

Damage Adaptation in Seconds for Architected Materials

James Avtges*, Jake Ketchum*, Helena Young, Taekyoung Kim, Ryan L. Truby, Todd D. Murphey
Northwestern University, Evanston, IL, USA

*Equal Contribution. Email: {javgtes, jketchum}@u.northwestern.edu




Fig. 1. **Latent Ensemble Adaptive Proprioception (LEAP)**: Modeling damage in architected materials enables simple yet robust ensemble methods to adapt to damage in seconds. We evaluate our approach on a soft wrist that completes a contour tracing task using proprioception alone. The wrist is damaged by melting actuators (left), cutting with scissors (right), and repairing them with glue. After adaptation, LEAP recovers up to all of the undamaged performance.

Abstract—Adaptation to damages and in-situ physical repairs is essential for long-term robot autonomy, yet challenging outside of narrowly defined and well-anticipated bounds. In this work we proprioceptively adapt to catastrophic damage in soft-actuated systems in under one minute. Architected materials are well equipped for adaptation: actuator failure occurs gradually rather than acutely, and damage can be described in a low-dimensional, discrete coordinate space. Surprisingly, latent damage representations plus a simple yet robust ensemble method is sufficient for adapting to unseen damage in real-time. Moreover, we identify conditions under which exponential sample complexity collapses to linear sample complexity for learned representations of architected materials, a concrete advantage over rigid components or continuum soft mechanisms. We demonstrate LEAP, our method for adaptive proprioception, via a tracing task for a 6DoF soft wrist based on Handed Shearing Auxetic (HSA) actuators. Our algorithm is able to adapt to cuts, burns, and actuator repairs, enabling simulation-free real-time adaptation that is critical for realizing the promise of soft robots outside the lab. Videos and more information are available at <https://murphey-lab.github.io/leap>.

I. INTRODUCTION

A major barrier to bringing soft robots out of the lab is the reliability of the compliant structural components, transducers,

and actuators. Soft robots use mechanical deformation to perform work, causing material failure through creep and fatigue. Actuators often have a continuous operating life of hours to even minutes, in contrast to brushless DC motors whose operating lives can be multiple orders of magnitude longer. All of these challenges are exacerbated in human-centric spaces such as kitchens, workshops, and hospitals where there are hazards that may lacerate, wear, or melt flexible devices.

Even within the nominal operating life of a soft actuator, the challenges associated with modeling and control have been well-documented [1]. In addition to the difficulties of modeling high-dimensional soft bodies, a soft actuator cannot be parametrically defined by constant terms that engineers typically have access to, as is the case with rigidly defined systems. Material properties, and therefore actuator models, are often hysteretic and vary with the robot’s configuration, payload, and temperature. Furthermore, actuator manufacturing is often done by hand and even when produced in parallel (i.e., on the same build plate of a 3D printer), current processes produce significant actuator-to-actuator variability [2].

Existing works in this field typically produce outputs that

implicitly assume “nominal” robot conditions: either immediately after the actuator is constructed or after some defined break-in procedure. As soon as a soft robot performs significant work, however, the robot’s constituent components inevitably change, causing a non-adaptive controller to be out-of-distribution with respect to its underlying model. Moreover, soft actuators often play a multifunctional role in an engineered system, providing structural integrity, motion, and contact surfaces, all of which are impacted upon experiencing damage.

Architected materials are advantageously positioned for damage adaptation. An architected material is a material whose structure and microscale properties are intentionally designed to produce desired functionalities and macroscale properties beyond material properties alone [2]. A hallmark of most architected materials is a lattice- or unit cell-based structure, such as the structures found in auxetics, kirigami, and active lattices. This provides a discrete mechanism to parameterize material damage, whereas soft body structures and rigid-linked systems have continuous parameterizations of structure damage. For example, a recent method to adapt to broken legs in a rigid quadruped [3] requires robust modeling of thousands of different damage states. The compliant nature of architected materials provides degeneracies—within each lattice member all damages are equivalent.

Since architected materials can replace multiple components in a traditional robotics setting such as gears, tendons, hinges, and joints, the multifunctional lattice of the material also provides redundancies, whereas the constituent components of a rigid analog are all single points of failure.

Furthermore, failure modes in architected materials are often non-acute. The materials themselves can continue operation, albeit with changes to the overall properties of the mechanism, before catastrophic material failure [4]. Architected materials in particular utilize redundancies in their discrete structure such that functionality can remain despite multiple links being severed.

In this work we introduce Latent Ensemble Adaptive Proprioception (LEAP). Our method utilizes these inherent advantages of architected materials to adapt to fully-unseen actuator damage in seconds on hardware. A key insight of our work is parameterizing the structure of architected materials into discrete coordinates for the purpose of latent modeling. This allows characterizing the damage state using a representation that is, importantly, independent of the robot state itself. While our method does not adapt to all types of damage in a single predictive step, we use this embedding as a means of adapting accurately and in near real-time. By modeling the meta-state of soft actuators as latent variables, “adaptation” and “damage” can be viewed both as functions of the same latent dynamics.

The contributions of this work are:

- 1) LEAP, a method for using an ensemble of meta-state encodings for damage adaptation in architected materials. In 6-DoF force/torque proprioception, LEAP requires only seconds of data and achieves a 90% reduction in proprioceptive force/torque prediction loss.

- 2) A novel 6-DoF soft wrist capable of proprioceptive force/torque sensing and rapidly adapting to changing dynamics.
- 3) We identify properties of architected materials that provide unique advantages for adaptation over rigid and continuum soft robots.

II. ON THE REPAIR OF SOFT ROBOTS

Well-defined calibration procedures in response to a motor or sensor repair have been essential elements of rigid mechanical systems for decades [5]. However, for soft systems, the concept of “repairing” a compliant actuator is typically overlooked, despite most soft robots being assembled by hand—and the ability to easily do so being one of their most transformative characteristics.

Physical repair is a vital prerequisite for the practical deployment of soft robots at scale: soft mechanisms function at a low enough degree of kinematic precision and net force output that many tools and methods used by hobbyists work well for repairing architected materials. Furthermore, repair of flexible, strain dependent actuators enables unique opportunities for in-situ mechanical adaptation when using direct replacements may be infeasible or impossible. Most repairs to flexible polymers meaningfully change the physical properties of the device; despite this, repair is generally not included in the class of settings where adaptation is necessary.

