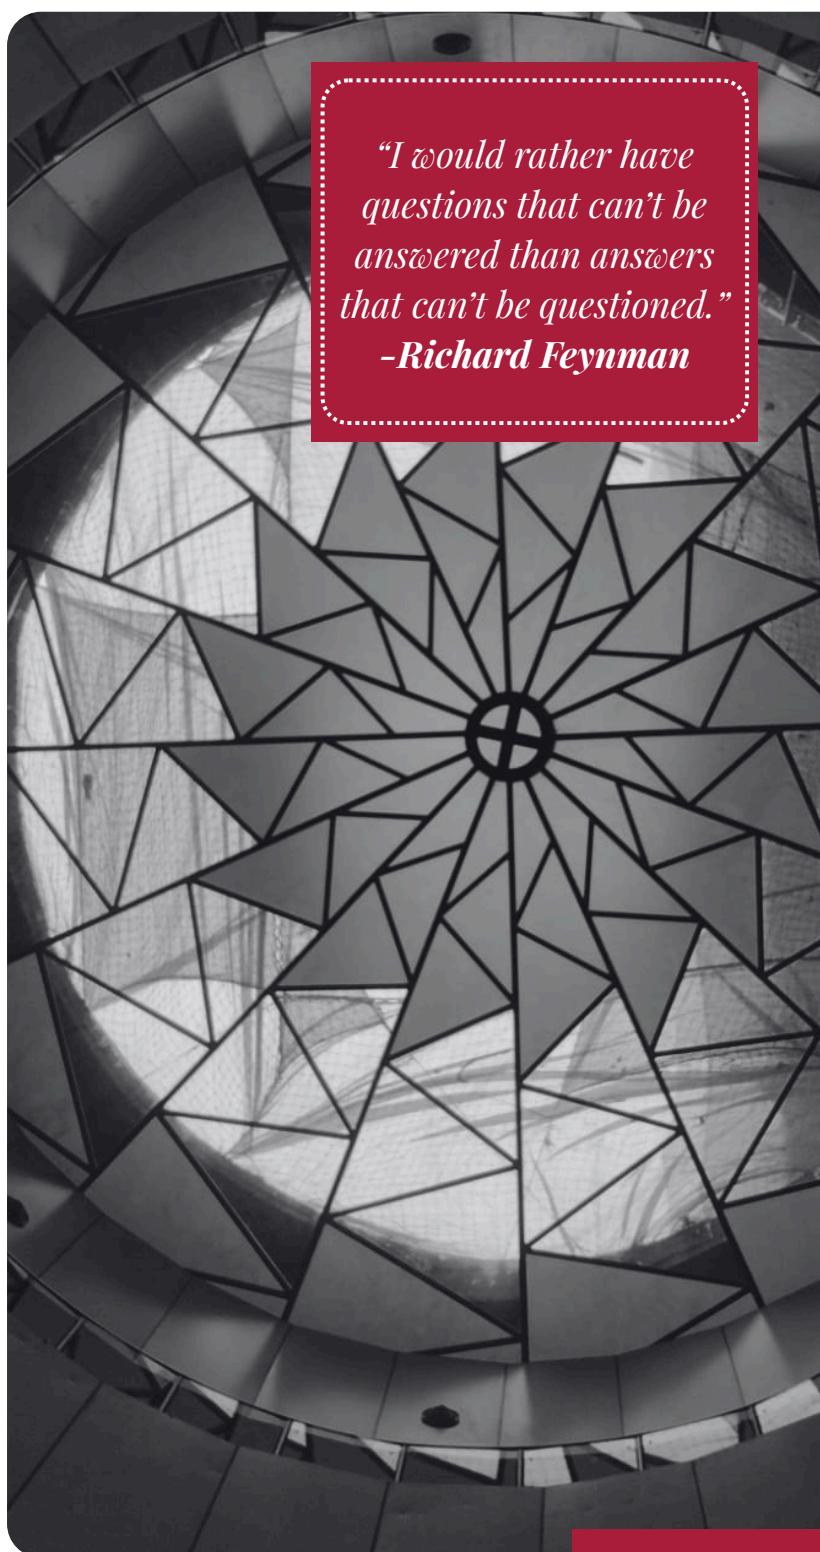


AI&DS INSIGHTS

**"CHEERS TO ONE YEAR OF AI&DS
INSIGHT! OUR MONTHLY AWARENESS
BULLETIN HAS OFFICIALLY
COMPLETED ITS FIRST SUCCESSFUL
YEAR."**

**we are now transitioning to a new
nomenclature for future issues."**

Page 02	Arealis Wins Google Cloud Agentic AI Hackathon 2025!
Page 03	Team noCodeBroCode Clinches Top Spot at INNOHACK 2.0 National Hackathon
Page 04	Publication
Page 05	Faculty Article: DocGenie – A Conversational AI Approach to PDF Exploration
Page 07	Student Article: Rain World's Emergent AI: A Milestone in Multi-Agent Intelligence and Its Future in Game AI Research



*"I would rather have
questions that can't be
answered than answers
that can't be questioned."*

-Richard Feynman

Achievement Spotlight

Arealis Wins Google Cloud Agentic AI Hackathon 2025!



