Australia's AI Advantage:
World-Class Research from Emerging Talent

A Snapshot of AI Research Capability and Talent in Australia

Kingston AI Group (KAIG)

August 2026,

Drawing on the inaugural Kingston AI Symposium (Adelaide Oval, 19 March 2026)

Main Authors: Adj Prof Sue Keay, Prof Shazia Sadiq, Dr Nisha Schwarz

In brief

Adjusted for population, Australia is one of the most productive AI research nations in the world, publishing at roughly six times the global per-person average, a higher research intensity than the United States, the European Union, or China. The 108 publications by predominantly early-career Australian researchers surfaced through the inaugural Kingston AI Symposium show this output is not peripheral: it is placed at the world's most selective AI venues, it tracks the global research frontier, and it reveals distinctive national strengths in trustworthy AI, machine unlearning, and multimodal understanding. These are precisely the capabilities a nation needs to accelerate AI in health, defence, and critical services on its own terms.

This paper was produced by the Kingston AI Group (KAIG), a national coordination and advocacy body comprising 14 of Australia's most senior AI researchers drawn from ten universities. KAIG was established in 2022 to align and accelerate Australia's national AI research effort and to advance the case for sovereign AI capability. Full details on KAIG and the inaugural Kingston AI Symposium that underpins this analysis appear in Section 7.

Key Findings

1. Australia ranks among the world leaders in AI research output per person (Stanford Institute for Human-Centered Artificial Intelligence (HAI), AI Index Report 2026). Australia produces an estimated 1.9% of global AI publications from just 0.33% of the world's population[1], a publication intensity of roughly six times the global average, ahead of the United States (~1.7×), Europe(~2.0×), and China (~1.0×)[2,3].

2. Our early-career researchers compete, and win, at the global frontier. The 108 publications in this snapshot were placed at venues such as CVPR, NeurIPS, ICLR, ICML, ACL, SIGMOD, and AAAI[4], where acceptance rates run between roughly 17% and 32%, and submission volumes have grown as much as 24-fold in a decade. Eleven (11) papers earned oral or spotlight designations, distinctions reserved for the top few percent of accepted work.

3. The research topics track the global frontier. Computer vision, multimodal learning, and language models dominate global AI research; they also dominate within this portfolio. Australia's emerging researchers are working on the problems the world considers most important, not in regional niches.

4. Australia shows distinctive strengths in responsible and trustworthy AI. Machine unlearning, anomaly detection, explainability, and privacy-preserving learning recur across institutions. With documented AI incidents rising 55% year-on-year globally, and right-to-erasure and clinical AI governance frameworks maturing, these strengths align with where global demand is heading.

5. The talent pipeline is deep in Australia. 50.5% of Australia's top AI researchers hold a PhD, the second-highest share of any country tracked, behind only the United Kingdom (51.1%), and 69% of the publications in this snapshot were led by PhD students. Australia's 30.1% female AI researcher representation is the second-highest of any country tracked, behind Saudi Arabia (32.3%). Without sustained investment, this talent is exportable.

[1] Australia's population officially passes 27 million | Australian Bureau of Statistics
[2] Australia's artificial intelligence ecosystem: growth and opportunities | National AI Centre
[3] The 2026 AI Index Report | Stanford HAI; ai_index_report_2026_chapter_1_research_development.pdf
[4] CVPR - IEEE/CVF Conference on Computer Vision and Pattern Recognition, NeurIPS - Conference on Neural Information Processing Systems, ICLR - International Conference on Learning Representations, ICML - International Conference on Machine Learning, ACL - Annual Meeting of the Association for Computational Linguistics, SIGMOD - ACM SIGMOD International Conference on Management of Data, AAAI - AAAI Conference on Artificial Intelligence.

A note on scope

This paper draws on 108 publications submitted to the inaugural Kingston AI Symposium (March 2026), of which 55 were selected for detailed review. The portfolio is a snapshot, not a census: it reflects the researchers who knew about and engaged with a first-year event, and it understates national activity in fields such as robotics, reinforcement learning, and generative modelling that are evidenced to be well represented in Australia. Future symposia are expected to broaden this representation as the venue becomes established. The full analysis underpinning this paper appears in the appendices.

Recommendation: The case for investing in sovereign capability is grounded in evidence. Frontier model production is concentrated in the United States and China; with a strong share of global talent, Australia can aspire to develop its own domain-specific AI models. The gap between Australia and other nations is not capability, it is scale, infrastructure, and coordination. That is the gap the Kingston AI Group was formed to close.