By representing the state of a damaged actuator entirely in a latent space, our method makes no distinctions between what is a “damage” and what is a “repair” to either the 6-DoF wrist or its constituent HSAs. Our method provides a new context for understanding repair—and generally the assembly of high variability components—in the fields of soft robotics and architected materials as an intermediate step requiring adaptation. Our method is faster and more computationally inexpensive than many traditional calibration procedures on servomotors, CNC spindles, and collaborative robot arms [6, 7, 8].

III. RELATED WORKS

A. Architected Materials

Architected materials are characterized by macroscale properties that are driven by the geometry of their internal structure, and which vary from the bulk material properties of their constituent materials. Architected materials can exhibit exotic behaviors such as negative Poisson’s ratios [9], very high strength/weight ratios, programmable responses [10], and multi-stability [11]. Architected materials are often built around a repeating unit cell, and come in a wide variety of types including kirigami, origami, auxetic, and lattice structures, among others [2]. Advances in additive manufacturing have helped architected materials become promising candidates for transmission and structural elements in soft systems [2, 12, 13].

In this work, we are concerned with lattice-type architected materials, and specifically Handed Shearing Auxetics (HSAs). HSAs are a type of rotary-to-linear soft-transmission first

described in [9]. They are composed of a 2D auxetic lattice, wrapped into a cylinder. The resulting structure will extend axially and diametrically when the two ends are rotated in one relative direction, and will correspondingly contract when the two ends are rotated in the other relative direction. HSAs can be manufactured from a variety of materials including steel [9], photopolymers [14], thermoplastics like TPU [12], or, as in this work, polypropylene.

Their high force, fast response time, and ease of control has made HSAs a promising transmission element for soft robots. HSAs have been used in a wide variety of soft robots including quadrupeds [15, 14, 16], parallel mechanisms [17], grippers [18], ball-kicking robots [19], and hoppers [20]. HSAs also have high mechanical reliability, making them a good candidate for real-world machine learning experiments [21, 22] which may involve prolonged actuation during training.

B. Soft Robot Proprioception

In biology, proprioception is defined as the subconscious sensation of one's own motion, position, and force [23]. In robotics, proprioception refers to the ability of the robot to sense its own configuration and the forces it is experiencing. Proprioception is a key enabler of dexterity in robotic systems [24]. This is particularly important for soft systems, which often exhibit time-varying dynamics and manufacturing variation that make precise open-loop control challenging [24]. As a result, significant work has been done on sensorizing soft actuators using pneumatic strain gauges [25], fluidic innervation [26], embedded magnetic elements [27], and direct resistance measurement [28]. Many of these methods have delivered promising results for single actuators, or small arrangements of actuators in a lab-bench context. However, in practice, methods like fluidic innervation that directly modify the actuators can greatly increase the cost, complexity, and fragility of engineered systems. Camera based approaches [29] bypass many of the practical drawbacks to more invasive methods and serve to move the sensing mechanisms away from contact surfaces, which may be subject to wear or damage.

The same viscoelastic, non-linear, and time-varying response that makes open-loop control of soft actuators challenging also complicates analytical approaches to proprioceptive sensing [24]. Machine learning, which can discover the required relationships from data, offers one solution to managing the complexity [29, 30]. Combined with camera-based sensing, learned proprioception promises to deliver robust and performant soft proprioception. However, at present, even learned proprioception systems generally lack mechanisms for handling damage or repair, and often sense forces in a maximum of two dimensions [29, 30]. By implementing our method for proprioception into a 6-DoF soft wrist, we are able to reliably sense 6-axis forces and torques at resolutions of up to 0.1 N in the F_x and F_y directions.

C. Damage Adaptation

The ability to detect and adapt to damage is essential for systems designed to operate in extreme environments,

in critical support roles, or in areas where human support is infeasible. This capability is particularly critical for soft mechanisms, which may be more vulnerable to damage and have more limited operating lives relative to their classical counterparts [31, 12, 32]. Fortunately, soft mechanisms lend themselves to gradual degradation, which facilitates learning [33, 34]. The majority of the damage adaptation canon to date can be split into two broad categories: Self modeling, in which the robot uses interactions with the environment to determine a model of its own dynamics, and control-adaptation, in which the robot uses interactions with the environment to determine an effective policy without first modeling its own configuration.

Control-adaptation methods, which directly modify the robot policy based on dynamics changes, can provide excellent adaptation performance for well-defined tasks. In a mobility context, simulation-based search based methods have demonstrated real-world adaptation for damaged robots [35, 36]. With quality simulation, control-adaptation methods can also handle very large dynamics changes which require new mobility strategies entirely [37]. However, control-adaptation methods do not directly model the system, and are thus unsuitable for use in proprioception.

In this work, we are concerned with self-modeling methods. Self-modeling, an extension of classical system identification, involves the robot building and maintaining a representation of its own dynamics. This representation can then be integrated into a variety of sensing and control methods. Self-modeling for real-time adaptation is characterized by relatively high sample and compute efficiency. Self-modeling has proven to be a reliable and effective technique in classical robotics, where it has enabled robust adaptation to catastrophic damage states in legged robots [38, 39] and serial manipulators [40].

Many self-modeling methods function on the robot's state-space [41] or on a lower-dimensional representation of the robot's geometry [42]. However, advances in machine learning and image processing have enabled the use of visual sensing for continuous self-monitoring and adaptation [43, 44]. Self-modeling is sample intensive, but can handle very large dynamics changes, in contrast with control-adaptation methods which often implicitly rely on similarity between a pre-damage simulation and post-damage robot dynamics [35, 36]. Where feasible, self-modeling methods are attractive because they enable the use of general purpose methods for control, guidance, and navigation.

IV. METHODS

Our method, Latent Ensemble Adaptive Proprioception (LEAP), utilizes ensembles of damaged architected materials that are parameterized by clustered latent states. By training a damage encoder and producing the ensemble members as a pre-training process, adaptation to fully-unseen types of damage and manufacturing variations in HSAs can be achieved in seconds.




