Towards a Robust Starfish Robot for AI Research

Mohannad Alhakami,[1,2] Dylan R. Ashley,*[1,3,4,5] Joel Dunham,*[6] Yanning Dai[1], Francesco Faccio,[1,3,4,5] Eric Feron,[1] and Jürgen Schmidhuber [1,3,4,5,7]

*Equal contribution.

[1]KAUST [2]SABIC [3]IDSIA [4]USI [5]SUPSI [6]optoXense, Inc. [7]NNAISENSE

Artificial intelligence has made great strides in many areas lately, yet it has struggled to achieve a similar level of success in general-purpose robotics. We believe that one of the reasons for this is the disconnect between the traditional robotic design needed by classic control algorithms and the properties needed by a robot to run open-ended, creativity-based AI systems. **To that end, we—taking selective inspiration from nature—have undertaken a journey of developing a robot from the ground up to handle free-form training of cutting-edge AI algorithms.** Specifically, the different designs we experimented with all aim to

- be highly resilient to the stresses needed to train modern neural network-based algorithms, enabling long periods of training without supervision;
- have extreme levels of redundancy to deal with any sensor or actuator loss that could occur when the robot damages itself;
- have a rich action space that enables many ways of reaching the same objective, reducing the likelihood of an algorithm falling into a strong local minima during training; and
- have a rich observation space with strong exploitable regularities (e.g., by having many cameras that can serve the role of many different sensors simultaneously).

The main contributions of the current version of this work are in

- designing and building an extremely robust robotic limb with rich sensory feedback and actuation potential to handle long periods of running advanced machine learning algorithms with minimal human oversight;
- evaluating the aforementioned robotic limb through experiments with two contemporary machine learning algorithms and under simulated sensor failure, demonstrating that our design is suitable to such algorithms; and
- designing and building two extensions of this concept, with the final one having a body and six limbs, with the limbs cumulatively having a total of 48 cameras and 48 motors.

First Prototype

Our first prototype focused on producing a single limb aimed at achieving the objectives mentioned above and particularly **(1)** used partially soft joints to resist repeated motion, **(2)** included a full electronics integration with several actuators and cameras, and **(3)** was able to perform proof-of-concept AI experiments with minimal supervision, despite us simulating sensor failures during execution.

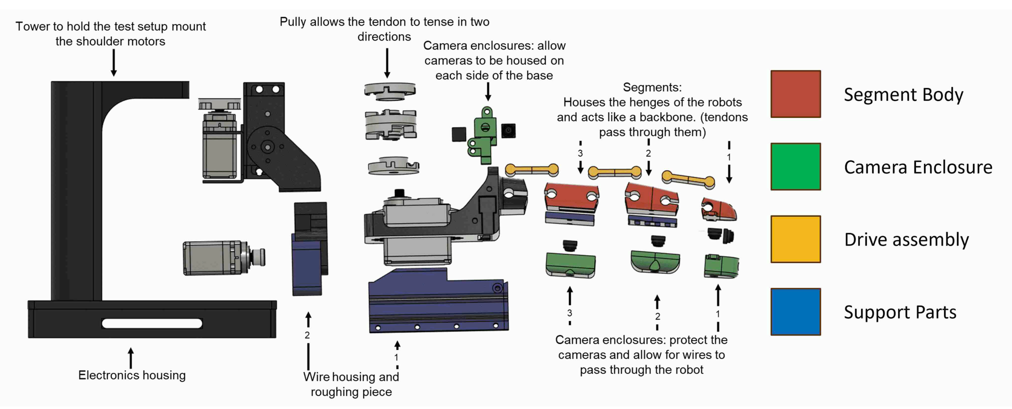

Figure 1: Exploded view of the first prototype with color-coding based on function---segments (red), brackets (black), camera enclosures (green), cable management (blue), and attachments(gray).

Figure 2: The effect on the field of view for the four front-facing cameras when the limb is extended and retracted.

