Lancaster University | DNS

Deep Reinforcement Learning for Network Pressure Management

Global Leakage Summit 2022

July 6, 2022

Ahmed Negm

Agenda

- Introductions
- Deep Reinforcement Learning
- Case Studies
- Closing
- Q&A

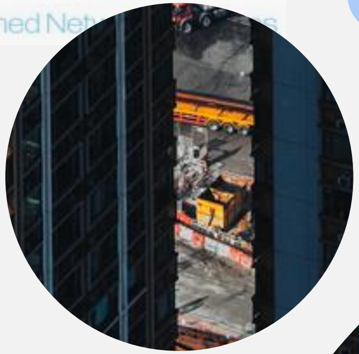

Introductions

Who am I?

A PhD candidate looking at novel AI technologies to develop the management of water distribution networks

Education:

- PhD Engineering (Currently)
- MEng Mechatronics (2020)

Previous Projects:

- Electric Formula Student – Lancaster Racing
- Assembly line optimisation – Oxley Development Ltd
- Conveyor belt optimisation – Butlers Farmhouse Cheese Ltd.

Achievements:

- Gold Employability Award – Lancaster University
- Engineering-in-business Fellowship Winner - Sainsbury Management Fellows

Who are we?

- Designing future proof networks enables network improvements to reduce water loss, increase efficiency, security and profitability.
- We create designs for new and existing mains and infrastructures.
- We help reduce leakage, burst events and overall maintenance costs.
- Offering a complete saving in production, carbon and environmental costs.

Supervisor: Craig Stanners (craig@dns-uk-ltd.co.uk)

Funded by CGE Eco-I NW and Partially Funded by Design Network Solutions Ltd – Contact: Georgia Faloone, cge-admin@lancaster.ac.uk, Project Administrative Support Officer, Centre for Global Eco-Innovation | Lancaster University , A17 Gordon Manley Building, Lancaster Environment Centre, Lancaster University, LA1 4YQ

http://www.globalecoinnovation.org/

Who are we?

- This research is being supported by Lancaster University's School of Engineering
- Access to state-of-the-art software and a high achieving research community
- The use of high-end computing capabilities

Supervisors: Prof. George Aggidis (g.aggidis@lancaster.ac.uk), Dr. Xiandong Ma (xiandong.ma@lancasster.ac.uk)

Funded by CGE Eco-I NW and Partially Funded by Design Network Solutions Ltd – Contact: Georgia Faloone, cge-admin@lancaster.ac.uk, Project Administrative Support Officer, Centre for Global Eco-Innovation | Lancaster University , A17 Gordon Manley Building, Lancaster Environment Centre, Lancaster University, LA1 4YQ

What are we doing?

The goal of this research project is to combine an understanding of the UK's water network system and water demand stresses, with an aim to develop a diagnostic tool to support a reduction in water waste and resulting carbon emissions.

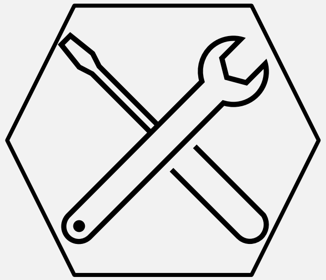

Technical

- Leakage and burst events
- Design and construction
- Aging infrastructure
- Varying topography

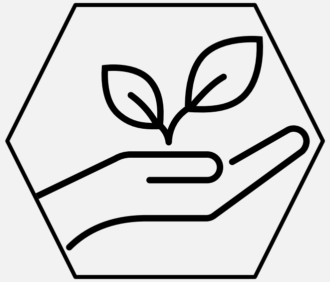

Environmental

- Water availability
- Carbon and GHG emissions
- Temperature fluctuations
- Droughts and floods