Fig. 2. **Architecture of LEAP:** Our proposed approach encodes damage states into a latent representation and uses the latent structure to produce ensembles—where online data collection and regression of ensemble weightings occurs within seconds. Our architecture contains two neural networks—the first produces latent encodings of the actuator damage state. The second uses the latent variables alongside the measured displacement of the actuator in $SE(3)$ to predict the 6-dimensional forces and torques at the actuator base. After training, the latent space provides clusters from which we produce an ensemble representing a sampling of possible damage states. In inference, we learn linear weightings per-component for the ensemble via least squares which provides adaptive force/torque predictions.

A. Representing Damage in Discrete Coordinates

The structure of HSAs and other lattice-based, architected materials allows for the assignment of discrete coordinates to each possible mechanism of damaging the material.

Figure 4A shows an undamaged HSA twisted to its maximum position of one full rotation. We consider in this work an HSA “cut” as fully severing either a vertical or horizontal auxetic link, the location of which is given by the vertical helix (A, B, C), and its height on the HSA (0-7 for vertical, 0-6 for horizontal). Burns, or melting auxetic links together, are segmented into vertical coordinates of (0-2) given the imprecise nature of melting HSAs. In total, this provides 54 different possible ways of damaging an HSA. We assume HSAs to be non-functional in their ability to generate forces and torques when more than 3 instances of damage are applied. The damage state of an HSA x is then represented as a sparse binary vector:

$$x \in \{0, 1\}^{54}, \quad \|x\|_0 \leq 3. \quad (1)$$

While this specific structure of vertical and horizontal struts is not universal to all architected materials, this coordinate system provides a physically grounded mechanism for engineering adaptive behaviors in auxetics and other architected materials. The goal of this encoding is not to explicitly and accurately model every combination of possible HSA damages. Rather, the goal is to produce an encoded representation such that the distribution of latent variables z produces ensembles of dynamically feasible force/torque predictions. Similar parameterizations are readily available for a wide range of architected materials [45, 46]. Since LEAP uses ensemble methods, full coverage of the latent space, while desirable, is not necessarily required. The goal is kinematically feasible and reasonably orthogonal basis clusters, not comprehensive coverage.

B. Modeling Damage via a Latent Space

We model the damage state of the architected material using a variational latent model, an architecture that produces smooth probabilistic latent spaces that act as a compressed representation of its input data. As described in Section IV-A, we transform the discrete coordinate representation of damaged links in the HSA into a sparse binary vector where the damaged coordinates are represented by ones. We then learn the latent space mapping (Eq. 2).

$$z = \mathcal{D}(x), \quad \mathbb{R}^{54} \rightarrow \mathbb{R}^3 \quad (2)$$

The final layer of $\mathcal{D}$ is parametrized as a multivariate Gaussian with means μ and a diagonal variance matrix whose non-zero elements are σ . This is then sampled to produce a latent vector z using the “re-parametrization trick” to allow proper gradient flow (Eq. 3).

$$z = \mu(x) + \sigma(x) \odot \epsilon, \quad \epsilon \sim \mathcal{N}(0, I) \quad (3)$$

The latent model is trained using a variant of the β -VAE formulation [47] of the Evidence Lower Bound (ELBO) objective (Eq. 4).

$$\mathcal{L}(\mathcal{W}, \mathcal{D}; y) = \mathbb{E}_{q_{\mathcal{D}}(z|x)}[\log p(y|z)] - \beta D_{KL}(q_{\mathcal{D}}(z|x) \parallel p(z)) \quad (4)$$

where y is a ground-truth force/torque measurement and $P \in SE(3)$ is the homogeneous transformation measured from the base to the top of the HSA. The first term in the objective is mean-squared error loss and D_{KL} represents the Kullback-Leibler Divergence (KL-D) using a β -annealing similar to [47].

The wrench predictor learns the function:

$$\hat{y} = \mathcal{W}(P|z), \quad \mathbb{R}^{12} \rightarrow \mathbb{R}^6, \quad (5)$$

where the input to the wrench predictor is the latent state z concatenated with a representation of the orientation component of P using a Gram-Schmidt-like process according to [48], which performs the mapping:

$$g_{\text{GS}}\left(\begin{bmatrix} | & & | \\ p_1 & \cdots & p_3 \\ | & & | \end{bmatrix}\right) = \begin{bmatrix} | & | \\ p_1 & p_2 \\ | & | \end{bmatrix} \quad (6)$$

This operation truncates the last column in p , which is the $SO(3)$ component of P , ensuring continuity of rotation. The wrench predictor is trained according to Eq. 4, which in practice simplifies to mean-squared error $\|y - \hat{y}\|^2$ assuming a unit variance.

Table I describes the distribution of the damaged HSA data provided to the pretrained LEAP model. Ground-truth data is collected via an OptiTrack fiducial tracker and a Robotous RFT60-HA01 six-axis force/torque sensor, both at 360Hz, by hand using a handle affixed to the end of an HSA. Training and testing data for each HSA includes bending, twisting, compression, extension, and buckling behaviors. Collecting a diverse set of training data is important for producing a structured latent space: while the number of combinations for 2+ types of HSA damage is large, we train on data from HSAs with one cut on each of the 21 “vertical cut” coordinates, and each of the 9 possible burn locations.

TABLE I
DATA QUANTITY FOR TRAINING

	Data (Minutes)	HSAs Used
No Damage	112	32
One Cut	146	32
Two Cuts	27	6
One Burn	41	9
2+ Cuts / Burns	45	10

C. Adapting to Damage with HSA Ensembles

Rapid adaptation is critical not only due to constraints on compute and time, but also on hardware—a partially-broken system is more likely to degrade during the adaptation process itself if it is sample-inefficient. LEAP overcomes this constraint by using a simple yet effective ensemble method, requiring very little data and consisting only of linear operators that regress in near real-time. Furthermore, practically functional adaptation methods must be robust to large variations in damage, rather than defining adaptable states into rigid bounds that can be easily predicted by first principles. As such, our method for adapting in-situ involves no training or fine-tuning of either the damage encoder $\mathcal{D}$ or force predictor $\mathcal{W}$.

