Final Report

AMPC
AUSTRALIAN MEAT PROCESSOR CORPORATION

Meat Industry Humanoids: Stage 1 Trial PUBLIC FINAL REPORT

Project code
2026-1051

Prepared by
D. Nie, A. Shahzeb, R. Vanzella, J. Cook

Date submitted
30/10/2026

Published by
Intelligent Robotics

Date published
2/11/2026

Contents

Contents 2

1.0 Abstract 4

2.0 Executive summary 5

3.0 Introduction 6

4.0 Project objectives 7

5.0 Methodology 7

6.0 Discussion 8

6.1 Robot Specification and Humanoid Ordering 8

6.2 Platform procurement, setup and operating baseline 10

6.4 Teleoperation stack and demonstration-data capture 11

6.5 High-level workflow overview 13

6.6 Experimental features explored 14

6.7 Bone pick-and-place 16

6.8 Knife handling 18

6.9 Meat Cutting 19

6.10 Processor Expression of Interest (EOI) 22

6.11 Office Demonstration to AMPC 22

6.12 Bone pick-and-place 23

6.13 Meat cutting 24

6.14 Milestone 5 planning and risk identification 25

6.15 Preparation at the Intelligent Robotics office 25

6.16 Site testing 26

Disclaimer The information contained within this publication has been prepared by a third party commissioned by Australian Meat Processor Corporation Ltd (AMPC). It does not necessarily reflect the opinion or position of AMPC. Care is taken to ensure the accuracy of the information contained in this publication. However, AMPC cannot accept responsibility for the accuracy or completeness of the information or opinions contained in this publication, nor does it endorse or adopt the information contained in this report.

No part of this work may be reproduced, copied, published, communicated or adapted in any form or by any means (electronic or otherwise) without the express written permission of Australian Meat Processor Corporation Ltd. All rights are expressly reserved. Requests for further authorisation should be directed to the Executive Chairman, AMPC, Suite 2, Level 6, 99 Walker Street North Sydney NSW.

6.17 Demonstration to site management 27

6.18 Demonstration to AMPC representatives 27

6.19 Next Stage Development Proposal 27

7.0 Conclusions 28

8.0 Recommendations 29

9.0 Project outputs 29

1.0 Abstract

The Australian red meat processing industry faces ongoing challenges with labour availability, workplace safety and the automation of tasks requiring human-like handling skills. This project investigated the potential for humanoid robots to perform these tasks, with the aim of developing more flexible automation solutions for meat processing facilities.

Intelligent Robotics selected and configured a humanoid robot and developed methods to train it by demonstrating tasks through human-controlled movements. The robot was trained and tested on two representative tasks: picking up and placing bones, and retrieving a knife and cutting meat.

The project successfully demonstrated the robot performing all tasks autonomously under controlled conditions and demonstrated to AMPC at Intelligent Robotics' office. The bone-handling capability was subsequently demonstrated at a processor site to site management and an AMPC representative, where the robot successfully completed 19 out of 20 trials despite differences from its training environment.

The results demonstrate the potential for humanoid robots to learn and perform tasks relevant to meat processing. However, further development is required to improve reliability, adaptability and safety before industrial deployment.

The project provides a foundation for future development of humanoid automation in the red meat industry, with the potential to reduce reliance on manual labour, improve workplace safety and enable automation of tasks that are difficult to address using conventional robotics.

The scope for a next stage development project was outlined based on the learnings from this project. A research proposal is currently being drafted based on this work to continue the development of humanoids for the red meat industry.

2.0 Executive summary

The Australian red meat processing industry faces ongoing challenges with labour availability, workplace safety and the automation of tasks that require human dexterity and adaptability. Conventional industrial robots have successfully automated many repetitive processes, but tasks involving irregular products, complex handling and the use of hand tools remain difficult to automate and require large footprints and site modifications. Humanoid robots, which are designed with human-like arms and hands, offer a potential alternative by providing the flexibility to perform a wider range of tasks within existing processing environments. This project, undertaken by Intelligent Robotics with funding from the Australian Meat Processor Corporation (AMPC), investigated whether humanoid robots could be trained to perform representative tasks relevant to red meat processing. The research was primarily targeted at Australian red meat processors and industry stakeholders seeking practical solutions to labour shortages, workplace safety risks and productivity challenges. The outcomes will help inform future research and investment decisions, providing AMPC members with evidence of the capabilities, limitations and potential commercial applications of this emerging technology.

The primary objective was to establish a working humanoid robot platform and demonstrate its ability to autonomously perform two tasks: picking up bones and placing them into a box, and retrieving a knife from different positions and placing a cut in a piece of meat. A further objective was to demonstrate these capabilities to AMPC and industry representatives, including at a meat processing facility. These objectives were substantially achieved, with all tasks successfully demonstrated under controlled conditions and bone handling subsequently demonstrated at an operating meat processing facility.

The project involved purchasing and configuring a humanoid robot, developing methods for controlling and training it, and conducting repeated trials to progressively improve performance. Training primarily involved a human operator demonstrating tasks through a headset that allowed the robot to follow their movements. These demonstrations were recorded and used to train artificial intelligence models, enabling the robot to subsequently perform the tasks independently. The team also investigated computer simulation, alternative training methods and additional sensors to identify opportunities for improving performance.

The project demonstrated that a humanoid robot could successfully learn and execute tasks relevant to red meat processing. In controlled office trials, the final bone-handling and meat-cutting programs each completed 20 out of 20 attempts successfully. The bone-handling program was subsequently tested at a processor site, achieving 19 successful attempts out of 20 despite differences in lighting, equipment and product presentation, without requiring additional training. Importantly, the project also identified limitations, including the robot's ability to adapt to unexpected situations, physical handling capacity, equipment reliability and safety requirements. These findings provide a valuable foundation for future development, but do not establish that the technology is yet ready for commercial deployment. For the red meat industry, humanoid robotics could ultimately provide a more flexible approach to automation, potentially reducing labour requirements and improving workplace safety for physically demanding or hazardous tasks. Further research is required to improve the robot's adaptability, reliability and safe operation in industrial environments, followed by more extensive trials under realistic production conditions.

The scope for a next stage development project was outlined based on the learnings from this project. A research proposal is currently being drafted based on this work to continue the development of humanoids for the red meat industry.

3.0 Introduction

The Australian red meat processing industry faces ongoing challenges associated with labour shortages, high staff turnover, workplace health and safety, and the physical demands of manual processing tasks. Many operations, particularly those involving cutting, boning and product handling, remain highly dependent on skilled manual labour. These tasks are often repetitive, physically demanding and require considerable dexterity and judgement. Increasing automation is therefore an important opportunity to improve productivity, workplace safety and the long-term competitiveness of the Australian red meat industry.

Conventional industrial robotics has been successfully applied to various meat processing operations. However, these systems are typically designed for specific tasks and often require specialised equipment, substantial engineering development and modifications to existing facilities. The natural variation in meat products, including differences in size, shape and positioning, presents additional challenges. As a result, many processing tasks remain difficult or commercially impractical to automate using conventional approaches.

Humanoid robots offer a potential alternative. Their human-like arms, hands and movement capabilities could allow them to use existing tools and operate within workspaces designed for people, reducing the need for specialised machinery and facility modifications. Furthermore, recent advancements in artificial intelligence have enabled robots to learn tasks through human demonstrations rather than relying entirely on manually programmed movements. This creates the possibility of more flexible automation systems that can be trained to perform multiple tasks, potentially reducing development costs and enabling wider adoption across the industry.

However, humanoid robotics remains an emerging technology, with limited practical evidence of its suitability for red meat processing. Significant questions remain regarding the ability of commercially available humanoid robots to handle irregular products, manipulate tools, perform precise movements and adapt to changing environments. Their reliability, safety and suitability for industrial applications also require investigation.

This project, undertaken by Intelligent Robotics (IR) with support from the Australian Meat Processor Corporation (AMPC), sought to investigate these questions and establish a foundation for future development. The project involved selecting and configuring a humanoid robot, developing methods to train it using human demonstrations, and evaluating its ability to autonomously perform two representative tasks: picking up and placing bones, and retrieving a knife and placing a cut into a piece of meat. Trials were conducted in a controlled office environment, followed by a demonstration at a meat processing facility.