Team Arealis has emerged as the champion of the **Google Cloud Agentic AI Hackathon 2025**, triumphing over 57,000+ innovators, 9,100+ solutions, and 700+ elite finalist teams from across India. Their groundbreaking innovation, **VELORA**, is a fully autonomous **City Data Management Platform** capable of identifying issues, tracing their origins, and deploying fixes instantly, that too without human intervention!

Built on the full power of Google Cloud's ecosystem, including Firebase, Pub/Sub, Vision AI, Vertex AI, Maps APIs, and Cloud Functions, VELORA stands out for its autonomous diagnostic agents and reactive multimodal architecture, demonstrating not just capability, but a leap toward true machine reasoning.

This victory underscores Arealis's vision of leading, not following, industry trends, and its commitment to redefining enterprise problem-solving with self-reliant AI systems. The team—Abhishek Shirasath, Avadhoot Ghewade, Darshana Ojha, and Neha Chaudhari—credits this success to months of innovation and tireless refinement.

This win is not just a milestone; it's a statement: the future of AI is autonomous, and Arealis is shaping it today.

Team noCodeBroCode Clinches Top Spot at INNOHACK 2.0 National Hackathon”

Thrilled to announce that Team noCodeBroCode from VIT Pune has secured 1st place at INNOHACK 2.0, a national-level hackathon hosted by Vishwakarma Institute of Technology in collaboration with BotBuddies, Microsoft Learn Student Club (MLSC), and the Department of Computer Engineering!

This achievement earned a trophy, certificates, ₹25,000 cash prize, and an internship opportunity. We're especially proud of Shankar Ramesh Jadhav (SY_B) and Samarth Sanjay Khadse (SY_D) from our department, alongside teammates Nirmit Hatti (CS-AI & ML) and Yash Wadhwani (CS), for this incredible accomplishment!

The sessions were led by expert representatives from Jamboree Institute – Ms. Rutuja



They developed an intelligent no-code data preprocessing platform that automates dataset cleaning, handles missing values, detects outliers, and performs advanced image processing (resizing, rotation, greyscaling, sharpening, brightness adjustments) to create model-ready data pipelines.

This solution compresses a task that would take four analysts four days into just five minutes for a non-technical user, making machine learning accessible without coding expertise or libraries like pandas, NumPy, OpenCV, or scikit-learn.

Publication

Dr. Ratna Nitin Patil, Aditya Pawar, Gitanjali Yadav, Atharva Nalawade – “Prostate Cancer Detection Using Large Vision Models on Apparent Diffusion Coefficient MRI”

Abstract:

Prostate cancer is among the most common cancers in men across the world, and early detection is considered crucial for enhanced health outcomes. This approach performs a computer-aided diagnosis of prostate cancer specifically on MRI-derived apparent diffusion coefficient (ADC) maps. Two proposed approaches are evaluated: the first uses feature extraction with Convolutional Neural Networks (CNNs), including GoogleNet, DenseNet-121, and DarkNet-19, and uses traditional Machine Learning Algorithms like SVM (including various kernels), Logistic Regression, and Random Forest for classification. The second approach leverages transfer learning on pre-trained CNNs, including GoogleNet and DenseNet121. Various metrics were evaluated to identify the best performance, including accuracy, precision, recall, F1 score, and ROC AUC. The results demonstrate that transfer learning and CNN-based feature extraction are indeed capable of achieving good results, which also points to a potentiality for deep learning in supporting clinical diagnostics. This comparative analysis demonstrates the feasibility of machine learning as a useful tool in the detection of various forms of cancers, the integration of which into radiological workflows can greatly enhance the reliability and quickness of diagnoses.

