

Robotic Dexterity – Handling our future

Programme thesis

v1.0

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 χεὶρ ἐστίν.

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: [Smart machines need smarter bodies](#).

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

Manipulation is the number one 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.

Several big tech companies are working on robots of unprecedented ability, building on recent developments in AI such as imitation and reinforcement learning^[3], and the use of multimodal LLMs to improve task planning^[4]. These efforts will doubtless be transformative, but ultimately these robots will reach their limit, because their hardware will not be sufficient to perform more dexterous, fine-grained, or high-torque tasks in a robust and energy-efficient manner. This programme aims to address this future bottleneck (Figure 1). We argue that 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 be definitive! [2]

What we expect to fund

We intend to fund work that can revolutionise manipulation by enabling the Darwin Paradigm, with the ultimate goal of demonstrating an ability to cost-effectively perform unskilled manipulation tasks that are undesirable for humans. We see this as requiring efforts across three broad areas, which we'll call Solutions, Design and Components.

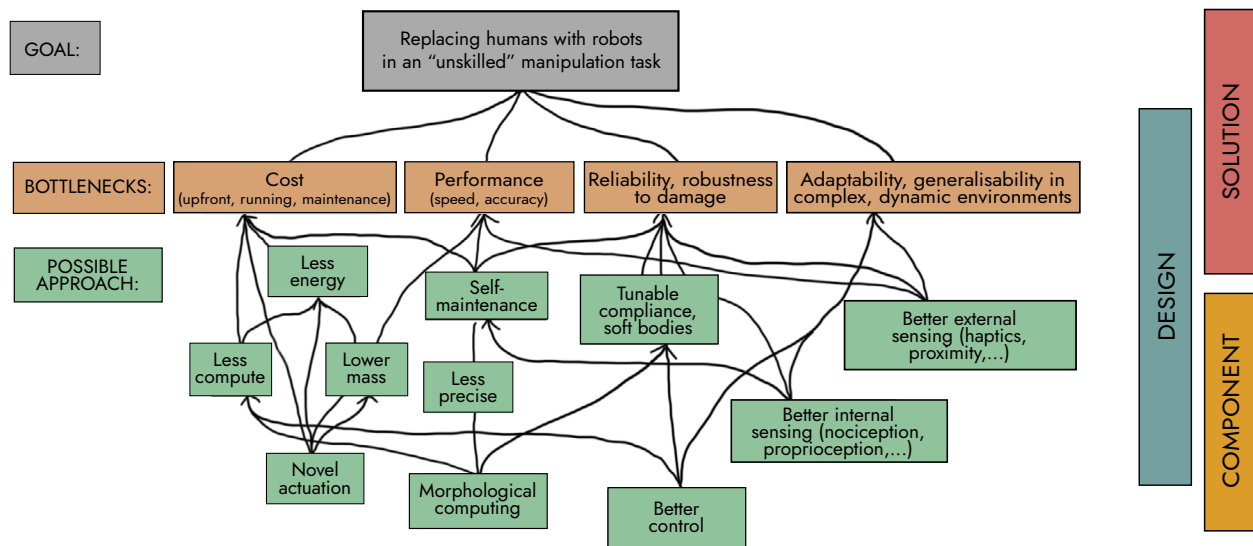


Figure 1. Bottlenecks preventing us from reaching the goal of widespread dexterous manipulation and possible routes to opening them up. On the right: where the three types of teams would mainly operate.

Solutions

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. We plan to invite Solution teams to define a high-value manipulation challenge in at least one domain which can inform the efforts of design and component development. Solutions teams will then be responsible for constructing a next-generation manipulator based on component and design advances, ultimately testing the proposed solution in an application-relevant environment. Solutions teams would include members who are either themselves end-users of the technology, or are working closely with those who are.

- + Solutions applicants will specify a set of tasks their desired manipulator should solve, along with any related constraints.
- + The tasks should have clear potential for social benefit, e.g. the ultimate application envisioned should not be "toy" or excessively niche.
- + The chosen tasks will determine the success metrics, e.g. how speed and accuracy should contribute to the definition of performance, as well as what key challenges must be overcome at both component and system levels.
- + Solution and Design teams will work together to design a manipulator to perform this task, using both standard components and those developed by Component teams.
- + In later stages of the programme, they will build and test a fully functional manipulator.

Design

At the heart of the Darwin Paradigm is the co-evolution of brains and bodies. Thus we plan to fund Design teams exploiting improvements in simulation and compute to advance techniques for co-designing hardware and software.

- + Design teams will work on an integrated approach to hardware and software design, e.g. simulated evolution that simultaneously learns how to design a body and how to control it to perform specific task(s) subject to constraints, e.g. minimising number of actuators needed, maximum force required, etc.
- + We expect these approaches will be applicable in robotics generally, although the focus here will be manipulation.
- + Design team efforts will draw information and inspiration from Solution teams to optimise their approach and toolsets.

Components

Biological evolution learns to produce new materials as well as body designs and control, but we can't rely on that process here. We intend to side-step literal evolution by funding Component teams to improve the materials and components available for Design and Solutions teams to draw on. While we have argued against copying particular manipulators, we do think it is appropriate to copy principles drawn from biology, whose ubiquity suggests that they have been found profoundly beneficial in every ecological niche.

- + Component teams will work on individual hardware components with potential to improve dexterity, such as haptic or other sensing, novel actuators, materials.
- + We see promising work in all these areas, but it is often not clear whether particular technical barriers, such as low durability or need for high voltages, can be overcome with more work, or are insuperable.
- + We anticipate making awards aimed at (a) answering these questions for individual technologies, and (b) if the answer is favourable, bringing components towards the point where they could be integrated into a novel manipulator.
- + Many of these technologies will have value in robotics generally and some are likely to have even broader application.