We ran a proof-of-concept target-finding experiment using both the proximal policy optimization and actor-critic algorithms. To test the added robustness offered by sensor failure, in some runs we randomly disabled cameras during execution. In all instances, the prototype was able to outperform Brownian-noise baselines---surviving the long periods of training required to do so.

Table 1: Comparison of Algorithm Performances Across Camera-Kill Settings

Camera-Kill Setting	Average Episode Lengths and Standard Deviation (SD)			
	BMSS	BMMS	PPO	AC
Kill 0	44.64 ± 47.44	16.43 ± 15.35	1.06 ± 0.17	15.65 ± 3.37
Kill 1	59.34 ± 71.98	20.89 ± 21.54	2.32 ± 0.57	18.06 ± 6.93
Kill 1-2	65.13 ± 76.57	26.47 ± 28.41	2.83 ± 0.74	20.23 ± 17.47

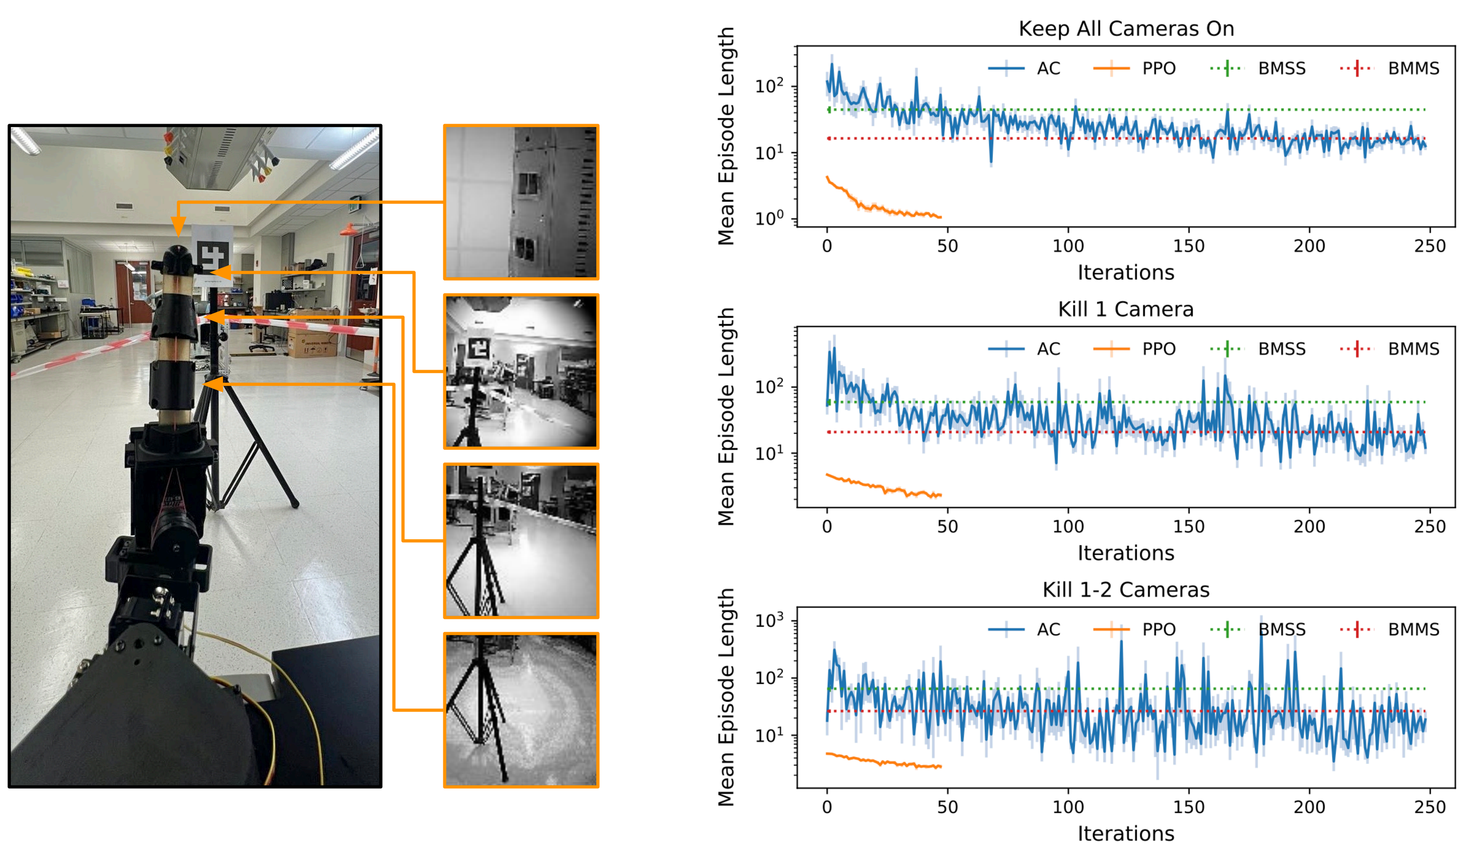

Figure 3: In our experiment, we use the four front-facing cameras and train the robot to move from the central position of the servos to a position where an AprilTag is in sight of one or more of the cameras.

