

Predictive Learning of Error Recovery with a Sensorised Passivity-based Soft Anthropomorphic Hand

Kieran Gilday¹, Thomas George Thuruthel¹, and Fumiya Iida¹

¹Affiliation not available

February 20, 2023

Abstract

Manipulation strategies based on the passive dynamics of soft-bodied interactions provide robust performances with limited sensory information. They utilise the kinematic structure and passive dynamics of the body to adapt to objects of varying shapes and properties. However, these soft passive interactions make the state of the robotic device influenced by the environment, making control generation and state estimation difficult. This work presents a closed-loop framework for dynamic interaction-based grasping that relies on two novelties: (i) a wrist-driven passive soft anthropomorphic hand that can generate robust grasp strategies using one-step kinaesthetic teaching and (ii) a learning-based perception system that uses temporal data from sparse tactile sensors to predict and adapt to failures before it happens. With the anthropomorphic soft design and wrist-driven control, we show that controllers can be generated robust to novel objects and location uncertainty. With the learning-based high-level perception system and 32 sensing receptors, we show that failures can be predicted in advance, further improving the robustness of the entire system by more than doubling the grasping success rate. From over 1000 real-world grasping trials, both the control and perception framework are also seen to be transferable to novel objects and conditions.

Corresponding author(s) Email: f224@cam.ac.uk

ToC Figure

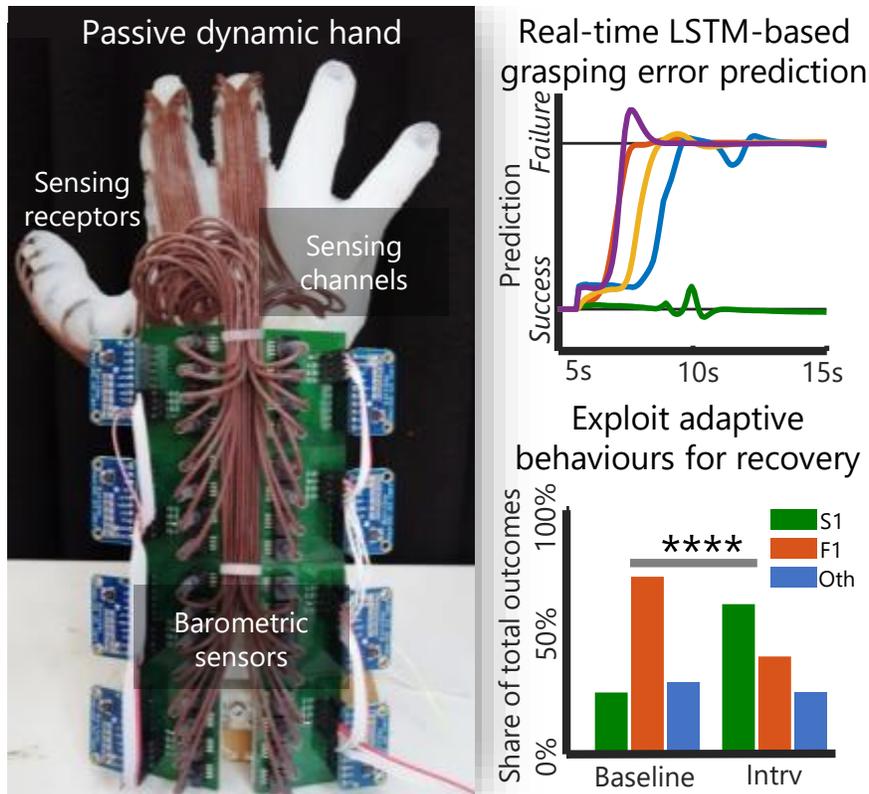


Figure 1: **ToC Figure.** A bio-inspired hand augmented with a novel soft-sensorised skin has been developed. Deep learning and large-scale physical experimentation allow robust and early predictions of task outcomes, such as grasping, with high-level intuitive feedback. The need to develop passive dynamic behaviours within robotic hands is showcased by exploiting emergent behaviours and early error predictions for 144% increase in grasping success rate.

Introduction

Grasping and manipulation using soft robotic hands has demonstrated notable advantages over more traditional rigid hands, especially in passive adaptive grasping (Piazza et al., 2019; Hughes et al., 2016; Bicchi, 2000; Mattar, 2013). Soft hands can readily conform to the environment, guaranteeing some robustness to inevitable sources of uncertainty in the environment (Eppner et al., 2015; Gilday et al., 2021; Hughes et al., 2018; Zhou et al., 2019). Hands such as the Pisa/IIT SoftHand (Catalano et al., 2014) utilise underactuation and muscle synergies (Alessandro et al., 2013) to greatly reduce control complexity for power grasping, even with substantial environmental uncertainty (Della Santina et al., 2018). Furthermore, soft hands open up additional functionalities that can be exploited for versatile manipulation strategies like the use of environmental constraints (Eppner et al., 2015) and programmable anisotropic stiffness properties (Hughes et al., 2018; Wang et al., 2022). When it comes to dynamic manipulation tasks, soft robotic designs have shown surprisingly robust behaviours. The RBO hand 3 (Puhlmann et al., 2022; Bhatt et al., 2022) displays robust manipulation capabilities when reorienting variable objects with a manually trained sequence of finger moves. This contrasts to other more rigid robotic designs which require millions of iterative learning episodes to generate robust control policies, e.g., with the Shadow hand (Andrychowicz et al., 2020).

Even with the robustness and adaptability that soft robotic hands provide, sensory feedback is still vital

for more complex and dexterous manipulation tasks (Dahiya et al., 2019; Yousef et al., 2011; Homberg et al., 2019; Zhao et al., 2016). In order to leverage the full capabilities of a soft hand’s underlying passive adaptive behaviours, the sensing and control systems must augment the existing soft structures and leverage the same behaviours to enhance perception (Della Santina et al., 2019; Shih et al., 2020). Soft robotic technologies for tactile sensing is a wide researched field with numerous technological solutions (Wang et al., 2018; Yousef et al., 2011; Kappassov et al., 2015). Tactile sensing in soft robotics generally suffers from issues of nonlinearity, low receptor density and/or lack of modularity and softness, preventing wide adoption of any single solution (Wang et al., 2018; Zhao et al., 2016). Distributed tactile sensing arrays go some way to solving the low-density problem, however are usually confined to flexible PCBs (Dahiya et al., 2009; Mouri et al., 2002) or soft sensors with highly nonlinear characteristics (Thuruthel et al., 2020; Yin & Müller, 2021). For tactile sensing with a soft interface and high effective density, vision-based sensors such as the gel-sight or tactip sensors achieve impressive tactile resolution, though in a constrained form factor (Yuan et al., 2017; Ward-Cherrier et al., 2018). Another novel solution is the use of sound-based sensors within the RBO hand 2 (Zöllner et al., 2018), which reads braille from the interference patterns of sound within its soft pneumatic fingers. Several recent works have shown how these soft sensors can be modelled for low-level state estimation (Thuruthel et al., 2019; Truby et al., 2020; Kim et al., 2018; Zhang et al., 2020; Wang et al., 2020). The main challenge in the field currently is the development of sensory models that can provide high-level sensory information for closed-loop control.

Execution of robust grasp policies require two components; a grasp planning algorithm and a reactive component that adapts to uncertainties in the environment. Model-driven approaches use models of objects and robot physics to optimally plan stable grasp strategies (Bicchi, 1995; Okamura et al., 2000). Due to the large variability and uncertainty in environment/robot parameters, data-driven approaches are more recently preferred (Bohg et al., 2013). Even then, most of these grasp strategies rely on quasi-static assumptions. The human grasp strategy heavily relies on the passive properties of the hand. They are characterized by dynamic interaction-based strategies that require little visual-feedback (Eppner et al., 2015). These morphological principles have been applied to develop soft robotic hands that require minimal grasp planning and visual processing to solve various tasks (Al-Ibadi et al., 2018; Krahn et al., 2017). Recent works have shown that these soft robotic designs also provide can leverage human knowledge to generate one-step grasp strategies, owing to their anthropomorphic nature (Gilday et al., 2021). With internal models and contact information, these open-loop strategies can be converted to sequential closed-loop strategies, making the controller more robust to varying objects (Pollayil et al., 2022).

Actions affect our cognition, so high-level decision-making should not be decoupled from low-level control (Lepora & Pezzulo, 2015). Error detection and recovery gives some coupling, which has potential to improve manipulation capabilities. For grasp adaptation, there has been promising investigation into detecting ‘anomalies’ during manipulations (Park et al., 2016; Zhang et al., 2020), training error detection using sequential neural networks with visual inputs (Lee et al., 2019), and frameworks for error recovery (Nakamura et al., 2018; Della Santina et al., 2019). Though, there is notable lack of investigation in this area with soft hands. Slip detection and prevention is one of the most common grasp adaptation strategies (Ajoudani et al., 2016; Dong et al., 2017; Chen et al., 2018; Veiga et al., 2018; Averta et al., 2022). However, failure due to slip is just one of the many failure modes that occur after grasping and requires dense tactile information and high control bandwidth for adaptation. For passive soft interaction-based designs, like in the case of passive prosthetic devices, such high sensory and control bandwidth is rarely available. Moreover, there are other modes of failure that can arise even before full closure with the object. Recent works have worked on identifying grasp quality using dense vision-based tactile sensors on static grasps (Hogan et al., 2018; Calandra et al., 2018).

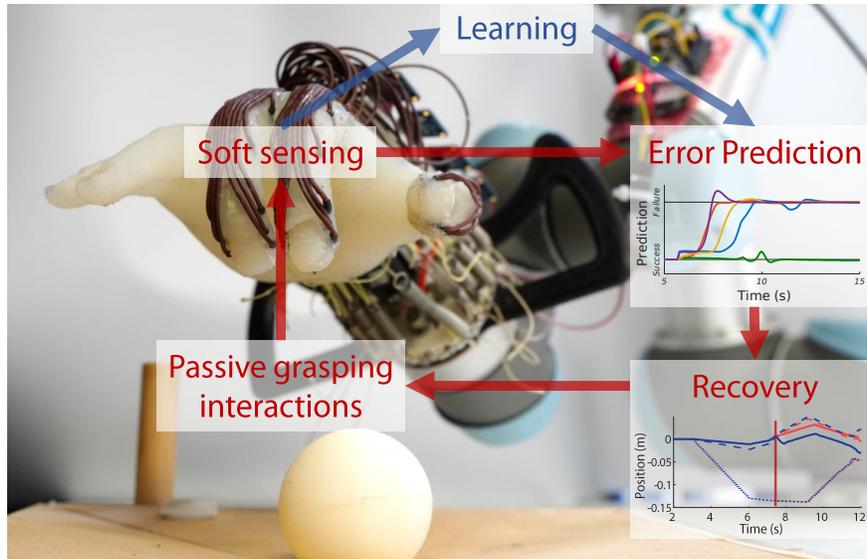


Figure 2: **Error detection and recovery from passive perception.** Demonstrated with a wrist-driven soft hand—which achieves grasping through sequential hand-environment interactions rather than any internal actuation—prediction of future errors in an open loop grasp can be learned using exteroceptive and proprioceptive information from a barometric sensing skin. With a self-resetting environment, large scale experiments can generate training data and evaluate error recovery performance.