1. The Global Context: AI Research Competition

Global AI research is scaling rapidly and becoming concentrated in certain nations. China now accounts for the largest share of AI publications worldwide, with 17.8% of publications and 20.6% of citations in 2024, while the United States retains leadership in high-impact output, producing 59 notable frontier models in 2025 to China's 35. Europe collectively holds roughly 11% of global publication share, though almost 40% of tracked publications carry no clearly identified national affiliation in the underlying data, a limitation of the source dataset rather than of any single country's output. Almost all frontier model development now occurs inside a small number of US and Chinese organisations. [5]

The competitive intensity of the field's publishing venues tells a consistent story. NeurIPS, the field's flagship machine learning conference, received 21,575 submissions in 2025 and accepted 24.5%. CVPR, the premier computer vision venue, accepted 22.1% of 13,008 submissions. ICML accepted 26.9% of 12,107 submissions. ICLR has grown roughly 24-fold since 2017 and now accepts around 32% of submissions. AAAI attracted close to 29,000 submissions for its 2026 edition and accepted 17.6%. Acceptance at any one of these venues places a paper in the top third, often the top quarter, of the world's AI research output for that year. Oral and spotlight designations sit a further order of magnitude higher: NeurIPS spotlights, for example, go to roughly 3% of submissions.[6]

This is the environment in which Australian research must be assessed. In a field where volume is dominated by economies twenty to fifty times Australia's size, we assess the quality of research output and where research effort is directed.

2. Australia's Position, Adjusted for Scale

The Stanford AI Index reports 2024 publication shares of 17.8% for China, 7.3% for the United States, and 11.1% for Europe, a grouping that may include non-EU countries such as the United Kingdom and Switzerland, so the EU-specific figure is likely modestly overstated relative to the EU-27 (27-member states of the European Union) population used below. Stanford does not break Australia out separately in this measure; for Australia's own share we draw on the Australian Government's National AI Centre,

[5] Stanford Institute for Human-Centered AI, The AI Index Report 2026, Stanford University, April 2026. Chapter 1 (Research and Development), Chapter 3 (Responsible AI), Chapter 1 and Chapter 7 (Education). https://hai.stanford.edu/assets/files/ai_index_report_2026_chapter_1_research_development.pdf

[6] CVPR 2025 Programme Statistics. 2,878 papers accepted from 13,008 submissions (22.1% acceptance rate). openaccess.thecvf.com; NeurIPS 2025 Conference Data. 5,290 papers accepted from 21,575 submissions (24.5% acceptance rate). Reflections on the 2025 Review Process, NeurIPS Blog, neurips.cc; Conference statistics compilations: NeurIPS and AAAI official programme data; OpenReview public data (ICLR, ICML); ACL Anthology (aclanthology.org); AAAI 2026 proceedings (aaai.org); ICLR 2017 Statistics - Paper Copilot; https://csconfstats.xoveexu.com/conferences/aaai/2026/

which reports Australia's global AI publication share at 1.9% in 2024[7]. Dividing each figure by its share of world population produces a comparable per-capita intensity index[8]. On this basis, Australia's research intensity is approximately 5.8 times the global per-person average, higher than the United States (~1.7×), the European Union (~2.0×), or China (~1.0×). This is an indicative comparison, not a definitive ranking, but it shows a clear and directionally consistent pattern of Australian strength.

Metric	China	United States	EU / Europe†	Australia
Population, 2025 (millions)	1,416	342	449	27
Share of world population	17.2%	4.2%	5.5%	0.33%
Share of global AI publications (2024)	17.8%	7.3%	11.1%	1.9%‡
Publication intensity (share ÷ population share)	~1.0×	~1.7×	~2.0×	**~5.8×**
Share of global AI citations (2024)	20.6%	12.6%	19.5%	~2%‡
Citation intensity (share ÷ population share)	~1.2×	~3.0×	~3.5×	**~6.1×**
Notable frontier models produced (2025)	35	59	Several	0 (Australia is not listed)
Papers in the world's 100 most-cited AI papers (2024)	41	46	n/a§	14

Table 1. Australia's AI research position adjusted for population. Publication and citation shares (China, US, EU) from the Stanford AI Index Report 2026, Chapter 1; Australia's publication share from Australia's National AI Centre (2026); population figures are 2025 UN estimates; intensity is the ratio of research share to population share. † Stanford's "Europe" grouping may extend beyond the EU-27 population figure used here. ‡ Not reported by Stanford for Australia specifically; sourced from the National AI Centre[9] (publications) or estimated proportionally (citations), treat as indicative. § Not reported as an EU-wide aggregate by Stanford; the highest-placed EU member state individually is Germany, with 8 papers. Several smaller advanced economies (for example Singapore, Switzerland, and South Korea, which leads the world in AI patents per capita) also show high per-person intensity; Australia's distinction is sustaining this intensity alongside top-tier venue placement and talent-pipeline depth.