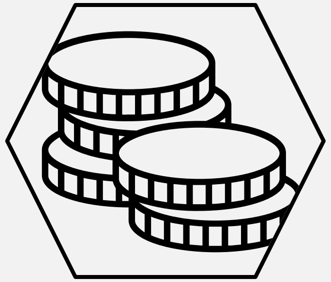

Economical

- Cost of investment
- Inadequate water networks
- Lack of cost optimisation
- Foresight

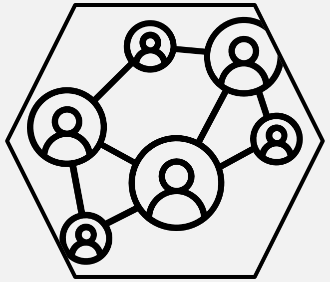

Socio-political

- Lack of political will
- Landmarks and legislation
- Social unrest during repairs
- Supply interruptions

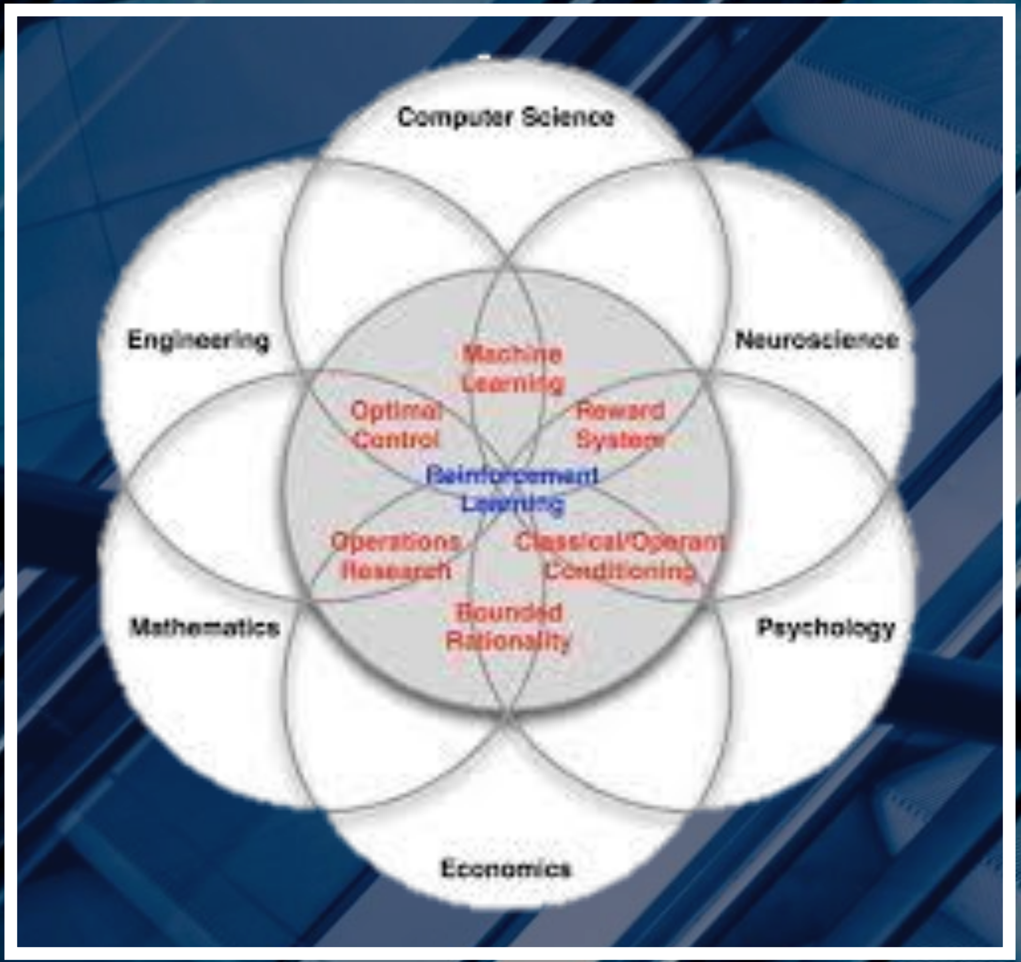
Computer Science
Engineering
Neuroscience
Mathematics
Psychology
Economics
Machine Learning
Optimal Control
Reward System
Reinforcement Learning
Operations Research
Classical/Operant Conditioning
Bounded Rationality

What is deep reinforcement learning?

