

TMR-VLA: Vision-Language-Action Model for Magnetic Motion Control of Tri-leg Silicone-based Soft Robot

Ruijie Tang*, Chi Kit Ng*, Kaixuan Wu, Long Bai, Guankun Wang,
Yiming Huang, Yupeng Wang, and Hongliang Ren†

Abstract—In-vivo environments, magnetically actuated soft robots offer advantages such as wireless operation and precise control, showing promising potential for painless detection and therapeutic procedures. We developed a trileg magnetically driven soft robot (TMR) whose multi-legged design enables more flexible gaits and diverse motion patterns. For the silicone made of reconfigurable soft robots, its navigation ability can be separated into sequential motions, namely squatting, rotation, lifting a leg, walking and so on. Its motion and behavior depend on its bending shapes. To bridge motion type description and specific low-level voltage control, we introduced TMR-VLA, an end-to-end multi-modal system for a trileg magnetic soft robot capable of performing hybrid motion types, which is promising for developing a navigation ability by adapting its shape to language-constrained motion types. The TMR-VLA deploys embodied endoluminal localization ability from EndoVLA, and fuses sequential frames and natural language commands as input. Low-level voltage output is generated based on the current observation state and specific motion type description. The result shows the TMR-VLA can predict how the voltage applied to TMR will change the dynamics of a silicon-made soft robot. The TMR-VLA reached a 74% average success rate.

I. INTRODUCTION

Miniature magnetic-actuated robots, with their wireless operation, are well-suited for biomedical applications, particularly in in vivo diagnosis and therapy [1]. They can reach deep inside the human body through narrow and complex pathways to perform highly precise clinical tasks. These include clearing cardiovascular blockages, delivering targeted drugs to specific lesions, and enabling minimally invasive diagnostics [2]. A key advantage is the compliance of soft materials, which significantly reduces the risk of collateral damage to surrounding soft tissues during internal operations. Among various actuation methods, actuation driven by an external magnetic field is widely recognized as an effective approach. This contactless mechanism enables flexible motion and precise attitude control without physical tethers—critical for navigating tortuous environments such as vascular networks. Through time-varying magnetic fields generated by external coils, forces and torques can be applied to magnetic particles embedded in the robot,

This work was supported in part by NSFC Young Scientists Fund — Category A T252500134, Ministry of Science and Technology (MOST) of China Key Project 2025YFE0122500, Hong Kong Research Grants Council (RGC) General Research Fund (GRF 14216022, 14206125), NSFC/RGC Joint Research Scheme (N_CUHK420/22); Guangdong Basic and Applied Basic Research Foundation (GBABF) #2021B1515120035 and Hainan Province Clinical Medical Center.

* These authors contribute equally to this work. All the authors are with the Chinese University of Hong Kong.

† Corresponding author: h1ren@ee.cuhk.edu.hk

achieving complex deformations and locomotion. This tether-free feature is key to unlocking the full potential of soft robots in biomedical applications.

Despite promising prospects, the development of micro-scale soft robots faces a bottleneck: at micron dimensions, the robot body is too small to integrate traditional onboard power, sensors, communication modules, or control [3]. Integrating such components increases the robot’s size, contradicting the goal of miniaturization. This limitation forces a separation between actuation (external magnetic fields) and perception (external observation devices). As a result, a critical gap emerges: the driver (magnetic field generator) and the robot (actuator) become physically decoupled. Unlike traditional electromechanical systems—such as motors, which use encoders for direct state feedback—external magnetic coils can only monitor their own current status. They cannot directly perceive the precise posture, deformation, or environmental interactions of the distant micro-robot.

With state information, controlling these systems is challenging. Soft robots are hyper-redundant, with infinite degrees of freedom. Their continuously deformable bodies exhibit complex nonlinear dynamics, defying accurate analytical modeling [4]. This modeling difficulty leads to a common reality: complex motions often depend on real-time human expert guidance. Experts manually adjust external magnetic fields using visual feedback, relying on experience and intuition. They become integral to the control loop [5]. A new approach is needed to add intelligence without increasing structural complexity. This would resolve the hardware limitations, eliminate reliance on human experts, and enable truly autonomous control. Therefore, we use the Vision-Language-action (VLA) framework for magnetically controlled soft robots.

The contributions of this paper are summarized as follows:

- **VLA for magnetic soft robotics.** We present the first end-to-end VLA framework for magnetic soft robots, achieving hybrid primitive control using only external vision feedback.
- **TrilegMR-Motion dataset.** We create a *TrilegMR-Motion dataset*, a vision–language–voltage dataset collected with a TMR to support learning direct voltage control from images and instructions.
- **Superior understanding and execution.** Compared with advanced general-purpose multimodal large language models, *TMR-VLA* attains consistently higher performance in both interpreting motion instructions and successfully executing decomposed action primitives,

providing a baseline for VLA in soft magnetic robots.