To adapt to out-of-distribution damage, our method relies on the weighted average of an ensemble of predictions driven by the latent structure. The ensemble is produced by clustering a large number of sampled encoder outputs to produce a computationally tractable ensemble of E latent coordinates $E = \{z_1 \dots z_K\}$ at the centroids of each k-means cluster. This approximates sampling from maximally differing HSA damages. These centroids are then used to produce an ensemble

Algorithm 1: Latent Ensemble Adaptation

Input: Networks $\mathcal{D}$, $\mathcal{W}$, Poses P , Ref. Wrenches y

Output: Ensemble Weights $\mathcal{A}$, Wrenches $\mathcal{F}$

Generate n binary vectors and latent coordinates:

$x_i \sim \mathcal{U}(\{0, 1\}^{54}), \quad i = 1, \dots, n$

$\hat{z}_{(i)} \leftarrow \mathcal{D}(x_i)$

Construct latent ensemble via clustering:

$E \leftarrow \text{K-Means}(\{\hat{z}_1, \dots, \hat{z}_n\})$ where $E = \{z_1, \dots, z_K\}$

Generate ensemble predictions:

for $i = 1$ **to** K **do**

$\hat{y}_{(i)} \leftarrow \mathcal{W}(P|z_i)$

$\hat{Y} = [\hat{y}_{(1)}, \dots, \hat{y}_{(E)}]$

Solve least-squares regression:

$\mathcal{A} \leftarrow \arg \min_{\mathcal{A}} \|\hat{Y}\mathcal{A} - y\|_2^2$

return $\mathcal{A}$

LEAP Inference:

$\mathcal{F} \leftarrow \mathcal{A}^T \cdot \mathcal{W}(P, E)$

return $\mathcal{F}$

of force/torque predictions by evaluating the entire adaptation window through the wrench predictor $\mathcal{W}(P, z)$.

Given either a ground truth force/torque signal or a reference weight picked up by a manipulator, weightings by-component for each member in the ensemble are calculated using a least-squares regression. Inference after obtaining ensemble weightings is given by:

$$y(P) = \sum_{i=1}^{i=30} a_i \mathcal{W}(p, z_i) \quad (7)$$

As all other elements of the ensemble are generated before runtime, regression of the ensemble weights is computationally lightweight and scales linearly with the number of actuators in the overall system. In practice, we find that for single HSAs, strong performance is achieved after 30 seconds of data collection at 100Hz, 5 seconds of data pre-processing and 2 seconds of regression, totaling 37 seconds for the whole process.

Our method does assume that the transform between HSA ends is known at runtime. For our evaluations on a 6-DoF soft wrist, this is achieved by a camera and AprilTag inside the mechanism as shown in Figure 3. For a single HSA, this ground truth is provided by OptiTrack motion capture cameras.

V. EXTENDING LEAP TO A 6-DOF SOFT WRIST

As a linear ensemble method, LEAP naturally extends to arbitrary hybrid soft robots provided that:

- Actuators originate and terminate exclusively between rigid elements,
- the relative positions of the rigid elements are known, and

- actuator models can be developed as described in Section IV-B.

If these conditions are met, the learned adaptation problem can be formulated as a linear least-squares regression over ensemble weightings—the only difference to the single-HSA case is the dimension of the linear regression.

To evaluate our method, we designed a 6-DoF soft wrist driven by six HSAs, as shown in Figure 3. Each HSA is driven by a servomotor that has 180 degrees of motion, resulting in a maximum linear travel distance of 7.2mm in Z . Contained within the wrist is a camera and AprilTag for proprioceptive pose sensing that is used for force/torque prediction during runtime.

We begin by establishing the kinematic relationships between each HSA actuator and the two rigid components of the system: the wrist chassis and the interface plate where the AprilTag is mounted. In this work those relationships are recorded in an XACRO file based on the system’s CAD model, but in principle could be determined experimentally or even learned. The internal camera establishes the position of the interface plate relative to the wrist chassis and the forward kinematics determines the transforms of each HSA relative to its base.

This transform is then augmented with servo rotation feedback to provide the inputs for LEAP. While it is feasible to model each actuator as its own ensemble with a separate set of damage seeds, we chose instead to initialize a single set of 30 latent clusters and determine different weightings for every constituent actuator on the wrist. The wrenches generated by each element of the ensemble can then be transformed to a common frame, weighted, and summed. The adaptation process, as with a single actuator, is equivalent to performing a least squares regression over those weights.

VI. RESULTS

Our experiments were structured to answer the following questions:

- How well does LEAP adapt to damage in a single HSA?
- Does LEAP extend to systems with multiple HSAs in parallel, such as a 6-DoF wrist?
- How well does LEAP generalize to out-of-distribution changes, including unseen damages and actuator repairs?
- Can LEAP adaptation produce task-level improvements?

A. Single HSA Damage Adaptation

In Figure 4, we show the impact of a vertical cut at every coordinate across an HSA. Figure 4A indicates the coordinate system used to describe where cuts and burns along an HSA are located. We validate LEAP by comparing a damage-adapted HSA to itself with LEAP weights learned for the fully-intact actuator (i.e., after damage occurs, but before adaptation). Each square in Figure 4C represents a cut at the corresponding coordinate in an HSA—we compare each cut to a prior evaluation of each HSA to control for wear and manufacturing inconsistencies.




Fig. 3. **HSA-Actuated Soft Wrist:** A. We demonstrate LEAP on a miniature, parallel 6-DoF soft wrist. Servomotors drive HSAs, and a proprioceptive camera determines platform pose via an AprilTag. B. The soft wrist mounted to a Franka Emika Panda robot arm. We demonstrate the wrist’s functionality via a tracing task where the objective is to traverse the contour of a 3D-printed surface while remaining in contact as much as possible.

Without adaptation, damaged HSAs have force/torque sensing errors of as much as $0.25 \text{ (N}^2, \text{N}^2\text{cm}^2)$ log-MSE, or $1.78 \text{ (N}^2, \text{N}^2\text{cm}^2)$ mean-squared error averaged per component. After adaptation, LEAP produces an average reduction in sensing error of 87% compared to the damaged and unadapted HSAs.

While LEAP improves performance across all axes, the largest raw improvements occur in T_x , T_y , F_z , and T_z . This is largely due to the geometry of the architected transmission: HSAs transduce torques and forces along the Z axis, and the beam bending behavior of HSAs along with the reaction forces due to the material’s handedness produces torques about X and Y . As shown in Supplemental Figure S1, unadapted predictions can be accurate in the X and Y axes but highly inaccurate in Z , the primary force/torque transmission axis.