Two further measures sharpen the picture. First, depth of expertise: the Stanford AI Index 2026 reports that 50.5% of Australia's top AI researchers hold a PhD, the second-highest share of any country tracked, behind only the United Kingdom (51.1%) and clearly ahead of the United States (42.0%). China is not included in this measure; the underlying dataset explicitly excludes it. Second, diversity: at 30.1%, Australia's female AI researcher representation is the second-highest of any country tracked, behind Saudi Arabia (32.3%) and just ahead of Canada (29.6%), modestly above the United States (27.7%) and above most individually tracked European economies, which range from roughly 18% to 30%[10]. The Index notes that no country approaches parity, and most have shown no meaningful improvement since 2010. Both measures matter increasingly to international funding eligibility, partnership criteria, participation in standards bodies, and the development of inclusive AI.

[7] Australian Government National AI Centre, Australia's Artificial Intelligence Ecosystem: Growth and Opportunities, 2026. ai.gov.au
[8] United Nations, World Population Prospects 2024, 2025 medium-variant estimates. population.un.org
[9] Australia's artificial intelligence ecosystem: growth and opportunities | National AI Centre
[10] Figure 1.8.4 in ai_index_report_2026_chapter_1_research_development.pdf

What the per-capita lens cannot show is whether the quality holds at the top of the field. For that, the evidence in this snapshot shows where Australian research is positioned.

3. What Australian Researchers Are Working On

The 108 publications gathered through the inaugural Kingston AI Symposium were authored predominantly by early-career researchers: 69% were led by PhD students, with postdoctoral researchers, research fellows, and early-career academics making up most of the remainder. They span 17 institutions and more than a dozen venues, and approximately 94% originated from researchers based in Australia at the time of publication[11].This is a window into the next generation of Australian AI research capability.

The venue distribution is striking for its selectivity. The three most represented venues, CVPR (21 papers), NeurIPS [19 papers]), and ICCV [9 papers], are among the most competitive publishing destinations in computer science, and together account for around 46% of the portfolio. The remainder spreads across ICLR, ICML, AAAI, IJCAI, ACL, EMNLP, ICRA, IROS, CoRL, SIGKDD, and SIGMOD, the recognised top-tier venues for machine learning, language, general AI, robotics, and data systems respectively. Eleven papers carried oral or spotlight designations, and one (Maincode) won the Best Special Theme Paper award at EMNLP 2025. This consistency of placement across a single twelve-month cycle reflects research competing at the global frontier.

3.1 Alignment with Global Research

The thematic shape of the portfolio mirrors the shape of global AI research activity. Across the world's major venues, computer vision, machine learning methods, and natural language processing are by a wide margin the largest research area; at AAAI 2026, for instance, the three most common keyword areas among nearly 29,000 submissions were computer vision (~10,000 papers), machine learning (~8,000), and NLP (~4,000)[12]. The Australian portfolio follows the same contour: of the 55 papers selected for detailed review, over half fall into vision and multimodal understanding or representation learning and generalisation, with language models and NLP forming the next largest cluster.

This means Australia's emerging researchers are not working at the periphery of the field but on its most contested ground: cross-modal grounding, foundation-model adaptation, instruction-following and alignment, and learning under distribution shift. These are the problems on which global AI progress currently turns, and on which Australian work is being accepted alongside that of the largest labs in the world.

3.2 Where Symposium Research Clustered

The 55 reviewed papers group into six thematically coherent clusters (full detail in Appendix C):

- Vision and multimodal understanding (~19 papers). Text–image–point-cloud pretraining, audio-visual segmentation, vision-language navigation, federated multimodal learning, and semantic LiDAR generation. Cross-modal alignment is the recurring sub-theme.
- Language models and NLP (~10 papers). LLM bias and fairness evaluation (including a pharmacovigilance study testing frontier models against structured FDA adverse-event data), conversational grounding, instruction-following, robustness, and alignment.
- Embodied AI, robotics and navigation (~7 papers). Vision-language navigation, language-grounded robotics, robot learning, and neuromorphic sensing hardware, a thread with clear relevance to advanced manufacturing and defence applications.
- Representation learning, adaptation and generalisation (~9 papers). Continual learning, domain generalisation, few-shot learning via causal disentanglement, attention efficiency, and machine unlearning.
- Anomaly detection, uncertainty and explainability (~5 papers). Multi-class anomaly detection, a major new large-scale anomaly benchmark dataset (MANTA, CVPR 2025), uncertainty quantification, and counterfactual explanation methods.

[11] Kingston AI Symposium Submissions Dataset. 108 submissions received; 55 selected for detailed review. Inaugural Kingston AI Symposium, Adelaide Oval, 19 March 2026. Internal dataset, Kingston AI Group, June 2026.
[12] AAAI-26 Review Process Update: Scale, Integrity Measures, and Pathways to Sustainability - AAAI

- Data systems and ML efficiency (~5 papers). Data centric AI, graph and tabular learning, federated learning over heterogeneous data, and efficiency-focused methods.