Unlike conventional automation projects focused on developing purpose-built systems for individual operations, this project investigated whether a commercially available humanoid platform could be trained to perform different meat processing tasks using a common approach. The research was intended to provide Australian red meat processors, AMPC levy payers and industry stakeholders with practical evidence of the technology's capabilities, limitations and potential applications.

The findings will guide future research and investment decisions, particularly regarding improvements in reliability, adaptability, safety and commercial feasibility. Ultimately, the project aims to support the development of more flexible and scalable automation solutions that could reduce reliance on manual labour, improve workplace safety and enable automation of tasks that remain challenging for conventional robotics.

4.0 Project objectives

- Procure and configure a humanoid robot for task training and operational control.
- Understand and program the humanoid robot to perform a specific human task.
- Demonstrate the humanoid can successfully:
    - Pick up a bone and place it in a box.
    - Grab a knife in different positions.
    - Cut a piece of meat with the knife.
- Demonstrate the humanoid robot's capabilities at AMPC offices and a meat processing facility to demonstrate its real-world application, safety and potential impact on automation in the red meat industry.

5.0 Methodology

Purchase Humanoid Robot and tools:

- Investigate the specifications and options of the humanoid robot.
- Purchase the humanoid robot.
- Purchase table and knife sharpener and knives and safety equipment.

Humanoid Programming:

- Imitation Learning to capture expert skills.
- Simulation Reinforcement Learning.

Humanoid testing:

- Real World Execution:
    - Run the program and measure the performance of the humanoid robot's performance in terms of blade angle and force exerted.
    - Place the bone in different places and see that the humanoid can find a place in a box.

Reporting, Demonstration and Future Planning:

- Demonstration to AMPC showcasing the concept and capabilities of humanoid robotics.
- Final report and summary video detailing project outcomes and system performance.
- Production of supporting videos to document training and/or demonstration results.
- Preparation of a potential Stage 2 proposal, outlining the next phase of development if trial outcomes support further progression.

6.0 Discussion

6.1 Robot Specification and Humanoid Ordering

The Unitree G1 was selected after doing a comparison with similar humanoids. The comparison was based on how open the humanoid code was and how open the robot code to training systems was. A comparison table is shown below. Humanoid options were assessed across the following dimensions:

- *Openness* – how "open" is the humanoid for the purposes of programming. A lot of suppliers currently have "closed" architectures, meaning that custom programming is not yet able to be applied to them. For the purposes of this project it is critical we have an ability to program the humanoid ourselves.
- *URDF / USD Models; Environment/Assets* – the availability of models and assets which can be used for simulation and training. These can be created, but the preference is to use models/assets which are already existing, have been created by the manufacturer, and have already been utilised by end-users. This will make development more efficient and enable the work on the project to be focused on the task building.
- *Synthetic Datasets; Policy training frameworks* – Similar to models and assets, utilising pre-existing synthetic datasets and policy training frameworks will aid with more efficient training to allow the focus of the project to be the task-specific data capture and training.
- *VR / Teleoperation* – a key element to training for this project will be imitation learning for the specific tasks selected. It is envisioned that VR/Teleop will be the most effective way to capture high-quality data for this training.
- *Additional development support; Community engagement* – As this is an R&D project, the access to support is important to assist with any sticking points. This is particularly important given the novel state of humanoid training and programming in general.
- *Onboard vs Cloud Compute* – onboard compute for the humanoid is preferred if possible to enable the humanoid to operate without internet connectivity. This will be particularly advantageous in the context of the site demonstration at a meat processing facility.

In order to run the simulation software a new computer was specified and purchased to run the Invidia Isaac Sim software. An AMD Ryzen 9 9950X with a 5090 GeForce RTX Graphics CPU was selected and purchased.

There are several hand options available, for the tasks we need to undertake we selected a five finger hands with touch capability.

Below is a list of the option items purchased.

- 1 off UTG1EDUULTI UnitreeG1HumanoidRobot– EDU Ultimate with 100 TOPS & 2 x RH56DFTP Multi-Touch Hands
- Includes:
    - 1 x 9000mAh Battery
    - 1 x Battery Charger
    - 1 x Handheld Remote Control
    - 1 x Suspension Rig
    - 2 x RH56DFTP Multi-Touch Hands with 5 Fingers











We have also purchased plastic knives and some plastic bones to start testing with.

We have also purchased a suspension rig as a safety aid for the robot.

6.2 Platform procurement, setup and operating baseline

Milestone 2 moved the project from planning into a working humanoid trial setup at Intelligent Robotics. The main objective for this stage was to assemble the physical platform, compute, task props and operating environment needed to begin teleoperation, policy training and office-based trials.

6.2.1 Hardware acquired and configured

The core hardware baseline is the Unitree G1 humanoid fitted with dexterous hands for object and tool handling trials. This gives the project a real humanoid embodiment for testing manipulation tasks that are closer to red meat industry workflows than a fixed industrial arm demonstration.

The project also established local compute for model training and inference, including RTX 5090 and RTX PRO 6000 workstations. Local compute is important because MS2 depends on fast iteration: collect demonstrations, train or fine-tune a policy, deploy it back to the robot, review failures, then collect more targeted data. Table 1 summarises the key hardware acquired for MS2. The hardware is also shown in Figure 1.

Table 1. Hardware acquired for Milestone 2

Hardware / setup item	Justification
Unitree G1 humanoid	Core embodiment for evaluating humanoid manipulation in this project.
Inspire RH56E2 dexterous hands	Core end effector for dexterous tasks, selected for versatility and compatibility with tasks designed around human hand use.
RTX 5090 workstation	Initial local compute purchased for the workflow; remains useful for policy inference and lighter experimentation.
RTX PRO 6000 workstation	Required to run augmentation models and for training more complex policies; reduces pod rental setup and data-transfer overhead.

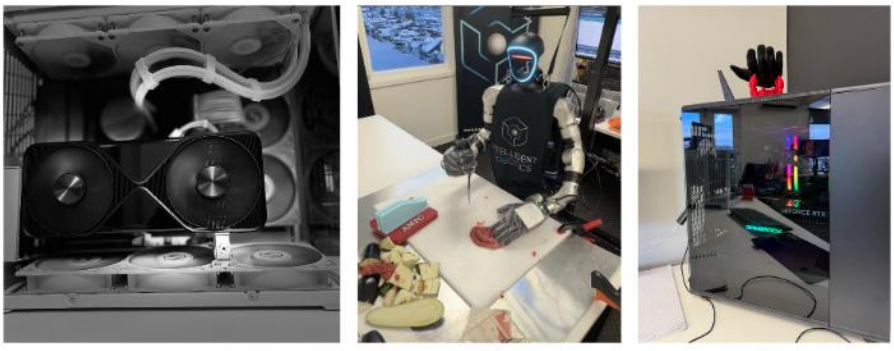


Figure 1. (a) The RTX PRO 6000 workstation (b) The Unitree G1 humanoid equipped with Inspire hands (c). The RTX 5090 workstation

6.2.2 Safety setup and practical operating constraints

The setup phase clarified that humanoid operation is not equivalent to commissioning a mature industrial arm. Industrial safety standards for humanoids are still developing, and the team believed further research, planning and validation are required before humanoid trials can be conducted on site. The lack of an industrial-standard E-stop is a key example.

Unitree has proposed an E-stop approach that removes energy from the system by initiating an electrical short through a port on the rear panel. However, this should not be presented as an industrial E-stop because it is single-channel and trips high rather than low. It also creates a practical operating issue: once triggered, the approximately 40 kg humanoid will lose support and collapse.

For the planned site trial, the team is planning to implement the E-stop approach proposed by Unitree. Although this is not an industrial-standard E-stop, it reduces some of the operational risk involved. Furthermore, the humanoid will be tethered to a frame, providing a physical limitation to its work envelope. The final approach will be re-evaluated before the site trial takes place including performing a risk assessment.

6.2.3 Challenges and lessons learned

During initial testing, the team gained valuable insights into the practical use of dexterous robotic hands for meat processing applications, particularly regarding their interaction with tools, products and work surfaces. These trials highlighted the importance of carefully managing contact forces and hand movements to minimise wear and the potential for damage during repeated operations. The team explored various operating techniques and protective measures to improve reliability and suitability for the intended applications. As humanoid robotics is still an emerging field, the underlying hardware continues to evolve, with ongoing improvements in durability, robustness and performance expected to further enhance its suitability for industrial applications.