In this work, we present predictive grasp quality estimation in a soft passive anthropomorphic hand with sparse sensory data, Figure 2. We trained a soft, passive, anthropomorphic hand with embedded soft sensors to detect and recover from errors even before they cause failures. Interaction-based wrist-control (Gilday et al., 2021; Gilday et al., 2021) is augmented with a form of active perception by a soft skin with 32 soft barometric sensors. Our framework demonstrates adaptive grasping capabilities while being able to reproduce multiple grasp classifications on a wide variety of objects when trained on a single object with a single trajectory. We study the grasping performance of the hand in a self-resetting environment, which allows large scale experimentation. We observe emergent grasping failure and success modes under real and artificial perturbations. A Long Short-Term Memory (LSTM) network is used to predict real-time failure and success from few trials, using exteroceptive and proprioceptive data from our soft, modular sensors. Real-time error predictions and a heuristic error recovery routine is implemented and compared to grasping with no feedback, resulting in an 144% improvement in success rate. In addition, we demonstrate the system’s generalizability when grasping and predicting errors with unknown objects and interference.

Rich media available at https://drive.google.com/file/d/1Cwcv27zuYqI_5ddNtxoo5-drToZp_Xec/view?usp=sharing

The main contributions of this work are:

- The anthropomorphic hand and sensor design for robust and adaptable object grasping with tactile feedback
- An early grasp error prediction algorithm using an LSTM network and the sparse tactile data for the described open-loop grasping strategy
- A method of error recovery which exploits early error predictions and passive hand behaviours to improve grasping success rate

Experimental Section/Methods

Passive hand design

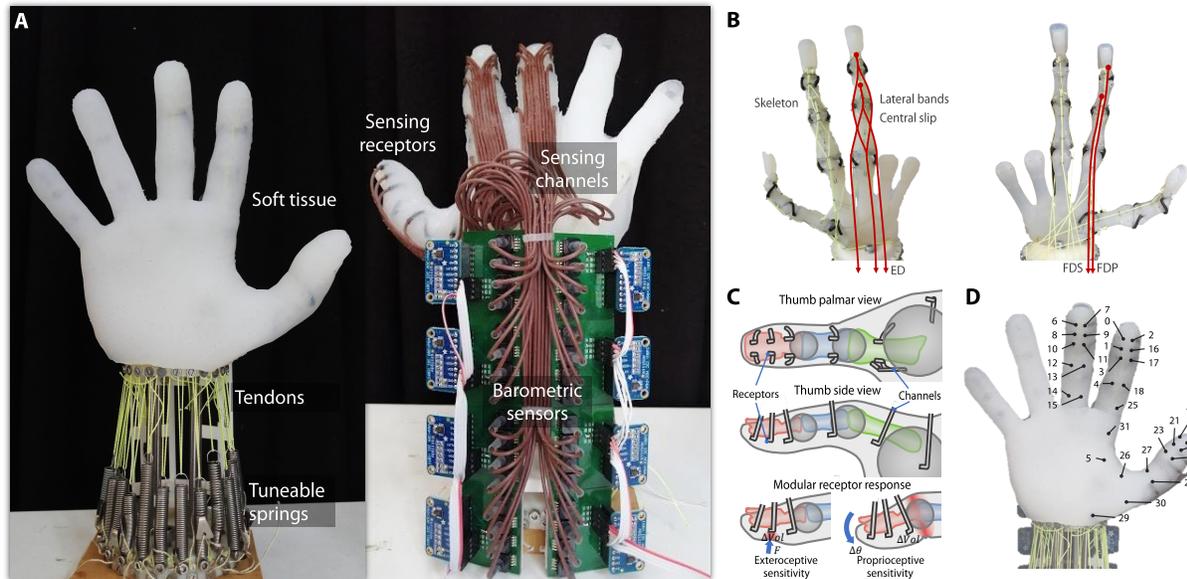


Figure 3: **Passive hand and modular sensor design.** (A) Passive anthropomorphic hand design with sensorised skin. Receptors are air chambers embedded throughout the soft skin tissue, where density, placement, and geometry are highly customisable. Receptors are coupled to surface mount barometric sensors through pneumatic channels. The hand is tendon-driven, with each tendon connected to re-configurable springs with tuneable pretension and stiffness. (B) Hand skeleton and tendon routing. Anatomical joints constrained by ligament pairs and 5 tendons per finger. (C) Modular receptor design. Modelled air channels connect receptor air pockets to ports on the rear of the skin for interfacing with pneumatic tubing. Sensitivity to exteroceptive force or joint deformation can be controlled by receptor placement and geometry. (D) Receptor distribution. 32 total receptors are placed in the thumb, index, middle finger and palm.

The anthropomorphic hand improves upon the design seen in (Gilday et al., 2021), primarily with the addition of sensing capabilities. The hand was adapted from a commercial 3D model purchased from TurboSquid (www.turbosquid.com) (Hughes et al., 2018). Muscles, tendons, and ligaments were removed, leaving only the skeleton. The bones are then modified for manufacturing of ligaments, tendon arrangement and skin molding. Collateral ligaments stabilise each joint (Knutson et al., 2000). For all joints, except the thumb carpometacarpal (CMC), the collateral ligaments form an ‘S’ shape over each side (Çulha & Iida, 2016), this allows rolling and prevents sliding. These ligaments allow some abduction/adduction for the metacarpophalangeal (MCP) joints. Though, the thumb CMC requires a greater range of motion (Degeorges et al., 2005; Liu et al., 2006), therefore an additional ligament is added for stability at larger deflections. Ligament mounting holes are modelled on each bone at the limits of the rolling motion. Volar plates, which prevent hyperextension (Minami et al., 1983), are omitted from this design for reduced complexity and ease of manufacturing (Figure 3B).

Rich media available at <https://3dviewer.net#model=https://dl.dropbox.com/s/3sww5y1bnhx1pqk/wholehand%20v26.step>

Tendon pulleys are modelled as rigid loops embedded into the bones for a five-tendon arrangement. Each finger has four degrees of freedom (DoF), therefore requires a suitable arrangement of five antagonistic

tendons for control (Pollard & Gilbert, 2002). The chosen arrangement is two flexor tendons, one to each of the intermediate and distal phalanx, and three extensor/abductor tendons, one to the distal phalanx while avoiding the proximal-interphalangeal (PIP) joint, the other two connecting to the intermediate phalanx while running laterally on the MCP for abduction/adduction (Figure 3B). This arrangement approximates the flexor tendons and extensor hood of a real finger (Wilkinson et al., 2003; Deshpande et al., 2008). Mounting holes for these pulleys are directly modelled into the bones.

A 3D scan of similar proportions to the skeleton is added, and the bones are aligned within. By subtracting the skeleton and a cavity around each joint (Gilday et al., 2021), the scan becomes a model of the soft tissues of a real hand (Wang & Hayward, 2007). This soft tissue model is where the sensing receptors are modelled. Each receptor is an 'L'-shaped tube with a curve in the stem to run around the bones. Four shorter receptors (2 mm) are placed in each fingertip where space is more constricted, these are aligned at a skin depth midway between the outer surface and bone to minimise any thin walls and are more sensitive to tactile pressure. Four longer receptors (4 mm) are arranged around each joint and placed deeper under the surface and close to the joint cavity, therefore are sensitive to both tactile pressure and skin deformation from joint bending. Two additional receptors are placed in the palm close to the skin surface, one in between the thumb and index finger, the other beneath the thumb. This gives a total of fifty receptor locations over the palm, thumb, index, and middle finger (Figure 3A, D).

One final modification to the skeleton is made, modelling fingernails. These serve dual purpose, one being improved grasping capabilities with fingernail pinching, the other being stabilisation during the skin molding process. For the molding process, the inner mold is modelled from the skeleton without ligament and tendon modifications, only joint cavities and fingernails. The outer mold is primarily a negative of the 3D hand scan, though individual sections for each of the thumb, index and middle finger containing higher detail receptor positives are generated to allow mold assembly.

The final hand component is the wrist-mounting hardware (Figure 3A). Consisting of a hand mounting plate with tendon routing holes, a UR5 robot arm mounting plate with tendon spring anchors and a sensor mounting shield on the rear.

Hand Manufacturing

Rich media available at https://3dviewer.net#model=https://dl.dropbox.com/s/617cltdyp15282h/wholehand_alt2_ani%20v3.step

The first manufacturing step is 3D printing components. Fused deposition modelling with PLA is used for the wrist-mounting hardware and the lower detailed inner and outer molds. For more detailed parts, the bones, and receptor molds, inkjet printing is utilised with a Stratasys Objet500. The bones are printed with Stratasys Durus for strength and toughness, the receptor molds are printed with a blend of Stratasys Rigur (high stiffness) and Agilus (low stiffness) for stability during molding and large elongations before breaking during mold assembly.

For the skeleton assembly process, the bones are printed with a scaffold keeping them aligned during ligament fabrication, these are snapped off once assembly is complete (Gilday et al., 2021). Ligaments are formed from Festo 2 mm flexible tube with shore hardness D52, cut to individual lengths for each joint, then bonded into the modelled ligament mounting holes using Araldite two-part epoxy. Flexible tubing provides a robust flexure joint with low rolling resistance and limited extensibility to reduce joint dislocation. Tendon pulleys are 'U' shapes with 2 mm diameter and 0.6 mm thickness formed of single core, copper, stripped 23AWG wire. These are inserted in the modelled mounting holes and bonded with Araldite epoxy. To resist higher forces, pulleys for the flexor tendons penetrate the thickness of the bone and hook onto the rear side. The final skeleton assembly step is attaching the wrist-mounting hardware and routing the tendons. Tendons tie to anchors at predefined termination points, then are routed through pulleys towards the wrist mounting hardware. Each tendon is looped through a spring and clamped to the hand mounting plate with a screw. Spring stiffnesses can be limited individually using hooks. With this mounting method, pre-tension and

stiffness of each tendon can be tuned for different starting postures and interaction force behaviours (Gilday et al., 2021).

The sensorised soft skin is cast into the 3D printed, assembled mold using Smooth-on Ecoflex 00-30 two part silicone. After mixing, the silicone is vacuumed to remove any air bubbles, then pored into the mold and allowed to cure for four hours. The cast skin is carefully removed from the mold, then placed over the pre-assembled skeleton like a glove. The skin is bonded to the distal phalanx of each finger, underneath the fingernail, using Smooth-on Silpoxy. The skin over the other joints is left unbonded and held in place via friction. For sensor connections, silicone tubing (BS2848, shore hardness: A40, inner diameter: 0.5 mm, outer diameter: 1.5 mm) is inserted to the chosen receptors, ten receptors on each of the thumb, index and middle finger and two in the palm. Tubing is routed on the rear of the finger and bonded to the skin with Silpoxy. As small a tube as possible is preferred, to reduce any added elasticity to the finger motions, allow for higher density of sensor placement and to increase sensitivity by maximising change in receptor chamber volume over total sensor volume.

The final step of hand assembly is the sensor readout. Pressure within each sensing receptor is transduced with NXP MPXH6300AC6U absolute pressure sensors. Analog voltages from each pressure sensor are measured with 16-bit ADS1115 analog to digital converters (ADC) capable of 860 samples/sec. These are I2C devices with four ADC channels and four possible addresses, hence a single I2C bus can support 16 sensor channels. Custom PCBs were manufactured for mounting and connecting the pressure sensors to the ADCs. Eight of these PCB are mounted on the sensor shield on the wrist-mounting hardware, where sensor tubing can be safely connected (Figure 3A, Figure 8). The total of 32 available sensing channels is divided onto two independent I2C buses of an FRDMKL25Z microcontroller. This controller configures each ADC and collates sensor readouts ready for a master device. With this configuration, the maximum sampling frequency of all channels together is 26.9 Hz, though to account for delays in communications and stabilising timings for processing, we sample all channels at 10 Hz.