4. Australia's AI Strengths

Three concentrations stand out as distinctively Australian, areas where the portfolio shows depth across multiple institutions rather than isolated papers, and where Australia has a plausible claim to international leadership.

4.1 Machine Unlearning and the Right to Be Forgotten

Multiple papers across at least two institutions address machine unlearning, the selective removal of learned information from trained models, including new methods published at CVPR, ICCV, and AAAI (one as an oral presentation), and a clinical-AI governance framing submitted to npj Digital Medicine. This is a coherent, growing research identity with direct policy traction: right-to-erasure obligations, clinical AI governance frameworks, and data-sovereignty requirements all depend on exactly this capability.

4.2 Trustworthy, Explainable, and Safety-Relevant AI

Anomaly detection, uncertainty quantification, counterfactual explanation, LLM fairness evaluation, and federated privacy-preserving learning recur throughout the portfolio. The timing is significant: the Stanford AI Index documents a 55% year-on-year rise in reported AI incidents (362 in 2025, up from 233 in 2024) and identifies a persistent global gap between AI capability and the governance frameworks needed to deploy it responsibly. Research of this kind is the raw material of safe deployment in health, finance, and critical infrastructure, and of the international standards leadership the Index suggests remains unclaimed.

4.3 Embodied and Physical AI

A smaller but distinct thread spans robot learning, language-grounded navigation, and neuromorphic sensing hardware. Given Australia's strengths in field robotics, mining automation, and agricultural technology[13], this thread is almost certainly under-represented in the present snapshot relative to national activity, robotics-focused researchers had less visibility of a first-year, ML-centred symposium. Even so, its presence signals the hardware–software perception capability that advanced manufacturing and defence applications require, and it is expected to strengthen in future symposia.

4.4 What does this signal for Australian AI Research

These clusters suggest a national research trajectory oriented toward deploying AI responsibly in high-stakes settings rather than toward frontier-scale model training. The strongest concentrations, unlearning, trustworthiness, multimodal perception, health-adjacent NLP, are the capabilities required to adapt, govern, audit, and certify powerful AI systems for domains such as clinical medicine, defence, and critical services. That is a rational specialisation for a nation of Australia's scale. It does not require hyperscale compute, and it compounds Australia's existing strengths in health and safety-critical industries.

The current portfolio contains comparatively little reinforcement learning, generative modelling as a primary focus, or formal verification. National activity exists in all three; their absence here reflects the snapshot's coverage rather than a national gap and tracking their representation at future symposia will indicate whether Australia's portfolio is broadening or specialising further.

5. Converting Talent to Sovereign Capability

Why does any of this matter? The evidence in this paper bears directly on whether Australia can credibly pursue sovereign AI capability, the ability to build, adapt, govern, and deploy AI systems for national purposes without depending entirely on technology controlled offshore.

The case rests on three observations. First, frontier AI is concentrating, not diffusing: notable model production sits almost entirely within the United States and China, and the Stanford AI Index records a

[13] Australia's robotics opportunity | | Department of Industry Science and Resources

collapsing performance gap between the two, a strategic competition in which Australia is currently a consumer, not a producer. For systems destined for health records, defence networks, and critical infrastructure, complete dependence on externally controlled models carries data-sovereignty, assurance, and supply-chain risks that no amount of procurement diligence fully removes.

Second, the binding constraint is not talent. The evidence of this snapshot is that Australia's early-career researchers already produce work the world's most selective venues accept, in the areas that sovereign deployment most requires, trustworthy AI, unlearning, and multimodal perception. The talent pipeline behind them is deep (a PhD share second only to the UK's) and diverse. What Australia lacks is not researchers; it is the scale, compute infrastructure, career pathways, and national coordination needed to retain them and translate their work.

Third, that constraint is addressable, and time-sensitive. The same Stanford data shows global AI talent flows reorganising rapidly, with researcher migration into the United States down 89% since 2017, and 80% in the past year alone. A talent base of this quality is exportable. The window in which Australia can convert demonstrated research excellence into durable sovereign capability, through research infrastructure, sovereign compute, and structured pathways for the international students who already make up part of this pipeline, is open now, but it will not remain open indefinitely.

The bottom line

Australia should have confidence that it possesses AI talent contributing to the defining technical challenges of the era, at world standard, in the fields that matter most for safe development, deployment and adoption of AI. This paper documents that the capability exists. The question for policymakers is whether Australia will invest to retain it, grow it, and put it to sovereign use.