- Trial and error learning
- Based on interaction and experience
- Learning how to map situations to actions to maximise a reward signal
- Helps us handles unstructured complex environments
- Provides a formalism for (human) behaviour

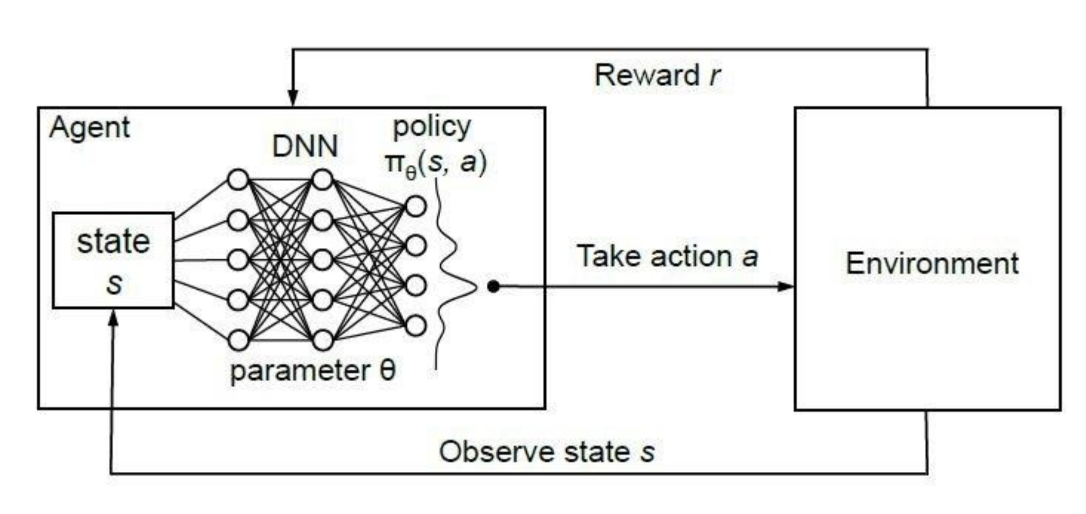

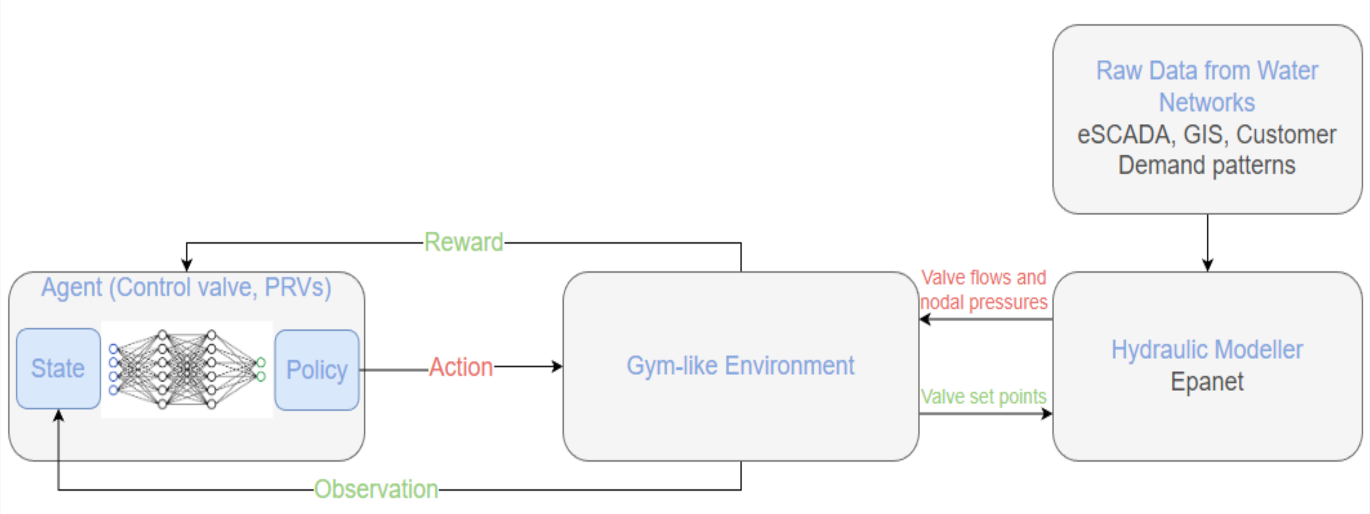

What is deep reinforcement learning?

To understand deep reinforcement learning it is important to understand its terminology and redefine our water network in the scheme of a Markov Decision Process (MDP).

Agent

The physical component that interacts with the environment through action, taking observations and receiving rewards.

Action

The steps the agent can take to affect the environment hence changing its observation Can be continuous or discrete actions

Environment

Everything that exists outside of the agent. In the most general sense – the universe.

Observation

Pieces of information that is relevant to the agent. Helps the agent identify the next action

Reward

The purpose of this numerical value is to grade the agent’s behaviour. Critical feedback to improve our agent.

Policy

A policy defines how the agent behaves. It takes an input state and outputs an action. Can be stochastic or deterministic.

What is deep reinforcement learning?

To understand deep reinforcement learning it is important to understand its terminology and redefine our water network in the scheme of a Markov Decision Process (MDP).

Agent