6.3 Software, simulation and training ecosystem explored

In parallel with the hardware setup, the team explored the software and training ecosystem required to support humanoid policy learning. This included evaluating specific tools and environments, as well as the broader workflow for turning robot demonstrations into policy training data.

The purpose of this work was not to make simulation the core milestone outcome. Rather, the intent was to understand which parts of the current Physical AI ecosystem are useful for this project now, which parts may become useful later, and which parts are not suitable for the Unitree G1 workflow at this stage.

6.4 Teleoperation stack and demonstration-data capture

Teleoperation is the current pathway for collecting robot demonstrations for imitation learning. The policy cannot learn the target behaviour without paired observations and actions from the real robot, so teleoperation quality directly affects the quality of the training data.

The team used Apple Vision Pro as the primary teleoperation interface. This was the most mature and tested path available for collecting real G1 demonstrations, and it allowed the team to begin validating the data capture, training and deployment workflow, as illustrated in Figure 2.



(a) (b)

Figure 2. (a) Robot reproducing the teleoperator's motion through teleoperation. (b) Apple Vision Pro headset acquired for teleoperation.

6.4.1 Demonstration data captured during teleoperation

The purpose of teleoperation is not only to move the robot manually. It is to capture demonstrations in a format that can later be used for policy training. During teleoperation, the system records visual observations, robot states and commanded actions so that the model can learn the relationship between what the robot sees and what action should be taken next. A summary is shown in Table 2.

Table 2. Data captured from the current teleoperation workflow

Data source	Purpose
RGB camera stream	Provides the visual observations used by the policy to reproduce learned behaviours from visual cues.
Joint position data	Provides robot state information so the policy can learn the relationship between visual observations and robot behaviour.

6.4.2 Apple Vision Pro strength and limitations

The Apple Vision Pro (AVP) is currently the most practical teleoperation solution evaluated for this project. It provides an intuitive interface for controlling the G1 and enables the team to capture demonstrations through an established and reliable workflow. Consequently, the AVP remains the primary method for teleoperation data collection.

As shown in Figure 3, the main limitation is that the current G1 teleoperation view presents a 2D RGB stream rather than a true spatial video feed. This makes depth judgement more difficult, particularly during fine manipulation, grasping and tool interaction. Noticeable latency also requires the operator to move more slowly and deliberately than when performing the task directly.

Occlusion presents another limitation. During grasping or tool handling, the robot hand, object or tool can obstruct the visual features needed to assess contact and alignment. This is particularly relevant to knife handling and cutting, where the point of contact may be partially hidden.

The AVP system also does not provide haptic feedback. Grip force, contact force and tool interaction must therefore be inferred visually, which can reduce operator confidence and the consistency of demonstrations during force-sensitive tasks.

Despite these limitations, their impact is minor within the context of this application. The AVP provides the best overall balance of usability, workflow maturity and data-capture capability among the methods evaluated, and remains suitable for achieving the objectives of this project. Future projects with the humanoid robot however may pivot towards alternative technologies, particularly as these continue to develop.

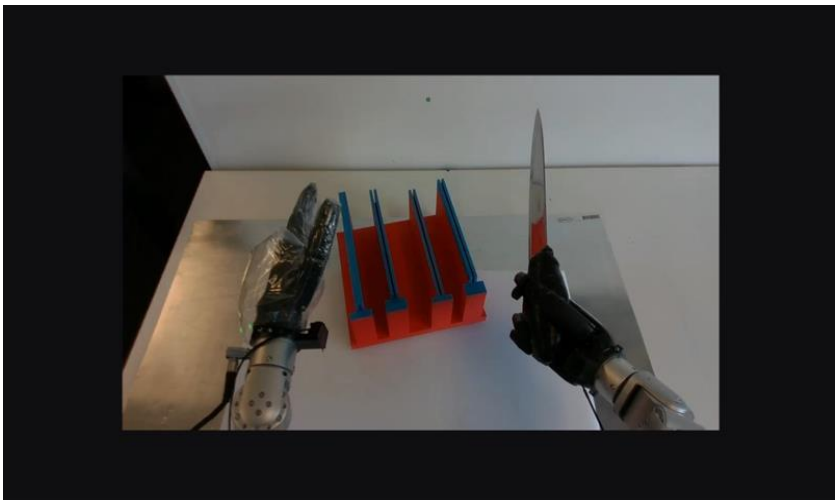


Figure 3. What the teleoperation operator sees in the headset: A 2D view surrounded by darkness

6.5 High-level workflow overview

The workflow starts with teleoperation, where the operator demonstrates the target task on the robot. During demonstration data collection, the system records the robot states, actions and video required for policy training.

The collected dataset is then formatted into a training-ready structure. Where useful, video augmentation can be applied to increase visual variation in the dataset.

Candidate policy models are then selected and trained using the prepared training dataset. Once a policy is trained, it is deployed back to the robot for evaluation. The result of deployment is used to identify gaps in the data, such as missing object poses, insufficient visual variation or failure cases that need more targeted demonstrations.

The important point is that this is an iterative workflow. The project does not expect one data collection round to produce a final policy. Instead, the team collects demonstrations, trains a policy, evaluates the result, identifies gaps, and then collects or augments more data for the next policy version.

Table 3. Summary of the workflow stages

Workflow stage	Role in the project
Teleoperation	Human operator demonstrates the task on the robot.
Demonstration data collection	Captures robot states, actions and video.
Dataset formatting	Converts collected data into a training-ready dataset.
Video augmentation	Adds visual variation where useful.
Policy training	Trains candidate policies using the prepared dataset.
Deployment and evaluation	Tests the trained policy on the robot and identifies gaps for the next iteration.

6.6 Experimental features explored

The team also tested several experimental features that may improve dexterous manipulation tasks. These include wrist cameras, tactile and force-torque data from the hand, and depth data from a RealSense camera. The purpose was to assess whether additional observations could help the policy with dexterous manipulation tasks such as secure grasping, knife pickup, adapting to object height, preventing slippage and controlling grip force.

The current finding is that these features are worth exploring but should not yet be presented as proven improvements. In general, adding a new sensor modality does not automatically improve policy performance. If the model has not been pretrained with that modality, it may require substantially more data before it can learn how to interpret the signal effectively.

6.6.1 Wrist cameras

The team 3D printed wrist camera mounts and tested wrist cameras on the humanoid, as shown in Figure 4. The intended benefit was to provide a closer view of hand-object contact, especially for dexterous tasks where the main camera view may not clearly show the grasp point or tool position.

Wrist cameras are widely used in robotics demonstrations and are a reasonable feature to test. The expected benefit is strongest for tasks such as secure grasping, knife pickup and adapting to different object or tool heights.

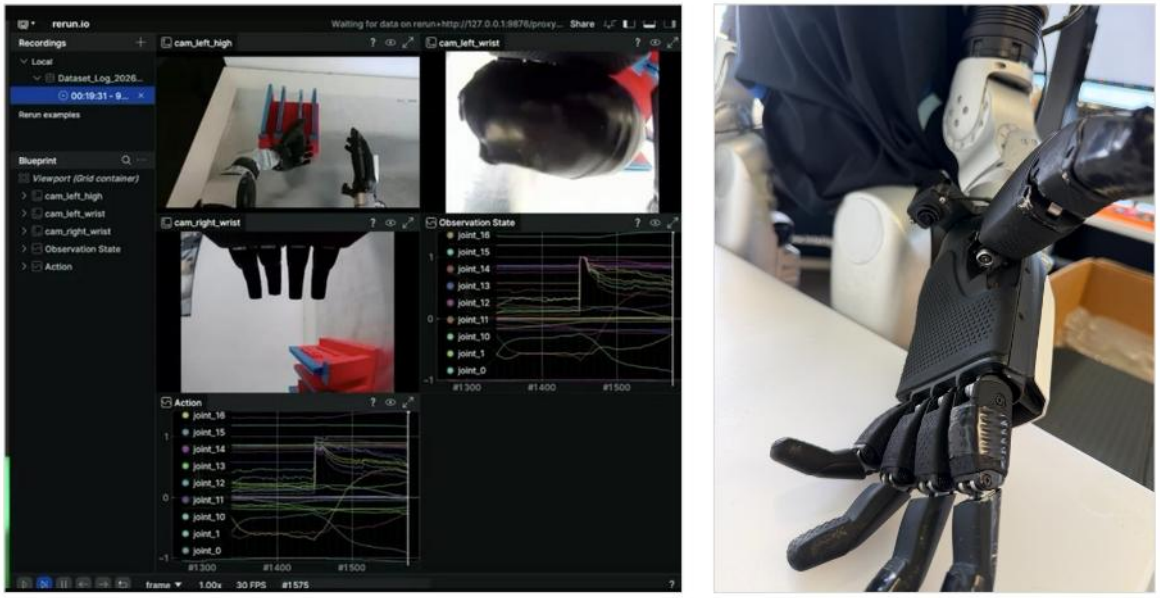


Figure 4. (a) Rerun Viewer showing footage from different camera sources; (b) Wrist camera mounted to the robot's wrist.