Published in: [2025 6th International Conference on Recent Advances in Information Technology \(RAIT\)](#)

Date of Conference: 06-08 March 2025

DOI: [10.1109/RAIT65068.2025.11089155](https://doi.org/10.1109/RAIT65068.2025.11089155)

Date Added to IEEE Xplore: 30 July 2025

Publisher: IEEE

► ISBN Information:

Conference Location: Dhanbad, India

Dr. Ratna Nitin Patil, Aditya Pawar, Gitanjali Yadav, Atharva Nalawade – “Prostate Cancer Detection Using Large Vision Models on Apparent Diffusion Coefficient MRI”

Abstract:

The increase in phishing websites reached 387% annually, and they became primary tools to conduct cyberattacks. This paper introduces a GraphSAGE-based detection framework that innovatively revolutionizes URL classification by modelling lexical, syntactic, semantic, and topological structures within four abstraction layers. The model achieved an accuracy rate of 99.73% with a latency of just 2.3 ms per URL and was tested against 1.6 million samples across 17 threat categories using Thor data. Significant innovations include hybrid NLP and graph-learning feature pipeline; adaptive neighbourhood sampling for robust inductive learning; and multi-head attention for interpretability in output. The framework consistently outperforms twelve state-of-the-art baselines concerning zero-day attacks (98.9% recall) and adversarial obfuscation (97.4% precision). When tested across three major ISPs, it was able to achieve an 83% reduction in false positives while processing more than 14,000 URLs per second on regular cloud infrastructure.

Published in: [2025 International Conference on Computing Technologies \(ICOCT\)](#)

Date of Conference: 13-14 June 2025

DOI: [10.1109/ICOCT64433.2025.11118684](https://doi.org/10.1109/ICOCT64433.2025.11118684)

Date Added to IEEE Xplore: 15 August 2025

Publisher: IEEE

► ISBN Information:

Conference Location: Bengaluru, India

Faculty Article

In the digital era, static PDFs often limit user interaction, making it difficult to search, summarize, or extract insights effectively. To overcome these challenges, DocGenie introduces an AI-powered solution that combines Retrieval-Augmented Generation (RAG) and Google Gemini Pro for dynamic PDF interaction.

DocGenie: A Conversational AI Approach to PDF Exploration

By integrating techniques like document chunking, FAISS-based embedding storage, and conversational AI, DocGenie allows users to query, summarize, and explore documents interactively. Libraries and research systems are shifting from traditional physical repositories to digital hubs. However, PDFs, despite being widely used, lack interactivity and searchability. DocGenie addresses this by providing a chat-based interface where users can upload documents, ask questions, and receive accurate, context-aware responses—bridging the gap between static content and intelligent interaction.

Research in question answering over PDFs, information retrieval with LLMs, and conversational agents has shaped DocGenie's foundation. Tools like LangChain and FAISS enable robust retrieval, while Gemini Pro enhances natural language understanding, summarization, and multilingual support.

The DocGenie system includes a user-friendly Streamlit-based conversational platform. PyPDF2 extracts and cleans text from uploaded PDFs, while Gemini Pro generates high-dimensional embeddings. FAISS enables efficient similarity search to quickly identify relevant document sections. The conversational chatbot then uses this information to deliver concise, natural answers. Additional features like adaptive summarization, real-time updates, and multi-document querying enhance overall usability, making the system suitable for a wide range of users.

This innovation has broad applications across sectors. In education, it can help students explore academic papers efficiently. In legal research, it enables quick retrieval of case references. In digital libraries, it transforms user experience by providing interactive access to large collections of documents. By combining advanced AI techniques with practical use cases, DocGenie demonstrates how intelligent systems can transform static documents into dynamic knowledge hubs.

Conclusion:

DocGenie illustrates the potential of AI-driven document interaction to revolutionize the way users engage with PDFs. By merging RAG, Gemini Pro, and FAISS, it delivers faster retrieval, better accuracy, and higher user satisfaction. It moves beyond passive reading to create an interactive, intelligent exploration process that benefits students, researchers, professionals, and organizations alike.

Future enhancements could include support for multiple file formats and expanded multilingual capabilities, but even in its current form, DocGenie stands out as a powerful tool for bridging the gap between static documents and dynamic knowledge discovery.



Prof. Palomi Gawli

Assistant Professor
AI & DS Department

Rain World's Emergent AI: A Milestone in Multi-Agent Intelligence and Its Future in Game AI Research

If VIT is analogous to a fruit basket and its vibrant community represents the fruits, then the gaming community would be an a berry. A berry, as we know, comes in many types, strawberry, blueberry, cranberry, gooseberry and many more.

A student looks at this basket doesn't immediately go, "Ah! Let me have just the berries instead of all the other delicious fruits." Similarly, As the same student in the beginner half on her learning journey, I have been questioning and fostering my curiosity in day to day life. Tasting different fruits and picking favorites is what I have been doing till now.

As we explore more, please take note that, as an aspiring Artificial Intelligence and Data Science engineer, I have written this from the perspective of a befitting student. I have written this article to marvel at and to appreciate the usage of Artificial Intelligence in Game Development by introducing a small part of the AI of the 2017 game Rain World.