B. Adaptation on the 6-DoF Soft Wrist

We used LEAP to adapt force/torque proprioception models to five different configurations of our HSA-actuated soft wrist system, as shown in Figure 5. Beginning with an undamaged robot, we incrementally introduce cuts, burns, and repairs to modify the dynamics of the system. In each case, we adapt a LEAP proprioception model to 60s of randomly selected data, a process which took 4.5s on an RTX3080 laptop GPU. The ground-truth for this experiment was provided by the six-axis force-torque sensor and motion tracking system described in Section V.

We then evaluate each resulting model on data from all five robot configurations and present the results in Figure 5. Along the diagonal of the confusion matrix in Figure 5, the LEAP ensemble weights are fit to the specific damaged embodiment. In the top-right, LEAP is evaluated on embodiments that are




Fig. 4. **LEAP Enables Robust, Rapid Damage Adaptation.** A.) **HSA Damage as Discrete Coordinates:** The periodic structure of architected materials introduces degeneracies: in an HSA, severed links can be either on a vertical strut (A, B, C) highlighted in the undamaged HSA, or on the horizontal links connecting the major struts. Torching HSAs occurs in one of three locations along each strut. B.) **Repair in Architected Materials:** Adaptation after repair in an architected material is operationally equivalent to damage adaptation and achievable by hobbyists with limited resources. Cut HSAs are repaired with glue and after adaptation their weighted latent coordinates are distinct despite all links being intact. C.) **Adaptation across the space of possible cuts.** We perform a vertical cut at every coordinate along an HSA. After evaluating LEAP, we observe a performance improvement of 87% in MSE loss across all components. Average component error decreases by 85% after adaptation.

more damaged than its weights were fit to, in the bottom-left the LEAP weights are fit to embodiments that are more damaged than what is evaluated.

We find that LEAP is able to model all five robot configurations, with mean error magnitudes of 1.5N in force and 0.074Nm in torque. LEAP delivers significant improvements—adapted models decrease error by an average of 49% compared to models trained on another robot configuration and deployed without adaptation. Figure 5 shows that unadapted models generally perform better in damage configurations that produce force/torque distributions that are closer to the support of their original LEAP adaptation data. However, it is interesting to note that the biggest drop in performance came from deploying models trained on the fully undamaged system on partially damaged datasets. We believe this reflects the additional stiffness imparted by having six undamaged actuators in a near-kinematically-constrained arrangement.

C. Adaptation to Repaired HSAs

Repair of architected materials such as HSAs introduces inconsistencies that must be accommodated with an adaptation procedure such as LEAP.

In Table II, we measure the Z-axis spring constant of five HSAs—as-printed, after a vertical cut, and after the auxetic links were repaired using glue. Arrows indicate whether the repair increased or decreased the spring constant compared

to the intact HSA. Repair is an imprecise process: despite repairing each HSA, their spring constants differed from the default stiffness by 9% in varying directions.

TABLE II
HSA SPRING CONSTANTS BEFORE AND AFTER REPAIR

	Default	One Cut	Repaired	Unit
HSA 1	0.36	0.30	↑ 0.39	N/mm
HSA 2	0.33	0.30	↑ 0.36	N/mm
HSA 3	0.34	0.30	↓ 0.32	N/mm
HSA 4	0.33	0.27	↑ 0.35	N/mm
HSA 5	0.34	0.29	↓ 0.30	N/mm

This behavior extends to LEAP’s latent representations—in Figure 4B we compare four undamaged HSAs to their repaired counterparts. While LEAP performs similarly in mean-squared error for force/torque prediction compared to the unadapted case, the weighted latent coordinates for the healthy and damaged ensembles differ significantly, indicating that a change occurs in the HSA dynamics. We explicitly made the choice to not include repaired HSAs as a means of testing out-of-distribution performance—in practical applications, actuator repair is an ad-hoc process in response to unpredictable damage, which occurs in the context of the environment and resources available.




Fig. 5. **Wrist Adaptation Performance:** Our whole system model fit to data from increasingly damaged wrist configurations. The confusion matrix diagonal shows that the system is able to adapt to a variety of system damage states using less than a minute of total data collection and computation time. The off-diagonal elements show a clear trend whereby increasing damage accumulation degrades model performance in the absence of adaptation.

D. Task Demonstration

To demonstrate LEAP’s performance in a contact-rich setting, we designed a tracing task wherein the HSA-actuated soft wrist is mounted to a Franka Emika Panda robot arm. The objective is to use the proprioception of the wrist alone (i.e., no vision or predefined motions) to slide along a curved contour, maintaining contact as much as possible.

In this setting, the wrist is passive and functions only as a sensor. For tracing, it is only required that the proprioceptive outputs are “force-like” and “torque-like” (i.e., have a monotonic relationship with respect to the real forces and torques experienced by the system). However, they must be directionally correct and maintain a consistent magnitude after adapting to damage, as the motion of the arm is driven by a PID loop about the wrist’s sensory output.

Before the wrist is damaged, the system is able to keep the end of the wrist in sliding contact with the contour for 85%




Fig. 6. **Tracing Task Demonstration:** Representative trajectories from the proprioceptive tracing task. Without adaptation, the damaged wrist is not task capable, it fails to follow the contour and cannot recover. After adapting with LEAP, the system recovers up to 100% of its original performance.

of the total trajectory. After adapting to burning one HSA and cutting a second (see the provided multimedia), the wrist is able to stay in contact with the contour for an identical 85% of the total path averaged over three trials, measured using a video analysis tool. Without LEAP adaptation, the damaged wrist maintains contact only briefly and catastrophically fails before reaching the end of the contour path—the damaged, unadapted system is not task capable.

VII. IN-DISTRIBUTION PROPERTIES

One of the insights from this work is that architected materials have a number of advantages over conventional continuum soft materials when it comes to characterizing and adapting to damage.

For truss-based architected materials, such as HSAs, cuts at different points along a discrete element of the structure are functionally degenerate, creating the same impact on the bulk material behavior. This observation forms the basis of first-principles truss analysis, used to characterize the behavior of infrastructure such as bridges, buildings, and towers. For linear elastic structures, the principle of superposition dictates that deflections caused by multiple forces can be summed independently of their ordering [49]—this implies that a finite number of destructive tests would cover the distribution of all damages.