6.6.2 Tactile and force-torque data

The team experimented with adding tactile and force-torque data from the hand into policy training, as shown in Figure 5. The intended benefit was to provide signals that RGB video alone cannot capture, such as contact force, slippage and grip stability.

For meat-handling tasks, this is relevant because the robot needs enough grip force to hold packaged product securely, but not so much force that it risks bag rupture. In principle, tactile and force-torque data could help the policy learn this balance.

There is also a teleoperation limitation: the operator does not have haptic feedback during demonstration collection. Even if tactile and force-torque data are added to the training dataset, the demonstrated behaviour was generated without the operator feeling those signals. This limits how much the current policy can learn from them. For this reason, tactile and force-torque data remain experimental rather than part of the proven training setup.

Future work may revisit this direction using models where tactile data is treated as a first-class modality rather than an added custom signal.

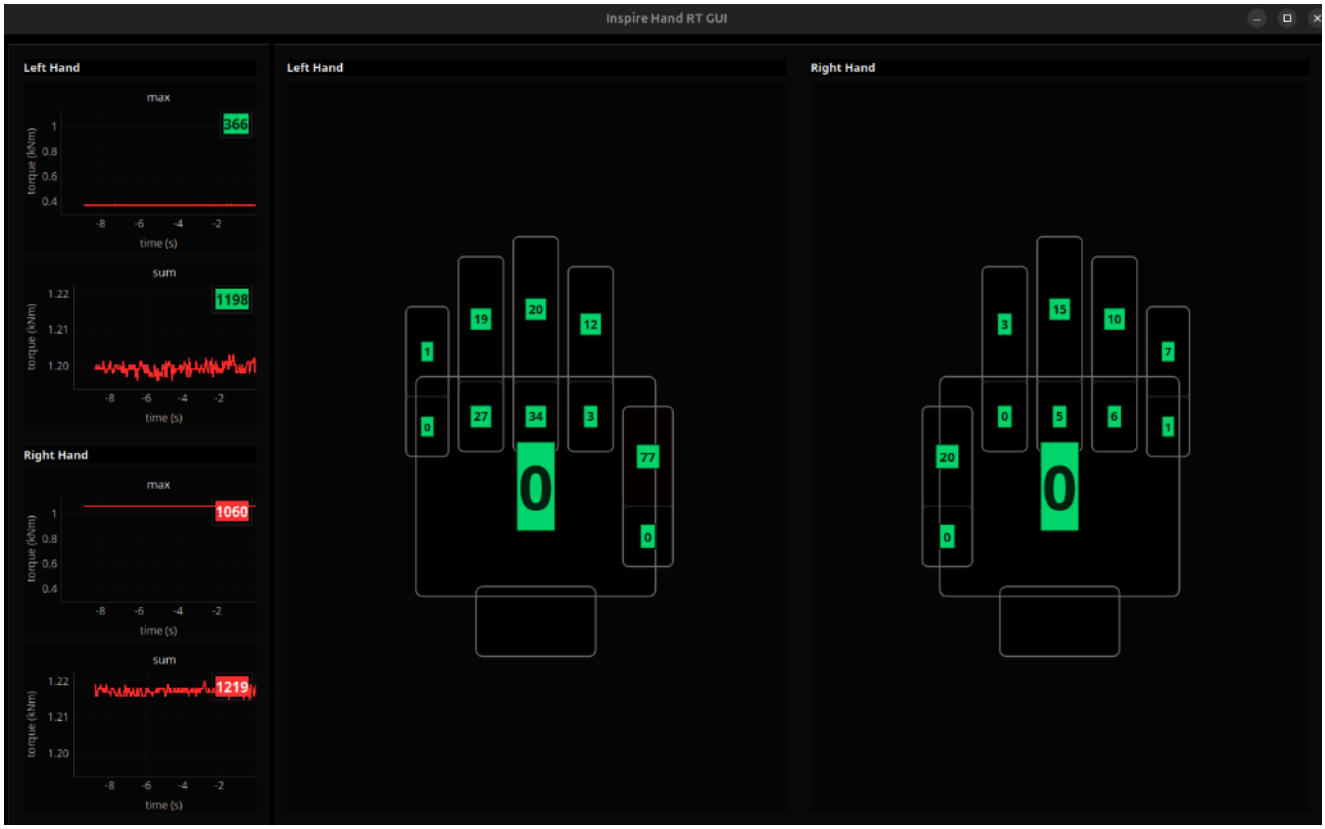


Figure 5. GUI designed by the team to visualise tactile and force data from the Inspire hands.

6.6.3 Depth data

The team also experimented with adding depth data from a RealSense camera to policy training. The intended benefit was to provide geometric information that RGB video does not directly capture, such as object height, approach distance and grasp pose.

Similar to tactile and force-torque data, current state-of-the-art policy models used in this workflow are not pretrained to treat depth as a native modality. The policy is also limited by the quality of the demonstrations: if the teleoperation operator does not have reliable depth perception during data collection, the policy cannot automatically learn depth-aware behaviour just because depth is added to the dataset.

6.7 Bone pick-and-place

6.7.1 Overview

The bone pick-and-place task is the most mature manipulation task in the current trial set. The policy can reliably pick up bones in the environment it was trained on, and it is reasonably robust to minor environmental changes, such as small lighting shifts or unexpected objects on the table, although these disruptions can still affect performance in some trials. The steps taken by the policy for each cycle are outlined in Figure 6.

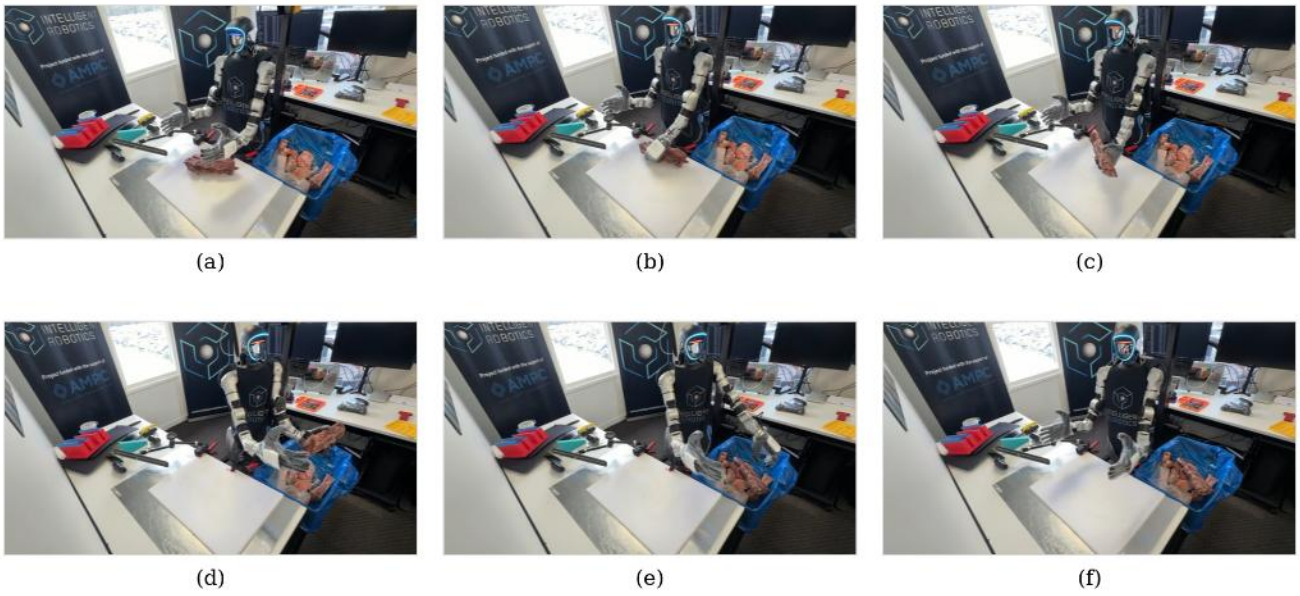


Figure 6. (a) starting position, bone in frame (b). gripping the bone (c). flipping the hand to better hold the bone (d). waist rotation (e). dropping off the bone (f). waist rotation back to starting pose

6.7.2 Task progression

The policy was developed through a series of versions, each targeting a specific weakness identified in the previous one. Table 4 summarises this progression.

