Ministry of Foreign Affairs

AI Research in Singapore
ESF Research Brief No.1

Commissioned by the Netherlands Enterprise Agency

>> Sustainable. Agricultural. Innovative. International.

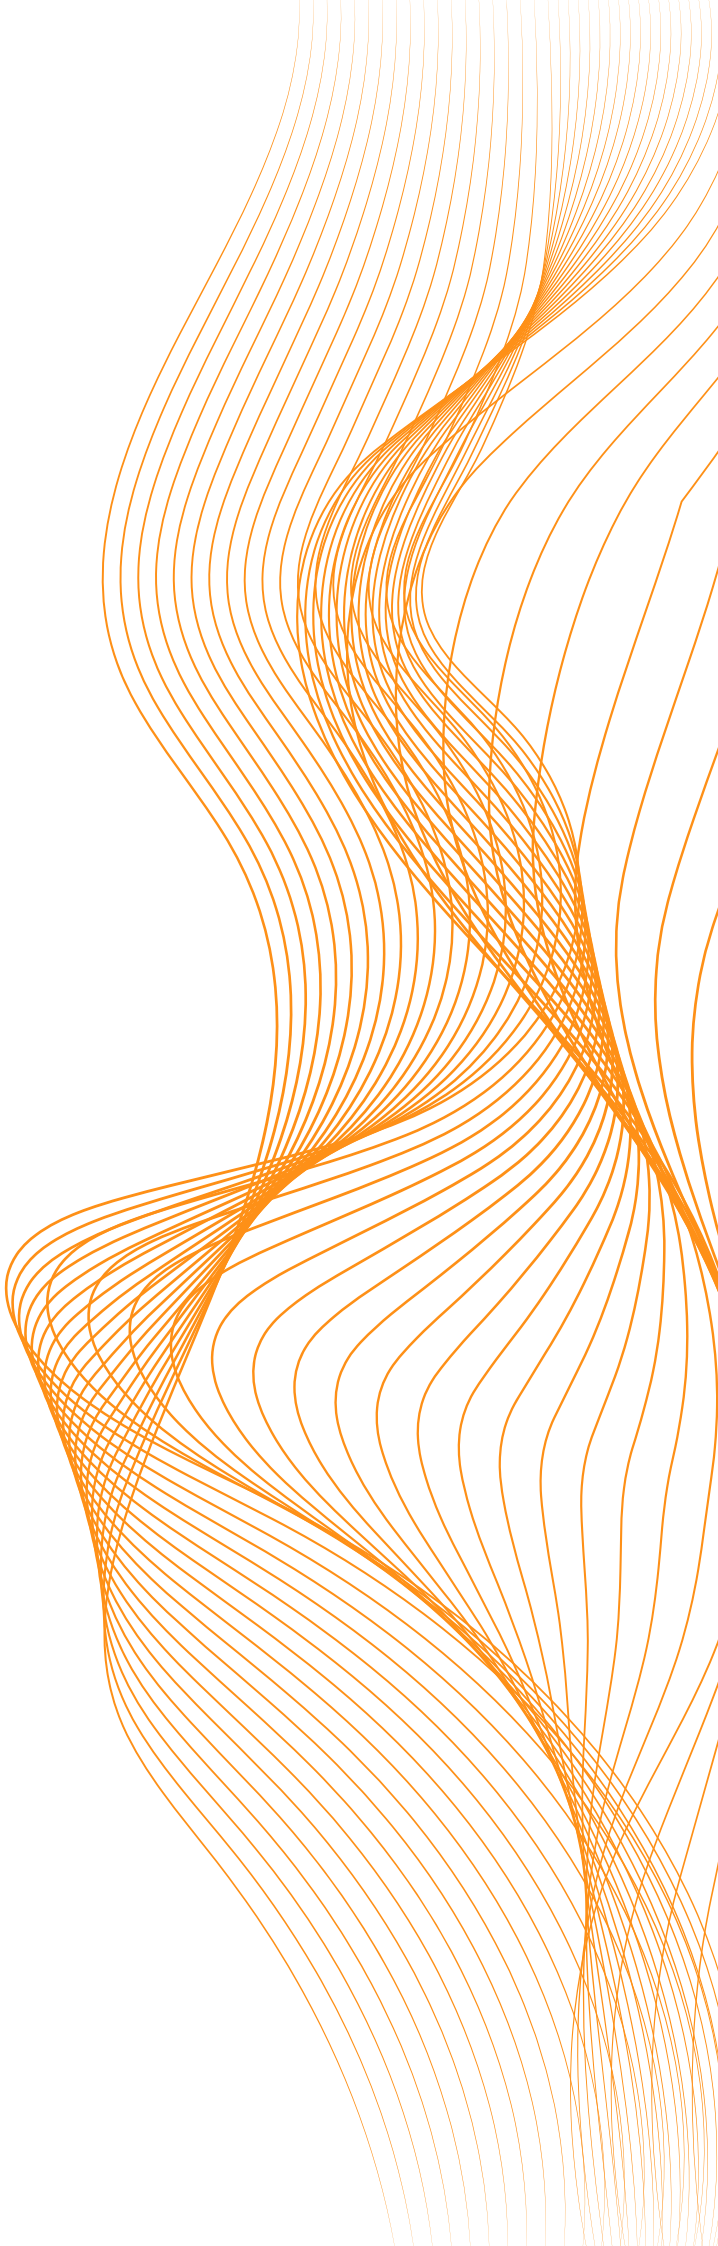

AI RESEARCH IN SINGAPORE

ESF Research Brief No.1

November 2025

ESF Research Brief No. 1

Jasper Verstappen
November 2025

Design: Ontwerp van de Buren

This publication was commissioned by the Netherlands Innovation Network Singapore and is part of the Embassy Science Fellowship programme, an initiative by NWO and the Ministry of Foreign Affairs of the Netherlands.

Content

1 Introduction

Singapore leverages AI technology to transform its economy and overcome labor and productivity constraints caused by its small geographic size and lack of natural resources. To this end, the city-state released its first National AI Strategy in 2019, charting a course to becoming a global hub for developing and deploying AI technologies. Currently, Singapore is broadly recognised as an AI leader and considered one of the most AI-ready countries in the world. This makes Singapore's industry, research institutions, and administrative bodies potentially valuable partners for collaboration. Moreover, Singapore has developed a thriving AI ecosystem that engages the national research community, the private sector, and the public sector. This experience in building a national AI ecosystem offers valuable learning opportunities.

This research project has therefore examined the priorities and strengths of Singapore's AI ecosystem by:

- examining the work done by Singapore's AI research community;
- exploring the types of AI applications being developed and deployed; and
- analysing Singapore's approach to AI governance.

2 Structure and method

This first Research Brief presents an overview of the findings with regard to the work done by Singapore's AI research community. These findings are based on semi-structured expert interviews, the Emerging Technology Observatory's (ETO) Country Activity Tracker and Map of Science, and a review of the relevant literature.[1] To this end, the Research Brief first places Singapore's AI research community within its broader AI ecosystem. It then evaluates the country's research output across five AI research areas, and finally highlights eight research clusters where the AI research community demonstrates particular strength.