Valves, Pumps, Resevoirs

Observation

Nodal pressures, Pipe flows, Energy Consumptions, Leakage values

Action

Open/close valves, Change pump speeds

Reward

Less pressure violations, Lower energy consumptions, Less leakage.

Environment

DMAs, Zonal networks

Policy

Advanced pressure management, Advanced pump control.

Deep Reinforcement Learning Success

Deep reinforcement learning is the true hope of artificial intelligence because of its immense potential ad its ability to adapt to multiple foreign environments and map the situations into beneficial actions. Below are a couple of real-life applications:

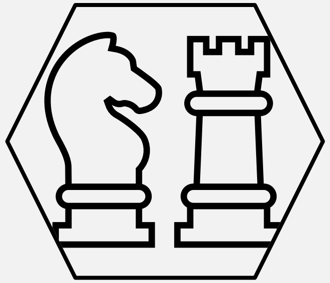

Gaming

- Go (AlphaGo Zero)
- Backgammon (TD-Gammon)
- Dota2 (Open AI 4)
- Atari (DQN)

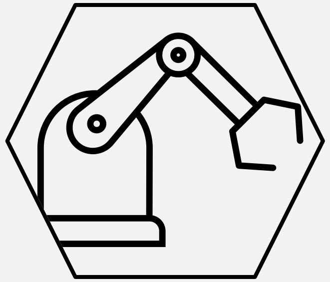

Robotics

- Grasping (QT-Opt)
- Learning dexterity
- Locomotion Behaviour
- Autonomous driving

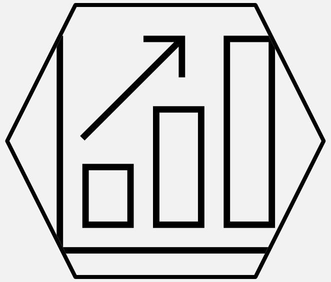

Finance

- Stock price prediction
- Financial trades
- Trading bots
- Risk optimisation

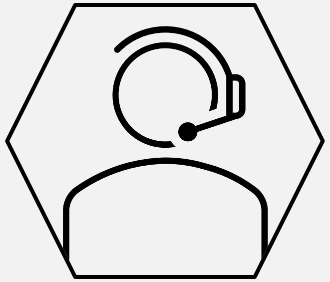

NLP

- Text summarisation
- Question answering
- Machine translation
- Chatbots

Deep Reinforcement Learning Challenges

Observation is a result of behaviour

An agent that performs bad actions will receive bad feedback (observations) from the environment that will not help them get closer to the desired reward. This could place a stubborn agent in a spiral under the impression that there is no path for positive reward.