Wrist-control framework

Grasping with a passive hand relies on wrist-driven interactions with the environment. Wrist control is essential in grasping and manipulation, especially in more constrained environments. Edge grasps are an example of wrist control being used to overcome limitations (Eppner et al., 2015; Della Santina et al., 2017). Techniques for generating similar environmental constraints’ exploitation are generally bespoke algorithmic solutions (Bimbo et al., 2019). These approaches require significant teaching effort, which the wrist-control framework (Gilday et al., 2021) attempts to solve. This is a method for training and adapting grasps based-on sequential interaction classification from only a single demonstration trajectory. The following process generates template trajectories for distinct grasping strategies on familiar objects:

- Kinaesthetic teaching is used to generate a trajectory for a particular grasping strategy on a training object.

$$\mathbf{D} = [x_i, y_i, z_i, rx_i, ry_i, rz_i]_{i=0, \dots, n-1} \quad (1)$$

- The trajectory \mathbf{D} is subsampled into k keypoints at interaction ‘inflection’ points, resulting in a simplified trajectory \mathbf{R} with a discrete set of interactions.

$$\mathbf{R} = [x_i, y_i, z_i, rx_i, ry_i, rz_i]_{i=[kp_0, \dots, kp_{k-1}]} \{ kp \in \mathbb{N} \mid 0 \leq kp \leq n \} \quad (2)$$

- Each interaction is given a classification which defines how it adapts to changing object geometry. Therefore, by inputting a familiar object with changes in its primitive geometry, this template trajectory \mathbf{T} can allow grasping adaptation (Figure 5A).

$$\mathbf{T} = \mathbf{R} + \Delta$$

$$\Delta = [\delta x_i, \delta y_i, \delta z_i, \delta r x_i, \delta r y_i, \delta r z_i]_{i=[kp_0, \dots, kp_{k-1}]}$$

(3)

Figure 4B shows an example trajectory with the wrist position, orientation, and sensor readout. The sensor readout enhances key point extraction and classification, which previously was processed manually, and with the additional sensor information has the potential to be automated.

Grasping and manipulation via sequential interactions can not only improve robustness and simplify control with a soft hand (Pollayil et al., 2022), it can add redundancy (Gilday et al., 2021) and extend behavioural diversity (Hughes et al., 2018). For successful robotic grasping, a minimum amount of knowledge is required about itself and the environment. With a passive adaptive hand, this information requirement is offset by the information gain through environmental interactions, therefore success is tolerant to significant deviations in self and environment state (Gilday et al., 2021). However, large uncertainty in the initial states can cause failures in later interactions, this is where additional information gain can keep the robot within the range of tolerated states.

Experimental setup

To facilitate large scale grasping experimentation, the hand is mounted on a UR5 arm and the grasping environment is self-resetting. Figure 2 shows the robot and environment. By grasping the sphere and releasing onto a plate with a shallow recession, the sphere always returns to the centre ready for the next trial. Since the hand is passive, releasing the object is achieved by pushing into a fixed dowel as part of the grasping trajectory. A grasp is considered a success if the object remains in the hand up until the point where it is knocked out by the dowel.

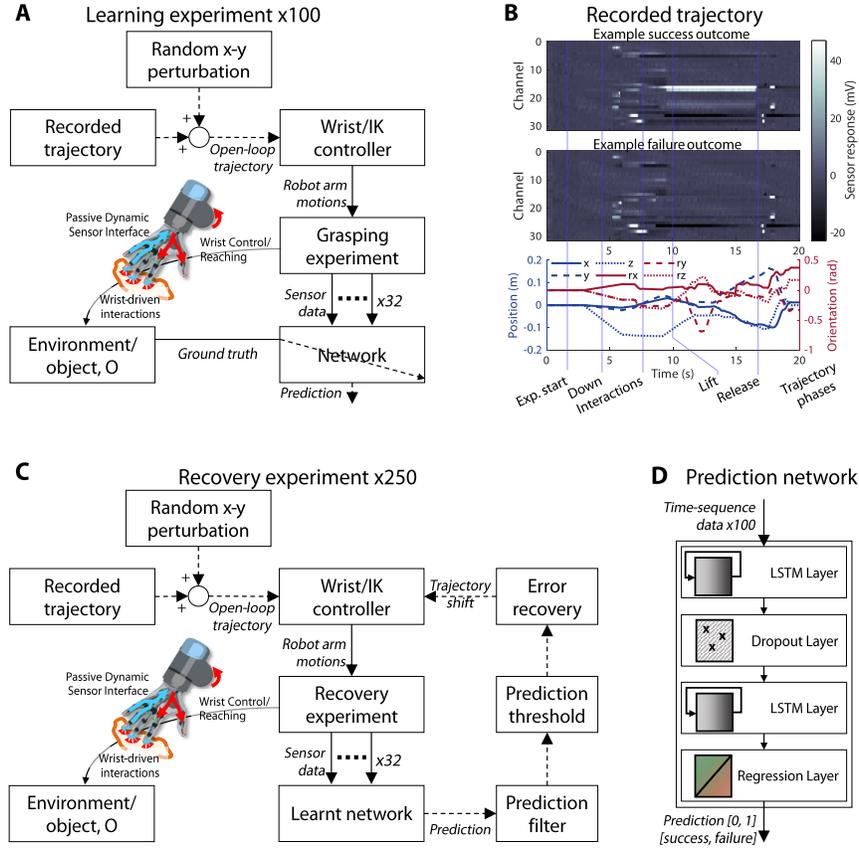


Figure 4: **Experimental setup and passive grasping/error recovery control architecture.** (A) System diagram for data collection. Learning experiment takes a recorded trajectory for wrist-driven grasping, adds perturbation, then records sensor data and grasp outcome. (B) Example recorded trajectory and sensor time series data through phases of reaching, interacting and releasing the object. (C) Modified system diagram for error recovery experiment, network predictions lead to error recovery actions once the prediction is ‘certain.’ (D) Prediction network architecture, trained on data from the learning experiment and outputting outcome predictions in real-time.

Many grasping strategies can be chosen for the same object (Gilday et al., 2021). Different strategies have strengths and weaknesses under different environmental constraints and uncertainties. The primary grasping strategy in this paper is chosen such that failures are relatively common (Figure 6C) but occur with little dependence on injected noise. The trajectory is recorded using the wrist-control framework, where kinaesthetic teaching gives a six DoF time series \mathbf{D} sampled at 10 Hz on a central PC which is them sub-sampled into seven keypoints (Figure 5B). The simplified trajectory can be replayed as a set of waypoints. Artificial noise is added in the form of a constant shift to the trajectory in the plane of the table, ($r \sim \mathcal{U}(0, 5mm)$, $\theta \sim \mathcal{U}(-\pi, \pi)$, $\mathbf{R}^n = \mathbf{R} + [\mathbf{r} \cos(\theta), \mathbf{r} \sin(\theta), \mathbf{0}, \mathbf{0}, \mathbf{0}, \mathbf{0}]$), essentially adding uncertainty to the starting location of the object. This noise is distributed radially, with uniform probability between 0 and 5 mm distance from the object origin.

Simultaneous to recording and replaying trajectories, sensor data \mathbf{S} is streamed to the central PC and video is recorded from a webcam. To ensure reliable sampling rate and data synchronisation, a central PC processes and collates all data. The result is a synchronised AVI file and tabulated data file (CSV) with timestamp,

robot pose and sensor readout.

$$\mathbf{S} = [s_0, s_1, \dots, s_{31}]_{i=[0, \dots, n-1]} \quad (4)$$

Grasping experiments can now be performed in rapid succession with minimal oversight. Six hundred fifty grasping trials are performed over 5 days, with different environmental conditions such as temperature, which offsets the barometric sensor readings. Footage of each trial is manually reviewed, and the final outcome is labelled. The label corresponds to the classification of outcome which is qualitatively evaluated, e.g., failure due to missing the object or failure due to the object getting trapped (Figure 6B). This dataset is used to train the network. With the network trained, sensor readout at 10 Hz is forwarded to the recurrent network, which returns a live prediction of grasping success. Trials can be run with or without error recovery enabled. The network and error recovery is tested with 250 trials (Movie S4), half have recovery disabled as a control, the other half error recovery is enabled.

Network Architecture and Training

The network for predicting failure is designed to perform a sequence-to-one regression. The architecture of the deep LSTM network is shown in Figure 4D. The network is made of two layers of LSTM units with a size of 60 and 15, with a dropout unit in between. The first LSTM layer performs a 'sequence-to-sequence' transformation and the second LSTM layer performs a 'sequence-to-one' transformation. The dropout unit has a dropout probability of 0.2. The network creation and training is done on the MATLAB Deep Learning Toolbox.

Six hundred and fifty sequence data is collected for training. All the sequence has its corresponding scalar outcome variable. Here, we give a value of 1 for failure and 0 for success, in order to convert the learning to a regression problem. The data set is augmented to a larger set, eighty times its size, by clipping the sequences randomly within a range of 5 to 15 s. Time-exponentially decreasing noise is also added to the sensor data. This data augmentation forces the network to predict the grasp outcome as early as possible without overfitting to the small dataset. Six hundred and ten samples are used for learning and the rest for testing.

Error recovery

The error recovery process allows us to test the real world implications of error detection. Figure 4C shows the modified wrist-control framework for error recovery. During an activity, the sensor readout is fed to a network which generates live outcome predictions. When the network is certain of failure, the output regresses to a steady '1', if the network is certain of success, the output regresses to a steady '0'. If there is uncertainty, the value may oscillate. The 'certain' states generally occur late in the trajectory, often times after failure or potential for failure has passed.

To perform an error recovery process, a decision has to be made whether to act or not during the trajectory. In order to obtain a prediction as soon as possible, analysis is performed to extract a prediction to act upon when network certainty is lower. The network output \mathbf{Z} is passed through a first order digital filter, with transfer function coefficients $a = [0.5 \ 0.25 \ 0.125 \ 0.0625 \ 0.03125 \ 0.03125]$, then the decision \mathbf{Z}^d is made if the filtered value is within 0.2 of each class and the derivative is within 0.2 of 0.

$$\mathbf{Z}(\mathbf{i}) = F(S_{0,\dots,i})$$

$$\mathbf{Z}^f(\mathbf{i}) = \sum_{j=0,\dots,5} a_j \mathbf{Z}(\mathbf{i} - \mathbf{j})$$

(5)

$$\mathbf{Z}^d(\mathbf{i}) =$$

$$\begin{cases} 1 & 0.8 < \mathbf{Z}^f(\mathbf{i}) < 1.2 \text{ and } -0.2 < \dot{\mathbf{Z}}^f(i) < 0.2 \\ 0 & -0.2 < \mathbf{Z}^f(\mathbf{i}) < 0.2 \text{ and } -0.2 < \dot{\mathbf{Z}}^f(i) < 0.2 \end{cases}$$

(6)

For the error recovery process, a heuristic strategy is chosen since. From the training data, one mode of failure dominated, with $\approx 50\%$ of total outcomes, therefore the recovery strategy targeted this. In this failure, mode ‘F1’ (Figure 6B), the sphere is trapped under the index finger and ejected before the grasp can be modulated by thumb contact. A simple hypothesis for recovery is to relieve pressure and bring the thumb into contact with the sphere before the elastic force within the index finger builds up too high. This can be achieved with a translation (away from the camera in Figure 7A). If this intuitive recovery technique is successful, then error detection is demonstrated to have great potential in robust grasping and manipulation when used in conjunction with human intuition or more advanced recovery processes which can account for types of failure, or even used as part of a training process itself.

