Role Definition
Job Title: Machine Learning Researcher (ARD)
Reporting to: Head of Applied Research and Development

The Applied Research + Development team produces cutting-edge technologies, including development and implementation of AI and machine learning algorithms, to support Architecture, Engineering, Construction and Operations (AECO) for the built environment. We take these technologies through all stages of the innovation lifecycle, from R+D prototypes through to robust production applications. We are looking for Machine Learning Researchers with an interest in spatial, interactive and generative design, computer vision and business insights to support this mission and deliver industry-leading solutions.

In this position you will be a member of the Applied R+D team. Your involvement will include strategic R&D around developing, implementing and evaluating new models, agents and techniques, as well as deploying standard models and methods towards prediction and automation - including building the relevant software to support wider deployment of these derivative applications. The role will include researching the latest models in Machine Learning and implementing proofs of concept for problem particular to the AEC industry, using standard frameworks such as PyTorch and JAX. Once a model is fully tested, the role will involve deployment to test and production on premises or the cloud. The role also involves identifying and leveraging opportunities in ML research for both architectural design innovation and operational business applications, training and grounding image and language-based ML models on proprietary data and investigating workflow automation through AI agents.

You will work alongside data scientist, front-end and back-end developers as well as domain specialists like material scientists and business analysts. You will contribute to applications that work across the practice.

Responsibilities

- Apply machine learning methods (including but not limited to Deep Learning, Reinforcement Learning, Probabilistic Programming, Gaussian Processes, Natural Language Processing, optimisation, computer vision and graph neural networks) for generative, automation and forecasting purposes.
- Research, develop and evaluate state-of-the-art ML models to solve complicated problems using standard frameworks like PyTorch and JAX.
- Take part in deploying and monitoring applications that make use of the developed ML models on prem and through cloud computing.
- Work with domain experts across the company to translate their business problems into Machine Learning solutions.
- Explain and defend analytical techniques and visualizations of relevant metrics to key stakeholders.
- Create trend reports that track the impact and effectiveness of processes and solutions over time.

Qualities and Skills required

Essential

- Master's degree in Computer Science, Machine Learning, Mathematics or equivalent experience.

- Deep knowledge and understanding of Deep Learning architectures.
- Ability to implement latest research ML papers.
- Strong programming skills, with a preference for Python.
- Excellent experience developing and training deep learning models using PyTorch and /or JAX.
- Excellent experience with Sci-Kit Learn, Pandas and Numpy.
- Experience with CUDA or Triton.
- Experience with Computer Vision, Natural Language Processing and Graph Neural Networks.
- Experience with a wide range of generative modelling techniques, including diffusion-based models, flow matching, and Latent Consistency Models for image synthesis, as well as large-scale transformer architectures for language modelling, with a focus on underlying algorithms and emerging industry trends.
- Knowledge of DataOps and MLOps concepts and tools (e.g. MLFlow).
- Experience with Huggingface libraries like Transformers, Diffusers and Accelerate.
- Effective communication skills, with the ability to deliver compelling presentations on how Machine Learning models can help the business.
- Effective listening skills in order to understand the requirements of the business.
- Able to manage sensitive and confidential information.
- A keen eye for detail.
- Excellent verbal and written communication skills
- Good interpersonal skills

Desirable

- Programming skills with C++
- Research track record through publications.
- Interest in Computer Graphics.
- Knowledge of Bayesian Machine Learning and Probabilistic Programming (PyMC, Pyro) – and Gaussian Processes in particular.
- A track record of using Deep Learning to solve complex problems.
- Familiarity with Python UI frameworks (e.g. Gradio or Streamlit).
- Experience utilizing multi-GPU nodes for training ML models.