

# Fourth Workshop on Multimodal AI

11-12 June 2026

**Social Media:** #MultimodalAI26, #UKOMAIN

**LinkedIn:** @UK Open Multimodal AI Network (UKOMAIN)

**Venue:** UCL East Campus, One Pool Street, E20 2AF ([direction](#))

## Programme

Workshop Day 1 – 11 June 2026	
Time	Event
09:30 - 10:00	Registration, coffee/tea, and biscuits
10:00 - 10:15	Welcome address: Amanda Brock, CEO at OpenUK
10:15 - 10:30	Opening talk: Haiping Lu, Director, UK Open Multimodal AI Network
<b>10:30 - 12:10</b>	<b>Session 1</b> - Chairperson: Nicola Morley, University of Sheffield
10:30 - 11:10	Keynote 1: Aron Walsh, CSO at CuspAI and Professor in Materials Design at Imperial College London <b>Materials on Demand</b>
11:10 - 11:45	Community talks 1 (Short talks and pitches)
11:45 - 12:10	Group photos
12:10 - 13:20	Lunch and posters
<b>13:20 - 14:50</b>	<b>Session 2</b> - Chairperson: Dezong Zhao, University of Glasgow
13:20 - 14:00	Keynote 2: Julia Hirschberg, Professor of Computer Science at Columbia University <b>Radicalization and De-Radicalization in Social Media</b>
14:00 - 14:30	Community talks 2 (Short talks)
14:30 - 14:50	Hackathon presentation
14:50 - 15:50	Coffee/tea, biscuits, and posters
<b>15:50 - 17:30</b>	<b>Session 3</b> - Chairperson: Richard Gilham, Bristol Centre for Supercomputing
15:50 - 16:30	Keynote 3: Chris Barnes, Head of Science of AI at National Physical Laboratory and Professor of Systems and Synthetic Biology at UCL <b>An Overview of AI at NPL</b>
16:30 - 17:20	Research Panel: From Models to Evidence: Making Multimodal AI Deployable (Panel Discussion 1)
17:20 - 17:30	Hackathon winner announcement
Networking Reception	
17:30 - 18:00	Arrival and networking
18:00 - 20:00	Networking reception: UKOMAIN connections corners
	Informal networking around three themes: <ul style="list-style-type: none"><li>• Funding, Proposals &amp; Collaborators</li><li>• Evidence, Evaluation &amp; Responsible Deployment</li><li>• Careers, Mentoring &amp; Inclusive Leadership</li></ul>

## Workshop Day 2 – 12 June 2026

Time	Event
08:30 - 09:00	Registration, coffee/tea, and biscuits
<b>09:00 - 10:25</b>	<b>Session 4</b> - Chairperson: Nataliya Tkachenko, Lloyds Banking Group
09:00 - 09:20	Invited talk: Andy Lawrence, Head of Engineering at EPSRC <b>Research and Innovation in AI</b>
09:20 - 10:00	Keynote 4: Henry Lee, CEO at Cultivarium <b>Scaling Biological Expertise</b>
10:00 - 10:25	OMAIB R1 showcases <ul style="list-style-type: none"><li>● CRITICAL-MM: Cross-Modal Reasoning with Integrated Clinical Assessment for Large Multimodal Models</li><li>● Towards Carbon-Neutral Living: An Open Dataset for Smart Social Housing</li><li>● SONAIR: Sim2Real Operational Benchmark for AI Robotics</li><li>● Neural Architecture Search Benchmark for Multimodal Self-Driving Car Test Data</li></ul>
10:25 - 11:25	Coffee/tea, biscuits, and posters
<b>11:25 - 13:00</b>	<b>Session 5</b> - Chairperson: Fei He, Coventry University
11:25 - 12:05	Keynote 5: Zoe Kourtzi, Professor of Cognitive Computational Neuroscience at University of Cambridge and CSO at Prodromic <b>Multimodal AI for Early Dementia Prediction: From Cloud to Clinic</b>
12:05 - 12:15	Invited talk: Jake Bowden, Bioinformatician at Twig Bio <b>Canopy: A Heterograph Foundation Model for Metabolic Engineering</b>
12:15 - 12:25	Invited talk: Farah Shamout, Assistant Professor of Computer Engineering at NYU Abu Dhabi <b>Multimodal AI for Patient-Centered Healthcare Technologies</b>
12:25 - 13:00	Community talks 3 (short talks and pitches)
13:00 - 14:00	Lunch and posters
<b>14:00 - 16:00</b>	<b>Session 6</b> - Chairperson: Peter Charlton, UKOMAIN Executive Board
14:00 - 14:40	Keynote 6: Tom Pollard, Research Scientist at Massachusetts Institute of Technology and Technical Director of PhysioNet <b>Data for Health AI in a Time of Change</b>
14:40 - 15:30	Career Panel: Building a Sustainable Research Career in Multimodal AI (Panel Discussion 2)
15:30 - 15:50	Presentation winner announcement
15:50 - 16:00	Closing remarks
16:00 - 17:00	Post-workshop networking