Statistical analysis

Statistical analysis is performed on the results to test the significance of the findings. A Chi-squared test is performed on the outcomes of the error recovery trials compared to the control trials (Figure 7B). The significance of the change in rates of the success and failure modes is tested to a 0.00001 significance with a single-ended p-value.

Results

The soft anthropomorphic hand

Embedding mechanical intelligence within a robotic hand increases greater behavioural diversity through mechanical complexity and redundancy (Pfeifer & Scheier, 2001; Pfeifer et al., 2007; Hughes et al., 2018). In addition, *intelligent* behaviours are readily exploited by cheap control (Montúfar et al., 2015), especially in common repetitive tasks such as grasping (Gilday et al., 2022). This is key for developing intelligent and general-purpose hands for use in industrial, social and prosthetic devices (Gilday & Iida, 2022; Deshpande et al., 2011). The human hand provides the ideal starting point for the design of such intelligent systems. Not only are they highly generalised and adaptable, but much of the modern world is designed with the specific ergonomics of human hand interactions (Mattar, 2013).

For this investigation into passive adaptive grasping and error prediction, we have developed a sensorised soft hand with an anatomically accurate structure. The replica synovial joints (Figure 3) ground the hand design with biological basis (Frost, 1994; Lepora et al., 2013; Gilday et al., 2021). These types of joints are the most common and most mobile joint seen in mammals, however, are exceedingly rare in robotics. In part due to the inherent instability of floating bones versus pin or socket joints. While synovial joints may have arisen due to biological constraints, they demonstrate significant advantages in terms of efficiency, customisability and resilience (Mow & Lai, 1980). Figure 3B shows the underlying skeletal structure of our passive hand. This design allows significant ranges of motion, including thumb abduction, resilience to impacts and dislocation, while retaining controllability and passive-adaptive behaviours (Figure 10 adaptive grasping).

Table 1: **Sensor range characterisation.** Sensor minimum and maximum forces observed with a 4mm indenting tool and two different sensor designs.

Sensor	$\Delta V_{min}(mV)$	$F_{min}(mN)$	$\Delta V_{max}(mV)$	$F_{max}(mN)$
High volume, shallow	1 ± 0.5	23 ± 6	172 ± 1	1600 ± 200
Low volume, deep	1 ± 0.5	99 ± 20	132 ± 0.5	5800 ± 90

The sensorised skin receptors act according to the ideal gas law. Air chambers located within the soft skin are deformed under tactile loads or joint bending (Figure 3C). The change in volume corresponds to a change in pressure measured by barometric sensors mounted on the wrist (Figure 3A). The skin has a resolution of 0.69 mN with a 16-bit ADC, a force range of 23 mN-5800 mN, and a response time up to 1 ms (Table 1).

Passivity and anthropomorphic design provides one-step stable adaptive grasps

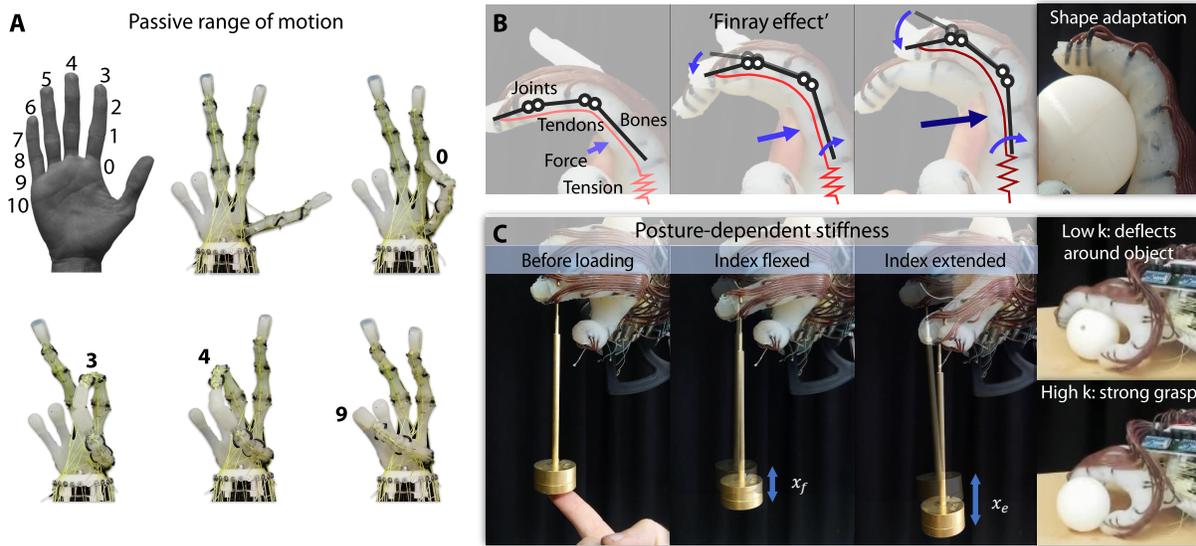


Figure 5: **Passive hand motions and adaptive behaviours.** (A) Skeleton range of motion in the Kapandji thumb test (Kapandji, 1986). Hand posture can be preset by tuning tendon springs (Fig. 3A). High test score excluding ring and little finger contributions. (B) Passive shape adaptation. Similar to the ‘finray effect’ (Pfaff et al., 2011), forces towards the base of a finger cause negative bending. As the MCP joint is extended, flexor tendon tension is increased, causing bending in the PIP and DIP joint. (C) Complex and anisotropic stiffness for diverse interaction, e.g., abduction/adduction stiffness varies with finger extension due to joint geometry (Gilday et al., 2022). Lower stiffness when extended allows the finger to deflect, higher stiffness when flex provides higher passive grasping force.

A previous study demonstrated adaptive grasping with the skeletal hand under a framework for passive grasping with sequential interactions through control of the wrist (Gilday et al., 2021). The skeletal hand is able to grasp spheres with diameter 25-75 mm from single demonstration with an overall success rate of 57.7%, with the wrist adaptations contributing to an 86% increase in success rate over simple open-loop control.

Figure 5 outlines some of the capabilities of our passive system. For use in a diversity of tasks, high posture control is desired. Using the Kapandji thumb test as an evaluation of thumb range of motion (Kapandji, 1986), Figure 5A shows successful posture setting to extreme positions in the test (some positions untestable). Additionally, intelligent behaviours in the hand can be exploited for more robust grasping or enabling new interactions. Figure 5B shows the curling behaviour of the index finger. Due to the tendon layout and joint design, when pressing a finger towards the base, the tip deflects towards the force. This negative bending, similar to the finray effect (Pfaff et al., 2011), enhances passive shape adaptation. In Figure 5C, posture dependent stiffness is demonstrated. Due to bone geometry at joint interfaces, abduction/adduction in the flexed position is restricted relative to the extended position, similar to a human hand (Degeorges et al., 2005). This can be exploited during grasping, where when forming grasps, lower stiffness is useful, then when holding, higher stiffness is preferred.

To evaluate adaptive grasping augmented by the soft sensorised skin, learned trajectories can be tested on everyday objects of similar size to the training object used during kinaesthetic teaching. The hand, in a partially closed position, is driven through a series of interactions with a 60 mm sphere (Movie S1). This trajectory has been simplified by interaction-based key point extraction (Gilday et al., 2021). One limitation

from the wrist-driven framework is the subjective manual trajectory processing, the presence of sensory information can enable automation and objectification of the manual steps such as keypoint extraction and interaction classification (Gilday et al., 2021).

Table 2: **Novel object grasping results.** Adaptive grasping trial objects, Figure 9, weights, dimensions, success rates and prediction accuracies. *Deformable.

	50mm sphere	70mm sphere	Peach	Battery	Block	DC charger	Grapes	Bubb
Weight/g	18	34	66	188	64	86	34	8
Dims./mm	50	70	67x77	30x46x86	35x43x94	50x59x74	18x140*	23x67
Success rate	1	0.4	0.1	0	0.3	0.5	0	0.4
Pred. accuracy	1	0.5	0.6	0.9	0.7	0.6	0.4	0.2

Table 2 shows the grasping success rates of each test object. Three of fourteen objects have no successes, either too heavy and small (battery), too small and deformable (grapes) or too large and concave (spool). These grasps may not be possible with this particular hand starting posture and trajectory, as grasping force in this case depends on elastic forces from index finger and thumb deformations. The remaining objects have variable success rate. Note that no information about the object and its pose is provided to the controller, showcasing the adaptability of the hand design. Success rates can be improved by visual feedback, wrist adaptations and error predictions, especially as some of the worst performing objects are orientation dependent unlike the training object. The wooden block, ground coffee and computer mouse geometries are all anisotropic and successes are only seen in favourable starting orientations, this is a problem that can be solved at the grasp planning stage through vision or exploration of the object with tactile feedback.

Notable objects which showcase the passive hands capabilities include the highly deformable bubble wrap, which has a high success rate when aligned with the hand; the smaller sphere, with a 100% success rate; the ground coffee bag, which is much heavier (550%) than the teaching spheres; and the wooden block, whose lack of smooth, curved surfaces doesn't prevent successful grasping in all cases (Grasping success: Movie S2, grasping failure: Movie S3).

Recurrent neural networks can predict grasping outcome before lift-off

The trajectory in Figure 5B is used for the majority of experiments. We expect the hand to be able to adapt and grasp from different starting states, such as under object position uncertainty. This recorded grasp policy is perturbed to generate training data and can be used for large-scale grasping experimentation, Figure 5A and C.

As the grasp policies are open-loop trajectories, the sensor data history from the onset of contact is indicative of future outcomes. We use an LSTM network (Figure 4D) to predict the grasp outcome in real time, assuming future actions are fixed. Even though the labelling of the dataset for training can only be done after the end of each grasp episodes (Figure 4A), our training framework drives the network to predict grasp outcomes as early as possible (See Materials and Methods). The speed of detection and its certainty depends on the nature of failure. An example of the network prediction for different kind of grasp outcomes are shown in Figure 4B.