While this experiment showed that the robot was able to achieve the aformentioned objectives to some degree, **(1)** cable routing caused more occasional camera failures, **(2)** the robot was limited in its "softness" and hence somewhat limited in its robustness, **(3)** the tendons provided an obvious point of eventual failure, and **(4)** the number of sensors (4/6 cameras) and actuators (4 servos) was still somewhat limited with little obvious angles of expansion.

Second Prototype

Our second prototype was designed to **(1)** eliminate the need for comparatively fragile tendons; **(2)** increase the number of useful cameras, all having useful lines-of-sight; and **(3)** improve the flexibility and maneuverability of the robot.

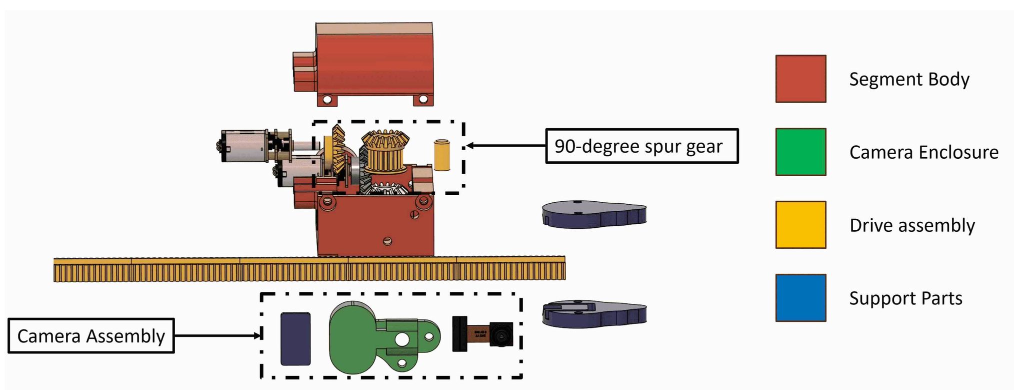

Figure 4: Exploded view of the second prototype showing the drive assembly (yellow), segment (red), camera enclosure (green), and the routing and management components (blue).

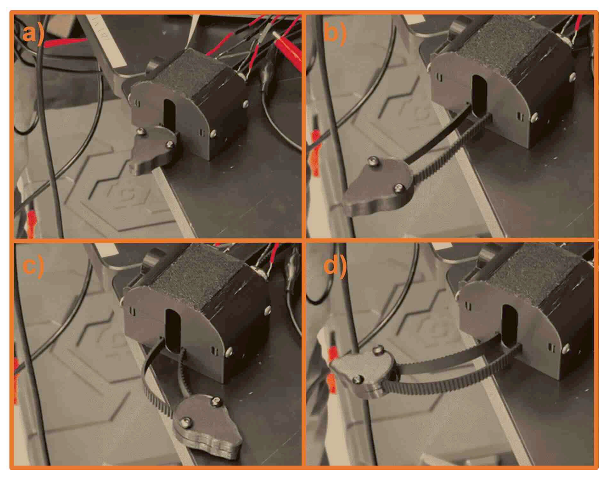

Figure 5: Examples of the second prototype's range of motion.