II. RELATED WORK

Magnetic Configurable Miniature Robots. Magnetically actuated soft miniature robots leverage programmed magnetization and global fields to realize rich, untethered gaits and task primitives [6]. Origami-inspired configurable designs show that field sequencing enables contraction, crawling, and steering in highly confined spaces. Programmable magnetic elastomers with multi-leg support functional integration (e.g., drug delivery, cargo transport) and robust state switching under field control [7]. Since global magnetic fields actuate all magnetized domains simultaneously—rendering these systems intrinsically underactuated—while embedding sensors and control electronics without compromising compliance or footprint remains difficult and magnetic feedback imposes fundamental limits on state estimation, the modeling and control of continuum magnetic soft robots confront underactuation, sparse onboard sensing, and stringent safety envelopes for medical use [3]. These challenges motivate us to enhance autonomy stacks that map task macros (e.g., squat/stand, leg-lift, heading correction, forward crawl) to time-varying field vectors/gradients with perceptual feedback—an approach we adopt for our silicon-based trileg robotic system that couples field actuation with voltage-induced morphology to realize sequential motion primitives.

Magnetically Actuated Robot Control. The autonomy of magnetic-driven robots is predominantly concentrated in the magnetic particles or capsules, focusing on positioning and navigation [8]. The application of reinforcement learning in eight-axis nonlinear magnetic drive systems achieves superior accuracy in controlling magnetic particle trajectories compared to conventional PID methods [9]. Magnetic continuum robots can leverage existing learning-based control methods developed for continuum robotics, provided datasets are generated through dual-arm magnetic actuation [10]. However, for mobile robots with multi-modal locomotion (e.g., jumping [11], crawling [12], rolling [13], and shape-shifting [14]), no intelligent solutions have been proposed to integrate these capabilities into task-specific frameworks systematically.

Vision-Language-Action Model. Vision-Language-Action (VLA) models represent a specialized category of multimodal architectures designed for embodied AI systems, integrating pre-trained vision-language foundations to jointly process visual observations [15], linguistic instructions [16], and action primitives [17]. While initial implementations [18], [19] demonstrate the feasibility of embedding Vision-Language Models (VLMs) into end-to-end visuomotor policies. Subsequent advancements [20]–[22] mitigate these issues through single-robot or simulated environment training paradigms, though such solutions sacrifice cross-platform generalizability and establish closed-loop systems resistant to efficient adaptation for novel robotic embodiments [23]–[25]. Parallel developments in visual imitation learning [26], [27] demonstrate alternative strategies by translating human behavioral videos into

executable robotic programs. Within clinical domains, emerging research efforts have focused on adapting VLA frameworks to medical applications. In contrast to prior end-to-end low-level control, we operationalize autonomy through a *library of instruction-constrained motion types* executed by a VLA policy. Each primitive (e.g., Squat, Leg-lift, Rotation, Forward) is parameterized by promptable task constraints to soft configurable robot [28], [29].




Fig. 1. The overall workflow of the VLA model for magnetic motion control of the trileg soft robot

III. TRI-LEG MAGNETIC SOFT ROBOT

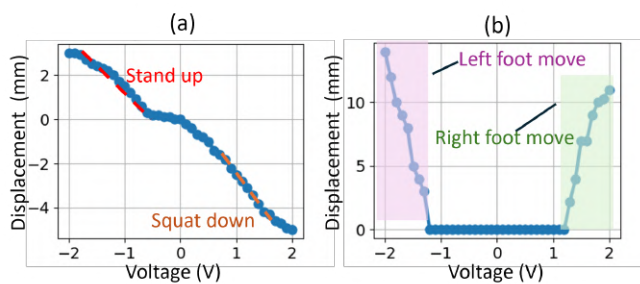


Fig. 2. Motion characterization of the Tri-leg Magnetic Soft Robot: (a) Squatting height (z -axis) vs. voltage (b) Stepping distance (alternating gait) vs. voltage

In [30], a three-legged robot has been designed, capable of multiple motion modes such as rolling, crawling, squatting, and leg lifting. In this paper, we preliminarily discretize these motion functions to enhance generalizability for VLA training. This approach avoids relying on hard-coded API calls, which could cause the robot to deviate significantly from intended paths in unfamiliar scenarios. In principle,

the robot operates within a torque generated by uniform magnetic field components along the x, y, and z axes, which counteract other possible external forces. Taking one leg as an example,

$$\tau_r = \mathbf{m} \times \mathbf{B}_r = \mathbf{m} \times |B| \frac{\vec{r}}{|r|}, \quad (1)$$

where the τ_r is the torque to move the leg, \mathbf{m} is the magnetic moment, \mathbf{B}_r is the external magnetic field to move the leg. In [30], to move forward in the x direction, the alternating torque ($\pm \tau_y$) application between the legs is achieved via a sinusoidal signal ($B_{(y)} = A_{(y)} \sin(2\pi ft)$). Similarly, movement in any direction requires activating this sinusoidal signal, whose amplitude and frequency are pre-programmed. To enhance generalization capability, these parameters must be discretized into the form shown in Fig. 2. This discretization enables multimodal large language models to interpret visual scenarios and precisely execute actions across diverse situations.

TMR motion characteristic. The state of a soft robot is influenced by the objects it interacts with, and its own structure also determines its steady state to some extent. Therefore, in Fig 2, we can observe that the squatting distance or lifting height does not increase linearly with the voltage. Only a specific segment shows a clear trend. For example, from -2 V to -0.8 V, the robot’s lifting height increases from 0 mm to 3 mm. Symmetrically, the squatting distance shows a more pronounced increase from 0 mm to 5 mm. The intermediate states, however, are not as evident. This occurs because the system surpasses a critical point in shape deformation, entering a range where bending becomes easier. For locomotion, the characterization is more complex. Fig 2(b) shows measurements of two-step movements in the x+ direction. It can be observed that the robot exhibits no significant movement between -1.2 V and 1.2 V, as the driving force is below the static friction of the ground. Once the critical point is surpassed, the displacement rapidly increases to its maximum—15 mm on the left and 12 mm on the right. However, if the friction between the ground and the feet increases, the critical points described in the aforementioned curves may no longer hold. Instead, the previous action—or a sequence of past actions combined with the current state—can serve as valuable information for making judgments. The robot can leverage the pattern of these curves to interpret its own current motion, rather than relying solely on rigid thresholds measured by human experts.