1 The ETO is a nonprofit public platform, aimed at collecting, linking, and transforming information from diverse source into data resources on the global emerging technology landscape. It is an initiative of the Centre for Security and Emerging Technology at Georgetown University; Emerging Technology Observatory, *Country Activity Tracker: Artificial Intelligence,* https://cat.eto.tech; Emerging Technology Observatory, *Map of Science*, https://sciencemap.eto.tech.

3 Singapore's AI research community

AI technology can support a wide variety of applications with a broad range of purposes. It has the potential to reshape the way people interact and express themselves, how knowledge is produced, and the way business is done. This has led to high expectations regarding AI's transformative impact on both societal and economic structures.

However, these technologies are still in their early stages of development. Today's AI-based enterprise solutions typically offer incremental gains in efficiency through business automation, customer support, knowledge management, and fraud detection. A similar pattern can be seen in consumer products, where generative AI assists with everyday tasks like finding information, checking texts, and providing recommendations. Such early applications indicate that ongoing scientific progress is necessary to unlock AI's expected transformative potential.

Singapore's AI research community is therefore pivotal in further unlocking the potential of AI technologies. In addition, its focus, activities, and output are relevant beyond Singapore's national AI ecosystem. The island nation is a leader in this domain and widely regarded as a major global player and a frontrunner in both AI development and deployment. Singapore's AI research community is therefore in a unique position to drive the national AI ecosystem whilst simultaneously shaping the development of globally deployed AI technologies.

4 Singapore's AI research

Given Singapore's leading role in AI development and the importance of continued scientific innovations in this respect, this study has explored the contributions of the nation's AI research community through the Emerging Technology Observatory 's (ETO) Country Activity Tracker and Map of Science. These tools contain detailed and deduplicated information on over 270 million academic articles, drawing from both commercial and open-access platforms such as Clarivate Web of Science, arXiv, and Semantic Scholar. Additionally, it includes data from The Lens's patent dataset, which enables the tracking of academic citations by patents. These tools offer insights into AI research activities of different countries and allow for detailed analyses of the research done on a national level.