# Community talks 1 (Day 1 | 11:10 – 11:45)

## Short Talks (5 mins each)

Name	Title
Dongsheng Wen	Kagome Intermetallic Database Assembled by Topology + Geometry Approach
Kewei Zhu	ReadMOF: Leveraging Systematic Chemical Nomenclature for Machine Learning Applications in Metal-Organic Frameworks

## Pitches (1 min each)

Name	Title
Xiang Li	From Lab to Market: Building Bridges for Multimodal AI Commercialisation
Francesco Murtas	Temporal Dynamics of Brain-Body Responses to Music: A Multimodal Analysis
Haolin Wang	Benchmarking Bandgap Prediction in Semiconductors under Experimental and Realistic Evaluation Settings
Zixuan Huang	DiffMG-Trans: Multimodal Safety-Critical Scene Generation
Arianna De Vecchi	A Multimodal Dataset for Cognitive Load, Stress, and Behaviour Analysis in a Collaborative Escape Room
Yanni Sun	Multimodal Analysis of 底层 ('bottom social stratum') Self-Representations on Kuaishou: A Scalable Video-Based Approach

## Community talks 2 (Day 1 | 14:00 – 14:30)

### Short Talks (5 mins each)

Name	Title
Cristina Palmero	UDIVA-HHOI: A Multimodal Dataset for Context-Aware Multi-Actor Interaction Understanding
Jonathan Carter	Next-Token Prediction Learns Generalisable Representations of Sleep Physiology
Baraa Al Jorf	How Effectively Can LLM Agents Synthesize Multimodal Clinical Data for Risk Prediction?
Misbah farooq	CMAT: Cross-Modal Attention-Weighted Triplet Loss for Multimodal Emotion Recognition
Sakib Ahammed	Towards Fair and Semantic Coverage-Aware Learning for Underrepresented Visual Concepts
Shrankhla Pandey	Speech Acoustics as Non-Invasive Biomarkers for Psychotic Disorder Severity Classification

## Community talks 3 (Day 2 | 12:25 – 13:00)


### Short Talks (5 mins each)

Name	Title
Jichun Li	An Effective Fragment-Based Dual Conditional Diffusion Framework for Molecular Generation
Longdan Hao	Myo-Guide: A Bidirectional Framework Bridging Multimodal AI and Rare Neuromuscular Disease Diagnostics
Giuseppe Tripodi	Uncertainty-Aware Multimodal Cancer AI: A Conformal Prediction Study across Pathology and Genomics

### Pitches (1 min each)

Name	Title
Aman Jat	Assessing Equity Gaps in Artificial Intelligence Tools for Mental Health Diagnosis and Triage: A CPRD Cohort Study of 7.19 Million Adults with Multiple Long-Term Conditions
Hazrat Ali	Mask-Conditioned Synthetic Prostate MRI Volumes using a Denoising Diffusion Probabilistic Model
Jiang Chang	Semantic-Guided Cross-Modal Interaction for Multimodal Depression Detection
Omotayo Omoyemi	Multimodal AI for Accessible Communication and Digital Health: The AURA Framework
Ricardo Valencia	Adapting Multi-Modal Variational Autoencoders for Microbiome Data
Tadiyos Hailemichael Mamo	Demographic-Aware Adversarial Transformer for Heart Failure Risk Prediction from ECG and Clinical Notes
Jingyu Hu	Probing and Fusion of Foundation Model Representations for Multimodal Cancer Analysis
Mohammad Suvon	Multimodal AI Improves Low-Cost Cardiovascular Hemodynamic Assessment from Chest X-Ray and Electrocardiography
Jiayang Zhang	Computational Design of Grafted Motifs on Protein Scaffold

# Keynotes



Keynote 1: **Aron Walsh**, CSO at CuspAI and Professor in Materials Design at Imperial College London

