Robot Dexterity — Handling our future
Programme thesis

v2

Jenny Read, Programme Director

CONTEXT

This document presents the core thesis underpinning a programme that is currently in development at ARIA. We share an early formulation and invite you to provide feedback to help us refine our thinking.

This is not a funding opportunity, but in most cases will lead to one. Sign up **here** to learn about any funding opportunities derived or adapted from this programme thesis.

An ARIA programme seeks to unlock a scientific or technical capability that

+ changes the perception of what's possible or valuable

+ has the potential to catalyse massive social and economic returns

+ is unlikely to be achieved without ARIA's intervention

ὥστε ἡ ψυχὴ ὥσπερ ἡ χείρ ἐστιν.

hence the soul is like the hand. — Aristotle, De Anima, III part 8, c. 350 B.C.E.

PROGRAMME THESIS, SIMPLY STATED

This programme thesis is derived from the ARIA opportunity space: Smarter Robot Bodies. *The original version of the thesis was published in February 2024 and has been updated following feedback from the community. To read the original version of the thesis please* **see here**.

Modern civilisation was built by human hands, the dexterity of which continues to underpin a great deal of the physical work in our lives and society. Until we create cost-effective dexterous robotic manipulators, general purpose automation of tiresome, dangerous, and otherwise unfavourable human labour will remain out of reach. This programme aims to create a novel robotic manipulator with world-leading performance, able to cost-effectively perform routine human tasks, leading to a step-change in human productivity and welfare.

Despite steady progress, general dexterous manipulation remains an unsolved problem in robotics. Key challenges include handling previously unseen objects, including delicate and deformable items, in a variety of lighting conditions, while avoiding error and damage over long periods of time. Advances in AI and machine learning are poised to produce significant improvements in robotics, but their impact on dexterity will be limited without comparable advances in hardware. Brute force, computationally intensive control of rigid structures can only get us so far.

This programme will focus on improving robotic dexterity primarily through advances in hardware. We plan to support development of new modes of sensing, transmission of sensory information, and actuation through hardware advances that benefit from co-design and integration with advanced software and controls.

In the early stages of the programme, we anticipate funding advances in individual components, e.g. actuation or sensing, in isolation. In later stages, we would combine advances made both within and beyond the programme to develop new manipulators, demonstrating a paradigm-shift in robotic abilities and establishing the basis for a powerful new industry that can help society better address the labour challenges of tomorrow.

PROGRAMME THESIS, EXPLAINED

A detailed description of the programme thesis, presented for constructive feedback.

Why this programme

Dexterous manipulation is a critical bottleneck to the wide adoption of robotics. The ability to deftly manipulate objects with a wide range of properties would enable automation of routine tasks across a wide range of sectors, as outlined in Table 1. Such tasks, often burdensome for human workers, range from repetitive and injury-prone to mundane and low-paying, and often occur in hazardous environments like sewers, factories, chemical plants or recycling facilities. Automating these tasks promises to reshape society by increasing economic productivity while freeing up humans for more rewarding tasks[1,2]. The benefits are particularly large in the UK due to our ageing population and low productivity. One study estimated that robot density in the UK warehouse logistics sector could grow from 3.3 robots per million hours worked in 2020 to 350 by 2035, increasing labour productivity by 25%[2].

Robots today are largely limited to highly controlled environments (e.g. factories, warehouses) or highly specialised tasks (e.g. vacuum cleaning, lawn mowing). Advances in AI, however, are enabling robots to venture into more complex and challenging environments, expanding the market and accelerating demand for robots capable of versatile tasks. In the context of manipulation, this will lead to demand for robots capable of handling objects with a wide range of properties — heavy, deformable, delicate, damp — with minimal damage and error rates, low power demand, and without noisy infrastructure such as air compressors. They will need to be robust and damage-resistant so that they can operate for long periods without requiring specialist repairs or maintenance. Figure 1 lays out some of the key requirements and bottlenecks.

Recent developments in AI, including in reinforcement learning[3] and the use of multimodal LLMs to improve scene understanding,[4] will help to increase generalisability and adaptability. However, truly unleashing the potential of robotics will require a paradigm shift from brute-force computation to more sophisticated hardware closely integrated with control. These ideas are already being explored in university labs and start-ups but need leadership, community-building and further investment to move from proof-of-interest at a component level to proof-of-value at a systems level.