6. Scope, Limitations, and Next Steps

This paper is deliberately framed as a snapshot. Its evidence base, 108 publications submitted to a first-year symposium, reflects who knew about the event and chose to participate, not the full breadth of Australian AI research, which the Kingston AI Group believes to be stronger still. Fields with established communities elsewhere, notably robotics, reinforcement learning, and generative modelling, are under-represented here relative to national activity. Quantitative claims drawn from the symposium dataset should be read as indicative of the participating cohort, not as representative of the country as a whole.

Three commitments follow for future editions. First, broader reach: as the symposium becomes established, targeted outreach to robotics, RL, and generative-modelling communities should bring those fields into the portfolio. Second, better measurement: future reports will track citation impact and venue designations longitudinally, so that claims about high-impact work can be quantified rather than asserted. Third, transparency of method: selection criteria for symposium presentation, and the meaning of multi-venue paper tags, will be documented so the evidence base can be independently assessed.

7. About the Kingston AI Group and the Inaugural Symposium

The Kingston AI Group (KAIG) is a national coordination and advocacy body comprising 14 of Australia's most senior artificial intelligence researchers, drawn from 10 universities. Established in late 2022 following a convening of leading professors in Canberra, KAIG was formed to align and accelerate Australia's national AI research effort and to advance the case for sovereign AI capability as a matter of national security and economic priority. Its members hold Laureate Fellowships, memberships of the learned academies, and sustained records of internationally competitive funding.

Recognising the calibre of research being produced across members' networks and the country more broadly, and a gap in how that capability and talent is recognised, KAIG convened its inaugural symposium to surface, assess, and communicate high-quality AI research from across Australia's research community.

The inaugural Kingston AI Symposium was held on 19 March 2026 at Adelaide Oval, bringing together Australia's leading AI researchers for a full day of presentations, poster sessions, and collaborative exchange. The program featured ten research presentations spanning reinforcement learning, explainability, optimisation theory, efficient learning, test-time adaptation, counterfactual reasoning, machine unlearning, and vision transformers, delivered by PhD researchers, postdoctoral fellows, and senior academics from institutions across the country. Two keynote addresses anchored the day: Professor Dana Kulic (Monash University) opened proceedings with a keynote on human-AI teaming with physical systems, and Professor Anton van den Hengel (Australian Institute for Machine Learning) closed the morning session with a provocation on the limits of data-centric pretraining. The symposium concluded with a gala dinner in the Ian McLachlan Room, where further paper presentations continued the day's exchange in a more informal setting.

Sponsors

The Symposium was made possible through the generous support of its founding sponsors.

Google Australia contributed foundational support reflecting its long-standing investment in Australian AI research and talent.

ResetData provided critical backing, consistent with its role as a key partner in Australia's and South Australia's sovereign compute agenda.

Commonwealth Bank of Australia demonstrated its commitment to responsible AI adoption in the financial sector and to building the research partnerships necessary to deploy AI safely.

Together, these sponsors signal the growing alignment between Australia's research community and the institutions, technology, infrastructure, and capital that will translate that research into national capability.

Appendices

The appendices provide the detailed data underpinning the analysis in the body of this paper.

Appendix A, Venue Distribution

The following table summarises publication counts by venue across the 108 total submissions. The 55 papers selected for symposium review are drawn from this set and are analysed in depth in Appendix C.

Venue	Papers	Domain
AAAI	14	AI (General)
ACL	11	NLP
CoRL	1	Robotics
CVPR	21	Computer Vision
EMNLP	3	NLP
ICCV	9	Computer Vision
ICLR	14	Representation Learning
ICML	1	Machine Learning
ICRA	4	Robotics
IJCAI	3	AI (General)
IROS	3	Robotics
NeurIPS	19	Machine Learning
SIGKDD	2	Data Mining
SIGMOD	2	Data Systems

CVPR (21), NeurIPS (19), and ICCV (9) together account for 49 of the 107 submissions (1 was ineligible), around 46%, roughly a third, confirming a strong centre of gravity in computer vision and perception. Robotics venues (ICRA, IROS, CoRL) represent an emerging thread distinct from the core vision work. The presence of SIGMOD and SIGKDD points to an active data systems research strand. These figures describe the submitters to the symposium and should not be read as a census of the country's overall research activity, which the Kingston AI Group believes to be stronger and broader still.

Appendix B, Institutional Profile of Submissions

The 108 submissions were received from researchers across 17 Australian and international institutions. Submitters were predominantly PhD students (75 submissions, 69%), with postdoctoral researchers (11), research fellows (9), and academic staff (7 lecturers) forming most of the remainder of the cohort. Two submissions came from industry-affiliated researchers at CSIRO/Data61. 101 submissions (93.5%) originated from researchers currently based in Australia.

108 Total submissions	**17** Institutions	**75** PhD students	**11** Postdoctoral

B.1 Submissions by Institution

The following table ranks institutions by number of submissions received. Where a submitter listed a joint affiliation (e.g. AIML / University of Adelaide), submissions are attributed to single main institution in Australia.

