

Indirect Thermal-Aware Supervision for Agonist–Antagonist SMA Ankle Actuation

Jácomo-Zavaleta, S.^{a,*}, Copaci, D.^a, Die-Pancorbo, S.^a, Blanco, D.^a, Moreno, L.^a

^aDepartment of Systems Engineering and Automation, Universidad Carlos III de Madrid, 28911 Leganés, Spain

Resumen

Las Aleaciones con Memoria de Forma (SMA) ofrecen alta densidad de potencia para rehabilitación portátil, pero su aplicación se ve limitada por la lenta relajación térmica y la acumulación de calor. Este trabajo presenta una unidad modular SMA agonista–antagonista para el tobillo, validada mediante Hardware-in-the-Loop. Se implementa un PID discreto coordinado por una lógica supervisora con supervisión térmica indirecta, capaz de favorecer la relajación térmica mediante intervalos intra-ciclo de “Enfriamiento Libre” e intervalos de recuperación entre lotes activos sin desactivar el lazo cerrado. En un modelo mecánico de tobillo con tres grados de libertad (3-DOF), el sistema mantuvo operación estable durante 40 minutos sin saturación térmica, logrando un RMSE global (fases activas) de 0.218° y un error máximo de 1.384°. La repetibilidad cíclica mostró una dispersión media de 0.068° sin deriva progresiva. La principal contribución es integrar explícitamente la recuperación térmica en la arquitectura de control mediante supervisión térmica indirecta, mejorando la consistencia operativa a largo plazo en sistemas SMA antagonistas.

Palabras clave: Sistemas mecatrónicos, Simulación Hardware-in-the-loop, Mecatrónica biomédica, Robótica blanda, Robótica de rehabilitación

Abstract

Shape Memory Alloys (SMAs) offer high power density for wearable rehabilitation devices, yet their practical use is constrained by slow thermal relaxation and heat accumulation. This paper presents a modular agonist–antagonist SMA ankle actuation unit validated in a Hardware-in-the-Loop framework. A discrete PID controller is coordinated by a supervisory state machine with indirect thermal-aware supervision, which favors thermal relaxation through intra-cycle “Free Cooling” intervals and recovery intervals between active batches without disabling closed-loop control. Experiments on an ankle surrogate with three mechanical degrees of freedom (3-DOF) demonstrate stable operation over 40 minutes without thermal saturation, achieving an active-only global RMSE of 0.218° and a maximum instantaneous error of 1.384°. Cyclic repeatability yielded a mean dispersion of 0.068° with no progressive drift. The main contribution is the explicit integration of event-driven thermal recovery into the control architecture through indirect thermal-aware supervision, enabling sustained and consistent long-duration operation of antagonistic SMA systems.

Keywords: Mechatronic systems, Hardware-in-the-loop simulation, Biomedical mechatronics, Soft robotics, Rehabilitation robotics

1. Introduction

Soft robotics and wearable rehabilitation devices demand actuation systems that are lightweight, compliant, and mechanically transparent to the user (Laschi et al., 2016). Conventional electric motors introduce inertia and mechanical impedance, while pneumatic systems require bulky external infrastructure (Reynaerts and Brussel, 1998). Thermally responsive Shape Memory Alloys (SMAs), particularly NiTi wires, offer high work density, silent operation, and bio-inspired linear contraction, making them attractive candidates for wearable joint as-

sistance (Hamid et al., 2023; Mohd Jani et al., 2014).

SMA actuation relies on a reversible thermo-mechanical phase transformation. Although Joule heating enables rapid contraction, cooling is governed by passive convection and exhibits nonlinear exponential recovery (Rao et al., 2015). This intrinsic asymmetry induces hysteresis, frequency-dependent displacement attenuation, and progressive heat accumulation across cycles. Under insufficient recovery time, peak delay and amplitude reduction emerge, eventually requiring forced rest intervals or thermal interruption (Copaci et al., 2019b; Ballester et al., 2023).

*Correspondence: sjacobo@pa.uc3m.es

Agonist–antagonist configurations partially mitigate relaxation limits by actively stretching the opposing actuator (Langbein and Czechowicz, 2012). However, without supervisory coordination, co-contraction may increase energy consumption and accelerate heat accumulation. Wearable implementations therefore often enforce temperature thresholds to prevent wire degradation, interrupting operation once critical limits are reached (Copaci et al., 2019b). Consequently, reported systems remain restricted to slow rehabilitation motions or require cooldown periods after limited cyclic operation.