Table 4. Version History for bone policy

Version	Results	Remaining challenges
v1	Robot reliably picked the bone from the fixed position and dropped it into the fixed box.	The policy memorised the fixed start and end positions rather than learning the task. Moving the bone even slightly caused the grasp to fail, as the robot did not adjust its approach.
v2	Reduced memorisation; the robot learned to pick the bone from anywhere within its work envelope.	Performance dropped when the robot was moved to a different desk, indicating the policy was still overfitted to the training environment.
v3	Significantly improved robustness to changes in work surface and visual appearance.	The work envelope remained limited, made worse by the narrow field of view of the egocentric RealSense camera at 480p. Teleoperation felt highly constrained.
v4	Waist-joint control and the wider 16:9 frame greatly improved visibility and expanded the work envelope. The box could now be placed in any direction around the desk, freeing space for the bone and giving the robot a much clearer view inside the box once it turned to face it. However, the added	The torque and tactile observations had to be removed to restore performance, addressed in v4.1.

Version	Results	Remaining challenges
	torque and tactile observations degraded the policy.	
v4.1	Restored stable behaviour, while keeping the waist-joint control and 720p resolution improvements.	Grasp position was sometimes inconsistent, and the excess plastic around the bagged bones made it difficult to identify bone shape and geometry.
v5	Unbagged bones made it much easier to perceive bone shape and adjust the grasp for a more stable grip. Wrist-camera contribution is inconclusive and still being evaluated.	The main improvements still required are policy stability across action chunks, robustness to stronger environmental changes, and recovery from policy drift over repeated pickups.

6.8 Knife handling

6.8.1 Overview

The knife handling trial requires the humanoid to pick up a knife handle from a custom 3D-printed knife block. This is a more dexterous task than bone pick-and-place because the robot must grasp the handle accurately and release the blade from a constrained holder. The steps taken by the latest policy for each cycle are outlined in Figure 7.

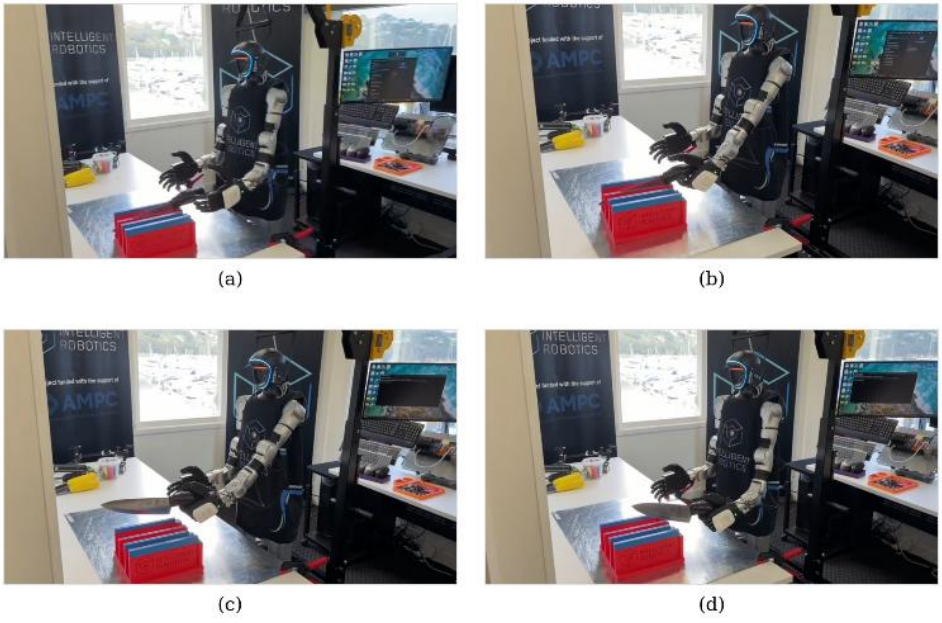


Figure 7. (a) policy convergence based on visual observations; (b) reaching the grasp pose; (c) pulling the knife from the knife block; (d) returning to the starting pose while holding the knife.

6.8.2 Task progression

The policy was developed over three versions, summarised in Table 5.

Table 5. Version history of the knife handling policy

Version	Results	Remaining challenges
v1	Robot successfully learned to pull the knife out of the block.	The yellow block suited only a narrow range of knife shapes and sizes, and the blade often caught on the plastic during removal.
v2	Robot learned the upward pull for the red block while retaining the outward-pull behaviour for the yellow block.	Grasp was not consistent on every attempt, particularly when the knife block was positioned further away.
v3	Grasp accuracy improved, but the gains persisted even when the wrist cameras were blocked, suggesting the improvement came from the additional demonstrations rather than the wrist cameras.	—

6.8.3 Results

The policy has been fine-tuned multiple times and can now handle two distinct knife blocks. For the yellow block, the robot learned to pull the knife backwards to draw the blade out; for the red block, it learned to lift the arm upward to release the blade.

The policy converges quickly and is robust to small environmental changes, supported by augmentation. For example, it performed comparably on a white chopping board despite the training data being captured on a wooden chopping board.

The policy was also able to grip knives when the knife tray was raised or lowered by approximately 50 mm. This may indicate a benefit from the wrist cameras, which give the policy a closer view of the hand and knife. However, the evidence is not yet conclusive, so it should be treated as a possible benefit rather than a proven result.

This trial is an important step for the project. It demonstrates generalisation beyond a simple pick-and-place task, shows more dexterous tool interaction, and provides a stronger foundation for the future cutting policy.

6.9 Meat Cutting

6.9.1 Overview

Cutting involves dexterous tool handling, force application, supporting-hand control and teleoperation fidelity. A more detailed summary of the challenges and proposed solutions is provided in Table 6. These issues are more complex than picking up a bone or drawing a knife from a block, and they directly affect whether the team can collect consistent demonstrations for policy training.