Artificial Intelligence (AI) is often discussed in terms of healthcare, finance, or autonomous vehicles, while its role in interactive media is undervalued as "non-serious" or "entertainment-driven." Yet game environments, as I see them, are uniquely powerful laboratories for multi-agent systems, emergent intelligence, and adaptive ecosystems. The 2D-game, Rain World (2017) stands as example of this potential. Though its AI is not revolutionary, it still has blown my mind. Its agent-based design represents a milestone in applied AI, demonstrating how relatively simple rules can generate emergent complexity at scale.



Vidhi Rahangdale
SY_E

Student Article



Though I am not a hardcore gamer, from the few games I have played, most non-player characters (called NPCs) in games operate on a set rules of instructions such as finite-state machines or behavior trees (thank you, Data Structures and Finite Automata). These methods, while efficient, are very predictable. Taking an example, a guard patrols the same hallway, reacts in a simple flowchart way, and resets once the player leaves. Such AI is player-centric and exists to serve the player's narrative or challenge.



Rain World diverges from this tradition. Its creatures are autonomous agents, each pursuing survival goals such as hunger, safety, and territory. The player-controlled "slugcat" (A mixture of cat and a slug) is not the focus but one participant in a complex food chain. The ecosystem continues to grow off-screen, with creatures hunting, relocating, or dying independently of player observation. Encounters thus become non-deterministic: no two playthroughs unfold identically.

When viewed from and compared with current research in Artificial Intelligence, this system uses principles long studied in AI such as usage of agents following simple rules and making them interact to generate complexity (called agent-based modeling), agents forming groups or their tribes and working in sync and with certain understanding with one another (called swarm intelligence and a very clear example are the orange lizards in Rain World), and ecological/real-world environmental simulations, which still comes as a surprise for me because I am still star-struck by such wondrous use.

A black and white version of the Rain World logo. The title 'RAIN WORLD' is in large, white, sans-serif capital letters. To the left of the text is a small, dark silhouette of the slugcat character. The background is dark and textured.

The predators and prey, both categories remembers the location of player/other agents in the world for a short amount of time (called persistent state maintenance). They simply do not remember, they take/plan/predict decisions based on their viewpoint (referred to world models).

Rain World also demonstrates probabilistic reasoning via its predator tracking behaviors. This is so because, predators don't know exactly where you are but track your last known location, moving probabilistically toward where you might be. Some agents even exhibit adaptive tendencies like repeating successful strategies against the player which creates the impression of updating strategies based on outcomes (very similar to experience-driven policy adjustment in reinforcement learning) without the computational cost of deep reinforcement learning. (Honestly, I still haven't dived into reinforcement learning, just splashing in shallow waters for now, so it comes off as daunting to me.)

Revolutionary in Practice, Not in Research

From an AI research perspective, Rain World does not break new ground. Its design is rooted in decades-old concepts from artificial life and multi-agent systems. However, its significance lies in implementation at scale within a commercial game. By balancing emergent complexity, real-time constraints, and player experience, Rain World demonstrates the feasibility of persistent agent-based ecosystems in an entertainment product—a rare achievement.

Thus, while it is not revolutionary in academic AI, it is revolutionary in game AI practice, offering a practical demonstration of emergent intelligence that bridges theoretical research and interactive application.

How Rain World Could Evolve with Modern AI

Had Rain World been built on contemporary AI research rather than decade-old paradigms, its ecosystem might look even richer. With deep reinforcement learning (DRL), agents could optimize survival policies over time rather than relying on hand-crafted heuristics. Large language models (LLMs) could enable dynamic communication between agents or player-facing dialogue. Neuroevolutionary methods could allow species-level adaptation across generations, yielding evolving ecosystems rather than static populations.

Such approaches are already emerging in modern games. *Middle-Earth: Shadow of Mordor* introduced the Nemesis System, where orc adversaries remember and adapt to players. More recently, experimental titles and research prototypes employ transformer, based NPCs capable of free-form dialogue or neural network–driven behaviors that learn from player interaction.

Compared to these, *Rain World*'s system appears computationally modest, but its elegance demonstrates that emergence does not require heavyweight algorithms.

Conclusion

Rain World's AI is best understood not as a scientific revolution but as a milestone in applied AI design. It bridges the gap between academic concepts of multi-agent intelligence and their practical realization in interactive environments. Its success demonstrates that emergent complexity, persistence, and autonomy can yield richer worlds than scripted NPCs. Looking forward, combining *Rain World*'s systemic elegance with modern advances in DRL, LLMs, and adaptive simulation could revolutionize not only game AI but also broader applications in artificial intelligence research.

For Feedback
and Query,
Scan this:



OUR TEAM



**Dr Minal
Barhate**



**Vidhi Rahangdale
SY AI&DS-E**



**Jatin Shende
SY AI&DS -E**



**Aarya Gondikar
SY AIDS-C**