Calling for a paradigm shift

Most robotics today follows what we will dub the **Genesis Paradigm**. In the Book of Genesis, God forms Adam's body from the dust of the ground, then animates him with the breath of life. Similarly in robotics, mechanical and electrical engineers design and build hardware, which is then animated either by human tele-operators or by algorithms designed by computer scientists[5]. In both cases, the body is treated as something quite distinct from the intelligence that controls it.

The Genesis Paradigm has encouraged the view that if a human can use a robot to execute an intended task, then the robotic hardware is demonstrably suitable for the task and the focus should shift to improving control software. However, this perspective overlooks a critical nuance: the human brain's remarkable power and adaptability. The human brain can navigate and overcome hardware limitations to accomplish desired tasks. This risks sending robotics down a rabbit hole of throwing ever more complex and expensive compute at a problem which likely has far simpler alternatives. In the past, this approach has paid off because compute was continually becoming cheaper and more powerful — but as Moore's Law comes to an end, it is looking increasingly unsustainable[6]. To truly unlock progress in robotic manipulation, we will need to move beyond the Genesis Paradigm.

Biology suggests that a different approach is possible. Biological organisms operate successfully with noisy, imprecise hardware and long, highly variable sensorimotor latencies (25ms for some proprioceptive reflexes, 200ms for saccades[7] in contrast to the high frequencies and low latencies (1ms) typical of robotic control. Biology is also far more economical with compute: even insect brains with around a million neurons can perform many complex tasks such as flight, foraging and object manipulation[8]. These results are possible because biology follows what we can term the **Darwin Paradigm**, the joint evolution of biological bodies and nervous systems optimised for a particular ecological niche[9]. This enables aspects of control to be effectively built into the mechanical properties of the body[10].

The Darwin Paradigm could be considered "bio-inspired robotics". However, it is important to point out that it does not mean copying animal bodies which have been selected for performance on a far wider range of tasks than even the most advanced robot and have been built from biological components not available to the roboticist. Instead, it is the biological approach of co-evolved design that holds promise. Just as evolution has produced a wide range of manipulator designs — from the single pincer-gripper used by many bird species to construct elaborate nest structures[11] to the sucker-cups used by octopuses to open clam shells — a similar approach in robotics could produce designs never seen in biology. For example, robotic manipulators could exploit reversible adhesion such as is used by geckos to climb walls, or have "eyes in their fingertips" to collect visual information locally, or use electromagnetic proximity sensors instead of whiskers. We now have unparalleled capabilities through advanced compute and generative AI to implement the Darwin Paradigm and succeed at novel co-evolution.

Sector	Example tasks requiring object manipulation	Environment complexity/ unpredictability	Robot density		Most important metrics	Comments
			Current	Potential increase via manipulation		
High-throughout manufacturing	Machining, welding, painting	Very low	High	Low (already highly automated)	Speed Accuracy Robustness	Limited scope for more automation
Warehousing/ logistics	Packing, shelf-stacking	Low	High	High	Cost Accuracy Adaptability Robustness	Much scope for more automation
Recycling	Sorting	Low	Low	High	Cost Robustness	Much scope for more automation
Horticulture (indoors)	Picking fruit, transplanting seedlings	Low	Low	Medium	Cost Adaptability Robustness	Seasonal, rural, low-paid
Food processing, preparation	Assembling sandwiches, butchering meat	Low	Medium	Medium	Cost Adaptability Robustness Speed	Limited scope for off-shoring
Extreme environments	Repairing nuclear reactor, retrieving satellite	Medium	Medium	Medium	Error rate Reliability Resistance to damage	Need to move from tele-operation to autonomy
Heavy industry	Turning valves in chemical plant	Medium	Low	Medium	Accuracy Robustness Adaptability	
Laboratory	Preparing samples	Medium	Medium	Medium	Speed Adaptability Robustness	Automation could enable AI to design and perform experiments
Surgery	Suturing, tissue dissection	High	Low	Medium	Accuracy Adaptability	Data available from human tele-operators, but regulatory approval challenging
Garment-making	Sewing shirts	Low	Low	Low	Cost Adaptability Robustness Speed	Hard to overcome cost advantage of low-paid off-shore labour
Household	Cleaning surfaces, tidying, cooking	Highly complex, variable and uncontrolled	Low	Low	Cost Adaptability Accuracy Robustness	Hard to move from specialised machines to generalised autonomy
Personal care	Lifting, dressing	Highly complex, variable and uncontrolled	Low	Low	Accuracy Adaptability Social acceptability	Even to assist rather than replace human care, exceptionally challenging