Title: **Materials on Demand**

Date and Time: 11 June 2026, 10:30 – 11:10

## Abstract

The inverse design problem (given a target property or function, identify the optimal material) represents one of the central challenges in materials science. The landscape of materials theory and simulation is addressing this problem through the integration of new techniques and tools from the artificial intelligence (AI) community. Progress in hardware, including classical supercomputers and emerging quantum computers, alongside software advancements incorporating advanced algorithms and statistical machine learning models, is expanding what is now possible. A particular opportunity lies in multimodal AI, which can bridge heterogeneous data streams spanning computation, synthesis, and characterisation to build richer and more transferable representations of materials. Recent developments, such as large language models and generative diffusion techniques, are unlocking application areas ranging from multimodal characterisation to integration with self-driving laboratories. I will survey the evolution of data-driven approaches to materials on demand, highlighting their potential to expedite the identification of compounds essential for the next generation of clean energy technologies. The talk will close with reflections on the translation of academic research into emerging industry, with particular attention to the growing AI-for-materials ecosystem.



Keynote 2: **Julia Hirschberg**, Professor of Computer Science at Columbia University

Title: **Radicalization and De-Radicalization in Social Media**

Date and Time: 11 June 2026, 13:20 – 14:00

## Abstract

Our prior work on radicalization focuses on increasing our understanding of online radicalization efforts from analyzing right- and left-leaning group videos in social media including groups from outside the USA. We studied what viewers found the videos appealing, what techniques were employed in producing radical content, and how we could computationally identify these techniques. More recently, we have worked to identify potentially effective de-radicalizing techniques and worked to build a unified multimodal processing pipeline toolkit to help us create conversational systems on social media that are useful in de-radicalizing viewers.



Keynote 3: **Chris Barnes**, Head of Science of AI at National Physical Laboratory and Professor of Systems and Synthetic Biology at UCL

Title: **An Overview of AI at NPL**

Date and Time: 11 June 2026, 15:50 – 16:30

### Abstract

The National Physical Laboratory (NPL) is the UK's national metrology institute. It develops and maintains the country's primary measurement standards and provides the measurement science, calibration, testing, and scientific expertise that underpin industry, innovation, and public services—helping ensure measurements are accurate, comparable, and trustworthy. In this talk, Chris Barnes (Head of Science for AI, NPL) will give an overview of NPL's AI activity and how it supports the UK's ambitions for trustworthy and deployable AI across industry.

The talk is structured into three parts covering the main aspects of AI@NPL. First, Trustworthy AI: how measurement thinking translates into AI practice, covering concepts such as uncertainty, robustness, interpretability, and data quality, and how these properties can be characterised, tested, and compared in a repeatable way. Second, AI for Metrology: examples of how modern machine learning can be used to enhance the scientific process. Third, AI assurance: the emerging methods, standards, and evaluation approaches needed to build justified confidence in AI systems for critical applications.



Keynote 4: **Henry Lee**, CEO at Cultivarium

Title: **Scaling Biological Expertise**

Date and Time: 12 June 2026, 09:20 – 10:00

### Abstract

Cultivarium builds scientific tools that turn biological discovery into real-world capability. Remarkable genetic capabilities and chemistries useful for beneficial biotechnologies can be found in natural biological systems, yet nearly all of them sit outside today's laboratory. To make the vast number of non-model organisms available for reproducible laboratory study, we have built a technology platform for domesticating them, compressing the typical time frame from decades to days. In this talk, I will trace our journey from deterministic software to AI tooling for accelerating scientific discovery, biological engineering, and the real-world deployment of knowledge systems. Hermes, our large language model (LLM) harness, helps scientists navigate existing knowledge, generate hypotheses, and run analyses across iterative discovery campaigns. PRISM, our vision-language model (VLM) harness, captures experimental work as audio and video feeds to generate traceable documentation and surface tacit knowledge that is otherwise earned one scientist at a time. These multimodal AI tools form the basis for a new software stack for embedding and distributing expertise, accelerating progress in frontier biology research programs.



Keynote 5: **Zoe Kourtzi**, Professor of Cognitive Computational Neuroscience at University of Cambridge and CSO at Prodromic

Title: **Multimodal AI for Early Dementia Prediction: From Cloud to Clinic**

Date and Time: 12 June 2026, 11:25 – 12:05

### Abstract

Early prediction of brain (e.g. neurodegenerative) disorders is key for clinical management and patient outcomes. Predicting whether individuals with mild cognitive impairment or people without symptoms will decline or remain stable is impeded by patient heterogeneity due to comorbidities, lifestyle and disease severity. Despite the importance of early diagnosis of dementia for prognosis and personalised interventions, we still lack robust tools for predicting individual progression. We propose a novel clinical AI predictive prognostic modelling approach that mines multimodal data to derive an individualised prognostic marker of cognitive decline at early stages of dementia or before symptoms occur. We validate our approach against routinely collected real-world patient data from memory clinics over time, showing that our clinical AI marker is more sensitive than the standard of care (cognitive tests, MRI scans). Our clinical AI approach has strong potential to facilitate effective patient stratification into clinical pathways and trials, reducing patient misdiagnosis and enhancing trial efficiency with important implications for clinical translation and drug discovery.