Exploitation/exploration dilemma

On the other hand, an agent can exploit an action that gives them a constant positive reward instead of actively exploring the environment for better route or vice versa.

Delayed Reward

It can be often unclear how one action can have future consequences so understanding and forecasting future value is essential.

Case Studies

Real-Time Optimization of Pumps in Water Distribution Systems

- The proposed technique relies entirely on live measurement data in the decision-making process. This property makes real-time optimal control of the pumps possible in a smart water network.
- The underlying algorithm is a duelling deep Q network (DDQN).
- Tested on Anytown and D-town test networks
- Outperforms conventional techniques in search speed

Contribution: Real-time pump optimisation of pump speeds solely through measurement data.

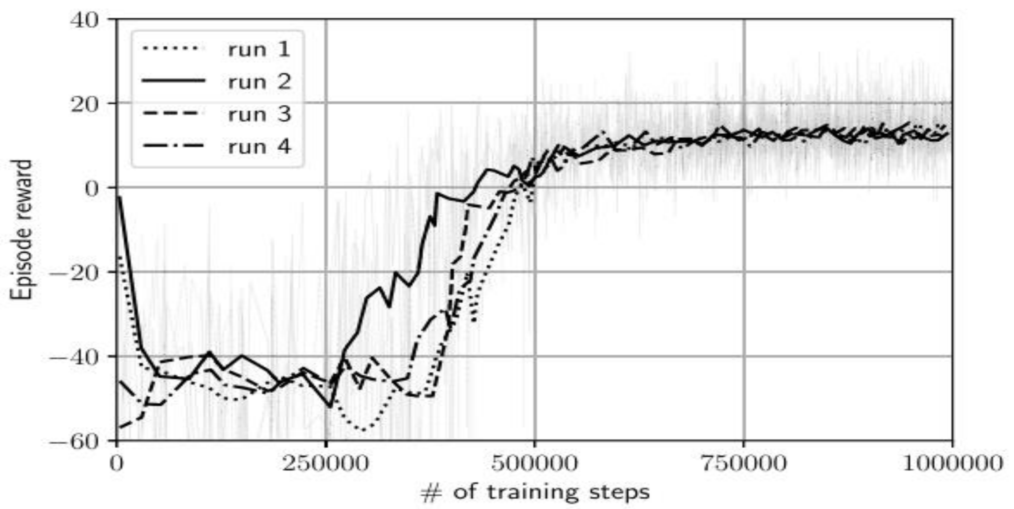

Run	Anytown-mod	D-Town-mod	D-Town-mod with one-shot RT
1	0.991	0.987	1.04
2	0.992	0.987	1.04
3	0.991	0.979	1.04
4	0.991	0.988	1.04

Zone scheduling optimization of pumps

- This article studies the pump scheduling optimization problem in water distribution networks (WDNs).
- The underlying algorithm is a Knowledge-Assisted Proximal Policy Optimisation
- Tested on Anytown network
- Compared to Nedler Mead and DDQN optimisations

Contribution: Real-time optimisation of pump scheduling outperforms conventional techniques