The ETO's Country Activity Tracker ranks countries based on the number of published AI-related research articles and how often these articles are cited. It covers five areas of AI research:

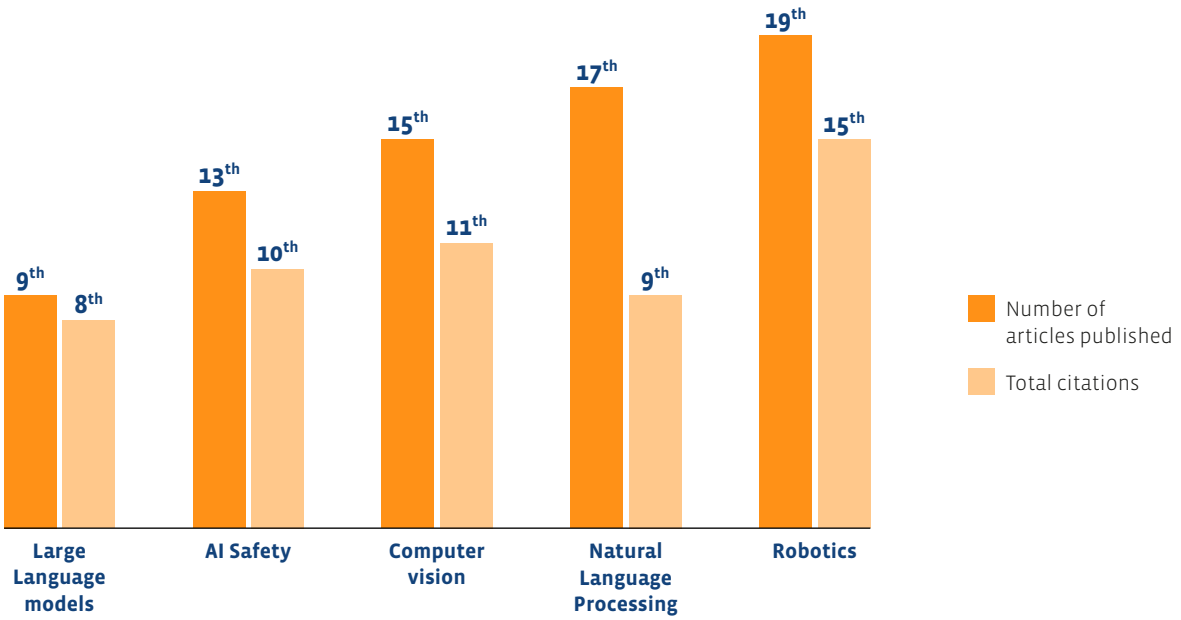

Overall, the Country Activity Tracker ranks Singapore 18th globally in the total number of AI-related research articles published and 13th in total citations.

ETO's Map of Science offers a more detailed view of the AI-related research clusters where the work of Singapore's AI research community is particularly active relative to other countries. By grouping research articles based on citations, identifying the key concepts within each group, and distilling a central theme from those concepts, the research focus of each cluster can be defined. Moreover, each research cluster is assigned an 'AI Percentage', a value that indicates the proportion of AI-focused articles within the cluster - making it possible to identify clusters primarily dedicated to AI research.[2] Lastly, ETO extracts the author's affiliated organization from each article and uses it to link the author to a specific country. This information is then used to determine the country affiliation of each cluster by measuring how often articles within a cluster are linked to specific countries. Based on this, each research cluster can be ranked according to the activity levels of authors from different countries. By identifying research clusters where Singapore ranks among the global top three, this study highlights eight AI-related areas in which the nation's research community is especially active relative to other countries.

2 An AI research cluster is one where at least 50% of research output is AI-related, based on ETO's 'AI percentage.

5 Research Areas

Human-object interaction and scene graph generation

Research in the context of this cluster explores methods for interpreting complex visual scenes in computer vision by detecting individual objects, their relationships, and their interactions. As such, this cluster is especially concerned with scene graph generation, which constructs image representations by identifying objects, their attributes, and their relationships. For example, in an image of a person with a surfboard on a rock in front of a beach, the system detects objects such as the person, the rock, the surfboard, and a wave. It then identifies attributes (e.g., the surfboard is red, the rock is large) and relations/predicates (e.g., the surfboard is being held, man standing on rock, person holding surfboard, wave near rock). Together, these form a structured representation of the scene as a graph.