Keynote 6: **Tom Pollard**, Research Scientist at Massachusetts Institute of Technology and Technical Director of PhysioNet

Title: **Data for Health AI in a Time of Change**

Date and Time: 12 June 2026, 13:40 – 14:20

### Abstract

We will explore how the health AI community is moving from an era focused mainly on model building to one that must confront deeper questions about data quality, access, governance, and representation. Using examples from recent controversies in medical AI and from the development of open research resources such as PhysioNet and MIMIC, the talk will argue that progress now depends on creating data ecosystems that are more trustworthy, more reusable, and more faithful to the realities of clinical care.

# Posters

## Day 1

Name	Title
Dongsheng Wen	Kagome Intermetallic Database Assembled by Topology + Geometry Approach
Kewei Zhu	ReadMOF: Leveraging Systematic Chemical Nomenclature for Machine Learning Applications in Metal-Organic Frameworks
Haoming Xiang	Data-Driven Resilience in Social Housing: Enhancing Retrofit Strategies via Physics-Informed AI
Francesco Murtas	Temporal Dynamics of Brain–Body Responses to Music: A Multimodal Analysis
Haolin Wang	Benchmarking Bandgap Prediction in Semiconductors under Experimental and Realistic Evaluation Settings
Zixuan Huang	DiffMG-Trans: Multimodal Safety-Critical Scene Generation
Cristina Palmero	UDIVA-HHOI: A Multimodal Dataset for Context-Aware Multi-Actor Interaction Understanding
Jonathan Carter	Next-Token Prediction Learns Generalisable Representations of Sleep Physiology
Baraa Al Jorf	How Effectively Can LLM Agents Synthesize Multimodal Clinical Data for Risk Prediction?
Misbah farooq	CMAT: Cross-Modal Attention-Weighted Triplet Loss for Multimodal Emotion Recognition
Ardita Shkurti, Ljiljana Stojanovic	Helios: Multimodal AI for Solar-Cell Materials Screening
Faith Sodipe	A Unified Multimodal Deep Learning Architecture for Cross-Modal Understanding and Cross-Modal Support
Hazrat Ali	Hyperglycaemic Event Prediction through Continuous Glucose Monitoring on Hardware Efficient Deep Learning
Nicola Morley	Materials Informatics for Magnetostrictive Materials
Xianyuan Liu	Towards Deployment-Centric Multimodal AI Beyond Vision and Language
Damith Senadeera	CoLoRSMamba: Conditional LoRA-Steered Mamba for Supervised Multimodal Violence Detection

## Day 2

Name	Title
Jichun Li	An Effective Fragment-Based Dual Conditional Diffusion Framework for Molecular Generation
Longdan Hao	Myo-Guide: A Bidirectional Framework Bridging Multimodal AI and Rare Neuromuscular Disease Diagnostics
George Searle	Integrating ECG and PPG Features for Hypertension Detection Using Multimodal Deep Learning Approaches
Giuseppe Tripodi	Uncertainty-Aware Multimodal Cancer AI: A Conformal Prediction Study across Pathology and Genomics
Aman Jat	Assessing Equity Gaps in Artificial Intelligence Tools for Mental Health Diagnosis and Triage: A CPRD Cohort Study of 7.19 Million Adults with Multiple Long-Term Conditions
Hazrat Ali	Mask-Conditioned Synthetic Prostate MRI Volumes Using a Denoising Diffusion Probabilistic Model
Jiang Chang	Semantic-Guided Cross-Modal Interaction for Multimodal Depression Detection
Omotayo Omoyemi	Multimodal AI for Accessible Communication and Digital Health: The AURA Framework
Ricardo Valencia	Adapting Multi-Modal Variational Autoencoders for Microbiome Data
Tadiyos Hailemichael Mamo	Demographic-Aware Adversarial Transformer for Heart Failure Risk Prediction from ECG and Clinical Notes
Jingyu Hu	Probing and Fusion of Foundation Model Representations for Multimodal Cancer Analysis
Wenrui Fan	MuSProt: A Multimodal Multistate Protein Dataset Towards Programmable Protein Design
Mohammad Suvon	Multimodal AI Improves Low-Cost Cardiovascular Hemodynamic Assessment from Chest X-Ray and Electrocardiography
Jiayang Zhang	Computational Design of Grafted Motifs on Protein Scaffold
Bulugahapitiye Mohottalalage Amantha Viraj Bulugahapitiya	A Synergistic Tri-Modal Framework for Alzheimer's Disease Diagnosis Using Self-Supervised 3D Swin Transformer and LSTM with Gated Fusion.
Haili Ye	Multimodal Hypergraph Modelling of Structural-Functional Connectivity Coupling for Interpretable ADHD Brain Network Analysis
Kewei Zhu	MOF in Water: Screening Water-Stable Metal-Organic Frameworks via ReadMOF diagram