Hardware setup. The TMR’s workspace is confined to the central region of a 3D Helmholtz coil system, with dimensions of $50\text{mm} \times 50\text{mm} \times 30\text{mm}$ ($L \times W \times H$). The host computer is the terminal for receiving visual information and transmitting magnetic field signal commands. The signal generator is used to collect voltage commands and convert them into magnetic field strength.

Data Collection. The commercial endoscope enters from the xz-side of the coil. It is fixed on a stand, providing a stable view. The endoscope system delivers a real-time

video stream at 60 FPS. The field of view resolution can be set to either 640×480 or 1280×720 . In this experiment, all data were fixed at 640×480 . Due to the lack of multimodal data for soft robots without magnetic control, we created the TMRVLA-Dataset. It contains 60 episodes, totaling 15793 image-action pairs. Based on movement tasks, it is divided into three categories:

- Moving from marker A to marker B on a generic white grid background;
- Moving from marker A to a white lesion in a medical intestinal endoscopy setting;
- Moving from the current position to a yellow lesion in the same medical setting.

Data was sampled at 10 Hz, pairing endoscope images with real-time magnetic voltage values. During post-processing, redundant static voltage values and corresponding images were removed to refine the dataset. An expert manually controlled the voltage along the x, y, and z axes, using motions such as squatting, leg lifting, turning, and advancing to complete the tasks.

IV. METHODOLOGY

A. Overview

We formulate TMR control as end-to-end *vision–language-to-voltage* prediction. At each control step t , the policy receives (i) a short history of overhead images $\mathcal{I}_{t-L:t} = \{I_{t-L}, \dots, I_t\}$ and (ii) a natural-language *motion-type instruction* I_m (e.g., SQUAT, STAND_UP, LIFT_BACK_LEG, ROTATE_LEFT/RIGHT, FORWARD). The policy outputs coil voltages $\Delta \mathbf{V}_t = [\Delta V_x, \Delta V_y, \Delta V_z]^T$ that generate three orthogonal magnetic fields and induce morphology changes realizing the requested primitive. Complex tasks are executed by chaining these primitives. Our controller fuses sequential vision and language to produce low-level voltages tailored to the current morphology and the intended motion type constrained by the prompt.

B. Problem Formulation

We cast control as a constrained partially observable MDP with observations

$$\mathbf{o}_t = (\mathcal{I}_{t-L:t}, I_m, \mathcal{O}_t^{\text{aux}}), \quad (2)$$

where $\mathcal{O}_t^{\text{aux}}$ denotes optional auxiliary cues (e.g., voltage state) available at training/inference time. The policy π_θ produces an *action token sequence* \mathbf{a}_t that represents a voltage increment $\Delta \mathbf{V}_t$; the applied voltage is

$$\mathbf{V}_t = \Pi_C(\mathbf{V}_{t-1} + \Delta \mathbf{V}_t), \quad (3)$$

$$\mathcal{C} = \left\{ \mathbf{V} : \|\mathbf{V}\|_\infty \leq V_{\max}, \|\Delta \mathbf{V}\|_\infty \leq \dot{V}_{\max} \right\}.$$

where Π_C projects onto hardware safety limits.




Fig. 3. Inference framework of the tri-leg magnetic robot VLA (TMR-VLA). The model consumes a short frame window and an instruction, then autoregressively emits quantized voltage increments that are dequantized and safety-projected before actuation.

Actuation model.: Voltages map to magnetic fields through a calibrated coil matrix $\mathbf{K} \in \mathbb{R}^{3 \times 3}$:

$$\mathbf{B}_t = \mathbf{K} \mathbf{V}_t, \quad \boldsymbol{\tau}_t = \mathbf{m} \times \mathbf{B}_t, \quad (4)$$

with effective magnetization \mathbf{m} of the robot. The induced torque drives leg lifting, centroid shifting, heading correction, and forward crawling. We operate in a quasi-static regime (small Δt) where viscous and inertial effects are dominated by field-induced shape changes, and we treat the morphology state as *implicitly observed* via $\mathcal{I}_{t-L:t}$.

C. EndoVLA-Initialized Policy and Fine-Tuning

We initialize the multimodal backbone from **EndoVLA**—an embodied VLA policy for endoluminal tasks [28]—preserving its vision encoder, text encoder, and cross-modal fusion blocks. We added a *voltage regressor adaptor* specialized to our coil-driven platform:

$$(\mathcal{I}_{t-L:t}, I_m, \mathcal{O}_t^{\text{aux}}) \mapsto \Delta \mathbf{V}_t, \quad (5)$$

and fine-tune by supervised behavioral cloning on 360 paired demonstrations, minimizing a voltage regression objective with regularization (Sec. IV-F) [31]. This retains EndoVLA’s strong vision–language grounding while adapting the action mapping to magnetic soft-robot actuation.