An important focus within this research is human-object interaction. Research in this area is closely related to scene graph generation, as it focuses on the actions or relationships between objects and persons in a visual scene. For instance, a system may identify the person walking down the rock whilst holding the surfboard. Such research moves beyond detection-classification pipelines and in many recent works employs transformer-based architectures to model interactions more holistically, though two-stage and hybrid designs also remain widely studied. This research addresses fundamental challenges in human-object interaction and scene graph generation, such as the need for contextual reasoning and the issue of imbalanced data caused by the overrepresentation of frequent interactions. Researchers active in this cluster may work on overcoming such challenges through the application of graph neural networks, query-based pair detection, or data rebalancing techniques.

Gradient-based minimax optimization

This research cluster concerns mathematical research into gradient-based minimax optimization problems such as nonconvex–nonconcave problems. This could underpin improved optimization techniques in adversarial learning and multi-agent reinforcement learning. This cluster could support the development of machine learning solutions that better reflect the real-world scenarios in which the technology is likely to be applied. Traditional methods like gradient descent perform best under specific conditions which often do not hold in practical applications, leading to suboptimal performance. Alternative optimization algorithms explored in this cluster include, for example extragradient and optimistic gradient methods, anchoring/accelerated schemes, and simple single-loop variants that aim for stable, often faster, and sometimes last-iterate convergence.

Applying machine learning to combinatorial optimization problems

Applying machine learning to combinatorial optimization problems could ultimately prove valuable in developing solutions to a wide range of logistical and planning challenges. These include, for example, manufacturing and production scheduling, infrastructure planning, and delivery optimization. Research in this cluster is mostly focused on routing problems, such as the vehicle routing problem, which seeks to determine the most efficient paths for a fleet of vehicles to deliver goods to multiple destinations. While such problems can be solved through traditional methods, these methods often scale poorly and are typically tailored to specific scenarios. This research cluster explores the integration of machine learning with traditional combinatorial algorithms to address such inefficiencies or to develop more general solutions.

Episodic memory

The research cluster on episodic memory is mainly focussed on enabling robots to encode, store, retrieve, and express experiences in a transparent and adaptable manner, with the goal of enabling more symbiotic human–robot interactions. Such research might explore brain-inspired memory models that use hierarchical clustering and Adaptive Resonance Theory (ART)–based networks to structure and retrieve memories contextually, thereby effectively encoding experiences as clusters of semantic elements (such as people, activities, places, times, and objects). Other research might instead focus on enabling robots to express and explain their behavior using natural language. By combining insights from research on robot verbalization and deep episodic memory architectures, this line of research could enhance a robot's ability to learn from experiences. This could also enable robots to express and explain those experiences in a meaningful way, thereby improving transparency and explainability.

Biohybrid insects

This final cluster works on integrating electronic systems into living animals to form cy-borg organisms. Most of the research done in this context focuses on insect–computer hybrid systems. Such biohybrids use the wings, muscles, or neural pathways of insects for movement and environmental interaction. By using the natural physiology of insects, this cluster aims to overcome the limitations related to size, power, or manoeuvrability of fully artificial micro-robots. For example, cyborg cockroaches equipped with sensors and wireless control units have been tested in controlled environments to navigate through obstacles, providing proof-of-concept data for search-and-rescue operations. Research in this cluster often focuses on neuromuscular stimulation for fine motor control as well as finite-state-machine–based navigation and closed-loop feedback control. This research could lead to practical applications of biohybrid robots in search and rescue operations or environmental monitoring.

Collaborative semantic mapping and SLAM

This research cluster focuses on enabling autonomous robots to collaboratively create maps of their environment that are both consistent and semantically rich. This cluster builds on traditional SLAM research by emphasising multi-agent settings. For example, multiple autonomous robots can be used to explore a disaster site. Each robot can explore a different section, and by sharing their data, they can create a single consistent map of the area that highlights key features such as collapsed buildings and blocked roads.

Key challenges focussed on by research within this cluster relate to perceiving, sharing, and merging geometric and semantic data in real-time. This cluster is particularly concerned with challenges that are unique to multi-agent systems, such as limited communication, unstructured or dynamic environments, and heterogeneous sensor configurations. A central theme of this research is the alignment and integration of maps produced by different agents. Because each agent contributes a unique perspective, achieving a consistent shared representation is inherently difficult. To address this, the research in this cluster develops hierarchical and probabilistic semantic mapping frameworks that support more reliable alignment across agents. This challenge is often exacerbated by differences in sensors, sensor inputs, and semantic interpretations. Research in this cluster explores solutions such as map fusion algorithms, deep learning–based segmentation, distributed SLAM architectures, and multimodal sensor fusion.