## Panel Discussion 1 (Day 1 | 16:30 – 17:20)

Research Panel: From Models to Evidence: Making Multimodal AI Deployable

Panel Members	Questions
<ul style="list-style-type: none"><li>• Chris Barnes</li><li>• Jonathan Carter</li><li>• Julia Hirschberg</li><li>• Henry Lee</li><li>• Tom Pollard</li><li>• Farah Shamout</li></ul>	<ol style="list-style-type: none"><li>1. What community outputs would most help multimodal AI move from research promise to real-world deployment?</li><li>2. What evidence should researchers provide before claiming a multimodal AI system is trustworthy enough for real-world use?</li></ol>

## Panel Discussion 2 (Day 2 | 14:40 – 15:30 )


Career Panel: Building a Sustainable Research Career in Multimodal AI

Panel Members	Questions
<ul style="list-style-type: none"><li>• Lu Gan</li><li>• Julia Hirschberg</li><li>• Henry Lee</li><li>• Tom Pollard</li><li>• Nataliya Tkachenko</li></ul>	<ol style="list-style-type: none"><li>1. How should early-career researchers decide which leadership, service or network opportunities to take on — and which to decline?</li><li>2. Beyond papers, what skills will matter most for the next generation of multimodal AI researchers?</li></ol>

# Event Guide

## Venue Information

UCL East Campus is easily accessible by public transport, with nearby tube, train, and bus connections. For more details, please see [UCL Find Us](#) or [Google Maps](#). No on-site parking is available. Closest paid parking is available at the [London Aquatics Centre](#).



Exterior of the UCL East



Interior of the UCL East



Cinema Auditorium

## Getting to the Venue

The UCL East campus is located on Queen Elizabeth Olympic Park in Stratford. The park is a vehicle-free site. It is therefore best to access our campus via public transport, the pedestrian paths, or cycle lanes.

- **Travel by train**
  - Stratford 10-min walk (Underground Central & Jubilee Lines, Elizabeth Line, Overground, DLR, Rail)
  - Stratford International 11-min walk (DLR, High-speed Rail)
  - Pudding Mill Lane 12-min walk (DLR)
  - Hackney Wick 19-min walk (Overground)
- **Travel by bus**
  - The closest bus stop is the London Aquatics Centre. For more information about transport links, accessible parking and cycle racks, please visit the [UCL Find Us](#) webpage.

## On-site Locations

All event activities will take place at UCL East Campus, One Pool Street. Locations are below:

- **Registration:** Ground Floor Atrium
- **Refreshments, lunch, and posters:** Ground Floor Atrium and First Floor
- **Talks:** Cinema Auditorium, Second Floor
- **Networking reception:** Teaching Space 4 Room 301, Third Floor

## On-site Help and Support

If you have any questions or need assistance during the workshop, please look for individuals wearing blue polo shirts with the **UKOMAIN** logo. They are available to provide support and help make your experience enjoyable.

## Internet Access

Connect to eduroam if you have an account, or connect to the UCLGuest Wi-Fi network and follow the registration process to create a guest account and access the internet.

## Emergency Information

The UCL East Campus, One Pool Street, is protected by an automatic fire alarm system. In the event of an emergency, please follow the UCL Hosts, who will lead you out to the relevant safe gathering place, remain with you, and advise you when it is safe to return. No scheduled tests or alarms are expected during the two days.

## Catering

Coffee, tea, water, and biscuits will be available during the coffee breaks. A Hot Buffet will be provided for lunch. If you have specific dietary needs or catering questions, please speak to our team at the registration desk for assistance.