D. Inputs and Action Space

Vision: A sliding window of L top-down RGB frames captures leg poses, posture, and contact state within the coil workspace. The overhead camera substitutes for onboard sensing, consistent with miniature robots’ payload limits.

Language: A concise instruction encodes the primitive and optional constraints (e.g., “rotate left towards the marker within 5° ”). Language constrains the otherwise high-dimensional soft-body behavior to a task-relevant subspace.

Actions: The controller outputs increments $\Delta \mathbf{V}_t = \mathbf{V}_t - \mathbf{V}_{t-1}$, bounded by hardware limits (Eq. (3)). Fields and torques follow Eq. (4).

E. Architecture

Visual encoder and temporal aggregation.: We use the native Qwen2.5-VL vision encoder—a dynamic-resolution ViT with *Window Attention* [32]. A short frame window is provided as a video clip to capture phase and contact transitions.

Cross-modal fusion.: Instructions are tokenized by the Qwen2.5-VL processor and interleaved with vision tokens using the model’s special delimiters (e.g., $\langle | \text{vision_start} | \rangle \dots \langle | \text{vision_end} | \rangle$). The concatenated stream is fed to the unmodified Qwen2.5-VL backbone to produce a fused representation \mathbf{o}_t .

c) Action Adaptor: Instead of a direct MLP regressor, we employ an *action adaptor* that lets the LLM *autoregressively* emit low-level voltage commands. Given \mathbf{o}_t and I_m , the adaptor outputs a short token sequence encoding $\Delta \mathbf{V}_t$.

We serialize the action as

$$\mathbf{a}_t = \langle \text{<SOS>, } \Delta \mathbf{V}_t^x, \Delta \mathbf{V}_t^y, \Delta \mathbf{V}_t^z, \text{<EOS>}, \quad (6)$$

The TMR-VLA p_θ autoregressively decoding:

$$p_\theta(\mathbf{a}_t | \mathbf{o}_t) = \prod_k p_\theta(a_{t,k} | \mathbf{o}_t, a_{t,<k}). \quad (7)$$

F. Training Objectives

Given demonstrations $\mathcal{D} = \{(\mathcal{I}_{t-L:t}, I_m, \mathcal{O}_t^{\text{aux}}, \mathbf{V}_t)\}$, we supervise both the *token* sequence and the *numeric* voltages. Let \mathbf{a}_t^* be the oracle action sentence. The loss is

$$\mathcal{L}_{\text{CE}} = - \sum_k \log p_\theta(a_{t,k}^* | \mathbf{o}_t, a_{t,<k}^*). \quad (8)$$

Calibration and domain randomization.: To improve transfer across coil rigs and minor morphology changes, we *abstract* robot constants (e.g., \mathbf{K}) into the post-processor and apply mild randomization to \mathbf{K} , and image appearance during training.

V. EXPERIMENTS

A. Tasks and Metrics

We evaluate five motion primitives that compose our navigation skill set: **Squat**, **Lift-leg**, **Rotation** (left/right), **Forward**, and **Recovery**. Let $p_t = (x_t, y_t)$ be the 2D centroid, ψ_t the heading, and h_ℓ a per-leg height proxy.

a) Success criteria: A trial succeeds if the task-specific criterion below is met within a budget T_{max} and sustained for T_s frames:

- **Squat**: vertical compression $\Delta h \leq -h_{\text{min}}$ (or projected area increase $\geq A_{\text{min}}$).
- **Leg-lift**: designated leg height $h_\ell \geq h_{\text{lift}}$ while the other two legs remain in contact.
- **Rotation**: $|\psi_T - \psi_0 - \psi^*| \leq \epsilon_\psi$ and centroid drift $< d_r$.




Fig. 4. (a) The illustration of the trileg robot. The magnetic field provided the torques in 3 dimensions. Experimental parameters description: (b) Maximum squat distance, (c) Leg lift height, (d) Anchor one foot, rotate body, (e) Forward distance (f) Reach target and recover.

- **Forward:** displacement along the commanded axis $\Delta x^{\parallel} \geq d_{\text{fwd}}$ with lateral error $|\Delta x^{\perp}| \leq d_{\perp}$ and heading error $\leq \epsilon_{\psi}$.
- **Recovery** designated leg returns to contact, i.e., $h_{\ell} \leq h_{\text{drop}}$ with posture stability (centroid drift $< d_s$) and no unwanted heading change ($|\Delta\psi| \leq \epsilon_{\psi}$).

A trial succeeds if the task-specific criterion (height/area change, heading change, or axial displacement with bounded lateral error) is met within a budget T_{max} and sustained for T_s frames. We report per-task *success rate* and the *macro-average* across tasks (Tab. II). A safety guard violation is strictly defined as exceeding the absolute voltage limit of 2.5V. Any voltage output beyond this threshold is considered a failure case. We also report the *per-axis MSE* of voltage prediction. MSE_x , MSE_y , MSE_z and the *overall MSE* for each motion type (Tab. I).

B. Hardware and Environment

Experiments are conducted in a three-axis coil rig producing orthogonal fields $\mathbf{B} = [B_x, B_y, B_z]^{\top}$ from coil voltages $\mathbf{V} = [V_x, V_y, V_z]^{\top}$ via a calibrated linear map $\mathbf{B} = \mathbf{K}\mathbf{V}$. The tri-leg silicone robot contains embedded magnetization; the interaction $\boldsymbol{\tau} = \mathbf{m} \times \mathbf{B}$ induces shape changes to realize the five primitives.