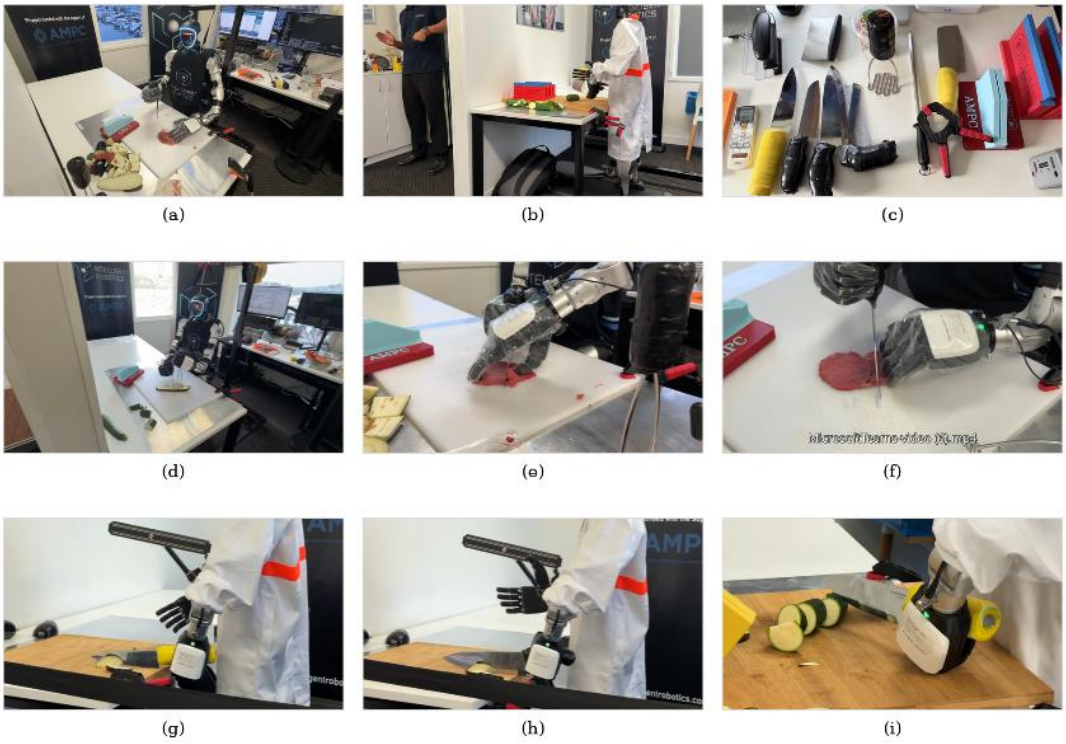


Figure 8. Footage from cutting trials conducted by the team thus far

Table 6. Challenges and Proposed solutions

Challenge	**Proposed solution**
Awkward chef's knife handling. The traditional chef's knife is difficult for the humanoid wrist configuration, causing floppy or abrupt motion during slicing.	Use a PETA Easi-Grip knife with a vertical handle.
AVP occlusion during slicing. When the operator curls a fist and performs an imaginary slicing motion, AVP can lose finger tracking, causing the knife to drop.	The PETA Easi-Grip knife helps because the operator does not need to move the hand as far downward.
Supporting hand cannot apply even pressure. The food can move during cutting if the supporting hand does not hold it consistently.	Use a potato masher or similar supporting tool to keep the object in place during cutting trials.
Inspire hand geometry limits fine slicing. The Inspire hand silhouette is large for highly dexterous slicing tasks; the knuckle can hit the table before the blade reaches the chopping board.	The PETA Easi-Grip knife helps reduce this issue by ensuring the hand always stays above the cutting board.
Inconsistent knife grip pose. Small changes in how the hand holds the knife can change the blade angle and result in a different cut path.	Iterative refinement of trained policies with more high-quality expert demonstrations showing optimal blade angle

6.9.2 Results

After extensive trialling with different knives and methods of securing the meat, the team selected the PETA Easi-Grip knife together with the potato masher, as this combination produced the most consistent motions across expert demonstrations.

The policy was trained using multiple episodes of vegetable cutting, steak cutting and slicing through beef mince. The resulting policy could complete the cutting task when the setup closely matched the conditions represented in the expert demonstrations, as shown in Figure 9.

However, the current policy remains sensitive to variations in the task setup and has several limitations that the team plans to address in MS4. These limitations and the corresponding proposed improvements are summarised in Table 7.

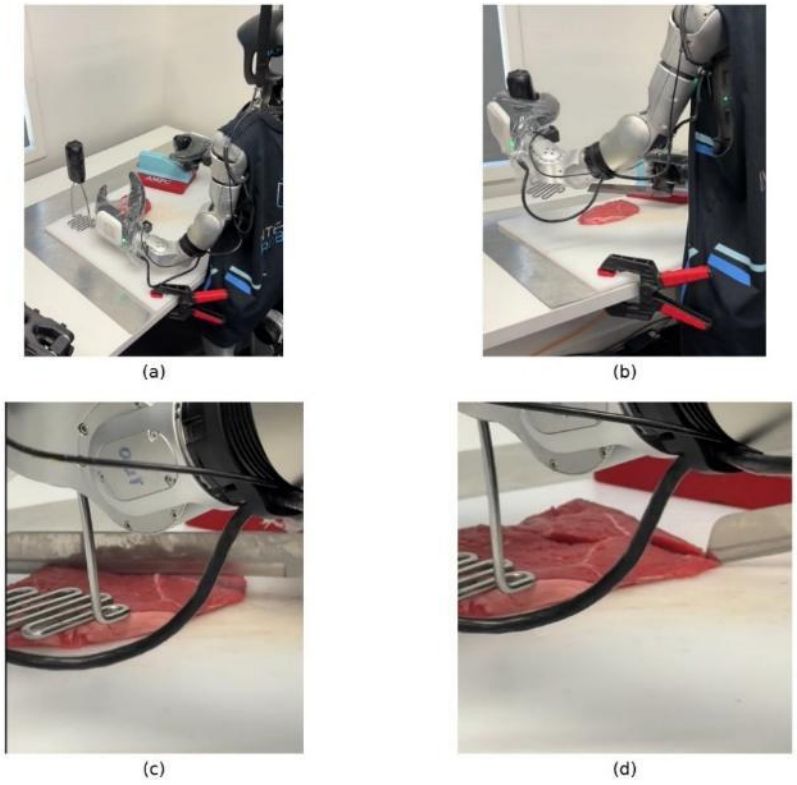


Figure 9. Meat Cutting Policy V1 rollout: (a) the robot reaches for the knife handle and pulls the knife out; (b) the robot grips the potato masher; (c) the robot uses the potato masher to support the meat and positions the tip of the knife for cutting; and (d) the robot completes the cut by moving the elbow joint backwards.

Table 7. Limitation and proposed solution

Limitation	Proposed solution
The policy is brittle and only works reliably under ideal conditions	Cutting is a complex, long-horizon task in which errors can compound over multiple stages. As a result, it requires substantially more training data than simpler tasks such as pick-and-place. Several factors may contribute to the current brittleness of the policy. The team plans to prioritise the areas expected to have the greatest impact by collecting additional demonstrations and introducing further constraints into the task space.

Limitation	Proposed solution
The policy converges slowly	In the current training dataset, the tools were placed in a wide range of positions and orientations to improve generalisation. However, this introduced substantial variation into the teleoperation motion patterns, making it more difficult for the policy to learn a consistent behaviour. For V2, the team plans to constrain the positioning and orientation of the tools so that generalisation is more targeted and the demonstrations contain more consistent motion patterns.
The policy struggles with tool handling	

6.10 Processor Expression of Interest (EOI)

Following the successful completion of Milestone 2, AMPC distributed an Expression of Interest to Australian red meat industry members to select a host site to perform the Site Demonstration for Milestone 5. From this process an Australian red meat processing plant was selected to host these trials.

6.11 Office Demonstration to AMPC

A demonstration was held at the Intelligent Robotics office on 21 September 2026 to AMPC representatives, fulfilling the key deliverable of Milestone 4.

The demonstration began with the bone pick-and-place policy, followed by the knife retrieval and meat-cutting policy. Together, these tasks demonstrated autonomous object handling, tool use and coordinated two-handed manipulation relevant to red meat processing. The demonstration presented the capabilities achieved within a controlled office environment, providing a baseline for the Milestone 5 site demonstration, where performance under processing-site conditions will be assessed.

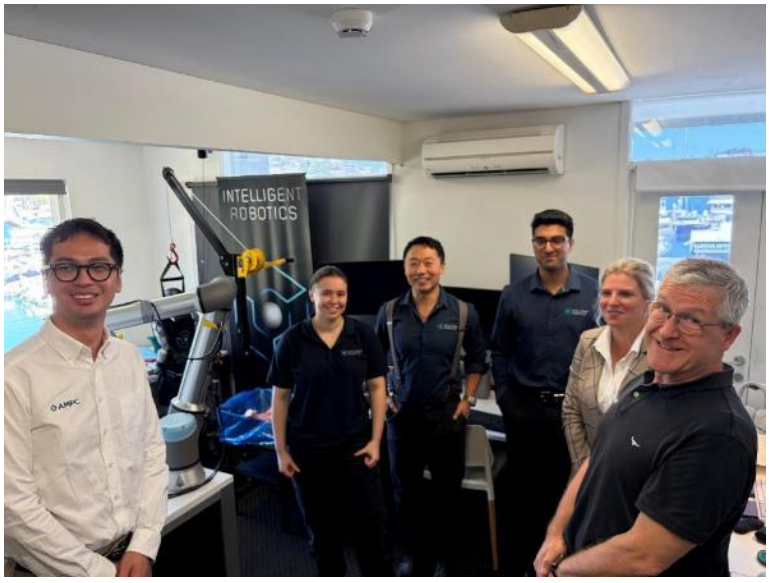


Figure 10. AMPC representatives and the Intelligent Robotics team at the office demonstration on 21 September 2026.

6.12 Bone pick-and-place

The bone policy was extended from the Milestone 2 v5.1 checkpoint. Version 6 was developed as an expert policy for leg bones, targeting reliable grasping across a wide range of starting orientations within a constrained workspace.