It may help at this point to provide an illustrative example of how the Darwin Paradigm might inform solutions that rely on advances in both design and component technologies. Let us simply consider the sensing functionality of robotic manipulators. At present, manipulation is guided almost solely by vision-based sensors. In biology, the key advantage of vision is its ability to relay information from a distance, whereas interactions with nearby objects are guided largely through somatosensory information obtained via skin touch or whisking. For example, human control of the arm and hand uses input from ~10 sensors for every 1 independently-controllable actuator^[12] – in today's robotics manipulators, this ratio is more like 2:1. The fact that evolution has universally produced animal bodies densely covered in a rich diversity of external and internal sensors suggests that we could enhance robotics' performance, robustness and adaptability by doing the same.

Maximising Social Benefit

Mechanisation and automation has freed billions from a life of toil. However, history also warns us of adverse consequences when these benefits are not shared equitably across society^[13]. AI-driven autonomous machines bring concerns beyond those of previous automations. Thus, in parallel with the technical funding described above, we intend to fund research in the social sciences and humanities to examine potential outcomes of this programme. Our aim will be to produce useful information for governments, researchers and the public regarding the likely societal risks and challenges of technology developed by this programme, and how these can be mitigated to minimise adverse impacts while maximising their benefit^[1]. Note: ARIA is running a separate programme dedicated to AI safety,^[14] so any research on social impacts will focus on robots operating as intended.

How we expect to fund

ARIA programmes primarily operate by funding people toward a clear objective, actively directing and coordinating projects towards this goal. Whether or not they formed part of the same group on application, it will be essential for all three types of teams described above to work closely during the programme (Figure 2). For example, Component teams will share with Design teams 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 / mm^2 will produce no further reduction in error rate". We will also consider the fit between team types when assessing applications. Applicants may come from all sectors including large companies, start-ups, SMEs, universities and/or other bodies such as Research Technology Organisations or Focused Research Organisations.

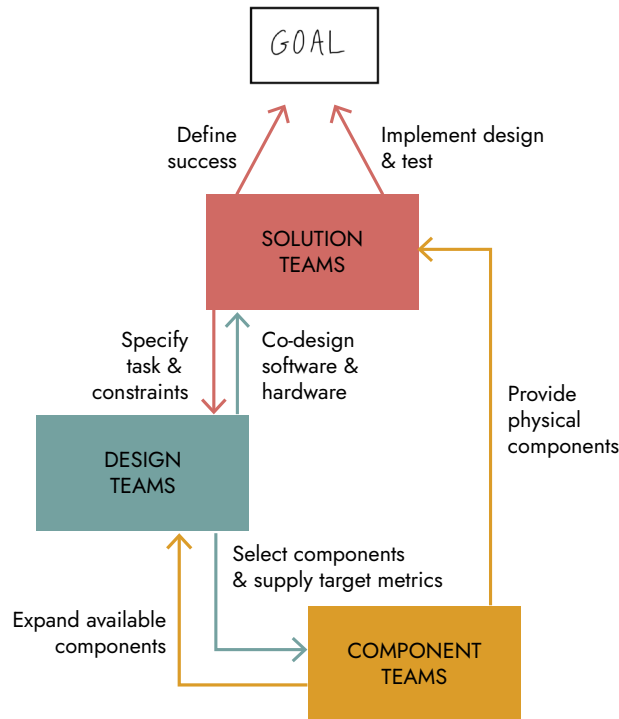


Figure 2. How the different team types will need to interact

To facilitate this close collaboration, we envisage launching calls for all three team types at once (Figure 3). Solution teams will be part of the programme from the beginning and will communicate with Component and Design teams throughout (Figure 2). Applications may contain only one type of team, or multiple, or applicants may indicate an intention to work initially as a Design and/or Component team and then subsequently as a Solution team.




Figure 3. One possible structure for an eventual programme. The columns represent 9 successful applications. Six groups work as Component Teams; one mainly as Component but with a Design element; one as a Design and Solution; and one works as Solution. Horizontal lines represent a go-no-go decisions.

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 therefore will require each application to identify a named project lead for whom this project will be their number one priority and occupy 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 multi-stage application process, beginning with a brief expression of interest, of which a subset of applicants will proceed to the development of a 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-reviewed journal articles are not discouraged, they are not a goal of the programme and will not be how we will assess success.

What we are still trying to figure out

The purpose of this document is to attract constructive feedback to guide programme structure. We invite readers to point out challenges we may not have thought of. Specifically, there are a number of questions that remain in shaping this programme, including:

- + How should we judge success in individual component technologies? What are meaningful metrics for actuators, force/torque sensors, haptic sensors etc, and what numbers are needed to be transformative? Is this even a question we can answer without designing them into a robot for a particular task?
- + To what extent should work focus on specialised manipulators that excel in one particular domain, versus a general manipulator capable of replacing human hands in a wide range of tasks?
- + Is simulation currently good enough to enable the co-design of hardware and software that we envisage? What AI techniques should be used? How much compute would be required to train and learn?
- + What are plausible budgets (£M/year) for the different types of project?
- + How can we best work with existing companies in this space to understand the needs and produce breakthroughs which will ultimately be applied in society?
- + How can we best engage people not currently working in robotics, e.g. in materials science or electrical engineering, to contribute?
- + How can we optimally engage and support women and other researchers who are woefully underrepresented in the robotics field?

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, Pruszyński. 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. Fuchslin, 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., Hruba, 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/AR-IA-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**.