Table 1: Possible applications of robotic manipulation now and in the future. Intended to stimulate discussion rather than to be definitive.[2]
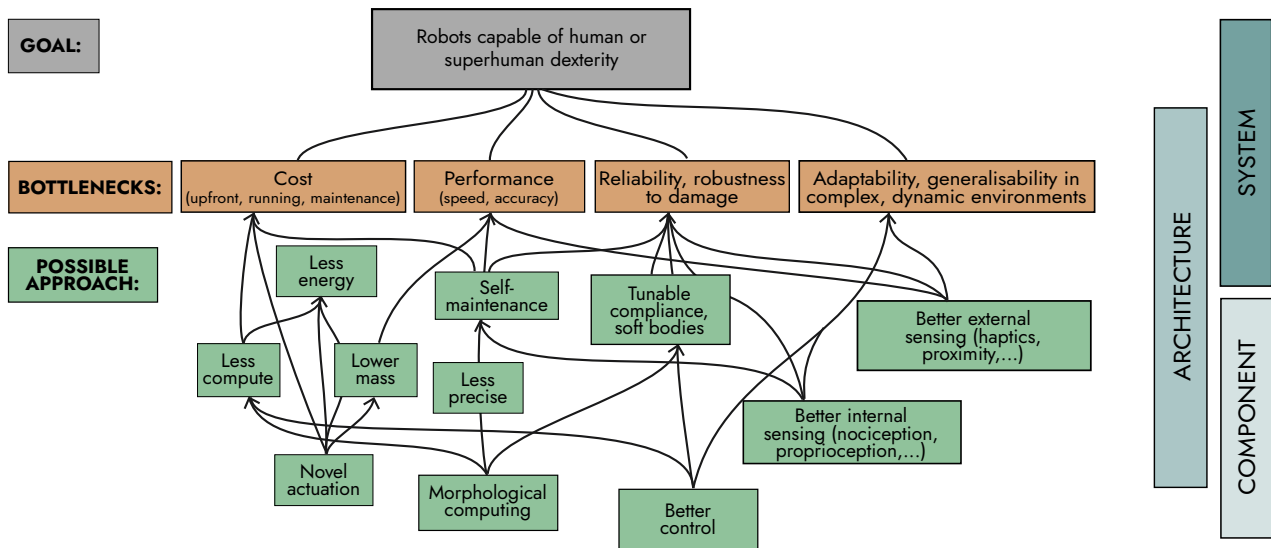
Figure 1: Bottlenecks preventing us from reaching the goal of widespread dexterous manipulation and possible routes to opening them up. On the right: mapped onto the three levels of Figure 2.

What we expect to fund

We expect to fund R&D Creators (individuals and teams who receive ARIA funding), who will

+ Create one or more novel robotic manipulators, demonstrating a dexterous ability that far exceeds what's possible today or likely to be achieved by existing approaches. Realise substantial improvements over the status quo in both performance and robustness, while not introducing any deal breakers in terms of cost, size, infrastructure or scalability.

+ Develop new techniques for designing robotic hardware and control software.

+ Produce advances in relevant technologies such as actuation and haptic sensing.

Figure 2 shows a systems-engineering analysis of the programme. When building a manipulator, it's important to be very clear about the goal, task requirements and any constraints on the solution (System level). However, we are also keen to fund novel components or materials which could have very broad applicability and are not necessarily limited to dexterous applications, or indeed to robotics at all (Component level). Accordingly, we expect to fund work addressing one or more of the Technical Areas labelled in Figure 2.
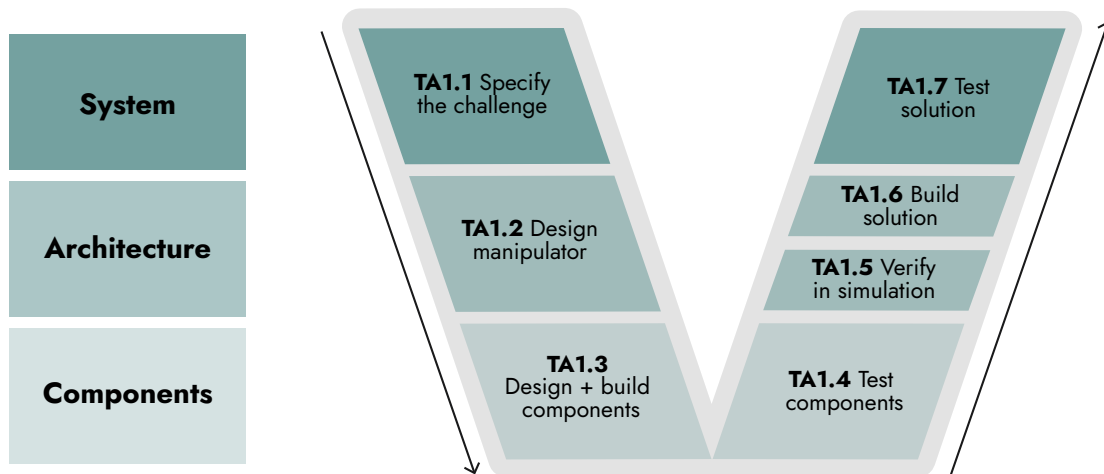
Figure 2: Systems-engineering V diagram of the programme. The seven Technical Areas represent the conceptually distinct Contributions. Many Creators will contribute to Several Technical Areas.