From a control perspective, nonlinear compensation schemes such as Bilinear Proportional–Integral–Derivative (BPID) have demonstrated effectiveness in SMA position regulation (Villoslada et al., 2015). However, most implementations focus on local tracking accuracy rather than long-duration thermal behavior. As a result, endurance and tracking stability are often treated as materials or cooling problems rather than control-architecture problems.

We address this gap by integrating an agonist–antagonist SMA architecture with a supervisory control strategy that introduces indirect thermal-aware supervision into the switching logic. Rather than relying on temperature thresholds or scheduled rests, the controller reduces unnecessary Joule heating through event-driven Free Cooling intervals and protocol-level recovery phases while maintaining closed-loop stability. The main contributions are:

1. **Modular antagonistic SMA ankle unit:** A long-wire Bowden-based agonist–antagonist actuation architecture designed for sustained cyclic operation.
2. **Indirect thermal-aware supervisory integration of BPID control:** A hierarchical scheme that prevents co-contraction and promotes passive thermal relaxation without explicit temperature sensing.
3. **Demonstration of long-duration stability:** 40-minute cyclic tracking without thermal saturation, achieving active-only $RMSE_{\text{active}} = 0.218^\circ$ and sub-degree cyclic repeatability.

2. Architecture

The proposed unit is a standalone module applicable to multiple joints (ankle, wrist, elbow). For validation, it was deployed on a three-degree-of-freedom (3-DOF) ankle surrogate.

2.1. Electro-Mechanical Design

The actuation mechanism comprises two agonist–antagonist muscle groups: Flexor (dorsiflexion) and Extensor (plantarflexion). Each group consists of three 2-way SMA wires (NexMetal®, diameter 0.5 mm, activation temperature 55°C , safe working contraction 4%) connected mechanically and electrically in parallel within a single Bowden cable of length $L_{SMA} = 3.0$ m.

The use of multiple SMA wires within a single Bowden sheath follows experimentally validated configurations (Copaci et al., 2019b), demonstrating improved force scalability, lower energy consumption, and enhanced tension uniformity compared to individually routed wires. Although separate Bowden channels improve convective cooling, the shared configuration

offers superior mechanical balance and reliability for antagonistic operation.

Actuators use a low-side switching topology: they share a common $+V_{DC}$ with independent, opto-isolated return paths.

Given $\lambda = 5.1 \Omega/\text{m}$ and $L_{SMA} = 3.0$ m, one wire has $R_{SMA} = 15.3 \Omega$ and three in parallel yield $R_{eq} = 5.1 \Omega$. With $I_{\text{wire}} = 3.9$ A (per 0.5 mm wire), $I_{\text{total}} = 11.7$ A and the required supply is $V_{DC} \approx I_{\text{total}}R_{eq} \approx 60$ V, corresponding to ≈ 702 W peak per channel during activation.

The complete electro-mechanical routing and ankle coordinate definition are illustrated in Figure 1.




Figure 1: Actuation unit overview. (a) Electro-mechanical configuration of the antagonistic SMA actuation system, illustrating the low-side switching architecture and Bowden-tendon routing. The supply polarity ($+V_{DC}$, GND) and current directions (I_{flex} , I_{ext}) are explicitly indicated. Red and blue lines denote the electrical connections and current flow orientation for each actuator channel. Dashed lines denote the mechanical coupling between the SMA wires and the mannequin foot. (b) Ankle coordinate frames and range of motion between x_f and x_0 , representing dorsiflexion–plantarflexion limits measured by a Hall-effect sensor. Here, x_f denotes the rotating foot reference axis with respect to the fixed ankle reference axis x_0 .

2.2. Embedded HIL Architecture

Experimental validation used a Hardware-in-the-Loop (HIL) setup with an embedded STM32F4 (Cortex-M4, 168 MHz) and a host PC (Caballero et al., 2016). The control algorithm runs at $T_s = 1$ ms and PWM is generated at 1 kHz to drive the MOSFET stage; board-level tasks execute at 0.5 ms (2 kHz) and angle sensing is acquired at 2 ms (500 Hz), using the most recent sample in the control loop. The host generates references and logs data via USB VCP (12 Mbit/s) using a binary protocol at 1 ms.