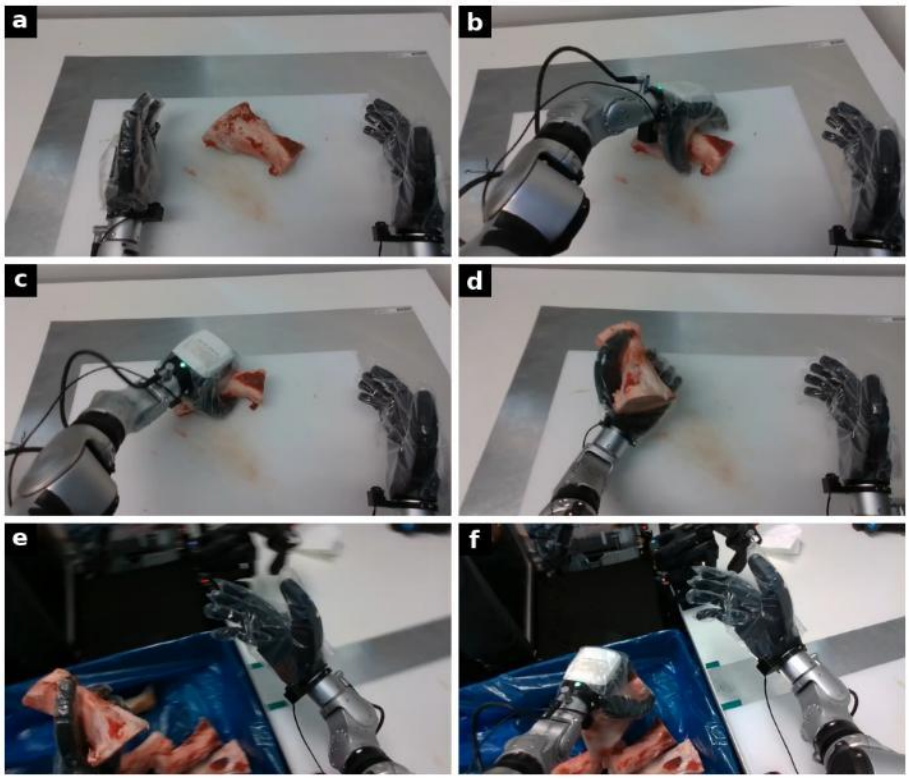


Figure 11. Bone pick-and-place v6.1: (a) home position; (b) reaching for an overhand grasp; (c) grasping phrase; (d) wrist roll after pickup; (e) transport to the box; and (f) drop-off and return to the home position.

Table 8. Bone policy development during Milestone 4

Version	Success rate	Observations
v6	Grasp 20 of 20; full sequence 10 of 20	Grasping was reliable and the policy selected its approach according to how the bone was presented. The sequence then stalled at the wrist roll following pickup, with the arm holding position rather than progressing to transport.
v6.1	20 of 20	The roll began immediately and decisively on grasp completion, and the full sequence ran to drop-off in every rollout.

6.12.1 Selecting the grasp from vision alone

Because bones were presented across a far wider range of starting orientations than in Milestone 2, no single grasp suited every presentation. A bone lying flat could be taken with a plain reach — fingers forward, thumb to the right. A bone standing closer to vertical required the wrist to be rotated clockwise before contact, bringing the fingers to the right and the thumb down.

Both approaches were demonstrated during capture, with no label distinguishing them and no instruction to the policy about which to use. The policy learned to select between them from the camera images alone, applying the rotated approach to upright bones and the plain approach to flat ones. This is the strongest evidence in this milestone that the policy responds to what it sees rather than replaying a single fixed motion.

6.12.2 Loss of hand tracking during capture

The Apple Vision Pro intermittently lost tracking of the operator's right little finger, recording a retracted position that did not reflect the operator's intent. These frames entered the training data and v6 reproduced the retraction during inference. The behaviour did not prevent task completion, since the task utilises the left hand, but it is a visible artefact and an instance of a general risk: teleoperation errors are learned as faithfully as teleoperation skill. Captured episodes will be screened for tracking dropouts before they enter the training set rather than after.

6.13 Meat cutting

The cutting task is a long-horizon, bimanual sequence requiring tool retrieval, product stabilisation, blade alignment and cutting. It is substantially harder than pick-and-place: four stages must complete in order, and a failure at any stage leaves the robot in a state that the later stages were never trained to recover from.

The *cut_all_v1* policy presented in Milestone 2 had an inconsistent cut path and did not use the supporting tool reliably. The team therefore trained *cut_meat_v1* using only steak demonstrations, with the robot pose, tool positions and product placement constrained. As with the bone policy, this traded generality for a more controlled expert policy.

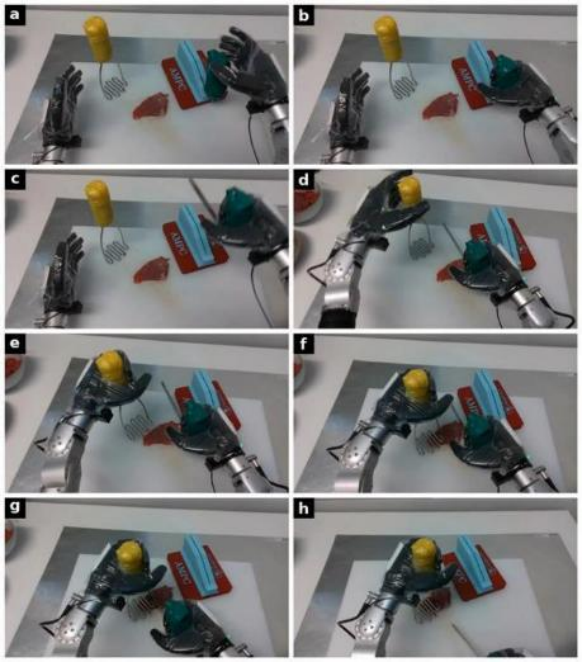


Figure 12. Meat-cutting v2: (a) home position; (b) reaching for and grasping the Easi-Grip knife; (c) pulling the knife out of the knife holder; (d) positioning the knife in the ready state and reaching for the potato masher; (e) grasping phase; (f) pressing down on the meat before cutting; (g) positioning the tip of the blade at the top of the meat, against the guard of the potato masher; and (h) sliding the blade through the meat to complete the task.

Table 9 summarises the policy evolution.

Table 9. Version history of meat cutting policy

Version	Success rate	Observations
v1	2 of 20	In most rollouts the right hand oscillated around the home pose and never began the reach to the knife block, as though waiting for a start signal. Where the task did begin it generally ran to completion, but masher placement and cut path were inconsistent.
v1.1	15 of 20	The policy now began the task immediately and converged directly on the knife handle. All remaining failures occurred at the knife grasp, where the sideways approach brought the hand in too low and the little finger met the knife block instead of the handle.
v2	20 of 20	Complete sequence reproduced in every rollout, with a secure five-finger grip on the handle. The policy will, however, run the same sequence over an empty board.

6.14 Milestone 5 planning and risk identification

Milestone 5 will demonstrate the trained policies at a processor site.

6.14.1 Risk

The principal risk is environmental difference. Both expert policies were trained under controlled conditions in the Intelligent Robotics office and performance outside the trained visual and geometric conditions has not been established. The processor environment will differ in lighting, background, work-surface height and finish, and object presentation. A degradation in performance on first deployment should be regarded as likely rather than exceptional.

6.14.2 Mitigation

The team will attend the site ahead of the formal demonstration and evaluate the current policies in place. If performance degrades, the team will:

- capture targeted demonstrations in the site environment, using the same teleoperation setup;
- preserve the general object placement area where practical; and
- post-train the current checkpoints on the site data, applying the curation practice established in this milestone, including screening for tracking dropouts before training

6.15 Preparation at the Intelligent Robotics office

The principal preparation task was to assess whether the policy could operate beyond the environment used for data capture. The team purchased a table 20 mm higher than the training table and set up the robot in a different setting, with changed lighting. Figure 13 shows the preparation setup.