In an architected material, these same degeneracies create a coordinate system by which damage can be encoded and damages can be superimposed. For example, in the HSAs used in this paper, any cut, burn, or combination thereof can be accurately described as part of a 54 dimensional binary vector.

We considered three types of damages: Removing a link in the structure (cutting), restoring a link in the structure (gluing), and fusing adjacent links (burning): all of these either increase or decrease the stiffness of the architected material. Since these changes are local, the damage accumulated in one part of the structure does not change the impact of additional damage elsewhere in the structure. Moreover, the impact of multiple damages is therefore order-independent, as long as they are non-overlapping [49].

Therefore, multiple instances of damage can be understood as a linear combination of their individual effects. This property is important for sample efficiency [50, 51] and, in effect, LEAP learns the impacts of individual damages along with residuals of their linear combinations. The approximate linearity of damage states on an HSA also motivates the choice of ensembles of wrench predictions. We rely on the support of the ensemble distribution to encompass the ground-truth signal and expect linear combinations of HSA damages to be supported under a data collection process where largely only single HSA damages are considered.

An architected material with a discrete damage description is intrinsically useful: a finite number of experiments can fully model the excitation of any damaged mechanism with respect to its learning inputs. For discrete damage, fully modeling the system requires N^m datasets, where N is the size of the binary damage vector and m is the number of times each transmission can be damaged.

However, superimposing linear combinations of damage collapses this exponential scaling. In a perfect linear elastic truss, only N datasets are required to fully model the system. In practice, we find that additional data is helpful for model stability—particularly in modeling the process variability of HSA manufacturing—though we consider $N \times m$ to be an approximate upper bound for the number of experiments required. Additionally, when designing HSA-actuated mechanisms, the quantity of HSA active units changes with either length or width (diameter) of the transmission, but not both. This means that the scaling complexity of N , and therefore the binomial distribution of possible damages, is linear with respect to changes in design.

This has implications for robots that expect to sustain and adapt to damage. In contrast to rigid or continuum soft mechanisms, in-distribution characteristics are guaranteed for learning representations of damaged architected materials given a finite and countable number of experiments. The ability to employ LEAP and LEAP-like algorithms reflects an inherent strength of architected materials relative to other classes of actuator.

VIII. CONCLUSION

In this work we introduce LEAP, a method of adapting to unseen damages in seconds in architected materials. We identify properties of architected materials that can be approximated with a truss-like structure that provides in-distribution guarantees for all local damages while using a comparatively small quantity of data and a countable number of physical

experiments—a distinct advantage over both rigid mechanisms and continuum soft actuators. We evaluate LEAP in single- and multi-HSA systems, introducing a 6-DoF soft wrist, and demonstrate that LEAP recovers proprioception performance in a tactile tracing task.

An additional empirical finding from our experiments in single- and multi-HSA systems, as shown in Figure 5 and supplementary Figure S1, is that improved proprioceptive results can be achieved when portions of the architected materials were broken. Given the nonlinearities present in HSAs, separating the energy-containing members of its auxetic structure provides an explanation for why this is: by removing nonlinearities, the architected material bears a closer resemblance to a linear spring or simple beam, which can be more easily parameterized. Architected materials possess an advantage in adaptability not only from their in-distribution properties as discussed in Section VII, but also from the selective removal of structure in order to improve a behavior or assume a new one altogether. Further investigation of the design tradeoff between force transmission and algorithmic observability is a direction of future work.

Limitations do exist to our method: one drawback is that despite approximately linear experimental complexity, a reasonable amount of actuator characterization is required in order to apply LEAP. Furthermore, there exists room for improvement in the ensemble selection process itself—our method presumes that selecting centroids of the latent clusters provides sufficient dynamical variation to solve the least-squares weighting. More intelligent selection of ensemble members or specific motion strategies for collecting adaptation data would both improve efficiency and performance. Another limitation is that we do not encode damages such as wear, abrasion, and localized hinging in this work, although anecdotally we find that LEAP is able to adapt well to all three, as well as HSA manufacturing variability.

This paper presents not only a method of adapting to damage and degradation of robotic components but also a new lens through which to view robot design. Representing the damage state of an architected material as a discrete encoding satisfies conditions that are useful for statistical learning beyond proprioception and self-modeling. **One interpretation of this work is that architected materials—because of their in-distribution guarantees—have advantages over rigid and continuum materials for mechanisms that will be required to learn.**

IX. ACKNOWLEDGMENTS

We thank Ayush Gaggar for his insightful feedback on this work. This research is based upon work supported by the US Army Research Office Grant no. W911NF-22-1-0286 and the National Science Foundation under Grant no. 2330040, the Engineering Research Center (ERC) for Human Augmentation via Dexterity (HAND). J.A. is supported by a National Defense Science and Engineering Graduate Fellowship.