Although successful, this prototype proved difficult to have move and this movement frequently resulted in interference between and tangling of the many wires.

Final Prototype

Our final prototype was designed to resolve the issues encountered in developing the previous prototypes. This prototype **(1)** is a full robot with 6 arms and a body, giving it the ability to move itself and manipulate objects; **(2)** has mostly flexible tentacles with a wide range of motion; **(3)** has an extreme number of simultaneous camera feeds (48) and motors (48); and **(4)** is equipped with an onboard Jetson Orin Nano 8GB to enable simple AI experiments and off-robot networking capabilities to enable more complicated ones. The maximum power draw of this version is 638.76 watts and the expected power draw is 226.92 watts. We noted that the startup sequence of such a complex prototype is paramount to prevent looping resets.

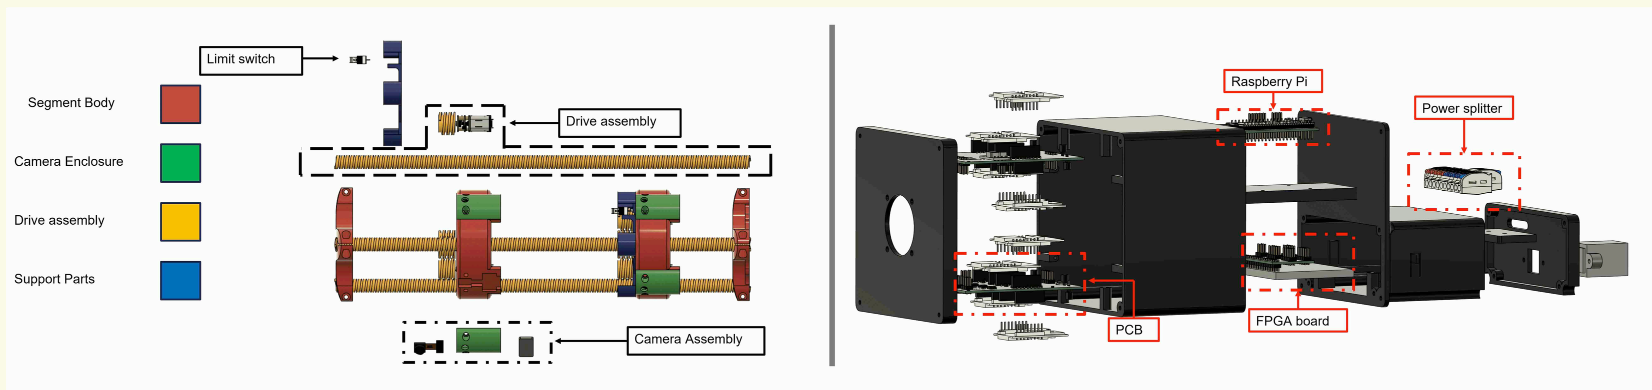

Figure 6: Exploded views of the final prototype showing the full tentacle assembly.