3. Control Strategy Formulation

The proposed control architecture is organized in two hierarchical layers. The lower layer computes the actuation effort required for trajectory tracking, while the upper supervisory layer determines which actuator is allowed to inject energy into the system. This structure simultaneously addresses three challenges inherent to SMA actuation: (i) nonlinear thermo-mechanical behavior, (ii) antagonistic coordination, and (iii) long-term thermal stability.

3.1. Problem Formulation

Let $\theta(t)$ denote the measured ankle angle and $\theta_{ref}(t)$ the desired trajectory. The tracking error is defined as

$$e(t) = \theta_{ref}(t) - \theta(t). \quad (1)$$

Angle measurements are obtained using a contactless magnetic rotary encoder (AS5040), enabling sagittal-plane motion tracking without constraining other degrees of freedom.

3.2. Supervisory Intention Logic

The supervisory layer determines which actuator channel is permitted to generate torque at each sampling instant. Its purpose is twofold: to prevent antagonistic co-contraction and to introduce indirect thermal-aware supervision into the switching logic by favoring intervals of reduced Joule heating whenever active actuation is unnecessary.

Here, *indirect thermal-aware supervision* does not rely on direct temperature measurement. Instead, it is achieved through the joint evaluation of the reference increment and the signed tracking error, allowing the supervisor to suppress unnecessary actuation when the error is already decreasing due to passive effects such as gravity, elastic recoil, or mechanical relaxation. In this way, the controller promotes zero-effort intervals that favor thermal relaxation without explicit thermal sensing.

Reference-based intention estimation. Motion intention is inferred from the discrete increment of the reference trajectory as:

$$\dot{\theta}_{est}[k] \triangleq \theta_{ref}[k] - \theta_{ref}[k-1], \quad (2)$$

where $\epsilon_{vel} > 0$ is used to avoid switching due to quantization noise. The discrete tracking error is defined as:

$$e[k] = \theta_{ref}[k] - \theta[k]. \quad (3)$$

Supervisory state law. The supervisory state $\mathcal{S}[k] \in \{\text{FLEX}, \text{EXT}, \text{IDLE}\}$ is defined as

$$\mathcal{S}[k] = \begin{cases} \text{FLEX}, & \dot{\theta}_{est}[k] > \epsilon_{vel} \wedge e[k] > \epsilon_{dead}, \\ \text{EXT}, & \dot{\theta}_{est}[k] < -\epsilon_{vel} \wedge e[k] < -\epsilon_{dead}, \\ \text{IDLE}, & \text{otherwise.} \end{cases} \quad (4)$$

where ϵ_{dead} is the tracking *deadzone* introduced to prevent chattering around the setpoint. Here, FLEX enables the flexor channel (dorsiflexion), whereas EXT enables the extensor channel (plantarflexion). Torque is therefore applied only when required in the intended direction of motion. The resulting supervisory behavior, including transitions to IDLE that favor thermal relaxation, is summarized in Figure 2.

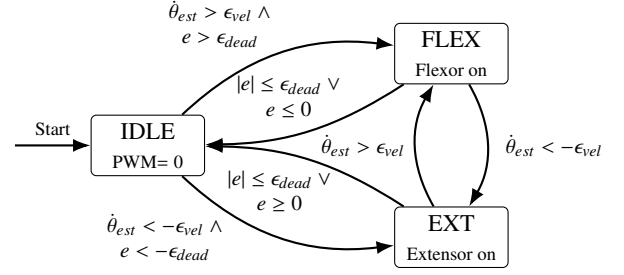


Figure 2: Supervisory intention logic. The reference increment $\dot{\theta}_{est}[k]$ selects the intended direction, while the signed error e enables corrective torque only when the joint lags. If passive effects already reduce the error ($e \leq 0$ in FLEX or $e \geq 0$ in EXT), or if $|e| \leq \epsilon_{dead}$, the supervisor returns to IDLE, increasing Passive Free Cooling and reducing antagonistic thermal load.