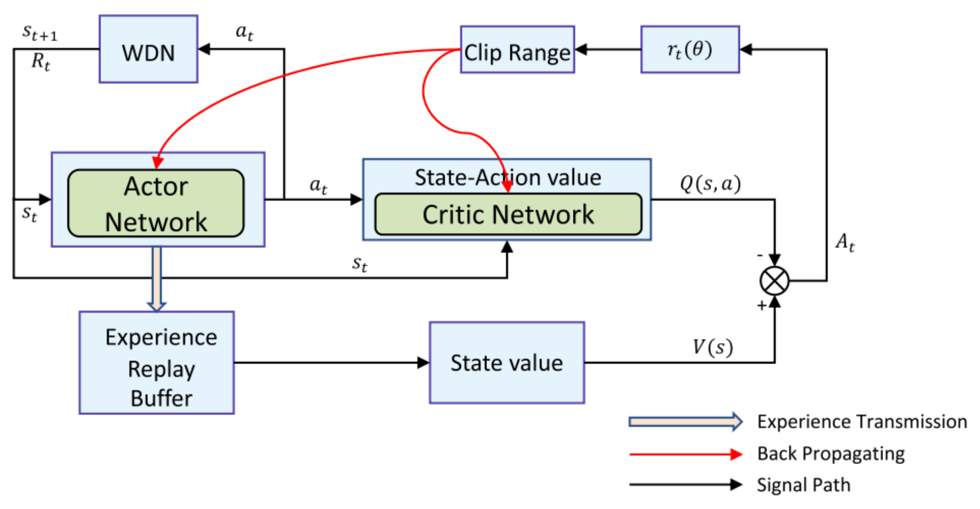

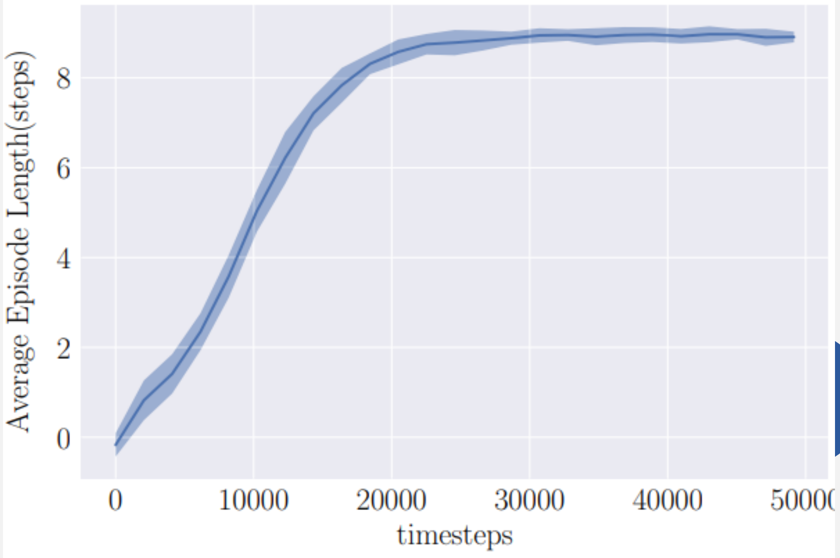

Pressure Valve Optimisation

- This experiment optimises the PRV set points to decrease the average pressure and Ofwat pressure violations.
- The underlying algorithm is Q-learning
- Tested on D town network