REFERENCES

- [1] M. Gazzola, L. H. Dudte, A. G. McCormick, and L. Mahadevan. Forward and inverse problems in the mechanics of soft filaments. *Royal Society Open Science*, page 171628, 2018. URL <https://royalsocietypublishing.org/doi/10.1098/rsos.171628>.
- [2] EunBi Oh, Taekyoung Kim, Pranav Kaarthik, and Ryan L. Truby. Architected Materials for Soft Robotics. *Journal of Materials Research*, 2026. URL <https://doi.org/10.1557/s43578-025-01778-2>.
- [3] The Skild AI Team. The case for an omni-bodied robot brain, 2025. URL <https://www.skild.ai/blogs/omni-bodied>.
- [4] Molly Carton, Jakub F. Kowalewski, Jiani Guo, Jacob F. Alpert, Aman Garg, Daniel Revier, and Jeffrey Ian Lipton. Bridging hard and soft: Mechanical metamaterials enable rigid torque transmission in soft robots. *Science Robotics*, 2025. URL <https://www.science.org/doi/full/10.1126/scirobotics.ads0548>.
- [5] Z. Roth, B. Mooring, and B. Ravani. An overview of robot calibration. *IEEE Journal on Robotics and Automation*, 1987. URL <https://ieeexplore.ieee.org/document/1087124/>.
- [6] Boston Dynamics. SpotCheck (joint and camera calibration), 2022. URL <https://support.bostondynamics.com/s/article/SpotCheck-Joint-and-Camera-Calibration-49935>.
- [7] Giovanni Franzese, Max Spahn, Jens Kober, and Cosimo Della Santina. MUKCa: Accurate and affordable cobot calibration without external measurement devices, 2025. URL <https://arxiv.org/abs/2503.12584>.
- [8] Zhibin Li, Shuai Li, and Xin Luo. Data-driven industrial robot arm calibration: A machine learning perspective. In *2021 IEEE International Conference on Networking, Sensing and Control (ICNSC)*, volume 1, pages 1–6, 2021. doi: 10.1109/ICNSC52481.2021.9702246.
- [9] Jeffrey Ian Lipton, Robert MacCurdy, et al. Handedness in shearing auxetics creates rigid and compliant structures. *Science*, 2018. URL <https://www.science.org/doi/abs/10.1126/science.aar4586>.
- [10] Chen Liu and Minh-Son Pham. Spatially programmable architected materials inspired by the metallurgical phase engineering. *Advanced Materials*, 2024. URL <https://advanced.onlinelibrary.wiley.com/doi/abs/10.1002/adma.202305846>.
- [11] Yi Yang, Marcelo A. Dias, and Douglas P. Holmes. Multistable kirigami for tunable architected materials. *Phys. Rev. Mater.*, Nov 2018. URL <https://link.aps.org/doi/10.1103/PhysRevMaterials.2.110601>.
- [12] Taekyoung Kim, Pranav Kaarthik, et al. A flexible, architected soft robotic actuator for motorized extensional motion. *Advanced Intelligent Systems*, 2024. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/aisy.202300866>.
- [13] Ahmad Rafsanjani, Yuerou Zhang, Bangyuan Liu, Shmuel M. Rubinstein, and Katia Bertoldi. Kirigami skins make a simple soft actuator crawl. *Science Robotics*, 2018. URL <https://www.science.org/doi/abs/10.1126/scirobotics.aar7555>.
- [14] Ryan L. Truby, Lillian Chin, and Daniela Rus. A Recipe for Electrically-Driven Soft Robots via 3D Printed Handed Shearing Auxetics. *IEEE Robotics and Automation Letters*, 6(2):795–802, 2021. doi: 10.1109/LRA.2021.3052422.
- [15] Jake Ketchum, Sophia Schiffer, et al. Automated gait generation for walking, soft robotic quadrupeds. *IEEE International Conference on Intelligent Robots and Systems (IROS)*, pages 10245–10251, 2023. doi: 10.1109/IROS55552.2023.10342059.
- [16] Taekyoung Kim, Alexander Torres-Soto, Haodong Wang, and Ryan L. Truby. Architected elastomers as load bearing actuators for untethered soft robot walking. *Advanced Materials Technologies*, 2025. URL <https://advanced.onlinelibrary.wiley.com/doi/abs/10.1002/admt.202500503>.
- [17] Jake Ketchum, James Avtges, et al. Force and speed in a soft stewart platform. In *2025 IEEE 8th International Conference on Soft Robotics (RoboSoft)*, 2025. doi: 10.1109/RoboSoft63089.2025.11020875.
- [18] Lillian Chin, Michelle C. Yuen, Jeffrey Lipton, Luis H. Trueba, Rebecca Kramer-Bottiglio, and Daniela Rus. A Simple Electric Soft Robotic Gripper with High-Deformation Haptic Feedback. In *International Conference on Robotics and Automation (ICRA)*, 2019. URL <https://ieeexplore.ieee.org/abstract/document/8794098>.
- [19] Taekyoung Kim, Eliot A. Dunn, and Ryan L. Truby. Clutchable soft actuators produce rapid, high-power movements in robotic artificial musculoskeletal systems. *IEEE Robotics and Automation Letters*, 2026. doi: 10.1109/LRA.2026.3655266.
- [20] Joseph Sullivan, Ian Good, Samuel A. Burden, and Jeffrey Ian Lipton. Spring-brake! handed shearing auxetics improve efficiency of hopping and standing, 2025. URL <https://arxiv.org/abs/2505.22898>.
- [21] James Avtges, Jake Ketchum, Millicent Schlafly, Helena Young, Taekyoung Kim, Allison Pinosky, Ryan L. Truby, and Todd D. Murphey. Real-Time Reinforcement Learning for Dynamic Tasks with a Parallel Soft Robot, 2025. URL <http://arxiv.org/abs/2509.19525>.
- [22] Zixin Zhang, James Avtges, and Todd Murphey. Sample-efficient online control policy learning with real-time recursive model updates. In *Proceedings of The 9th Conference on Robot Learning*, Proceedings of Machine Learning Research. PMLR, 27–30 Sep 2025. URL <https://proceedings.mlr.press/v305/zhang25d.html>.
- [23] John C. Tuthill and Eiman Azim. Proprioception. *Current Biology*, 2018. URL <https://www.sciencedirect.com/science/article/pii/S0960982218300976>.
- [24] Hongbo Wang, Massimo Totaro, and Lucia Beccai. Toward perceptive soft robots: Progress and challenges. *Advanced Science*, 2018. URL <https://advanced.onlinelibrary.wiley.com/doi/abs/10.1002/advs.201800541>.