Event extraction

The process of understanding text involves converting it from an unstructured form into a structured representation of events. This research cluster develops techniques for detecting, identifying, and describing such events. This process, known as event extraction, is essential in natural language processing because it allows AI to answer questions, retrieve information, and build knowledge bases. Research conducted within this cluster aims to develop more flexible, scalable, and generative approaches that often operate at the document level as well as the sentence level, and aim to reduce but not completely remove dependence on annotated data. For instance, from the sentence ‘she is figuring out if the waves are decent for longboarding,’ an event extraction system might identify the event (evaluating surf conditions), the trigger word that signals the event, and the roles such as Agent (she), Object (waves), and Context (longboarding), with the understanding that roles vary by event type. Document-level event extraction could combine contextual information, linking who is evaluating the conditions, what kind of waves are being considered, and which type of board, to build a richer, structured account of the event while reducing rather than eliminating reliance on annotated training data. This research explores generation-based and sequence-to-structure methods, as well as the use of prompt learning and prefix tuning to incorporate type-specific or contextual information into pre-trained models.

Sentiment analysis

The research cluster on sentiment analysis focuses on computationally identifying and classifying the subtleties of human emotions. It does so by moving beyond traditional three-class models that categorize text as positive, neutral, or negative. This line of research explores solutions inspired by human cognition, enabling a more nuanced, human-like understanding of emotions that can recognize the intensity of emotions and incorporate contextual nuances.

A common theme in this research is the development of explainable sentiment analysis tools. One approach involves using a multi-level modular structure that mirrors the process by which humans understand language. The integration of multiple knowledge bases, which serve as structured repositories of information for guiding models' decisions, makes it possible to perform fine-grained, multiclass sentiment classification. By using identifiable components, this method makes it possible to explain the reasoning behind a particular sentiment assignment. For instance, when analysing the sentence 'he was looking forward to surfing, but almost slipped when walking down the rock,' the system can separate the positive sentiment associated with looking forward to surfing from the negative sentiment tied to almost slipped. Drawing on knowledge bases, it identifies both emotions, classifies the overall sentiment as mixed, and can explain that the positive classification comes from anticipation of surfing while the negative classification arises from nearly slipping.

An alternative approach generalizes language into a form of 'proto-language' by converting input into semantic primitives. These primitives are then linked to emotion and polarity labels via a commonsense knowledge graph using the Hourglass model. This neurosymbolic approach ensures explainability by making the generalization process transparent and linking final classifications to both emotional states and the underlying concepts that express them.

6 Key insights

- Singapore's AI research community, a key pillar in driving the nation's AI ambitions, is particularly active in a select yet diverse set of research clusters.
- The output of Singapore's AI Research Community in terms of the number of published AI-related articles and citations scores particularly well in the areas of Large Language Models and AI Safety, while ranking top-20 globally for all 5 areas of AI research.
- Singapore's AI research is particularly active in 8 research clusters, seven of which concern computer science and one of which concerns mathematics.

References

Emerging Technology Observatory, *Country Activity Tracker: Artificial Intelligence*, https://cat.eto.tech.

Emerging Technology Observatory, *Map of Science*, https://sciencemap.eto.tech.

Nestor M et al, Artificial Intelligence Index Report 2025 (2025), available at: https://arxiv.org/pdf/2504.07139.

Oliver Wyman Forum, *Global Cities AI Readiness Index* (Oliver Wyman Forum, 2019), available at: https://www.oliverwymanforum.com/city-readiness/global-cities-ai-readiness-index-2019.html (last accessed 20 August 2025).

Oxford Insights, Government AI Readiness Index 2022 (2022), available at: www.unido.org/sites/default/files/files/2023-01/Government_AI_Readiness_2022_FV.pdf (last accessed 20 August 2025).

This is a publication of

Netherlands Enterprise Agency
Prinses Beatrixlaan 2
PO Box 93144 | 2509 AC The Hague
T +31 (0) 88 042 42 42
Contact
www.rvo.nl

This publication was commissioned by the ministry of Foreign Affairs.

Publication number: RVO-029-2026/RP-INT

NL Enterprise Agency is a department of the Dutch ministry of Economic Affairs that implements government policy for Agricultural, sustainability, innovation, and international business and cooperation. NL Enterprise Agency is the contact point for businesses, educational institutions and government bodies for information and advice, financing, networking and regulatory matters.

Netherlands Enterprise Agency is part of the ministry of Economic Affairs.