C. Models and Training

a) *Training and Inference Implementation:* The **TMR-VLA** model was fine-tuned based on EndoVLA. Training was conducted using multiple GPUs, including four NVIDIA RTX A6000 (48 GB), and took approximately three hours to complete. We use AdamW, cosine decay, 5% warmup. Parameter-efficient fine-tuning (LoRA/PEFT) fits a single high-memory GPU. Inference was performed on a single NVIDIA RTX 5090 (32 GB) GPU, achieving a speed of approximately 2 Hz.

b) *Baselines:* We compare against strong open models trained under the same recipe and head: Gemma-3 (vision variant) [33], LLaVA-1.6 [34], Qwen2.5-VL-3B/7B [32], and MiniCPM-V [35]. All baselines receive the same inputs (sequential frames + instruction) and are trained to regress voltages with identical losses, schedules, and augmentations. For backbones lacking native temporal fusion, we insert a small causal transformer over per-frame tokens.

D. Evaluation Procedure

First, we compared the action error of the baseline model and our three models against the ground truth on the test set. Then, following the previously defined success criteria, each action was executed 10 times under random initial states, and the average was taken as the success rate.

E. Prompt Design

Table II presents the results of (TMR-VLA (p_0), TMR-VLA (p_1), TMR-VLA (p_2)) with motion-type success rates. The system prompts for TMR-VLA (p_0) and TMR-VLA (p_1) are relatively concise in task description. Their common part is defined as follows: "You are a voltage controller. You can control the voltage to move the trileg robot as shown in the images. Current image is <Im0>. Past images are <Im1>, <Im2>. Current state is s0. Previous states are s1 and s2. What is the next action?" The key difference is that TMR-VLA (p_0) includes additional historical action information: "Previous actions are a1 and a2." TMR-VLA (p_2) extends TMR-VLA (p_1) (which excludes historical actions) by providing a detailed task description. Its prompt is as follows: "The TMR is white with three legs, each equipped with a magnet. You control the voltage to manipulate the magnetic field, which in turn moves the robot. The robot's state represents the current strength of the magnetic field (in x, y, z directions). An action involves adjusting the magnetic field in a specific direction. For the robot to walk, it must first crouch, then lift a leg, and subsequently alternate between legs to move. You can observe the state of the robot from the images and correlate it with the numerical values."

F. Per-Motion Voltage Prediction Error (MSE)

Table I details per-axis voltage MSE (x, y, z) and overall MSE for each motion type and baseline backbone. These results complement success rates by quantifying low-level control fidelity and highlighting axes/motions with higher actuation difficulty.

G. Baseline Comparison (Success Rate)

Table II reports motion-type success rates (%) and the macro-average across the five primitives for baselines and our model. We observe the relative ranking among public backbones under identical training and evaluation conditions. It should be noted that while MSE serves as a metric for evaluating action accuracy, we observed a distortion effect during experiments: when the actual action change is very small, even if the robot's output is zero, the resulting average error remains low. In contrast, the output generated by VLA

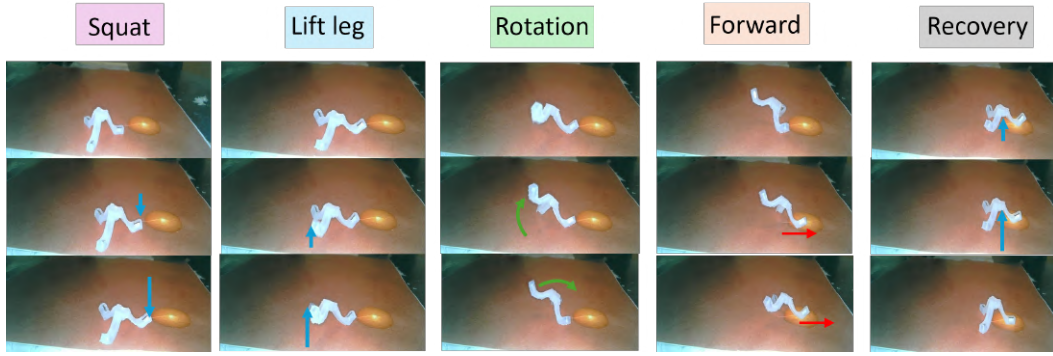


Fig. 5. Experimental results demonstrate that TMR-VLA achieves higher success rates in executing multi-step actions. The arrow indicates the magnitude and direction of the action output.