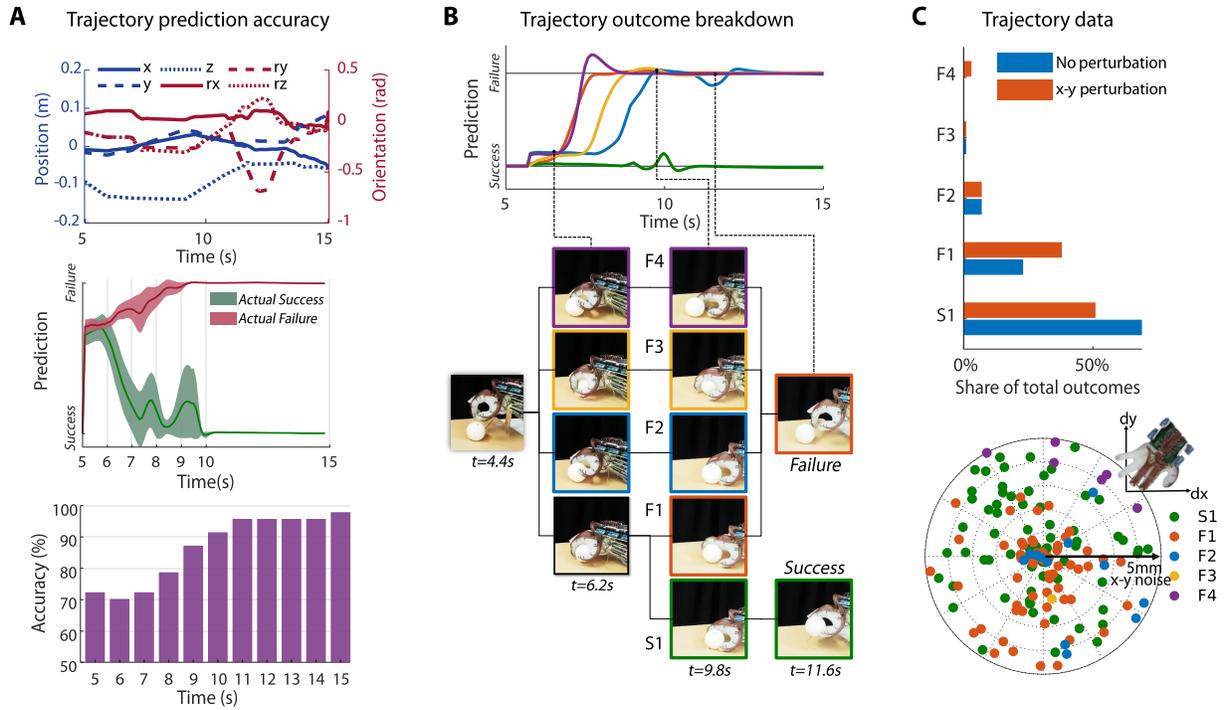


Figure 6: Learning experiment results. (A) Prediction accuracy over trajectory phase (upper). Further into the trajectory, predictions stabilise (middle) and become more accurate (lower). (B) Different grasping outcomes are classified manually by the observed interactions. A single success type was observed, four failure types were observed. (C) Effect of perturbations on the learning dataset. With no perturbation, success outcomes were the majority (upper). With perturbation added, outcome variance increases, some failures types are dependent on the direction of perturbation (lower).

The prediction accuracy with respect to the length of sensory data on the test set provides an insight on the highest accuracy achievable with the current setup and training data (Figure 6A (lower)). Given the whole sequence of sensory data after lift-off (time > 10 s), the learned network can predict the final outcome with very high accuracy. Note that this is not a trivial problem because the sensors are significantly affected by non-linearities like hysteresis and drift in the material and the temperature surrounding the skin. The prediction accuracy reduces as the length of the sensory data is reduced, as expected. However, most of the failure cases are still predicted early and with higher confidence, which is vital for error recovery and robustness. Figure 6A (middle) shows the average predictions for the test set, labelled with the actual outcome, where it can be observed that failure predictions are more accurate and robust. Even though convergence of the predictions happens at a later stage, by observing the trends of the prediction, we can make reasonable predictions of the outcome early enough that it can be used for error recovery using low-bandwidth controllers.

Large scale grasping experimentation reveals multiple failure modes with simulated noise

The passive anthropomorphic hand has a vast number of possible grasps. Not only are there multiple grasp type possible (Feix et al., 2015; Cutkosky & others, 1989), but also multiple possible trajectories to achieve the same grasp type. Limiting the hand to a single starting posture (Figure 3A) narrows down the grasping possibilities, however there is still a multitude to choose from.

Table 3: Effect of noise on experimental trajectory (Figure 6). Success and failure mode counts for experiment trajectory, with and without added noise.

	S1	F1	F2	F3	F4
Count (no noise)	69	23	7	1	0
Count (noise)	51	38	7	1	3

An exploration of grasping trajectories is performed, with the criteria that we find a trajectory with a significant proportion of failure modes, therefore less training and test data is needed to observe failures and recovery. Figure 4B shows this trajectory. Figure 6C and Table 3 show the results from a trial of 200 grasps; the first 100 trials follow the same trajectory without artificially added perturbation, and the last 100 have a random perturbation applied in the plane of the table to simulate uncertainty in grasping location. The first 100 trials have a success rate of 69%, which drops to 51% with perturbation. Failures are distributed over four distinct modes (Figure 6B, Figure 12 left), with the first failure mode being the mode common. Figure 6C (top) shows the outcomes of these 100 trials over the object starting locations. The presence of failures under the no noise condition suggests there are other sources of uncertainty, the most likely being the hand starting posture, where joint and tendon friction can result in position hysteresis.

Table 4: Effect of noise on alternate trajectory (Figure 11). Success and failure mode counts for test trajectory, with and without added noise.

	S1	S2	F1	F2
Count (no noise)	100	0	0	0
Count (noise)	64	3	32	1

During exploration, a second trajectory of note was found. This trajectory (Figure 12 right) demonstrated a much higher success rate than the trajectory used for the remaining experiments (Figure 12 left). Table 4 shows the outcomes of 200 trials, 100 with added noise. With no noise, this trajectory achieved a 100% success rate, signalling that it is robust to uncertainties in the starting hand posture and that trajectories can be optimised for particular environments (Movie S5). With an appropriate trajectory, the passive hand framework can achieve highly robust grasp performances under slight environmental uncertainties. As we are interested in error prediction and recovery, this grasp trajectory is not investigated in detail. The trajectory did see a significant drop in success rate under the presence of added noise, down to 67%. Within the 67 successes, three were observed as distinct new grasps, where the teaching sphere ended in a power grasp against the palm, rather than a pinch between the thumb and index finger. This gives a passive boost in performance which can be exploited in more novel and challenging tasks. In addition, this demonstrates the behavioural diversity enabled by passive design which can be exploited by minor changes in control (as this mode is observed to be starting position dependent, Figure 11A right).

Error recovery is possible via heuristics-based wrist trajectory adaptation

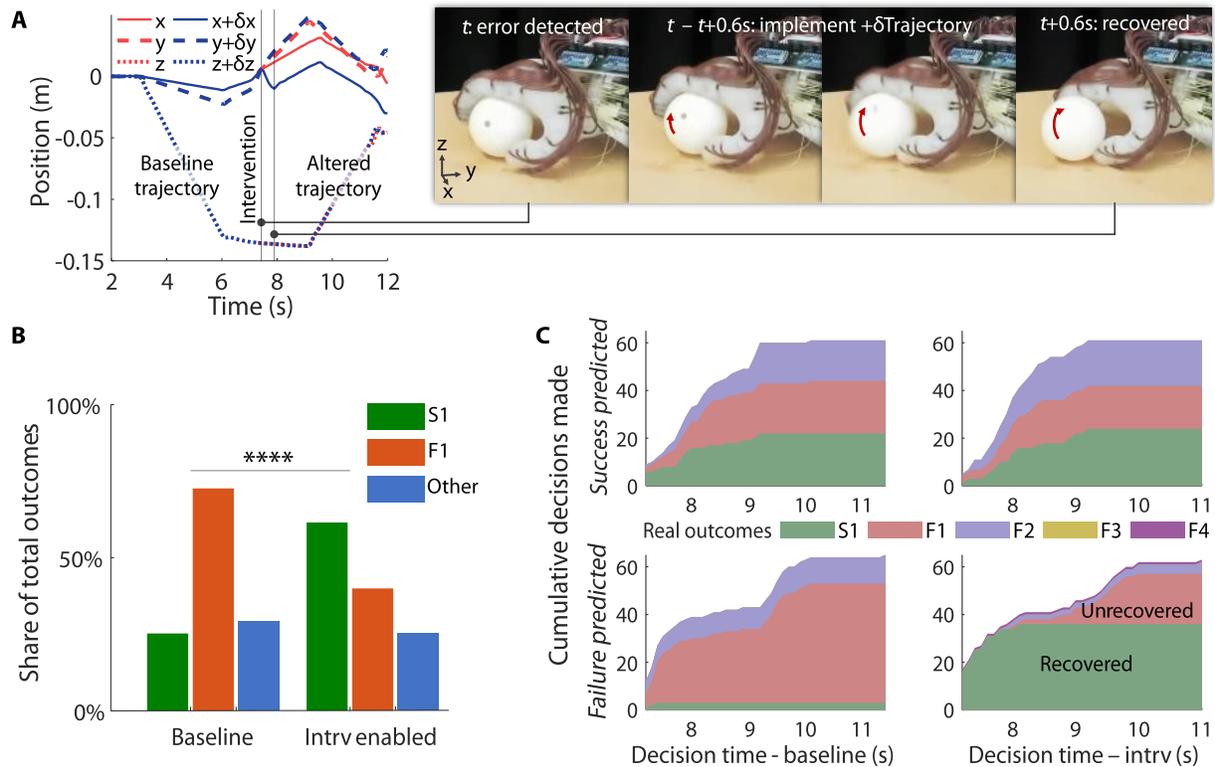


Figure 7: **Error recovery experiment results.** (A) Heuristic error recovery procedure. Prediction is monitored, as soon as error is ‘certain,’ the recovery intervention is implemented. The routine corrects for the most common error (F1: Fig. 6), relieving the trapped object with a shift, primarily in the x direction. (B) Change in outcomes when implementing error recovery, successes increase and failure mode 1 decreases significantly (to 0.001%). Other failures are not significantly affected. (C) Real grasping outcomes broken down by time of prediction, whether success was predicted (upper) or not (lower) and intervention disabled (left) as a control or enabled (right). Change in outcome distribution when failure is predicted and intervention enabled (lower right) demonstrates the recovery—early decisions are critical for successful recovery.

Passive grasping relies upon successful completion of a series of interactions. There is tolerance in these interactions. Passivity in design aims to increase these tolerances, meaning less information is required about the self and environment (Figure 5). In this way, the task of information gathering is exported from perception systems to the physical dynamical system. This means for many simple or familiar tasks, much lesser burden is placed on the perception system, though for more complex or unfamiliar tasks, perception is required.

The addition of a perception system can allow reaction to surprise during grasping (Park et al., 2016; Kieu et al., 2019). This way, expensive attention is required less frequently (Hogan et al., 2018). If this reaction comes too late, the grasp may have already failed, generally requiring an expensive rerun of the grasp (Laursen et al., 2015). Therefore, if potential failures can be detected early, they can be accounted for without either expensive attention or regrasping (Movie S6).

From Figure 6C, we see failure mode one is the most common. In this mode, the teaching sphere is trapped

under the index finger and the table, rather than in a more stable contact with the thumb (Figure 6B). This leads to a build-up of pressure, which when released knocks the sphere out of the grip. A heuristic recovery method is introduced, which intervenes with a shift to the trajectory when the failure prediction is ‘certain.’ This shift relieves pressure by rolling the sphere from under the finger and bringing it into contact with the thumb. Figure 7A illustrates the error recovery process. Of note, the recovery can intervene at any moment in time, so can react immediately when failure prediction exceeds a threshold (Figure 4C) (Movie S7).

The heuristic recovery intervention only operates as intended for a single failure mode, though acts as a proof of concept for improving manipulation robustness with error prediction. With a cheap control technique and exploitation of the anthropomorphic hand’s passive dynamics, intuitive and natural skills can be observed.

Table 5: Error recovery experiment results. Real outcomes against predictions (correct predictions in grey, incorrect in red). Predicted failures with intervention enabled are more likely to have real outcome ‘Success 1’.