3.3. Discrete Control Law

Low-level position regulation is performed using a BPID controller, widely adopted in previous SMA-based actuation systems within our laboratory and the literature (Villoslada et al., 2015; Copaci et al., 2019a; Serrano et al., 2023) to mitigate nonlinearities in hysteretic plants while preserving implementation simplicity in embedded platforms. Compared with a conventional PID, the BPID introduces a state-dependent bilinear scaling that improves regulation under amplitude-dependent nonlinearities and hysteresis while preserving a low-complexity structure suitable for embedded implementation.

3.3.1. Control Structure

Given the tracking error $e[k] = r[k] - y[k]$, the linear PID effort is computed as:

$$u_{lin}[k] = K_p e[k] + K_i I[k] + K_d D[k], \quad (5)$$

where K_p , K_i , and K_d are the proportional, integral, and derivative gains, respectively, while $I[k]$ and $D[k]$ denote the integral state and the filtered derivative term.

The bilinear compensator is defined as:

$$\Gamma[k] = \frac{1 + K_0 r[k]}{1 + K_0 y[k-1]}, \quad (6)$$

where K_0 defines a bilinear scaling factor. Therefore, the final control effort becomes:

$$u_{PID}[k] = u_{lin}[k] \Gamma[k]. \quad (7)$$

The output $u[k] \in [0, 1]$ represents the normalized PWM duty cycle applied to the SMA actuator after saturation.

3.3.2. Antagonistic Routing and Reset

Given that only one actuator is enabled at a time, we define the routing of the unified control effort $u_{PID}[k]$ according to the supervisory state $\mathcal{S}[k]$:

$$u_{flx}[k] = \begin{cases} \text{sat}(u_{PID}[k]), & \mathcal{S}[k] = \text{FLEX} \\ 0, & \text{otherwise} \end{cases} \quad (8)$$

$$u_{ext}[k] = \begin{cases} \text{sat}(|u_{PID}[k]|), & \mathcal{S}[k] = \text{EXT} \\ 0, & \text{otherwise} \end{cases} \quad (9)$$

where $u_{flx}[k]$ and $u_{ext}[k]$ denote the routed and saturated PWM duty cycles applied to the flexor and extensor channels,

respectively. The absolute value in the extensor channel ensures non-negative PWM modulation while preserving directional information through the supervisory state.

To prevent integral windup and switching transients, each BPID channel is reset through a level-type supervisory signal whenever its actuator is inactive. This reset returns the integrator and derivative-filter states of the inactive channel to their initial conditions (set to zero). Thus, no residual integral action is retained during IDLE intervals, and the newly activated channel starts without residual bias (Ballester et al., 2023).

Control Architecture Overview. The control strategy defines a two-layer hierarchical logic consisting of a supervisory decision layer and a low-level BPID regulation layer. From an implementation standpoint, this logic is executed across three distinct layers within the HIL workflow: *Host*, *Target*, and *Plant*, as illustrated in Figure 3. The supervisory logic determines the active actuator channel, while the BPID controller computes a unified effort subsequently routed and saturated to drive the antagonistic SMA actuators. This separation between decision-level energy management and embedded actuation control enables indirect thermal-aware supervision without modifying the underlying regulation law.

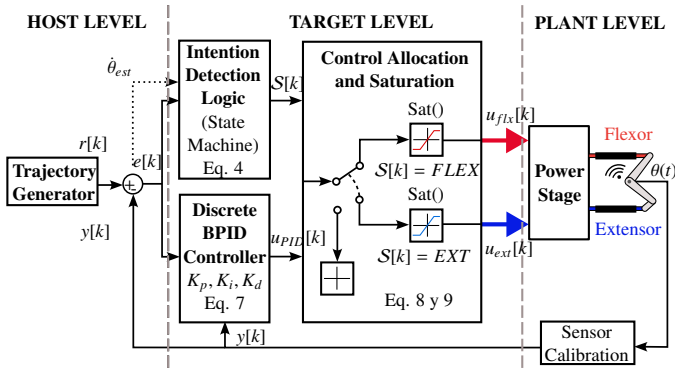


Figure 3: Control execution layers in our HIL architecture. The Host generates the reference trajectory and logs data. At the Target (embedded controller), the supervisory logic determines the active state $S[k]$, and the controller computes the unified effort $u_{PID}[k]$. The control allocation and saturation block routes the signal to $u_{flx}[k]$ or $u_{ext}[k]$, which drive the SMA actuators at the Plant level.