TABLE I

PER-MOTION MSE OF VOLTAGE PREDICTION (LOWER IS BETTER). COLUMNS SHOW x , y , z , AND OVERALL MSE FOR EACH MOTION TYPE.

| Model | Squat | | | | Lift Leg | | | | Rotation | | | | Forward | | | | Recovery | | | |
|-------------------|------------------|------------------|------------------|--------------|------------------|------------------|------------------|--------------|------------------|------------------|------------------|--------------|------------------|------------------|------------------|--------------|------------------|------------------|------------------|--------------|
| | MSE _x | MSE _y | MSE _z | Overall | MSE _x | MSE _y | MSE _z | Overall | MSE _x | MSE _y | MSE _z | Overall | MSE _x | MSE _y | MSE _z | Overall | MSE _x | MSE _y | MSE _z | Overall |
| Gemma-3 (V) | 0.150 | 0.000 | 0.040 | 0.070 | 0.015 | 0.005 | 0.010 | 0.010 | 0.080 | 0.403 | 0.000 | 0.158 | 0.100 | 0.100 | 0.000 | 0.066 | 0.110 | 0.000 | 0.050 | 0.053 |
| LLaVA-1.6 | 0.1 | 0.005 | 0.176 | 0.093 | 0.095 | 0.105 | 0.000 | 0.066 | 0.093 | 0.1 | 0.000 | 0.064 | 0.0381 | 0.108 | 0.000 | 0.049 | 0.050 | 0.085 | 0.050 | 0.061 |
| Qwen2.5-VL-3B | 0.004 | 0.028 | 0.128 | 0.054 | 0.095 | 0.005 | 0.000 | 0.033 | 0.078 | 0.021 | 0.000 | 0.033 | 0.001 | 0.100 | 0.000 | 0.033 | 0.050 | 0.050 | 0.000 | 0.033 |
| Qwen2.5-VL-7B | 0.009 | 0.000 | 0.100 | 0.036 | 0.010 | 0.005 | 0.005 | 0.006 | 0.1788 | 0.021 | 0.000 | 0.066 | 0.100 | 0.100 | 0.000 | 0.066 | 0.055 | 0.000 | 0.050 | 0.035 |
| MiniCPM-V | 0.057 | 0.100 | 0.143 | 0.100 | 0.795 | 0.105 | 0.010 | 0.303 | 0.078 | 0.303 | 0.000 | 0.127 | 0.368 | 0.313 | 0.000 | 0.227 | 0.150 | 0.190 | 0.052 | 0.134 |
| TMR-VLA (p_0) | 0.005 | 0.000 | 0.005 | 0.003 | 0.010 | 0.005 | 0.005 | 0.007 | 1.160 | 1.833 | 1.934 | 1.644 | 0.000 | 0.009 | 0.000 | 0.003 | 0.11 | 0.007 | 0.012 | 0.043 |
| TMR-VLA (p_1) | 0.005 | 0.000 | 0.005 | 0.003 | 0.879 | 0.000 | 1.158 | 0.678 | 0.230 | 0.209 | 0.121 | 0.187 | 0.008 | 0.076 | 0.000 | 0.028 | 0.157 | 0.01 | 0.095 | 0.087 |
| TMR-VLA (p_2) | 0.004 | 0.004 | 0.009 | 0.006 | 0.060 | 0.005 | 0.045 | 0.036 | 0.233 | 0.2 | 0.060 | 0.164 | 0.000 | 0.099 | 0.000 | 0.033 | 0.050 | 0.005 | 0.060 | 0.038 |

control may involve random trial-and-error. Therefore, it is essential to also consider the actual experimental success rate for a comprehensive evaluation.

TABLE II
MOTION TYPE SUCCESS RATE (%)

| Model | Squat | Lift Leg | Rotation | Forward | Recovery | Mean |
|-------------------|------------|-----------|------------|------------|-----------|-----------|
| Gemma-3 (V) | 0 | 0 | 0 | 10 | 0 | 2 |
| LLaVA-1.6 | 0 | 0 | 20 | 0 | 0 | 4 |
| Qwen2.5-VL-3B | 10 | 0 | 0 | 0 | 0 | 2 |
| Qwen2.5-VL-7B | 50 | 0 | 10 | 30 | 0 | 18 |
| MiniCPM-V | 0 | 10 | 0 | 0 | 0 | 2 |
| TMR-VLA (p_0) | 100 | 0 | 0 | 60 | 30 | 38 |
| TMR-VLA (p_1) | 100 | 20 | 80 | 100 | 70 | 74 |
| TMR-VLA (p_2) | 100 | 50 | 100 | 0 | 20 | 54 |

H. Robustness and Generalization

As shown in Table II, providing additional system descriptions does not improve the performance of the large language model in scenarios completely lacking common-sense understanding. Instead, it introduces a certain degree of interference. Comparing the results of TMR-VLA (p_0) and TMR-VLA (p_1), our experimental observations indicate that due to the repeated occurrence of certain directional voltage adjustments (e.g., $-0.1V$ in the $x+$ direction) in the dataset, TMR-VLA tends to imitate previous actions, while TMR-VLA (p_1) infers the next action based on the current state values. As a result, TMR-VLA (p_1) achieves higher success rates across varying states.

VI. DISCUSSION AND CONCLUSION

Typical failures include cases where the robot returns to its initial state after completing squatting and leg lifting, without proceeding to locomotion. Additionally, when human intervention forcibly alters the model's output action, two or more modifications can cause the model to persistently

follow the newly introduced action, resulting in failure. This indicates that in the absence of visual auxiliary labels such as bounding boxes or masks, the comprehension ability of large models deteriorates in non-universal scenarios. It is necessary to enable the model to genuinely understand why the robot should generate such deformations.