Outcome mode		S1	F1	F2	F3
Baseline	Success predicted	[HTML]EFEFEF22	[HTML]F7D7D522	[HTML]F7D7D517	[HTML]F7D7D
	Failure predicted	[HTML]F7D7D53	[HTML]EFEFEF50	[HTML]EFEFEF12	[HTML]EFEFE
Intervention	Success predicted	[HTML]EFEFEF24	[HTML]F7D7D518	[HTML]F7D7D519	[HTML]F7D7D
	Failure predicted	[HTML]EFEFEF36	[HTML]F7D7D521	[HTML]F7D7D55	[HTML]F7D7D

With both the error predictions and recovery functioning, an experiment is run to evaluate the combined performance. Two hundred fifty grasping trials are run, alternating between open-loop control as a controlled baseline and with intervention enabled (Figure 4C). Table 5 shows all real outcomes and predictions for these trials. Observing the baseline results, the true positive rate of successes is 36% and for failures is 95%. Figure 7B shows the change in success and failure modes between the baseline and trials with intervention enabled. With intervention enabled, the number of successful outcomes increases from 25 to 60 and the number of failures decreases from 72 to 39. As the distribution of predictions is approximately equal between the two cases, this would suggest the intervention is impactful (with statistical significance $p < 0.00001$). The remaining failure modes are grouped and do not change significantly, this is expected due to the targeted nature of the error recovery.

Average predictions over time show patterns in the different outcomes. Figure 13 shows mean and standard deviation for the baseline trials (left), success (top) and failure (bottom) outcome chosen. Incorrect decisions happen early on in the trajectories, when the predictions are less distinguishable. The more common occurrence being failure outcomes looking similar to successful outcomes before 8 s. By waiting to make a decision, more information can be gathered, and higher accuracy can be achieved (with an upper limit from the regression accuracy, Figure 6A). However, this results in delayed reaction to recover from errors. There is high prediction variance within the different outcome modes, especially in failure mode two. This is partly because of our coarse classification of outcomes. Within each success and failure mode, there is some variance in terms of catastrophic failure point. Especially in failure mode 2, in which the sphere is dropped any time after lifting and before being reset. Therefore, significant variance in sensor information is expected, potentially causing overlaps in sensory information which can be resolved by gathering more information (either from further hand-environment interactions or with placement of additional sensing receptors).

Observing the decision times for the different predictions and real outcomes confirms the effect of the recovery intervention (Figure 7C). Firstly, the shapes of the total cumulative predictions counts are near identical between the baseline and intervention modes. Secondly, for the success predictions, when no recovery happens in either the baseline or in the intervention case, the distribution of real outcomes is on average within 13%. Therefore, the difference in real outcomes when failure is predicted is evident.

There are patterns in both the successful predictions and failure predictions based on the time the decisions are made. For the cases where failure outcomes were misidentified as successes (Figure 7C top, modes F1 and F2), 87% are decided before 9 s into the trajectory. Failure prediction accuracy increases if they can be

identified early (before 8.0 s) or later (after 9.5 s) (Figure 7C bottom left). The first peak is critical for the recovery routine. From the predicted failures in the intervention case, the successful outcomes significantly increase over the baseline case. However, after 8.1 s into the trajectory, there are no additional success outcomes. The failure decision has to come before this time for a chance of recovery. At 7.8 s, 35 failures are identified, 34 of them are successful outcomes (assuming 3 misclassified successful outcomes from the baseline failure predictions), the recovery success rate is $\approx 97\%$. At 8.4 s, the recovery success rate drops to $\approx 87\%$ and at 9.5 s drops to $\approx 66\%$.

Grasp predictions are also generalizable

While the error predictions are only trained on the grasping of spheres with position uncertainty, grasping success potentially is indicated by underlying patterns in the sensor readout which are not unique to spheres. Patterns such as net forces, enclosure of the object or stability of each contact (e.g., lack of slip), have the potential to be learned and transferred by the network to wider situations.

Table 2 shows the prediction accuracy on the set of everyday objects. The accuracy compares the decision point of the prediction to the real outcome. The most accurate predictions are seen with the 50 mm sphere (100%), the battery (90%) and the spool (80%). For the sphere, the grasp is near identical to the grasp of the training object (60 mm sphere), the battery and spool both have no successes and fail early in the trajectory, giving more information for the predictions. The lowest accuracy is seen in the 70 mm sphere (50%), bubble wrap (20%) and bottle (20%). The 70 mm sphere grasp is successful only in a different grasping mode which utilises the middle finger, hence the predictions are less accurate for the successes. Additionally, both the bubble wrap and bottle succeed with slightly different emergent grasps (Movie S2). This would suggest the network is overfitting on the success outcomes. This is not unexpected due to the lower diversity in grasping success in the training data versus the diversity in failures.

The predictions demonstrate some intuition for more general cases. This intuition is observed in predictions for familiar objects, in addition to predictions with unfamiliar environmental interference. For example, when grasping and the object is knocked out of the grasp during lift, the prediction rapidly updates to a failure (Figure 14, Movie S8). This exact interference case is not present in the training data, though failure mode two can look similar. Other live prediction updates do back up the intuition and generalised performance, particularly the ability of the prediction to update after the heuristic error recovery. The final prediction of the network (at 15 s) has 100% accuracy (50 of 50) for the baseline cases where failure was correctly predicted, for the intervention cases where failure was predicted but not recovered the final prediction was correct 100% (21 of 21), for cases where the intervention succeed, the prediction correctly updated to success 97% (35 of 36) of the time. The failure predictions were made on average 7.403 s into these trajectories, then by 8.2 s, the average prediction became uncertain (regression output 0.5) and by 8.5 s, the prediction average (0.143) passed the threshold to a success (< 0.2). Therefore, there is potential for closed-loop error recovery and self-supervised learning of recovery behaviours.

Discussion

Performance of robotic manipulators in general-purpose tasks is lacking when compared to human manipulation. Two significant problems we see in robotic manipulation include the development of appropriate designs that can exhibit adaptive dynamic behaviours and the development of control strategies for adaptation and learning in novel and niche situations. The passive dynamics of the hand have shown to be essential for generating diverse adaptive behaviours through interactions (Eppner et al., 2015; Hughes et al., 2018), ensuring low control effort through strategies including passive shape adaptation (Catalano et al., 2014) and stiffness control (Hughes et al., 2018). In passive soft grasping, emergent behaviours allow robustness to uncertainties in object and environment properties, though these behaviours can result in complex sensory responses which makes perception difficult. The described soft anthropomorphic hand is the ideal platform

for testing passive dynamic behaviours and investigating sensing strategies for intuitive manipulation. In total, 1240 grasping trials were performed (650 training, 250 heuristic recovery testing, 140 everyday object trials, 200 alternate trajectory grasps), providing diverse manipulation data.

The passive anthropomorphic hand demonstrates significant adaptive grasping behaviours. Trained in one-shot on a regular sphere, the hand is able to successfully grasp 11 of 14 everyday objects chosen at random with similar size to the sphere. These objects include highly irregular objects such as a computer mouse, and highly deformable objects such as a roll of bubble wrap. The soft sensing skin demonstrates a low-cost, highly sensitive and modular design (sensor response is customisable and receptors can be placed in any geometry of hand). These sensors provide tactile and proprioceptive information, though the current solution has limited receptor density and scalability is challenging with individual barometric pressure gauges for each sensing channel. The sparse data is informative for studying the hand-environment interactions, Figure 4B shows sensor channel peaks corresponding to the extracted keypoints, which mark changing interactions (Gilday et al., 2021). The anatomically correct design approach coupled with the soft sensing capabilities allows intuitive training from non-expert users.

Our error prediction approach is able to overcome traditional soft sensing problems including non-linearities, drift, and temperature dependence by using recurrent learning architectures. Low-level, high-bandwidth control is performed in a pure mechanical manner by the passive-adaptive behaviours (Figure 5) and emergent grasps (Figure 10). Hence, sensory feedback is only required for higher-level monitoring and predictive reactions. When incentivising early predictions, prediction accuracy can be as high as 79% with data length of 80 (8 s), improving to 98% with data length 150. Extracting the trends of the prediction, we are able to make a decision about the outcome in advance of catastrophic grasp failure (greater than 1 stable contact to one or fewer, e.g., object is ejected from grasp), therefore allowing us to correct for any error and recover the grasp. Implementation of a heuristic recovery routine targeted at the most common failure mode improves the success rate from 19.8% to 48.4%. Of the 35 earliest failure predictions, 34 were successfully recovered. The recovery strategy is a simple translation added based on intuition, showing the power of this predictive approach only requiring simple interventions, which is promising for scalability and developing recovery routines for other manipulations and failure modes, potentially through self-supervised learning (Pinto & Gupta, 2016; Wong, 2016).

To go beyond just improving grasping robustness and adaptation, we hope to see the emergence of intuitive behaviours. The first of these is being able to recognise grasping failures on unfamiliar objects. When only trained on grasping of a regular sphere, the network is able to detect failures with 68% accuracy (successes with 41%). This is promising and highly likely to be improved with a more diverse training set. Secondly, is being able to react to external sources of interference. Even when the object is knocked out of the hand post-lift, the network output reacts within 0.2 s to update the prediction and smaller sources of interference. Additionally, by being similar to a human hand, the achievable grasp types are well understood (Feix et al., 2015; Sundaram et al., 2019; Cutkosky & others, 1989), the hand matches ergonomically with the environment, and human input during teaching is intuitive due to familiarity (Ravichandar et al., 2020). This has great implications for making robust and intuitive manipulation systems, for example in flexible industrial or logistic robotic applications where the environment can constantly change, but highly robust and adaptable systems are required. Or for prostheses, where the anthropomorphic design is highly desirable, both for aesthetics (Biddiss et al., 2007; Pylatiuk et al., 2007) and for ease of exploitation (Kashef et al., 2020; Biddiss & Chau, 2007). An additional requirement for prostheses is low-cost, simplicity and sensory feedback during use (Cordella et al., 2016; Hermens et al., 2011; Zhao et al., 2016). Our hand achieves all of these, in particular the passive nature greatly simplifies the system and usability in the real world, and the high-level sensory feedback is easily interpretable by the user and can give real-time feedback of grasp quality.

Acknowledgements

This work was funded by UKRI grant [RG99055] and Arm Ltd. For the purpose of open access, the author has applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising.

Conflict of interest

The authors declare no conflict of interest for this work.

Supplementary materials

- Figure 8. Robot system diagram.
- Figure 9. Novel objects.
- Figure 10. Grasping of novel objects.
- Figure 11. Alternate grasping trajectory.
- Figure 12. Grasping outcome classifications.
- Figure 13. Breakdown of average network predictions.
- Figure 14. Prediction reaction.
- Movie S1. Teaching demonstration.
- Movie S2. Successful novel object grasping.
- Movie S3. Unsuccessful novel object grasping.
- Movie S4. Trajectory 1 (experimental) time-lapse.
- Movie S5. Trajectory 2 (alternate) time-lapse.
- Movie S6. Trajectory 1 failure examples.
- Movie S7. Heuristic recovery demonstration.
- Movie S8. Live predictions and reactions showcase.