3.4. Thermal Recovery Mechanisms

In the experimental protocol, a *cycle* denotes one flexion–extension period of the sinusoidal reference, whereas a *batch* denotes one contiguous active segment composed of several consecutive cycles. Successive batches are separated by recovery intervals, whose exact durations are specified in the endurance experiment.

Within this framework, two complementary thermal recovery mechanisms naturally arise from the hierarchical control architecture:

Passive Free Cooling (intra-cycle): intervals where $u_{flx} = u_{ext} = 0$ while the joint continues converging due to gravity, elastic recoil, or mechanical relaxation.

Recovery interval (inter-batch): intervals in which the controller remains closed-loop but the reference is temporarily aligned with the measured angle ($\theta_{ref} \approx \theta$), minimizing energy injection while preserving stability.

4. Experimental Results

4.1. Agonist–Antagonist Coordination

Two mechanically distinct configurations were evaluated under a sinusoidal reference ($A = 5^\circ$, $f = 0.03$ Hz):

High pre-tension: both SMA tendons are tightened to maintain elevated baseline tension, reducing slack but increasing antagonistic stiffness.

Relaxed initial state: no additional tightening is applied, yielding lower antagonistic stiffness and greater passive compliance.

Table 1 summarizes the main tracking metrics.

Table 1: Performance trade-off: mechanical pre-tension vs. thermal robustness. RMSE: root mean square error; MAE: mean absolute error; Max Error: peak absolute deviation; Tail Mean Error: mean error over the last 20 s.

Metric [°]	High Pre-Tension	Relaxed State
RMSE	0.388	0.346
MAE	0.212	0.259
Max Error	2.440	1.170
Tail Mean Error	0.067	0.034
Thermal Behavior	Saturation Onset	Free Cooling

High pre-tension reduces mechanical backlash, which explains the lower MAE during the initial cycles. However, it also increases antagonistic stiffness, so the active channel must overcome a larger internal counter-torque. This raises the required PWM duty, accelerates heat accumulation, and progressively amplifies peak deviations, leading to earlier saturation effects.

In contrast, the relaxed state increases passive compliance, allowing gravity and elastic recoil to assist the return phases. Consequently, the supervisor more frequently selects IDLE, increasing zero-effort intervals ($u_{flx} = u_{ext} = 0$) and reducing antagonistic effort density. This translates into lower global RMSE and smaller peak error.

Figures 4 and 5 illustrate the qualitative difference between both configurations. These results indicate that thermal robustness is not solely determined by controller tuning, but emerges from the interaction between mechanical pre-tension and supervisory coordination.

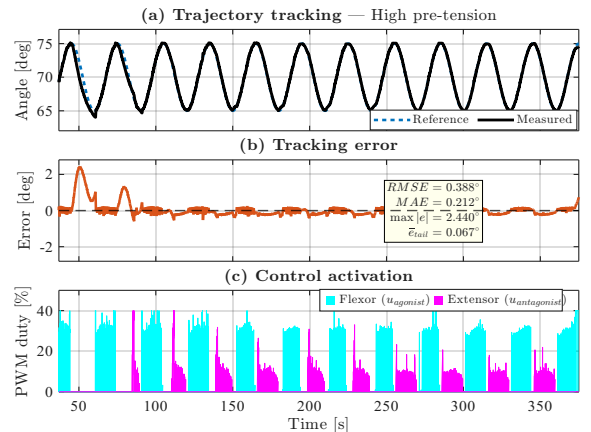


Figure 4: **High pre-tension configuration.** Reduced backlash yields lower instantaneous error in early cycles; however, increased antagonistic stiffness elevates energy injection, leading to progressive thermal accumulation, peak amplification, and onset of saturation effects after ~ 11 cycles.

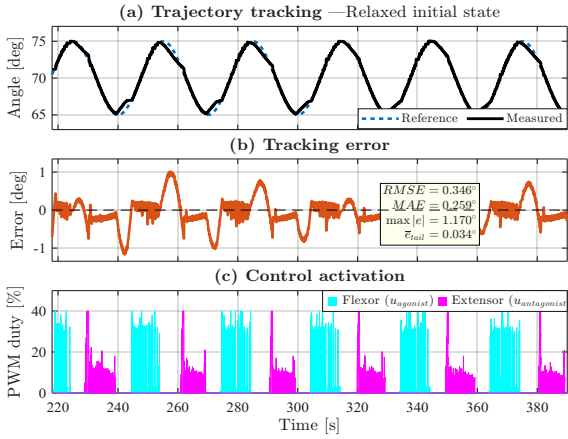


Figure 5: **Relaxed initial state configuration.** Increased passive compliance promotes recurrent Passive Free Cooling intervals ($u = 0$), reducing antagonistic effort density and delaying thermal accumulation.