Institution	Submissions	Notes
Adelaide University / AIML	24	Strong across NeurIPS, ICCV and ICLR, with AAAI, ICML and CVPR also represented
University of New South Wales	16	Largest single submitter after Adelaide/AIML; strongest at CVPR, with NeurIPS, ACL, ICCV and ICLR also present
University of Melbourne	15	Broad coverage led by ICLR and NeurIPS, plus ACL, CVPR, EMNLP and SIGKDD
Monash University	13	Strongest robotics output (ICRA, IROS) alongside CVPR, AAAI, ICCV, ICLR and ACL; includes one IITB-Monash Research Academy joint submission
Deakin University	9	AAAI-heavy, with NeurIPS, IJCAI, ICML, ACL and SIGKDD also represented
The University of Queensland	5	ICCV and CVPR-led, plus ICLR, NeurIPS and one SIGMOD paper
Australian National University	5	NeurIPS-led, with CVPR also represented
The University of Western Australia	5	AAAI and NeurIPS-led, plus CVPR; 3D object recognition and generative modelling
RMIT University	4	IJCAI-led, with AAAI, ACL, IROS, ICLR and ICML also represented
Queensland University of Technology	3	CVPR-led, plus IROS (oral), ACL and ICCV; robotics and vision
Maincode	2	Two NeurIPS/EMNLP papers, including a Best Special Theme Paper at EMNLP 2025
CSIRO / Data61	2	Industry-affiliated; NeurIPS and ICLR
Macquarie University	1	EMNLP
Dalian University of Technology (China)	1	CVPR; China-based affiliate
Ho Chi Minh University of Education (Vietnam)	1	ICLR; Vietnam-based affiliate
Federation University Australia	1	ICLR
Murdoch University	1	CVPR

B.2 Presentation Quality Indicators

Of the 108 submissions, 57 were published (camera-ready) and 44 had been accepted at the time of submission. Four (4) papers were under review. Eleven (11) papers had received oral or spotlight presentation designations, which at venues like NeurIPS indicate exceptional peer recognition.

57 Published	**44** Accepted	**94%** Published or accepted	**11** Oral / Spotlight

Institution	Venue	Presentation Type
AIML / Adelaide University	AAAI / ICLR / ICML / NeurIPS	Spotlight poster (published at TU Munich)
ANU	NeurIPS	Spotlight
University of Queensland	CVPR / ICCV / ICLR / NeurIPS	Oral
Monash University	ACL	Oral
Monash University	ICRA	Oral
Monash University	AAAI	Oral (×2)
Deakin University	IJCAI	Oral

QUT	IROS	Oral
RMIT University	AAAI	Oral
Monash University	IROS	Oral

The EMNLP 2025 Best Special Theme Paper award (Maincode) is a further quality signal beyond the oral/spotlight category.

Appendix C, Research Strengths: Cluster Analysis of Selected Papers

This section analyses the 55 papers selected for presentation at the Kingston AI Symposium. Papers are grouped into six thematically coherent research clusters based on topic, method, and domain signals. 11 papers were submitted with multiple venue tags, indicating either simultaneous workshop/main-track submissions or extended versions across conferences.

55 Papers reviewed	**6** Research clusters	**14** Venues

C.1 Research Clusters

Cluster	Papers	Description
Vision & Multimodal Understanding	~19	Text-image-point cloud pretraining, audio-visual segmentation, text-to-video retrieval with uncertainty quantification, federated multimodal learning, multi-view change localisation, vision-language navigation (SAME), semantic LiDAR generation via diffusion, text-to-image with concept control, multi-view anomaly detection (MANTA dataset).
Language Models & NLP	~10	LLM bias and fairness in pharmacovigilance (FAERS/persona-based evaluation), conversational grounding, instruction-following, robustness and alignment. Includes the Best Special Theme Paper of EMNLP 2025 (Maincode).
Embodied AI, Robotics & Navigation	~7	Vision-language navigation (SAME adaptive mixture-of-experts), event-driven neuromorphic hardware, ICRA/CoRL robot learning, IROS language-grounded robotics.
Representation Learning, Adaptation & Generalisation	~9	Continual learning with mixture-of-adapters, domain generalisation, few-shot learning via causal disentanglement, machine unlearning via Nash bargaining (MUNBa), attention efficiency, self-supervised methods.
Anomaly Detection, Uncertainty & Explainability	~5	Multi-class anomaly detection with class-aware contrastive learning, MANTA large-scale anomaly dataset (CVPR), text-to-video retrieval uncertainty, counterfactual explanations, IoU-based detection (MGIoU, Monash).
Data, Systems & ML Efficiency	~5	SIGMOD data systems (UQ), SIGKDD graph/tabular learning (Deakin/ANU), federated learning with heterogeneous multimodal data, efficiency-focused NeurIPS/ICLR work.