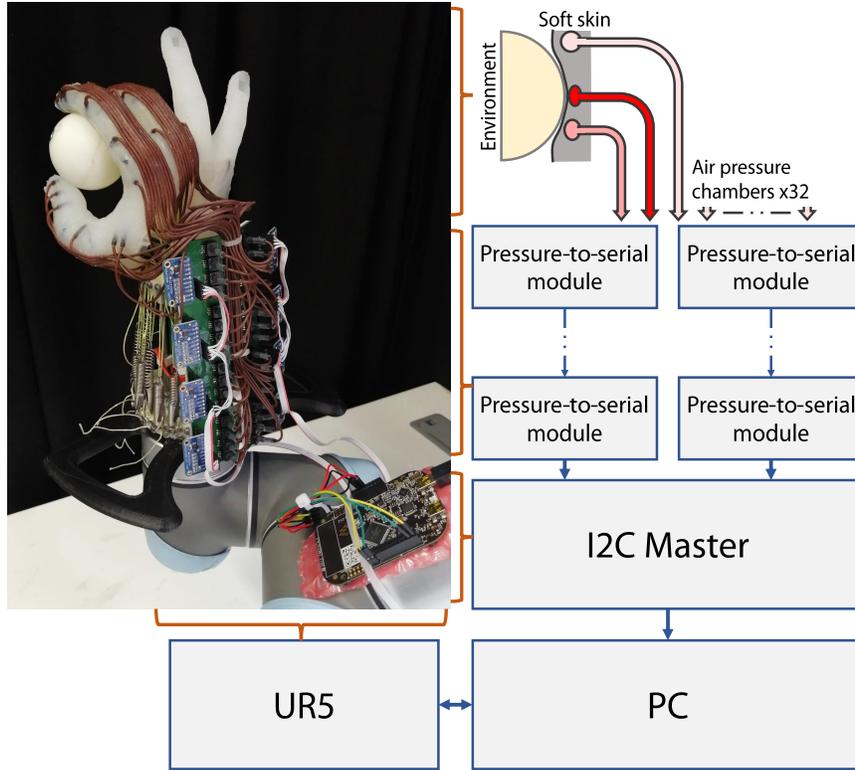


Figure 8: **Robot system diagram.** Barometric sensors are read by a bank of analog-to-digital converters and transmitted to a local microcontroller over two I2C buses. This I2C master device converts the information to serial data for the master PC which also controls the robot arm.



Figure 9: **Novel objects.** Randomly selected objects for the adaptive grasping experiment.

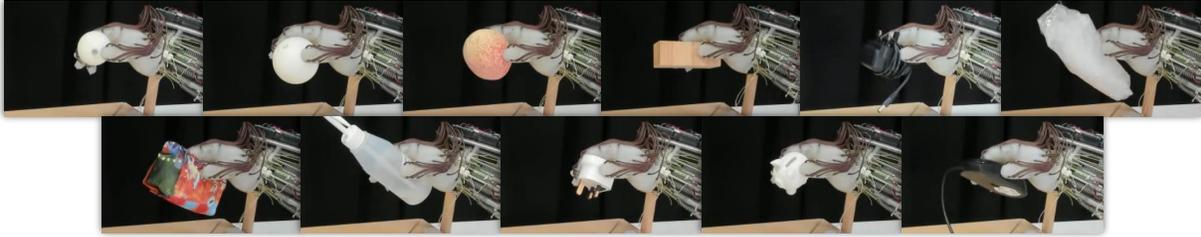


Figure 10: **Grasping of novel objects.** 11 of 14 randomly selected objects, Figure 9, can be grasped successfully with the original open-loop trajectory.

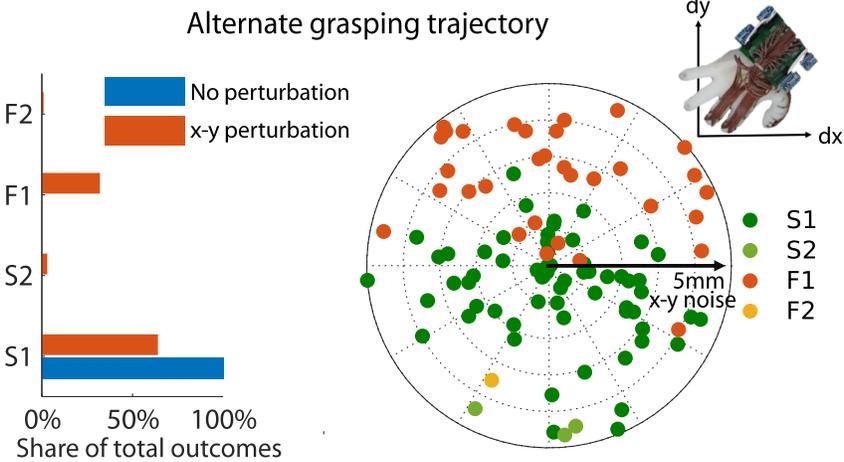


Figure 11: **Alternate grasping trajectory.** Share of grasping outcomes (left). Multiple unique success and failure modes emerge (Figure 12 right), dependent on orientation of perturbation (right).

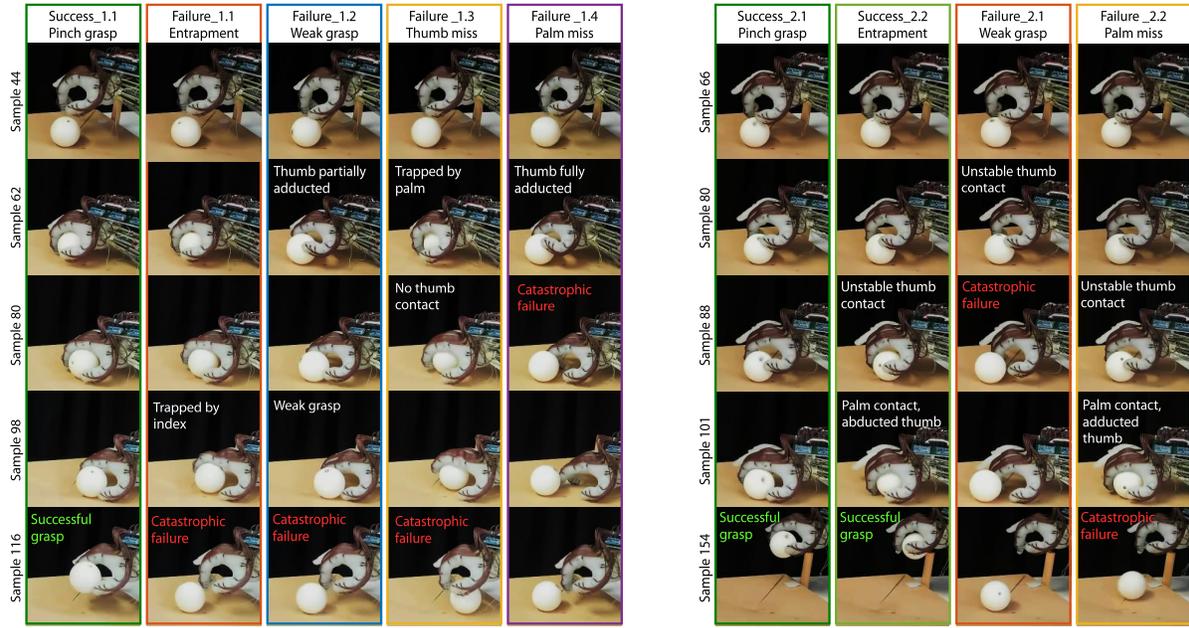


Figure 12: **Grasping outcome classifications.** Grasping outcomes with main (left) and supplementary (right) trajectory (1 and 2). Multiple unique failure modes emerge in both trajectories. An emergent success mode is seen in trajectory 2 (S2.2).

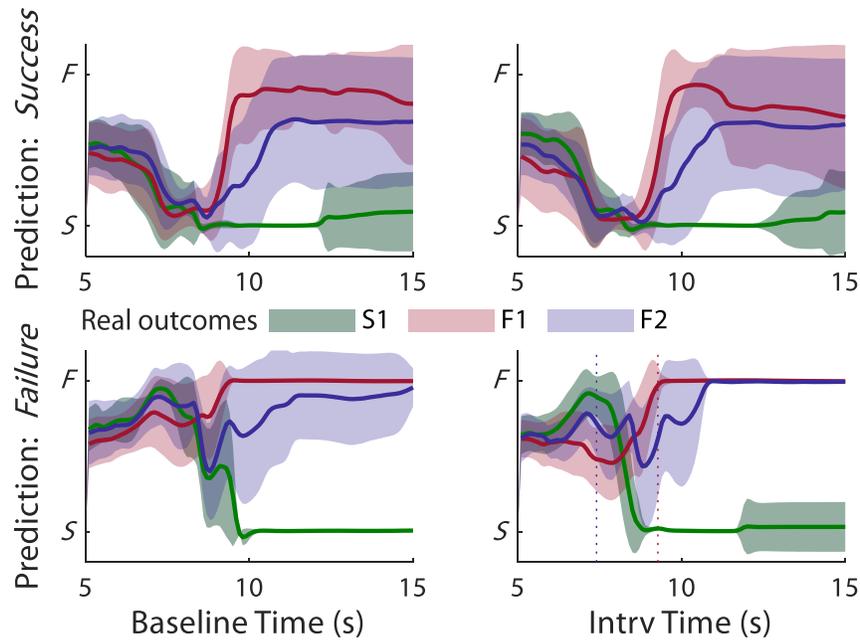


Figure 13: **Breakdown of average network predictions.** High variance reduces prediction certainty, i.e., there is a tradeoff between prediction accuracy and predicting early enough for successful recovery.

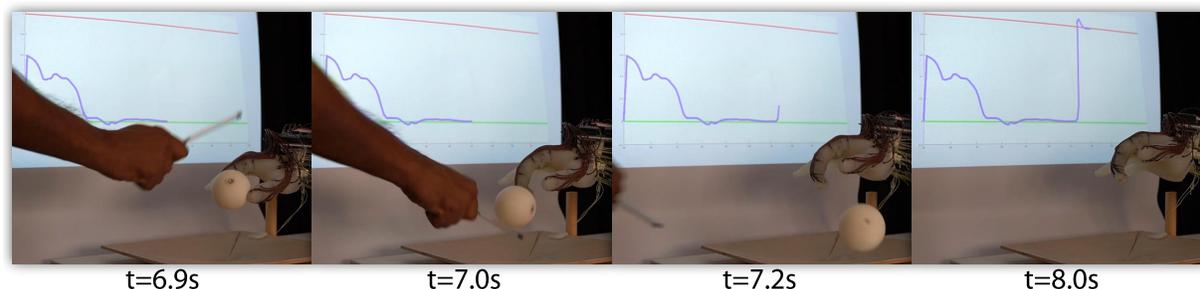


Figure 14: **Prediction reaction.** Real-time prediction reaction to external interference. Training sphere is knocked from the grasp, prediction begins to update after 0.2 s.