Contribution: Reinforcement learning to optimise pressure reducing valves

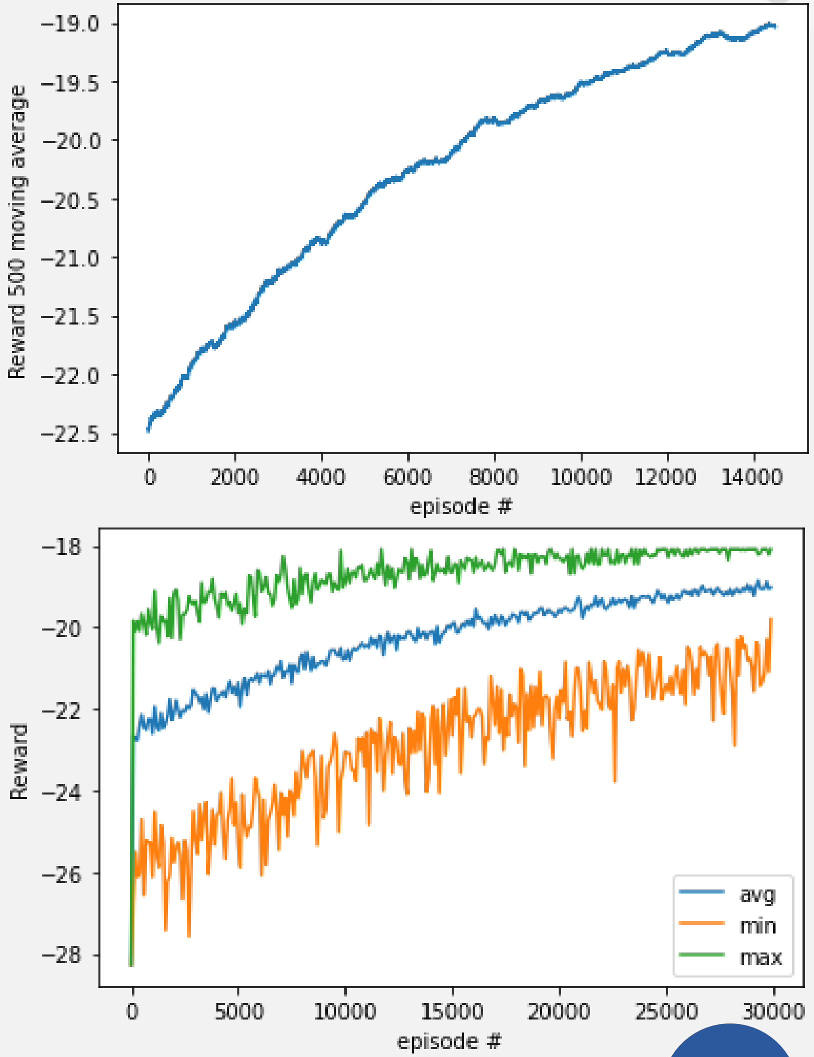

Closing

Conclusion

The use of this technology in the water sector is an avenue of great potential and could lead to:

- Increased water preservation
- Accounted for the rising customer demands
- Provided a framework for battling the persisting leakage issues
- Promoted net-zero carbon emissions

Ahmed Negm

a.negm@lancaster.ac.uk

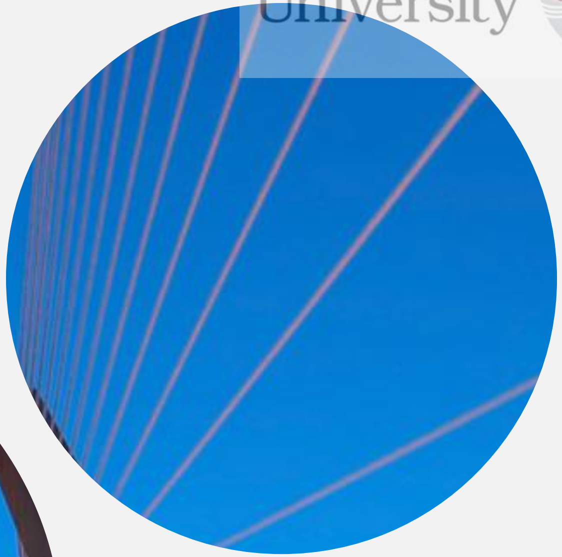

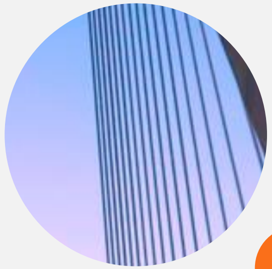

Lancaster University | DNS

Thank You

July 6, 2022

Ahmed Negm