C.2 Key Observations

Dominant research focus.
Over half the papers fall into vision and multimodal understanding (Cluster 1) or representation learning and adaptation (Cluster 4) [19+9], consistent with CVPR and, NeurIPS being the two most represented venues. This reflects a group with deep computer vision infrastructure and a growing interest in cross-modal grounding and generalisation.

Embodied AI as a distinct thread.
The ICRA, IROS, and CoRL papers, primarily from Monash University, AIML and QUT, including neuromorphic hardware work from AIML/Adelaide, point to a hardware-software embodied perception

thread analytically distinct from the core vision work, with potential relevance to advanced manufacturing and defence applications.

Policy relevance of research.
Multiple papers across Monash University and AIML address machine unlearning, including MUNBa (Nash Bargaining, CVPR/ICCV) and an AAAI oral on data-diet unlearning. A separate npj Digital Medicine submission (AIML) frames machine unlearning as a clinical AI governance imperative (not part of the 55 reviewed however).

Health AI signal.
Cluster 2 includes a pharmacovigilance/LLM bias paper from Adelaide University evaluating frontier language models against structured FAERS data using a persona-based framework. Combined with the machine unlearning/clinical AI governance work, this positions the group at the intersection of AI safety, health, and NLP.

Notable gaps.
The current participation portfolio contains almost no work in reinforcement learning, generative modelling as a primary focus, or formal verification. Broader evidence exists of these areas being studied elsewhere by Australian researchers, and their inclusion will be an important consideration for future symposia.

High-impact indicator.
The Fan et al. paper (UNSW) is the standout high-citation-potential paper in the set. The NeurIPS Spotlight (ANU) is a further high-visibility output worth tracking.

Appendix D, Relevance to Sovereign AI Strategy and Talent Policy

The comparative data in Table 1 (Section 2) is the primary evidence base for this appendix; the points below draw out its policy implications and add two further data points on Australian venue placement not shown in Table 1: Australian researchers, including multiple submissions from Adelaide University, placed papers at both CVPR 2025 (22.1% acceptance) and NeurIPS 2025 (24.5% acceptance) in this snapshot alone.

D.1 What Distinguishes Australia's Position

Venue selectivity, not volume.
CVPR 2025 accepted just 22.1% of 13,008 submissions. NeurIPS 2025 accepted 24.5% of 21,575 submissions. ICLR and ICML operate at broadly similar thresholds (26–32%). Placement across all four venues in a single twelve-month cycle, alongside ACL, EMNLP, ICRA, and IROS, reflects a research operation competing at the global frontier, not in lower-tier regional venues.

PhD talent depth, second globally.
The Stanford AI Index 2026 reports that 50.5% of Australia's top AI researchers hold a PhD, the second-highest share in the world, behind only the United Kingdom (51.1%) and ahead of the United States (42.0%). The Kingston AI Symposium submission data reinforces this: 75 of 108 submissions (69%) came from PhD students. This is the talent base that produces frontier venue placements. It is not replicated at scale, and it is at risk of export to industry or overseas institutions without sustained investment in academic research infrastructure.

Gender representation, second globally.
At 30.1% female representation, Australia ranks second among tracked countries, behind Saudi Arabia (32.3%) and just ahead of Canada (29.6%), modestly above the United States (27.7%). The Stanford AI Index notes that no country approaches gender parity, and that representation has shown no meaningful improvement since 2010 in most nations. Australia's position here is a genuine, evidence-based differentiator for funding, partnership, and international collaboration. However, the Kingston AI Symposium did not capture this data, we do not believe this to deviate from the consensus.

D.2 Relevance to Sovereign AI Strategy

- Trustworthy AI and explainability. Aligns with national AI safety priorities at a moment when the Stanford Index notes documented AI incidents rose 55% year-on-year (362 in 2025 vs 233 in 2024).

- Federated and privacy-preserving systems. Supports health and defence data constraints, directly relevant to Australia's sovereign data obligations.
- Machine unlearning. Directly supports the right-to-erasure and clinical AI governance frameworks currently under development domestically and internationally.
- 3D perception and robotics. Supports advanced manufacturing, defence, and agricultural applications.
- Clinical AI governance. Positions Australia for international standards leadership, a domain where the Index identifies a global gap between capability and policy.
- Non-US/China research pipeline. Model production remains concentrated in the US and China; Australia's coordinated academic research pipeline is one of the few non-US/China sources of frontier-quality ML output globally.

D.3 Implications for AI Education and Talent Strategy

The following observations are drawn from the Stanford AI Index 2026, Chapter 7 (Education), and from the Kingston AI symposium submission demographics.