- [25] Michelle C. Yuen, Rebecca Kramer-Bottiglio, and Jamie Paik. Strain sensor-embedded soft pneumatic actuators for extension and bending feedback. In *2018 IEEE International Conference on Soft Robotics (RoboSoft)*, 2018. doi: 10.1109/ROBOSOFT.2018.8404920.
- [26] Ryan L. Truby, Lillian Chin, et al. Fluidic innervation sensorizes structures from a single build material. *Science Advances*, 8(31):eabq4385, 2022.
- [27] Vani Sundaram, Khoi Ly, Brian K. Johnson, Mantas Naris, Maxwell P. Anderson, James Sean Humbert, Nikolaus Correll, and Mark Rentschler. Embedded magnetic sensing for feedback control of soft hasel actuators. *IEEE Transactions on Robotics*, pages 808–822, 2023. doi: 10.1109/TRO.2022.3200164.
- [28] Yusen Zhao, Chiao-Yueh Lo, Lecheng Ruan, Chen-Huan Pi, Cheolgyu Kim, Yousif Alsaied, Imri Frenkel, Rossana Rico, Tsu-Chin Tsao, and Ximin He. Somatosensory actuator based on stretchable conductive photothermally responsive hydrogel. *Science Robotics*, 2021. doi: 10.1126/scirobotics.abd5483. URL <https://www.science.org/doi/abs/10.1126/scirobotics.abd5483>.
- [29] Annan Zhang, Ryan L. Truby, Lillian Chin, Shuguang Li, and Daniela Rus. Vision-based sensing for electrically-driven soft actuators. *IEEE Robotics and Automation Letters*, 7(4):11509–11516, 2022. doi: 10.1109/LRA.2022.3201604.
- [30] Yifan Zhu, Mei Hao, Xupeng Zhu, Quentin Bateau, Alex Wong, and Aaron M. Dollar. Forces for free: Vision-based contact force estimation with a compliant hand. *Science Robotics*, page eadq5046, 2025. URL <https://www.science.org/doi/full/10.1126/scirobotics.adq5046>.
- [31] Dickson Chiu Yu Wong, Yucheng Luo, and Rob Bernardus Nicolaas Scharff. Durability of soft pneumatic actuators: A review and benchmarking protocol. *Advanced Robotics Research*, 2026. URL <https://advanced.onlinelibrary.wiley.com/doi/abs/10.1002/adrr.202500172>.
- [32] Abhirup Sarker, Tamzid Ul Islam, and Md. Robiul Islam. A review on recent trends of bioinspired soft robotics: Actuators, control methods, materials selection, sensors, challenges, and future prospects. *Advanced Intelligent Systems*, 2025. URL <https://advanced.onlinelibrary.wiley.com/doi/abs/10.1002/aisy.202400414>.
- [33] Thomas George Thuruthel, Benjamin Shih, Cecilia Laschi, and Michael Thomas Tolley. Soft robot perception using embedded soft sensors and recurrent neural networks. *Science Robotics*, 2019. URL <https://www.science.org/doi/abs/10.1126/scirobotics.aav1488>.
- [34] Hedan Bai, Young Seong Kim, and Robert F. Shepherd. Autonomous self-healing optical sensors for damage intelligent soft-bodied systems. *Science Advances*, 2022. URL <https://www.science.org/doi/10.1126/sciadv.abq2104>.
- [35] Antoine Cully, Jeff Clune, Danesh Tarapore, and Jean-Baptiste Mouret. Robots that can adapt like animals. *Nature*, 2015. URL <https://doi.org/10.1038/nature14422>.
- [36] Konstantinos Chatzilygeroudis, Vassilis Vassiliades, and Jean-Baptiste Mouret. Reset-free trial-and-error learning for robot damage recovery. *Robotics and Autonomous Systems*, 100:236–250, 2018.
- [37] Sam Kriegman, Stephanie Walker, Dylan S. Shah, Michael Levin, Rebecca Kramer-Bottiglio, and Josh Bongard. Automated shapeshifting for function recovery in damaged robots. In *Robotics: Science and Systems XV*. Robotics: Science and Systems Foundation, 2019. URL <http://dx.doi.org/10.15607/RSS.2019.XV.028>.
- [38] J.C. Bongard and H. Lipson. Automated damage diagnosis and recovery for remote robotics. In *IEEE International Conference on Robotics and Automation (ICRA)*, volume 4, pages 3545–3550, 2004. doi: 10.1109/ROBOT.2004.1308802.
- [39] Josh Bongard, Victor Zykov, and Hod Lipson. Resilient machines through continuous self-modeling. *Science*, 2006. URL <https://www.science.org/doi/abs/10.1126/science.1133687>.
- [40] Robert Kwiatkowski and Hod Lipson. Task-agnostic self-modeling machines. *Science Robotics*, 2019. URL <https://www.science.org/doi/abs/10.1126/scirobotics.aau9354>.
- [41] Mengyuan Yan, Yilin Zhu, Ning Jin, and Jeannette Bohg. Self-supervised learning of state estimation for manipulating deformable linear objects. *IEEE Robotics and Automation Letters*, 5(2):2372–2379, 2020. doi: 10.1109/LRA.2020.2969931.
- [42] Kuai Yu, Crystal Su, Xiang Liu, Judah Goldfeder, Mingyuan Shao, and Hod Lipson. LyTimeT: Towards robust and interpretable state-variable discovery, 2025. URL <https://arxiv.org/abs/2510.19716>.
- [43] Boyuan Chen, Robert Kwiatkowski, Carl Vondrick, and Hod Lipson. Fully body visual self-modeling of robot morphologies. *Science Robotics*, 2022.
- [44] Anuj Pokhrel, Aniket Datar, Mohammad Nazeri, Francesco Cancelliere, and Xuesu Xiao. Zero-shot adaptation to robot structural damage via natural language-informed kinodynamics modeling, 2026. URL <https://arxiv.org/abs/2602.12385>.
- [45] Alexandra Otto, Max Rosenkranz, Karl Alexander Kalina, and Markus Kästner. Data-driven inverse design of spinodoid architected materials. *GAMM-Mitteilungen*, 48, 2025. doi: 10.1002/gamm.70008.
- [46] Zuyu Li, Wei Gao, Nicole Kessissoglou, Sebastian Oberst, Michael Yu Wang, and Zhen Luo. Multifunctional mechanical metamaterials with tunable double-negative isotropic properties. *Materials & Design*, 232: 112146, 2023. doi: 10.1016/j.matdes.2023.112146.
- [47] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-VAE: Learning basic visual concepts with a constrained variational framework. *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=Sy2fzU9gl>.
- [48] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the Continuity of Rotation Representations in

- Neural Networks, 2020. URL <http://arxiv.org/abs/1812.07035>.
- [49] William S. Slaughter. *The Linearized Theory of Elasticity*. Birkhäuser, Boston, MA, 2002. doi: 10.1007/978-1-4612-0093-2.
- [50] Robert Kohn, Michael Smith, et al. Nonparametric regression using linear combinations of basis functions. *Statistics and Computing*, 11(4):313–322, 2001.
- [51] Milan Korda and Igor Mezić. Optimal construction of Koopman eigenfunctions for prediction and control. *IEEE Transactions on Automatic Control*, 65(12):5114–5129, 2020.