4.2. Cyclic Consistency and Inter-Cycle Dispersion

Repeatability was evaluated over $N = 9$ consecutive cycles ($t = 360\text{--}660$ s, period $T = 30$ s) as shown in Figure 6. Each cycle was normalized to a 0–100% phase axis and stacked to quantify dispersion. The inter-cycle standard deviation profile yielded:

$$\bar{\sigma} = 0.068^\circ, \quad \sigma_{max} = 0.202^\circ.$$

Normalized with respect to the reference range, the mean dispersion corresponds to:

$$\bar{\sigma}_{norm} = 0.684\%, \quad \sigma_{max,norm} = 2.020\%.$$

The worst-case inter-cycle peak-to-peak envelope was:

$$\Delta_{pk,max} = 0.637^\circ \quad (6.37\% \text{ of reference range}).$$

Tracking metrics across all stacked cycles were:

$$RMSE_{all} = 0.169^\circ, \quad MAE_{all} = 0.147^\circ.$$

The low dispersion confirms that the BPID combined with asymmetric supervisory logic effectively mitigates thermo-mechanical hysteresis during sustained operation. Immediately after the ninth cycle, increased actuation effort was required for relaxation, indicating residual heat accumulation. This observation motivated the structured batch recovery strategy implemented in the endurance protocol.




Figure 6: Cyclic repeatability analysis ($N = 9$, $T = 30$ s). The mean response (solid) tracks the reference (dashed). Shaded region represents $\pm 1\sigma$ dispersion.

4.3. Thermal Endurance and Long-Duration Stability

A 40-minute endurance experiment evaluated long-duration robustness under thermal constraints. Cyclic tracking at $f = 0.03$ Hz was executed in $N = 6$ active batches, each composed of multiple consecutive cycles and separated by recovery intervals on the order of 53–163 s.

During recovery, the controller remained closed-loop while the reference was aligned with the measured angle ($r(t) \approx y(t)$), allowing thermal relaxation without destabilizing the system. The overall temporal structure of the experiment, including active batches and recovery intervals, is shown in Figure 7.

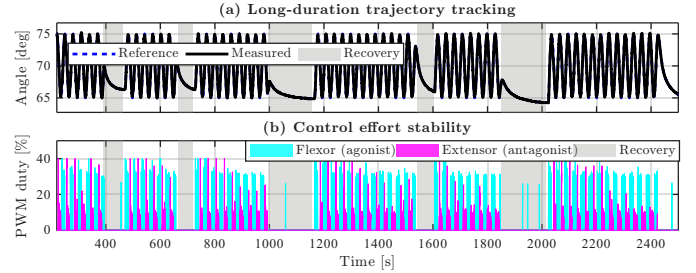


Figure 7: **40-minute endurance experiment.** (a) Angle tracking across six active batches separated by recovery intervals (shaded regions). (b) PWM duty cycles. The absence of persistent saturation confirms effective thermal recovery.

Performance results are summarized in Table 2, which also reports the number of cycles executed in each batch.

Across all active data, tracking accuracy remained within sub-degree range ($RMSE = 0.218^\circ$, $MAE = 0.163^\circ$), with a worst-case instantaneous error of 1.384° . No persistent saturation was observed during the 40-minute session.

Table 2: Thermal endurance performance over active batches. Error metrics are reported in degrees ($^\circ$). Bias: mean signed tracking error within the batch.