- PhD-to-industry flow reversal. New AI PhD graduates in the US and Canada are increasingly staying in academia rather than moving to industry, a structural window for a research-centre model. The KAIG submission pool (69% PhD students) shows the pipeline is active.
- A policy opportunity in AI education. There is scope for Australian institutions to lead curriculum development frameworks at the national level, drawing on research depth across the KAIG network.
- Diversity as a competitive differentiator. Australia's 30.1% female researcher representation becomes increasingly valuable as diversity metrics become central to funding eligibility and international partnership criteria.
- Applied research demand. Growing AI literacy without engineering depth in the broader workforce reinforces demand for Australia's translation and applied research capabilities, particularly relevant to the CSIRO/industry-affiliated submissions in this cohort.
- International student pipeline. Of the 108 submissions, 7 were from researchers based outside Australia at the time (5 China, 1 India, 1 Vietnam), several affiliated with Australian institutions. This international pipeline is a talent funnel worth nurturing through structured pathways.

Appendix E, Abbreviations and Full Names

E.1 Conference and Journal Venues

Abbreviation	Full Name	Field / Location
ACL	Annual Meeting of the Association for Computational Linguistics	Natural Language Processing
AAAI	Association for the Advancement of Artificial Intelligence (Conference)	AI (General)
CoRL	Conference on Robot Learning	Robotics / Robot Learning
CVPR	IEEE/CVF Conference on Computer Vision and Pattern Recognition	Computer Vision
EMNLP	Conference on Empirical Methods in Natural Language Processing	Natural Language Processing
ICCV	IEEE/CVF International Conference on Computer Vision	Computer Vision
ICLR	International Conference on Learning Representations	Machine Learning / Representation Learning
ICML	International Conference on Machine Learning	Machine Learning

ICRA	IEEE International Conference on Robotics and Automation	Robotics
IJCAI	International Joint Conference on Artificial Intelligence	AI (General)
IROS	IEEE/RSJ International Conference on Intelligent Robots and Systems	Robotics
NeurIPS	Conference on Neural Information Processing Systems	Machine Learning
SIGKDD	ACM SIGKDD Conference on Knowledge Discovery and Data Mining	Data Mining / ML
SIGMOD	ACM International Conference on Management of Data	Database / Data Systems
npj Digital Medicine	Nature Partner Journal Digital Medicine (Nature Portfolio)	Clinical / Digital Health
ICCVW	ICCV Workshop track (co-located workshops at ICCV)	Computer Vision

E.2 Institutions

Abbreviation	Full Name	Field / Location
AIML	Australian Institute for Machine Learning, Adelaide University	Adelaide, SA
ANU	Australian National University	Canberra, ACT
A2I2	Applied Artificial Intelligence Institute, Deakin University	Geelong, VIC
CSIRO	Commonwealth Scientific and Industrial Research Organisation	National (Australia)
Data61	CSIRO's Data Science and Artificial Intelligence Research Group	National (Australia)
FEIT	Faculty of Engineering and Information Technology, University of Melbourne	Melbourne, VIC
IITB-Monash	IITB-Monash Research Academy (joint: IIT Bombay and Monash University)	Mumbai, India / Melbourne, VIC
KAIG	Kingston AI Group	National (Australia)
QUT	Queensland University of Technology	Brisbane, QLD
RMIT	RMIT University (Royal Melbourne Institute of Technology)	Melbourne, VIC
TU Munich	Technical University of Munich (Technische Universität München)	Munich, Germany
UNSW	University of New South Wales	Sydney, NSW
UQ	University of Queensland	Brisbane, QLD
UWA	University of Western Australia	Perth, WA

E.3 Technical Terms and Paper Names

Term	Full Name / Description
AI	Artificial Intelligence
FAERS	FDA Adverse Event Reporting System, US pharmacovigilance database used in drug safety surveillance
LLM	Large Language Model, a neural network trained on large text corpora to generate and understand language

MANTA	Multi-view and Visual-Text Anomaly Detection Dataset, a large-scale benchmark introduced by Fan et al. at CVPR 2025
MGIoU	Modified Generalised Intersection over Union, an improved bounding-box detection metric
ML	Machine Learning
MUNBa	Machine Unlearning via Nash Bargaining, a method for selectively removing learned information from trained models
NLP	Natural Language Processing
SAME	State-Adaptive Mixture of Experts, a vision-language navigation architecture
SAC	Senior Area Chair, a senior program committee role at NeurIPS and similar conferences
AC	Area Chair, a mid-level program committee role coordinating reviews at major AI conferences
DIET	Data Diet Machine Unlearning, a method using selective data pruning
SG-LDM	Semantic-Guided LiDAR Diffusion Model, a method for LiDAR point-cloud generation

End of document