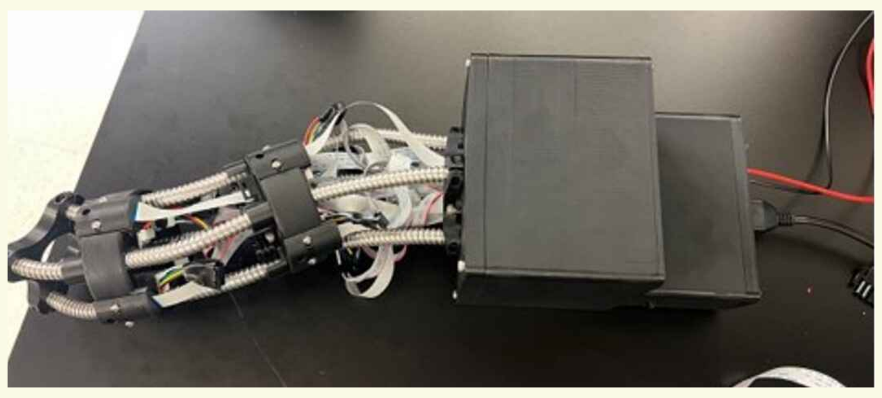

Figure 7: A fully assembled tentacle.

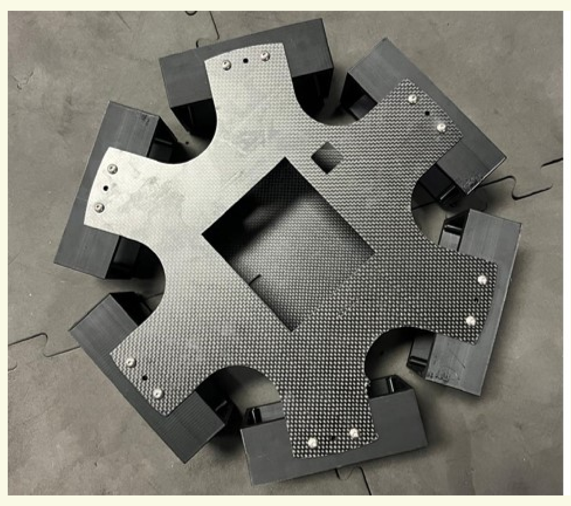

Figure 8: The body of the robot that serves to connect the tentacles, perform networking, and house the Jetson.

Figure 9: The principal biological inspiration for the final prototype's design.

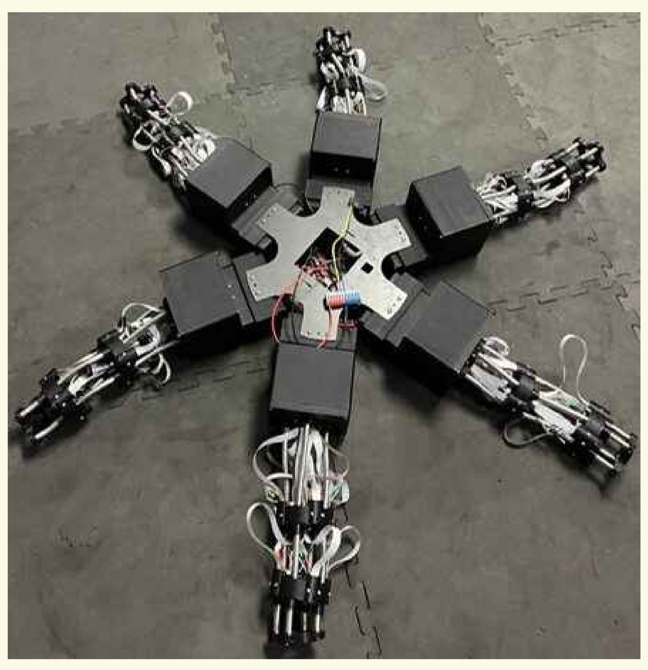

Figure 10: The full assembled final prototype.

This fully-functional prototype achieves the aforementioned objectives better than the previous prototypes. Future work will **(1)** use this as a platform to experiment with creativity-driven AI algorithms, **(2)** enclose the robot with a soft skin-like cover to further increase the robustness of the robot, and **(3)** further increase the richness of the sensor and actuation space by elongating the tentacles.

This work was supported by the European Research Council (ERC, Advanced Grant Number 742870) and the Swiss National Science Foundation (SNF, Grant Number 200021 192356).