In conclusion, this work addresses a challenge in the development of miniature magnetically actuated soft robots: the hardware limitation that decouples external actuation from robot perception, creating a reliance on human experts for control. To bridge this gap and enable autonomous, intelligent operation without increasing structural complexity, we introduced TMR-VLA, the first end-to-end Vision-Language-Action (VLA) framework for magnetic soft robotics. Our system successfully translates high-level natural language commands and visual observations into low-level voltage control signals, predicting how applied voltages alter the state of a silicone-made tri-leg robot to execute diverse and flexible motions. Supported by the novel TrilegMR-Motion dataset, TMR-VLA demonstrated a superior ability to interpret instructions and execute hybrid motions compared to general-purpose models, achieving a 78% average success rate. This approach provides a foundational baseline for embedding intelligence into magnetic soft robotic systems and is a significant step toward their autonomous application in complex in-vivo environments.

REFERENCES

- [1] Z. Chen, S. Yang, J. Law, X. Du, and Y. Sun, "Magnetic soft robots in medicine: material and structural designs for organ-specific applications," *Soft Science*, vol. 6, no. 1, pp. N–A, 2026.
- [2] R. Dreyfus, Q. Boehler, S. Lytle, P. Gruber, J. Lussi, C. Chautems, S. Gervasoni, J. Berberat, D. Seibold, N. Ochsenbein-Kölblle, *et al.*, "Dexterous helical magnetic robot for improved endovascular access," *Science Robotics*, vol. 9, no. 87, p. eadh0298, 2024.
- [3] X. Li, X. Xiao, X. Xiao, Z. Liu, J. Gong, Z. Lin, B. Xue, S. Liu, X. Wu, W. Zhang, *et al.*, "Magnetic field-enhanced vertical integration enables embodied intelligence in untethered soft robots," *Science Advances*, vol. 11, no. 37, p. eadv9572, 2025.
- [4] D. A. Haggerty, M. J. Banks, E. Kamenar, A. B. Cao, P. C. Curtis, I. Mezić, and E. W. Hawkes, "Control of soft robots with inertial dynamics," *Science robotics*, vol. 8, no. 81, p. eadd6864, 2023.
- [5] Y. Kim, E. Genevriere, P. Harker, J. Choe, M. Balicki, R. W. Regenhardt, J. E. Vranic, A. A. Dmytriw, A. B. Patel, and X. Zhao, "Telerobotic neurovascular interventions with magnetic manipulation," *Science Robotics*, vol. 7, no. 65, p. eabg9907, 2022.
- [6] Q. Ze, S. Wu, J. Nishikawa, J. Dai, Y. Sun, S. Leanza, C. Zemelka, L. S. Novelino, G. H. Paulino, and R. R. Zhao, "Soft robotic origami crawler," *Science advances*, vol. 8, no. 13, p. eabm7834, 2022.
- [7] Y. Dong, L. Wang, N. Xia, Z. Yang, C. Zhang, C. Pan, D. Jin, J. Zhang, C. Majidi, and L. Zhang, "Untethered small-scale magnetic soft robot with programmable magnetization and integrated multifunctional modules," *Science advances*, vol. 8, no. 25, p. eabn8932, 2022.
- [8] Y. Liu, H. Wang, X. Wu, J. Qu, X. Liu, and Q. Fan, "Autonomous navigation of magnetic microrobots with improved planning and control in complex environments," *IEEE Transactions on Automation Science and Engineering*, 2024.
- [9] S. A. Abbasi, A. Ahmed, S. Noh, N. L. Gharamaleki, S. Kim, A. M. B. Chowdhury, J.-y. Kim, S. Pané, B. J. Nelson, and H. Choi, "Autonomous 3d positional control of a magnetic microrobot using reinforcement learning," *Nature Machine Intelligence*, vol. 6, no. 1, pp. 92–105, 2024.
- [10] M. Brockdorff, T. da Veiga, J. Davy, P. Lloyd, J. H. Chandler, G. Pittiglio, R. K. Mathew, and P. Valdastrì, "Hybrid trajectory planning of two permanent magnets for medical robotic applications," *The International Journal of Robotics Research*, p. 02783649241264844, 2024.
- [11] D. Tang, C. Zhang, C. Pan, H. Hu, H. Sun, H. Dai, J. Fu, C. Majidi, and P. Zhao, "Bistable soft jumper capable of fast response and high takeoff velocity," *Science Robotics*, vol. 9, no. 93, p. eadm8484, 2024.
- [12] Z. Ren and M. Sitti, "Design and build of small-scale magnetic soft-bodied robots with multimodal locomotion," *Nature protocols*, vol. 19, no. 2, pp. 441–486, 2024.
- [13] Z. Yang, C. Xu, J. X. Lee, and G. Z. Lum, "Magnetic miniature soft robot with reprogrammable drug-dispensing functionalities: Toward advanced targeted combination therapy," *Advanced Materials*, vol. 36, no. 48, p. 2408750, 2024.
- [14] L. Yang, M. Sun, M. Zhang, and L. Zhang, "Multimodal motion control of soft ferrofluid robot with environment and task adaptability," *IEEE/ASME transactions on mechatronics*, vol. 28, no. 6, pp. 3099–3109, 2023.