Batch	Time[s]	Dur.[s]	Cycles	RMSE	MAE	MaxE	Bias
1	220–392	172	5	0.293	0.217	1.373	-0.062
2	462–665	204	5	0.220	0.165	0.870	-0.009
3	719–1001	282	8	0.188	0.152	0.741	-0.029
4	1156–1544	387	12	0.232	0.170	1.384	-0.017
5	1597–1850	253	7	0.167	0.141	0.651	-0.049
6	2013–2448	435	12	0.212	0.155	1.348	-0.020
Global	–	–	–	0.218	0.163	1.384	–

4.3.1. Batch-to-Batch Trend and Safety Margin

Because the protocol includes recovery intervals between active segments, time variation is evaluated at the batch level. Linear regression of batch RMSE and mean signed error (Bias) over time yielded:

$$RMSE = -0.120^\circ/h, \quad Bias = +0.022^\circ/h.$$

The negative RMSE trend indicates slight improvement consistent with mechanical bedding-in, while the positive bias remains negligible relative to the control deadzone. With a conservative safety bound $\epsilon_{safe} \approx 5^\circ$ and $\max(|e(t)|) = 1.384^\circ$, the resulting safety margin is $M_{safety} = 3.62^\circ$, well within the safety limit. No emergency stop or thermal runaway was observed.

4.4. Thermal Endurance in Antagonistic SMA Systems

SMA actuators are increasingly employed in wearable rehabilitation systems due to their compactness, compliance, and high work density (Hamid et al., 2023). Nevertheless, cyclic operation is fundamentally constrained by cooling dynamics: under passive convection, recovery follows a nonlinear profile, leading to peak delay, displacement attenuation, and progressive heat accumulation as the actuation frequency increases.

Most reported approaches mitigate this limitation through hardware-oriented thermal management. Forced convection can extend actuation bandwidth, but it increases system complexity and power consumption (Affandi and Wahid, 2026). Similarly, temperature thresholds are frequently enforced to prevent overheating, interrupting operation once critical limits are reached. In rehabilitation prototypes operating at 0.03 Hz, thermal cut-offs (e.g., 281°C) are explicitly used to guarantee safe actuation (Affandi and Wahid, 2026).

At the opposite end of the dynamic spectrum, antagonistic architectures using ultra-thin SMA bundles have demonstrated actuation frequencies up to 200 Hz by combining aggressive current pulses with active airflow assistance (Barakat et al., 2025). While dynamically remarkable, this performance depends on enhanced cooling conditions and elevated power input to avoid thermal saturation.

Table 3 summarizes representative approaches across architectures and thermal strategies. The comparison indicates that endurance is typically achieved via thermal interruption, external cooling augmentation, or high-power dynamic operation, rather than being embedded within the control structure.

Table 3: Comparison of antagonistic SMA joint actuation studies ^a: (Ballester et al., 2023), ^b: (Copaci et al., 2019b), ^c: (Mansilla Navarro et al., 2024), ^d: (Affandi and Wahid, 2026), ^e: (Barakat et al., 2025). Antag. + sup.: antagonistic architecture with supervisory control; Arch.: control architecture.

Work	Arch.	Thermal Strategy	Interruption	Freq. (Hz)
Bench ^a	Single	Passive + smoothing	Heat accumulation	0.03
Elbow ^b	Antag.	48°C threshold	Forced stop	0.06
Ankle ^c	Antag.	Free cooling	1–2° peak error	0.08
Facial Rehab. ^d	Single	Forced air + 281°C cut-off	Thermal shutdown	0.03
High-Freq. ^e	Antag.	Active airflow + high current	Cooling assisted	200
This work	Antag. + sup.	Event-driven free cooling	No interruption	0.03

In contrast, the present work embeds endurance at the control-architecture level through indirect thermal-aware supervision. The supervisory law suppresses antagonistic actuation when passive mechanical recovery is already reducing the tracking error, creating intra-cycle zero-effort intervals (Passive Free Cooling). Batch-level reference alignment further preserves closed-loop stability during extended recovery phases without interrupting operation.

Experiments demonstrate stable cyclic tracking for 40 minutes at 0.03 Hz under free convection, without thermal shutdown, emergency stops, or observable performance drift. Although the operating frequency is intentionally low, the results support the view that endurance in antagonistic SMA systems can be extended through indirect thermal-aware supervision rather than through hardware cooling augmentation alone. **Overall, the results support the hypothesis that thermal robustness in SMA-driven joints is not solely a material limitation, but also a control-architecture problem.**

Open Data and Code. All experimental datasets and MATLAB scripts required to reproduce the reported results and figures are publicly available at <https://github.com/SJacoboZavaleta/RBVM2026>.