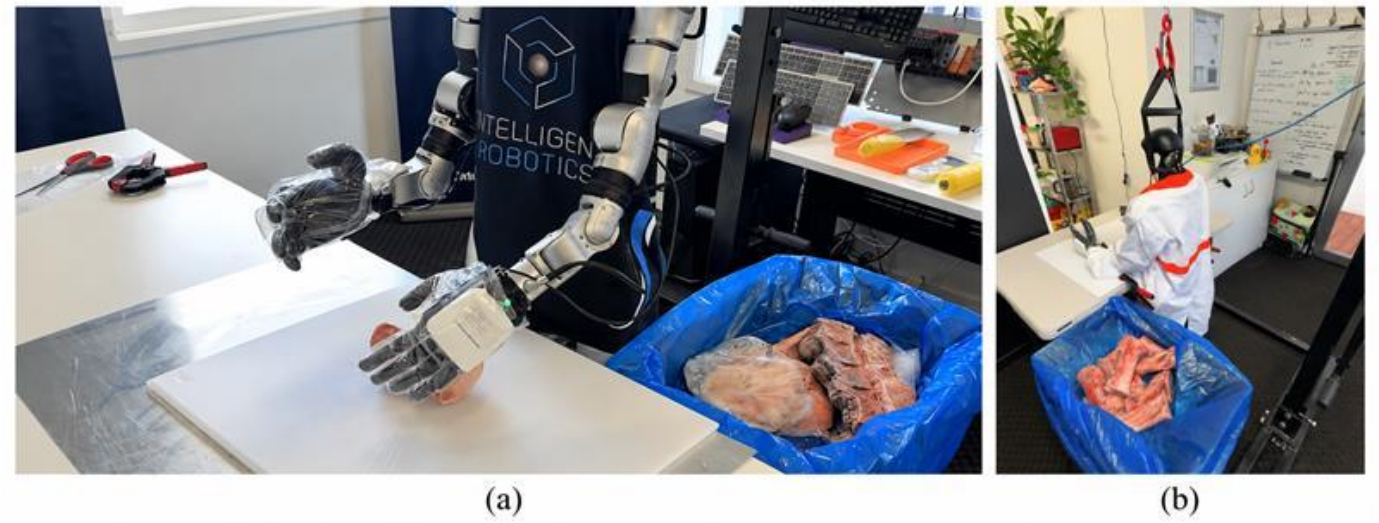


Figure 13 (a). training environment for the policy; (b) modified environment in preparation for the site demonstration.

The policy occasionally struggled to converge but completed 23 of 25 full sequences. This provided confidence to test the unchanged checkpoint on site. Table 10 summarises the evaluation results, including the Milestone 4 baseline for context.

Table 10. Evaluation results

Evaluation	Conditions	Successful sequences	Success rate
Milestone 4 baseline	Configured office setup	20 of 20	100%
Milestone 5 office preparation	Different setting and lighting; table 20 mm higher than the training table	23 of 25	92%
Milestone 5 site testing	Site setup; changed visual and geometric conditions; smaller bones	19 of 20	95%

Note that these results describe separate evaluations under their respective conditions. They demonstrate successful transfer in the tested setups; they do not establish a general success rate across processing environments.

6.16 Site testing

The team arrived at the processor site on 13 October 2026, one day before the AMPC demonstration.

The site presented several differences from the training environment:

- A different background and low-frequency ceiling lighting that appeared as flicker in the camera images.
- Different work-surface texture and height.
- Changes in arm appearance and wrist-camera angles.
- Different bone size and appearance: the site-supplied femur mid-sections were substantially smaller than those used for training and were vacuum-packed.

Despite these differences, pick_bone_v6.1 completed 19 of 20 full sequences. The checkpoint and runtime settings were unchanged, and no additional demonstrations or post-training were required. This fulfilled the intended zero-shot deployment: the policy trained at the Intelligent Robotics office performed the task at the processor site using the existing model.

The result also addressed the environmental-transfer risk identified in Milestone 4. For this task and setup, the changes were within the policy's operating capability. The remaining failure involved an edge case associated with contact physics and out-of-distribution bone placement, consistent with the limitations observed during office preparation.

6.17 Demonstration to site management

On 13 October 2026, the team successfully demonstrated autonomous bone pick-and-place to site management. The robot picked up the bone and placed it into the box using the policy trained at the Intelligent Robotics office.

6.18 Demonstration to AMPC representatives

On 14 October 2026, the team successfully demonstrated the same bone pick-and-place capability to AMPC representatives. This demonstration completed the key deliverable of Milestone 5 and presented the policy's capability at a location different from its training environment.

6.19 Next Stage Development Proposal

Based on the successful outcomes in this project, a next stage development proposal is currently being drafted by Intelligent Robotics in consultation with AMPC. During the site visit, a number of potential applications were investigated and a target application for the next stage was selected.

The project will also incorporate an extended site trial, where the humanoid is run in production for a target period of four weeks. This will be under the supervision of Intelligent Robotics, but will give significant learnings around deployment, longer term operation, and the production realities of the application and operating environment. Having a humanoid performing a task for an extended period of time, at a processor site, will also represent a big milestone for the technology and the industry.

The research proposal for this work is currently being drafted between Intelligent Robotics and AMPC.

7.0 Conclusions

The project has successfully demonstrated the potential for commercially available humanoid robots to learn and autonomously perform tasks relevant to the Australian red meat processing industry. The project established a working humanoid robotics platform and developed methods for training the robot through human demonstrations. The successful completion of bone handling, knife retrieval and meat cutting tasks represents an important initial step towards assessing the feasibility of humanoid robotics for meat processing applications.

The results demonstrated that the humanoid robot could learn increasingly complex tasks, progressing from basic product handling to tool manipulation and coordinated two-handed cutting operations. Following successive improvements to the training process, both the bone-handling and meat-cutting tasks achieved success rates of 100% across 20 trials under controlled conditions. Importantly, the bone-handling capability was subsequently demonstrated at a processor site, achieving 19 successful trials out of 20 despite differences in lighting, work surfaces and product presentation. This was achieved without additional training, providing encouraging evidence that learned capabilities can transfer beyond the original training environment.

A significant finding was the importance of training data quality and the methods used to teach the robot. The project demonstrated that carefully preparing human demonstrations could substantially improve performance without requiring extensive additional training. These findings provide valuable insights into how humanoid robots can be trained more efficiently and reliably for future applications. The research also identified that simulated environments and additional sensing technologies did not necessarily improve performance, reinforcing the importance of practical testing and selecting training methods appropriate to the task.

Despite these achievements, the technology remains at an early stage of development. While the robot demonstrated reliable performance under defined conditions, its ability to respond appropriately to unexpected situations remains limited. For example, the meat-cutting program could reproduce the cutting sequence even when no meat was present, demonstrating that successful task execution does not necessarily indicate an understanding of the broader operation. The project also identified limitations relating to optimal use of hand hardware, adaptability and safety systems. These challenges must be addressed before humanoid robots can be considered suitable for continuous operation within commercial meat processing facilities.

For Australian red meat processors, the findings demonstrate a potential new approach to automating tasks that have traditionally been difficult to address using conventional industrial robotics. Humanoid robots may eventually provide greater flexibility by performing multiple tasks using existing tools and workspaces, potentially reducing the need for specialised machinery and extensive facility modifications. This could create opportunities to reduce reliance on manual labour, improve workplace safety and increase operational flexibility. However, the project has not yet established the reliability, productivity or commercial viability required to demonstrate these benefits in a production environment.

The outcomes provide a foundation for further research and development. Future work should focus on improving the robot's ability to adapt to changing conditions, recognise unexpected situations, recover from errors and safely interact with its environment. Extended trials under realistic processing conditions will be necessary to establish achievable operating speeds, reliability and commercial feasibility.

Overall, the project has demonstrated that humanoid robotics is a promising technology for future red meat processing automation, while also identifying significant challenges that remain before practical adoption. The results provide AMPC, levy payers and industry stakeholders with an evidence-based understanding of the technology's current capabilities and limitations, supporting informed decisions regarding future research investment and the development of potential commercial applications.

8.0 Recommendations

The project has demonstrated promising capabilities for humanoid robotics in red meat processing, while identifying several technical and practical challenges that must be addressed before commercial adoption.

A research proposal is currently being drafted between Intelligent Robotics and AMPC on the next stage of development, as outlined in this report. It is recommended that this project proceed to continue the development of humanoids for the red meat industry.

9.0 Project outputs

Outcome: Demonstration of the humanoid robot performing red meat industry-specific tool and product handling

- Validated motion policy developed through imitation and reinforcement learning, enabling consistent and adaptive knife and product handling by the humanoid robot.
- Demonstration to AMPC showcasing the concept, training approach and capabilities of humanoid robotics.
- Site trial at a red meat processors site.
- Submission of a Stage 2 project proposal (subject to trial outcomes supporting further progression), building on learnings from this trial.