- [15] S. Nair, A. Rajeswaran, V. Kumar, C. Finn, and A. Gupta, "R3m: A universal visual representation for robot manipulation," *arXiv preprint arXiv:2203.12601*, 2022.
- [16] S. Karamcheti, S. Nair, A. S. Chen, T. Kollar, C. Finn, D. Sadigh, and P. Liang, "Language-driven representation learning for robotics," *arXiv preprint arXiv:2302.12766*, 2023.
- [17] D. Driess, F. Xia, M. S. Sajjadi, C. Lynch, A. Chowdhery, B. Ichter, A. Wahid, J. Tompson, Q. Vuong, T. Yu, *et al.*, "Palm-e: An embodied multimodal language model," *arXiv preprint arXiv:2303.03378*, 2023.
- [18] M. Shridhar, L. Manuelli, and D. Fox, "Cliport: What and where pathways for robotic manipulation," in *Conference on robot learning*, pp. 894–906, PMLR, 2022.
- [19] A. Stone, T. Xiao, Y. Lu, K. Gopalakrishnan, K.-H. Lee, Q. Vuong, P. Wohlhart, S. Kirmani, B. Zitkovich, F. Xia, *et al.*, "Open-world object manipulation using pre-trained vision-language models," *arXiv preprint arXiv:2303.00905*, 2023.
- [20] J. Huang, S. Yong, X. Ma, X. Linghu, P. Li, Y. Wang, Q. Li, S.-C. Zhu, B. Jia, and S. Huang, "An embodied generalist agent in 3d world," *arXiv preprint arXiv:2311.12871*, 2023.
- [21] X. Li, M. Liu, H. Zhang, C. Yu, J. Xu, H. Wu, C. Cheang, Y. Jing, W. Zhang, H. Liu, *et al.*, "Vision-language foundation models as effective robot imitators," *arXiv preprint arXiv:2311.01378*, 2023.
- [22] H. Zhen, X. Qiu, P. Chen, J. Yang, X. Yan, Y. Du, Y. Hong, and C. Gan, "3d-vla: A 3d vision-language-action generative world model," *arXiv preprint arXiv:2403.09631*, 2024.
- [23] N. Dorka, C. Huang, T. Welschhold, and W. Burgard, "What matters in employing vision language models for tokenizing actions in robot control?," in *First Workshop on Vision-Language Models for Navigation and Manipulation at ICRA 2024*, 2024.
- [24] J. Liu, M. Liu, Z. Wang, P. An, X. Li, K. Zhou, S. Yang, R. Zhang, Y. Guo, and S. Zhang, "Robomamba: Efficient vision-language-action model for robotic reasoning and manipulation," *Advances in Neural Information Processing Systems*, vol. 37, pp. 40085–40110, 2025.
- [25] P. Ding, J. Ma, X. Tong, B. Zou, X. Luo, Y. Fan, T. Wang, H. Lu, P. Mo, J. Liu, *et al.*, "Humanoid-vla: Towards universal humanoid control with visual integration," *arXiv preprint arXiv:2502.14795*, 2025.
- [26] Y. Wang, G. Gonzalez-Pumariega, Y. Sharma, and S. Choudhury, "Demo2code: From summarizing demonstrations to synthesizing code via extended chain-of-thought," *Advances in Neural Information Processing Systems*, vol. 36, pp. 14848–14956, 2023.
- [27] G. Chen, M. Wang, T. Cui, Y. Mu, H. Lu, T. Zhou, Z. Peng, M. Hu, H. Li, Y. Li, *et al.*, "Vlmimic: Vision language models are visual imitation learner for fine-grained actions," *arXiv preprint arXiv:2410.20927*, 2024.
- [28] N. C. KIT, L. Bai, G. Wang, Y. Wang, H. Gao, K. yuan, C. Jin, T. Zeng, and H. Ren, "EndoVLA: Dual-phase vision-language-action for precise autonomous tracking in endoscopy," in *9th Annual Conference on Robot Learning*, 2025.
- [29] S. Wang, J. Zhang, M. Li, J. Liu, A. Li, K. Wu, F. Zhong, J. Yu, Z. Zhang, and H. Wang, "Trackvla: Embodied visual tracking in the wild," *arXiv preprint arXiv:2505.23189*, 2025.
- [30] R. Tang and H. Ren, "A frequency-modulated tripod soft magnetic robot with diverse motion modalities for ingestible applications," *IEEE Robotics and Automation Letters*, vol. 9, no. 3, pp. 2168–2175, 2024.
- [31] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, W. Chen, *et al.*, "Lora: Low-rank adaptation of large language models," *ICLR*, vol. 1, no. 2, p. 3, 2022.
- [32] S. Bai, K. Chen, X. Liu, J. Wang, W. Ge, S. Song, K. Dang, P. Wang, S. Wang, J. Tang, *et al.*, "Qwen2. 5-vl technical report," *arXiv preprint arXiv:2502.13923*, 2025.
- [33] G. Team, A. Kamath, J. Ferret, S. Pathak, N. Vieillard, R. Merhej, S. Perrin, T. Matejovicova, A. Ramé, M. Rivière, *et al.*, "Gemma 3 technical report," *arXiv preprint arXiv:2503.19786*, 2025.
- [34] H. Liu, C. Li, Q. Wu, and Y. J. Lee, "Visual instruction tuning," *Advances in neural information processing systems*, vol. 36, pp. 34892–34916, 2023.
- [35] Y. Yao, T. Yu, A. Zhang, C. Wang, J. Cui, H. Zhu, T. Cai, H. Li, W. Zhao, Z. He, *et al.*, "Minicpm-v: A gpt-4v level mllm on your phone," *arXiv preprint arXiv:2408.01800*, 2024.