Acknowledgments

This research is part of the project *Robótica blanda para la rehabilitación del tobillo (SRAR)*, I+D+i grant PID2023-149141OB-I00, funded by MCIN/AEI/10.13039/501100011033 and by FEDER/UE.

References

- Affandi, R. I. R., Wahid, A. N., 2026. Mechanical Performance Evaluation of an Enhanced-Cooling SMA Actuator for Multidirectional Assistive Facial Rehabilitation Device. *IEEE/ASME Transactions on Mechatronics*, 1–9. DOI: 10.1109/TMECH.2026.3656626
- Ballester, C., Copaci, D., Arias, J., Moreno, L., Blanco, D., Apr. 2023. Hoist-Based Shape Memory Alloy Actuator with Multiple Wires for High-Displacement Applications. *Actuators* 12 (4), 159. DOI: 10.3390/act12040159
- Barakat, R., Kirsch, S.-M., Welsch, F., Motzki, P., Mar. 2025. Bidirectional rotational antagonistic shape memory alloy actuators for high-frequency artificial muscles. *Scientific Reports* 15 (1), 9108. DOI: 10.1038/s41598-025-93209-9
- Caballero, A. F., Copaci, D. S., Peciña, Á. V., Rojas, D. B., Lorente, L. M., Jul. 2016. Sistema Avanzado de Protipado Rápido para Control en la Educación en Ingeniería para grupos Multidisciplinares. *Revista Iberoamericana de Automática e Informática Industrial RIAI* 13 (3), 350–362. DOI: 10.1016/j.riai.2016.05.004
- Copaci, D., Blanco, D., Moreno, L. E., Aug. 2019a. Flexible Shape-Memory Alloy-Based Actuator: Mechanical Design Optimization According to Application. *Actuators* 8 (3), 63. DOI: 10.3390/act8030063
- Copaci, D., Martin, F., Moreno, L., Blanco, D., 2019b. SMA Based Elbow Exoskeleton for Rehabilitation Therapy and Patient Evaluation. *IEEE Access* 7, 31473–31484. DOI: 10.1109/ACCESS.2019.2902939
- Hamid, Q., Wan Hasan, W., Azmah Hanim, M., Nuraini, A., Hamidon, M., Ramlı, H., Jun. 2023. Shape memory alloys actuated upper limb devices: A review. *Sensors and Actuators Reports* 5, 100160. DOI: 10.1016/j.snr.2023.100160
- Langbein, S., Czechowicz, A., 2012. Adaptive resetting of SMA actuators 23 (2), 127–134. DOI: 10.1177/1045389X11431741
- Laschi, C., Mazzolai, B., Cianchetti, M., 2016. Soft robotics: Technologies and systems pushing the boundaries of robot abilities 1 (1). DOI: 10.1126/scirobotics.aah3690
- Mansilla Navarro, P., Copaci, D., Arias, J., Blanco Rojas, D., 2024. Design of an SMA-Based Actuator for Replicating Normal Gait Patterns in Pediatric Patients with Cerebral Palsy 9 (7), 376. DOI: 10.3390/biomimetics9070376
- Mohd Jani, J., Leary, M., Subic, A., Gibson, M. A., 2014. A review of shape memory alloy research, applications and opportunities 56, 1078–1113. DOI: 10.1016/j.matdes.2013.11.084
- Rao, A., Srinivasa, A. R., Reddy, J. N., 2015. Design of Shape Memory Alloy (SMA) Actuators. *SpringerBriefs in Applied Sciences and Technology*. Springer International Publishing. DOI: 10.1007/978-3-319-03188-0
- Reynaerts, D., Brussel, H. V., 1998. Design aspects of shape memory actuators 8 (6), 635–656. DOI: 10.1016/S0957-4158(98)00023-3
- Serrano, D., Copaci, D., Arias, J., Moreno, L. E., Blanco, D., Sep. 2023. SMA-Based Soft Exo-Glove. *IEEE Robotics and Automation Letters* 8 (9), 5448–5455. DOI: 10.1109/LRA.2023.3295994
- Villoslada, A., Flores, A., Copaci, D., Blanco, D., Moreno, L., Nov. 2015. High-displacement flexible Shape Memory Alloy actuator for soft wearable robots. *Robotics and Autonomous Systems* 73, 91–101. DOI: 10.1016/j.robot.2014.09.026