How we expect to fund

We envisage that, as a minimum, a Creator team will usually work on both right and left arms of the V at a given level, and so cover pairs of Technical Areas at the same level (Systems/Architecture/Components, e.g. TA1.1+TA1.7). Some larger groups may also wish to work across multiple levels.

ARIA programmes primarily operate by funding people toward a clear objective, actively directing and coordinating projects towards this goal. Thus whether or not they formed part of the same group on application, it will often be essential for teams at different levels to work closely during the programme. For example, Component teams (TA1.3+1.4) will share with Design teams (TA1.2+1.5) the metrics they hope to be able to achieve, while Design teams may be able to inform the metrics targeted by Component teams, e.g. "our simulations suggest that increasing receptor density beyond x / mm2 will produce no further reduction in error rate".
We will also consider the fit between team types when assessing applications.

Given this, we expect that applications will usually combine Technical Areas in one of the four patterns A-D laid out below. We are open to other possibilities, except that we do not expect to fund Creators focusing solely on the right-hand side of the V diagram (TA1.4, TA1.5, TA1.6 or TA1.7).

### Pattern A. Challenge specification (TA1.1 + 1.7)

The engine of evolution is selection, and that requires a definition of fitness. Thus, we will need a clear definition of the particular manipulation problem to be solved, along with any constraints. In TA1.1, Creators define a high-value manipulation challenge in at least one domain which can inform the efforts of design and component development. In TA1.7, they test out proposed solutions to this challenge.

At this stage we are not ruling out any sectors or use-cases, other than defence/military. The challenge should have clear potential for social benefit, e.g. the ultimate application envisioned should not be "toy" or excessively niche. Challenges may be very specific, e.g. a surgical end-effector for a particular operation, or very general, e.g. a manipulator capable of all household tasks. However, even a "general purpose manipulator" will necessarily have limits (e.g. on maximum load, power or tolerance to environments) which will bound its applicability, and these should be specified.

Teams who apply in Pattern A do not themselves propose or create the solution. Such Creators will participate in the programme by developing specifications and benchmarks for the task, and/or through testing solutions built by other Creators. TA1.1+1.7 teams will include members who are themselves end-users of the technology, or at minimum be working closely with end-users.

Pattern B. Integrated solutions (TA1.1 - 1.7)

Some applicants may wish to propose an integrated solution encompassing all the technical areas - from the challenge and also the design and build of a novel manipulator to solve it. We expect that this would usually involve novel components as well.

Pattern C. Novel techniques for robotic design (TA1.2+1.5, perhaps also 1.6)

At the heart of the Darwin Paradigm is the co-evolution of brains and bodies. Thus we plan to invite proposals for exploiting improvements in simulation and compute to advance techniques for co-designing hardware and software, and use these to design highly capable robotic manipulators. Such Creators may work entirely in silico; if they do, we will expect to team them with other Creators who can realise their designs.

Alternatively, Creators may wish to include rapid prototyping and/or final build (TA1.6) as part of the process. We will expect Creators working on TA1.2+1.5 to work closely with those at other levels, and will help them collaborate where they were not part of an integrated team at application, as in Pattern B.

Pattern D. Novel components (TA1.3+1.4)

Applicants may also request funding only to build and test novel components which will aid robot dexterity, whether because they are useful for dexterity specifically or for robotics in general. These could be haptic or other sensors, novel actuators or artificial muscles, novel materials or anything else, provided that a case can be made for how they will benefit dexterity.

