



Trinity College Dublin
Coláiste na Tríonóide, Baile Átha Cliath
The University of Dublin

AI Research In Action Research Programme

Funded PhD Positions in Artificial Intelligence

CKDelta have recently established a research programme in collaboration with the School of Computer Science & Statistics, Trinity College Dublin. The *AI Research In Action Research Programme* will establish a globally excellent research team which will assist in the application of leading edge AI techniques to a range of real-world problems within the utilities sector.

Trinity College Dublin, located in the heart of Dublin, is Ireland's number 1 university and ranked 87th internationally (QS Rankings 2024). The School of Computer Science and Statistics, Trinity College Dublin is the highest ranked School of Computer Science on the island of Ireland (QS World University Subject rankings 2024) and delivers a collegiate, friendly, and a world-class research-intensive centre for academic study and research excellence attracting scholars of global calibre from across the globe.

CKDelta is headquartered in Dublin, Ireland, and is part of CK Hutchison Holdings Limited, a multi-national conglomerate employing 300,000 people in over 50 countries around the world. CKDelta is an AI software business within the CK Hutchison Holdings Ltd group, leveraging data, expertise and intellectual property on a conglomerate scale. The company has a mission to empower customers with insights that transform decision-making and deliver impact, leveraging assets from across CK Hutchison group companies. CKDelta develop and offer innovative solutions across utilities, retail, financial services, transport and logistics solving digital transformation challenges to promote a safe and sustainable future.

A fully funded PhD scholarships are available providing a tax-free bursary of €25,000 per annum for 4 years together with fee payment. PhD(s) in the following area:

Graph-based Reasoning in Artificial Intelligence

Utility organisations maintain large, heterogeneous corpora of technical documentation — engineering specifications, maintenance records, inspection reports, operational procedures, and regulatory filings — that are tightly coupled to complex physical asset networks. Despite the widespread adoption of digital document management systems, retrieval remains largely keyword- or metadata-driven, and critically, documentation normally exists in isolation from the asset hierarchies, operational states, and maintenance histories it describes. This result in persistent operational problems, where engineers need a reliable standard procedures, and documentation that

reflects the current state of an asset, whether a maintenance event has introduced inconsistencies across related documents, or how a change in one part of the network propagates through its associated knowledge base.

While Large Language Models (LLMs) have demonstrated strong capability for natural language reasoning over enterprise knowledge, and Retrieval-Augmented Generation (RAG) has shown promise in grounding model outputs in document corpora, these approaches treat documents as independent, semantically similar artefacts rather than as nodes in a structured relational network. This is a fundamental mismatch with the operational reality of utility asset management, where the relationships between documents — and between documents and physical assets — are as important as the content of the documents themselves.

The core contribution of this research is:

- Design, implementation, and empirical evaluation of a domain-specific graph database schema and reasoning framework for utility asset-document networks;
- Investigating how unstructured and semi-structured technical documentation can be parsed and linked into a knowledge graph that explicitly represents asset hierarchies, maintenance event chains, document provenance, and cross-document dependencies.;
- Addressing the disconnection between maintenance history and asset specifications, while preserving documents currency and consistency.
- Development of a graph-grounded LLM reasoning architecture that enables natural language querying across this structure, with explicit provenance tracking and temporal validity;
- Exploration of the extension of the graph databases and schemas to incorporate a digital twin representation of asset state, including telemetry, lifecycle data, and real-time operational metadata. This extension raises distinct research questions about how dynamic state information can be integrated into a reasoning framework designed primarily around static documentation, and how temporal inconsistencies between document state and asset state can be detected and surfaced to users.

We are seeking a highly motivated PhD candidate in AI focusing on transformative AI solutions focused mainly in the utilities sector. This role blends cutting-edge research with real-world application, driving innovation in energy, water, and infrastructure systems through GenAI technologies.

Applications ought to have a 2.1 or 1 Degree in Computer Science or cognate discipline and/or an M.Sc in Computer Science/cognate discipline. Applications will be received until the positions are filled.

Prospective candidates would be expected to have an interest in several of the following areas: Artificial Intelligence, Multi-Agent Systems, Machine Learning, Data Analytics, IoT systems & Ubiquitous Sensing.

Interested parties should send via email a detailed academic Curriculum Vitae together with a letter of application indicating their interest in research,, to Professor Gregory O'Hare Gregory.OHare@tcd.ie
Please ensure the subject Heading is ARIA Application.