References

- A century of robotic hands. (2019). *Annual Review of Control, Robotics, and Autonomous Systems*, 2, 132.
- Soft manipulators and grippers: a review. (2016). *Frontiers in Robotics and AI*, 3, 69.
- Hands for dexterous manipulation and robust grasping: A difficult road toward simplicity. (2000). *IEEE Transactions on Robotics and Automation*, 16(6), 652662.
- A survey of bio-inspired robotics hands implementation: New directions in dexterous manipulation. (2013). *Robotics and Autonomous Systems*, 61(5), 517544.
- Exploitation of environmental constraints in human and robotic grasping. (2015). *The International Journal of Robotics Research*, 34(7), 10211038.
- Wrist-driven passive grasping: interaction-based trajectory adaption with a compliant anthropomorphic hand. (2021). *Bioinspiration & Biomimetics*, 16(2), 026024.
- An anthropomorphic soft skeleton hand exploiting conditional models for piano playing. (2018). *Science Robotics*, 3(25).
- A soft-robotic approach to anthropomorphic robotic hand dexterity. (2019). *Ieee Access*, 7, 101483101495.
- Adaptive synergies for the design and control of the Pisa/IIT SoftHand. (2014). *The International Journal of Robotics Research*, 33(5), 768782.
- Muscle synergies in neuroscience and robotics: from input-space to task-space perspectives. (2013). *Frontiers in Computational Neuroscience*, 7, 43.
- Toward dexterous manipulation with augmented adaptive synergies: The pisa/iit softhand 2. (2018). *IEEE Transactions on Robotics*, 34(5), 11411156.
- Data-driven Simulation Framework for Expressive Piano Playing by Anthropomorphic Hand with Variable Passive Properties. (2022). *2022 IEEE 5th International Conference on Soft Robotics (RoboSoft)*, 300305.
- RBO Hand 3—A Platform for Soft Dexterous Manipulation. (2022). *ArXiv Preprint ArXiv:2201.10883*.
- Surprisingly robust in-hand manipulation: An empirical study. (2022). *ArXiv Preprint ArXiv:2201.11503*.
- Learning dexterous in-hand manipulation. (2020). *The International Journal of Robotics Research*, 39(1), 320.
- Large-area soft e-skin: The challenges beyond sensor designs. (2019). *Proceedings of the IEEE*, 107(10), 20162033.

- Tactile sensing for dexterous in-hand manipulation in robotics—A review. (2011). *Sensors and Actuators A: Physical*, 167(2), 171187.
- Robust proprioceptive grasping with a soft robot hand. (2019). *Autonomous Robots*, 43(3), 681696.
- Optoelectronically innervated soft prosthetic hand via stretchable optical waveguides. (2016). *Science Robotics*, 1(1), eaai7529.
- Learning from humans how to grasp: a data-driven architecture for autonomous grasping with anthropomorphic soft hands. (2019). *IEEE Robotics and Automation Letters*, 4(2), 15331540.
- Electronic skins and machine learning for intelligent soft robots. (2020). *Science Robotics*.
- Toward perceptive soft robots: Progress and challenges. (2018). *Advanced Science*, 5(9), 1800541.
- Tactile sensing in dexterous robot hands. (2015). *Robotics and Autonomous Systems*, 74, 195220.
- Tactile sensing—from humans to humanoids. (2009). *IEEE Transactions on Robotics*, 26(1), 120.
- Anthropomorphic robot hand: Gifu hand III. (2002). *Proceeding of International Conference ICCAS2002*, 12881293.
- Drift-Free Latent Space Representation for Soft Strain Sensors. (2020). *2020 3rd IEEE International Conference on Soft Robotics (RoboSoft)*, 138143.
- Integration of deep learning and soft robotics for a biomimetic approach to nonlinear sensing. (2021). *Nature Machine Intelligence*, 3(6), 507512.
- Gelsight: High-resolution robot tactile sensors for estimating geometry and force. (2017). *Sensors*, 17(12), 2762.
- The tactip family: Soft optical tactile sensors with 3d-printed biomimetic morphologies. (2018). *Soft Robotics*, 5(2), 216227.
- Acoustic sensing for soft pneumatic actuators. (2018). *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 69866991.
- Soft robot perception using embedded soft sensors and recurrent neural networks. (2019). *Science Robotics*, 4(26), eaav1488.
- Distributed proprioception of 3D configuration in soft, sensorized robots via deep learning. (2020). *IEEE Robotics and Automation Letters*, 5(2), 32993306.
- Deep full-body motion network for a soft wearable motion sensing suit. (2018). *IEEE/ASME Transactions on Mechatronics*, 24(1), 5666.
- Towards learning to detect and predict contact events on vision-based tactile sensors. (2020). *Conference on Robot Learning*, 13951404.
- Real-time soft body 3d proprioception via deep vision-based sensing. (2020). *IEEE Robotics and Automation Letters*, 5(2), 33823389.
- On the closure properties of robotic grasping. (1995). *The International Journal of Robotics Research*, 14(4), 319334.
- An overview of dexterous manipulation. (2000). *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No. 00CH37065)*, 1, 255262.
- Data-driven grasp synthesis—a survey. (2013). *IEEE Transactions on Robotics*, 30(2), 289309.
- Active soft end effectors for efficient grasping and safe handling. (2018). *IEEE Access*, 6, 2359123601.

- A soft-touch gripper for grasping delicate objects. (2017). *IEEE/ASME Transactions on Mechatronics*, 22(3), 12761286.
- Sequential contact-based adaptive grasping for robotic hands. (2022). *The International Journal of Robotics Research*, 02783649221081154.
- Embodied choice: how action influences perceptual decision making. (2015). *PLoS Computational Biology*, 11(4), e1004110.
- Multimodal execution monitoring for anomaly detection during robot manipulation. (2016). *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 407414.
- Making sense of vision and touch: Self-supervised learning of multimodal representations for contact-rich tasks. (2019). *2019 International Conference on Robotics and Automation (ICRA)*, 89438950.
- Technique of Recovery Process and Application of AI in Error Recovery Using Task Stratification and Error Classification.. (2018). *J. Robotics Netw. Artif. Life*, 5(1), 5662.
- Reflex control of the pisa/iit softhand during object slippage. (2016). *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 19721979.
- Improved gelsight tactile sensor for measuring geometry and slip. (2017). *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 137144.
- Tactile sensors for friction estimation and incipient slip detection toward dexterous robotic manipulation: A review. (2018). *IEEE Sensors Journal*, 18(22), 90499064.
- Grip stabilization of novel objects using slip prediction. (2018). *IEEE Transactions on Haptics*, 11(4), 531542.
- Learning to Prevent Grasp Failure with Soft Hands: From Online Prediction to Dual-Arm Grasp Recovery. (2022). *Advanced Intelligent Systems*, 4(3), 2100146.
- Tactile regrasp: Grasp adjustments via simulated tactile transformations. (2018). *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 29632970.
- More than a feeling: Learning to grasp and regrasp using vision and touch. (2018). *IEEE Robotics and Automation Letters*, 3(4), 33003307.
- Jamming Joints for Stiffness and Posture Control with an Anthropomorphic Hand. (2021). *2021 IEEE 4th International Conference on Soft Robotics (RoboSoft)*, 134140.
- Intrinsic and extrinsic contributions to the passive moment at the metacarpophalangeal joint. (2000). *Journal of Biomechanics*, 33(12), 16751681.
- Enhancement of finger motion range with compliant anthropomorphic joint design. (2016). *Bioinspiration & Biomimetics*, 11(2), 026001.
- Three-dimensional rotations of human three-joint fingers: an optoelectronic measurement. Preliminary results. (2005). *Surgical and Radiologic Anatomy*, 27(1), 4350.
- The application of three dimensional analysis system in measuring rotation angles of thumb carpometacarpal joints. (2006). *Zhonghua Wai Ke Za Zhi [Chinese Journal of Surgery]*, 44(16), 11151118.
- Ligamentous structures of the metacarpophalangeal joint: a quantitative anatomic study. (1983). *Journal of Orthopaedic Research*, 1(4), 361368.
- Tendon arrangement and muscle force requirements for human-like force capabilities in a robotic finger. (2002). *Proceedings 2002 IEEE International Conference on Robotics and Automation (Cat. No. 02CH37292)*, 4, 37553762.

- An extensor mechanism for an anatomical robotic hand. (2003). *2003 IEEE International Conference on Robotics and Automation (Cat. No. 03CH37422)*, 1, 238243.
- Understanding variable moment arms for the index finger MCP joints through the ACT hand. (2008). *2008 2nd IEEE RAS & EMBS International Conference on Biomedical Robotics and Biomechanics*, 776782.
- In vivo biomechanics of the fingerpad skin under local tangential traction. (2007). *Journal of Biomechanics*, 40(4), 851860.
- Postural hand synergies during environmental constraint exploitation. (2017). *Frontiers in Neurobotics*, 11, 41.
- Exploiting robot hand compliance and environmental constraints for edge grasps. (2019). *Frontiers in Robotics and AI*, 135.
- Understanding intelligence*. (2001). MIT press.
- Self-organization, embodiment, and biologically inspired robotics. (2007). *Science*, 318(5853), 10881093.
- A theory of cheap control in embodied systems. (2015). *PLoS Computational Biology*, 11(9), e1004427.
- Sensing, Actuating, and Interacting Through Passive Body Dynamics: A Framework for Soft Robotic Hand Design. (2022). *Soft Robotics*.
- Intelligent Soft Hands and Benchmarking towards General-Purpose Robotic Manipulation. (2022). *IOP Conference Series: Materials Science and Engineering*, 1261(1), 012010.
- Mechanisms of the anatomically correct testbed hand. (2011). *IEEE/ASME Transactions on Mechatronics*, 18(1), 238250.
- Perspectives: a vital biomechanical model of synovial joint design. (1994). *The Anatomical Record*, 240(1), 118.
- The state of the art in biomimetics. (2013). *Bioinspiration & Biomimetics*, 8(1), 013001.
- Recent developments in synovial joint biomechanics. (1980). *Siam Review*, 22(3), 275317.
- Clinical test of apposition and counter-apposition of the thumb. (1986). *Annales De Chirurgie De La Main: Organe Officiel Des Societes De Chirurgie De La Main*, 5(1), 67.
- Application of fin ray effect approach for production process automation. (2011). *Annals of DAAAM & Proceedings*, 22(1), 12471249.
- The grasp taxonomy of human grasp types. (2015). *IEEE Transactions on Human-Machine Systems*, 46(1), 6677.
- On grasp choice, grasp models, and the design of hands for manufacturing tasks.. (1989). *IEEE Transactions on Robotics and Automation*, 5(3), 269279.
- Outlier Detection for Time Series with Recurrent Autoencoder Ensembles.. (2019). *IJCAI*, 27252732.
- Automatic error recovery in robot assembly operations using reverse execution. (2015). *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 17851792.
- Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours. (2016). *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 34063413.
- Towards lifelong self-supervision: A deep learning direction for robotics. (2016). *ArXiv Preprint ArXiv:1611.00201*.
- Learning the signatures of the human grasp using a scalable tactile glove. (2019). *Nature*, 569(7758), 698702.

Recent advances in robot learning from demonstration. (2020). *Annual Review of Control, Robotics, and Autonomous Systems*, 3(ARTICLE), 297330.

Consumer design priorities for upper limb prosthetics. (2007). *Disability and Rehabilitation: Assistive Technology*, 2(6), 346357.

Results of an Internet survey of myoelectric prosthetic hand users. (2007). *Prosthetics and Orthotics International*, 31(4), 362370.

Robotic hand: A review on linkage-driven finger mechanisms of prosthetic hands and evaluation of the performance criteria. (2020). *Mechanism and Machine Theory*, 145, 103677.

Upper limb prosthesis use and abandonment: a survey of the last 25 years. (2007). *Prosthetics and Orthotics International*, 31(3), 236257.

Literature review on needs of upper limb prosthesis users. (2016). *Frontiers in Neuroscience*, 10, 209.

Myoelectric forearm prostheses: State of the art from a user-centered perspective. (2011). *J. Rehabil. Res. Dev.*, 48(6), 719.