Additional principles that will guide our approach to funding in this programme

Our goal of producing transformative change in robotic manipulation within a five-year timeframe is the yardstick against which we measure all aspects of the programme. This has a number of implications:

1. Delivering ambitious projects on a short timescale requires a high level of commitment and focus. We will want to be convinced that the right structures are in place to drive the project forward. For example, each workstream on an application could have a named project lead, for whom the ARIA workstream is their number one priority and occupies the overwhelming majority of their time. To enable this for university-led applications, we are keen to explore structures not typical in academic research such as supporting early career researchers as project leads, exploring secondments or funding >80% of senior academics' time so that they can focus fully on their ARIA project.

2. Instead of a system where applications are reviewed in isolation and then either awarded or not, we envisage a more iterative process where researchers propose ideas and we help shape projects so they make the optimum contribution to the programme goal. This includes taking into account other projects within the programme. Thus we envisage a two-stage application process, in which applicants initially submit a 3-page concept paper. Following evaluation of these, a subset of applicants will be encouraged to submit a more detailed proposal incorporating ARIA feedback.

3. In terms of outputs, we will value results over academic papers. We will strongly encourage open publication of methods, results and code where this is consistent with researchers' own IP strategy. ARIA staff will review these and offer assistance with, e.g., documentation and packaging of data and code. While traditional peer-review or journal articles are not discouraged, they are not a goal of the programme and will not be how we will assess success.

SOURCES

References cited in this document.

1. Oxford Economics (2019). How robots change the world. Available at: https://www.oxfordeconomics.com/wp-content/uploads/2023/07/HowRobotsChangetheWorld.pdf

2. Department for Business, Energy & Industrial Strategy (2021). The economic impact of robotics & autonomous systems across UK sectors. Available at: https://assets.publishing.service.gov.uk/media/6193996bd3bf7f055b293381/ras-final-report-nov-2021.pdf

3. Ibarz. J, Tan. J, Finn. C, Kalakrishnan. M, Pastor. P, Levine. S (2021). How to train your robot with deep reinforcement learning: lessons we have learned. The International Journal of Robotics Research 40(4-5):698-721. doi:10.1177/0278364920987859

4. Google Deepmind (2023). RT-2: Vision-Language-Action Models Transfer Web Knowledge to Robotic Control. Available at: https://robotics-transformer2.github.io

5. Ackerman E (2024). That Awesome Robot Demo Could Have a Human in the Loop. IEEE Spectrum. Available at: https://spectrum.ieee.org/amp/robot-teleoperation-autonomy-2667060864

6. Shalf J (2020) The future of computing beyond Moore's Law. Phil. Trans. R. Soc. A378. doi:20190061

7. Codol. O, Kashefi. M, Forgaard. C, Galea. J, Pruszynski. J, Gribble. P (2023) Sensorimotor feedback loops are selectively sensitive to reward. eLife, 12:e81325 doi: https://doi.org/10.7554/eLife.81325

8. Science Magazine (2018). Bees have more brains than we bargained for. [Online video]. Available at: https://www.youtube.com/watch?v=fJyHeUaSMVw.

9. Pfeifer, R. Bongard, J. (2006). How the body shapes the way we think: a new view of intelligence. The MIT Press

10. Hauser, H. Füchslin, R. Pfeifer, R. (2024) E-book on Opinions and Outlook on Morphological Computation [Online]. Available at: https://www.morphologicalcomputation.org/e-book

11. Sheard. C, Street. S E, Evans. C, Lala. K N, Healy S D, Sugasawa S. (2023) Beak shape and nest material use in birds. Phil. Trans. R. Soc. B378: doi:20220147.

12. Gesslbauer, B., Hruby, L.A., Roche, A.D., Farina, D., Blumer, R. and Aszmann, O.C. (2017). Axonal components of nerves innervating the human arm. Annals of Neurology, 82(3), pp.396–408. doi: https://doi.org/10.1002/ana.25018

13. Allen, R. (2017) Lessons from history for the future of work. Nature. 550, 321–324. doi: https://doi.org/10.1038/550321a

14. Dalrymple, D. (2024) Safeguarded AI: constructing safety by design. Advanced Research and Invention Agency. Available at: https://www.aria.org.uk/wp-content/uploads/2024/01/ARIA-Safeguarded-AI-Programme-Thesis-V1.pdf

ENGAGE

Our next step is to launch a funding opportunity derived or adapted from this programme thesis.

Click **here** to register your interest, or to provide feedback that can help improve our thinking.

If you require an accessible version of this document and/or form, please contact us at **info@aria.org.uk**.