## Filming and Photography

Please be advised that there will be media coverage, including filming and photography, during the workshop. The images and videos taken may be used for promotional purposes on the UK Open Multimodal AI Network website, social media, or YouTube channel. If you do not wish to appear in any videos or photographs, please inform the organisers by emailing [multimodalai26-group@sheffield.ac.uk](mailto:multimodalai26-group@sheffield.ac.uk), speaking to the registration desk, or informing the photographer/videographer.

## Accessibility Information

Please contact the organisers at [multimodalai26-group@sheffield.ac.uk](mailto:multimodalai26-group@sheffield.ac.uk) if you have any accessibility requirements that you would like to discuss, and we will endeavour to meet your requirements. A prayer room is available at UCL East Campus, and the staff can direct delegates to its location.

## Partners

This workshop is hosted by the [UK Open Multimodal AI Network \(UKOMAIN\)](#), a national initiative funded by the Engineering and Physical Sciences Research Council (EPSRC) and led by the University of Sheffield, University College London, and the University of Glasgow. This workshop is also supported by the Alan Turing Institute's Interest Group on [Meta-learning for Multimodal Data](#) (welcome to [sign-up and join](#)).



## Sponsors

We gratefully acknowledge the support of our sponsors, whose generosity has ensured the success of this workshop. Their contributions have been valuable in enabling students and researchers to attend and present their work.

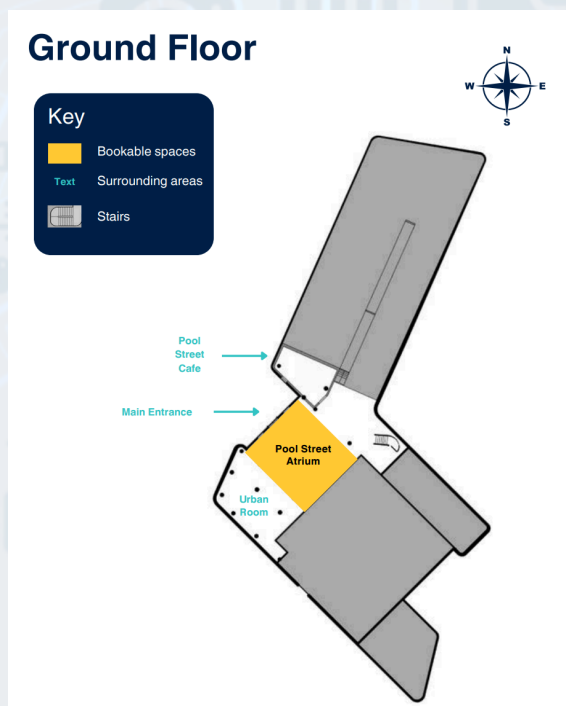


Engineering and  
Physical Sciences  
Research Council

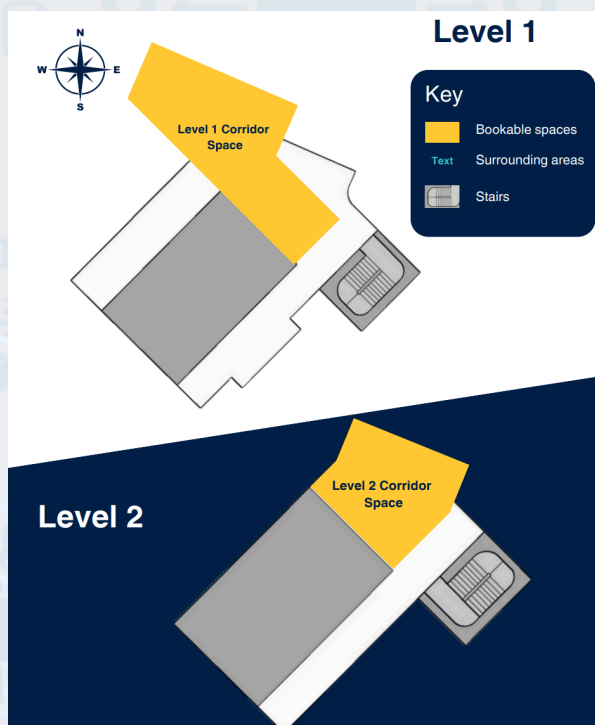
HENRY . . . .  
ROYCE . . . .  
INSTITUTE

## Floor Plan

The floor plan below shows the locations of catering points, poster boards, and other workshop facilities.



Registration & information desk, poster board, and catering are within the **Yellow** area.



Poster board and catering are also within the **Yellow** area on Level 1.