TOWARD CONVINCING BIOLOGY WITH LANGUAGE: OFFLINE LEARNING OF PROMPT-CONDITIONED INTERVENTIONS FOR CELLS, ORGANOIDS, AND BIOBOTS

Anonymous authors
Paper under double-blind review

ABSTRACT

Artificial intelligence increasingly serves as a natural-language interface to complex technical systems, letting people accomplish sophisticated tasks by describing what they want rather than specifying how to do it. Extending this interface to living systems is harder: unlike code or images, a biological intervention has no closed-form linguistic meaning, and the paired language-intervention-outcome data needed to learn such a mapping is expensive to collect, since each example requires its own wet-lab experiment. One way around this is to treat an existing archive of interventions and their already-observed outcomes as a fixed, offline dataset, and use a vision-language model to judge, without any new experiments, whether an archived outcome matches a natural-language description. But whether that judgment is reliable enough to train a language-to-intervention mapping on – without new experiments and without human validation – has remained unclear. Here we show that a natural-language interface for a living organism – a xenobot, a synthetic multicellular construct with no nervous system – can be learned entirely offline this way, using a vision-language model's own judgment as the sole training reward: an instruction is mapped to the intervention already on record as producing the described behavior. This mapping generalizes to entirely new instructions, evaluated against archive data withheld from training (80.0% held-out accuracy vs. a 66.7% chance baseline, matching a network trained directly on ground-truth labels). These results indicate that offline archives, combined with automated judgment, can substitute for costly new experiments when learning to ground natural language in a living system's own recorded behavior – an intermediate step toward directing living matter with the same ease as code, though whether these predicted interventions actually reproduce their intended behavior on live, newly-stimulated cells remains to be tested.

1 INTRODUCTION

Imagine a clinician or bioengineer simply describing what they want – calm this inflamed tissue, direct these cells to close a wound, steer this synthetic biobot toward a target – and having cells, organoids, or synthetic biobots (Kriegman et al., 2021; Blackiston et al., 2021a; Levin, 2021) carry out that intention directly, the way a modern software system already responds to a natural-language request. Controlling living systems this way is a long-standing goal across biology, medicine, and synthetic biology, yet the interface remains fundamentally low-level: practitioners specify interventions parameter by parameter, through trial and error informed by domain expertise. In other technical domains, natural language has emerged as a high-level interface to complex systems: users specify a goal – generate this code (Fried et al., 2023), design this molecule (Bengio et al., 2021), plan this trajectory (Huang et al., 2022) – and a learned model translates intent into low-level action. This raises a question central to applied machine learning (ML): can a similar interface be learned for biological systems, mapping a natural-language description of desired behavior to the physical intervention that elicits it?

Two properties make this hard. First, collecting paired (prompt, intervention, outcome) data online is prohibitively expensive: each datum is a wet-lab experiment, ruling out approaches that need a fresh experiment per prompt or per objective. Second, the objective space is inherently multi-task and

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

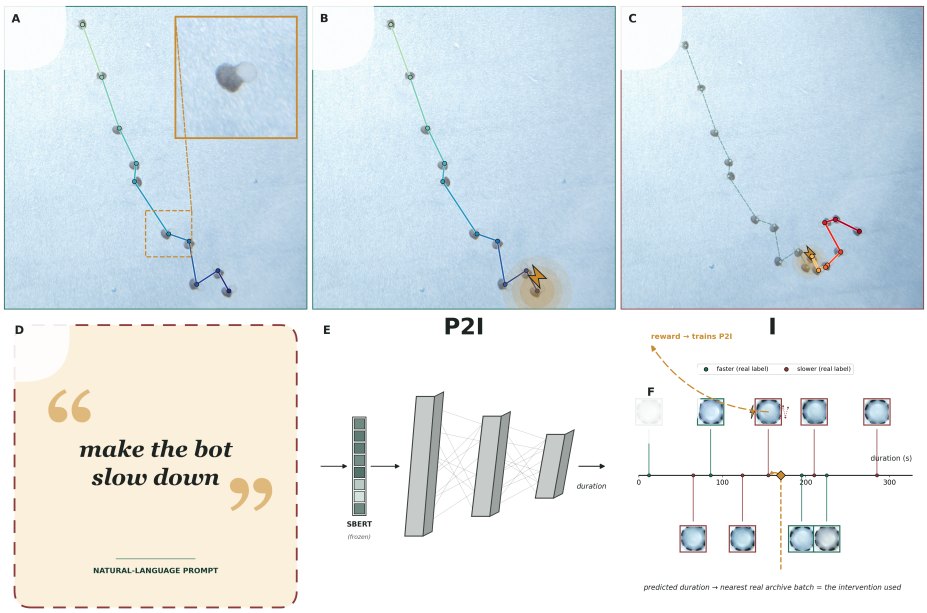

Figure 1: Overview of the pipeline. (a) Tracked trajectory of a xenobot before stimulation (PRE), with a zoomed view of the construct. (b) The same trajectory, marking the moment the electrical stimulus is applied. (c) The trajectory continuing after stimulation (POST), showing the resulting change in behavior. (d) An example natural-language prompt. (e) The Prompt-to-Intervention (P2I) network: a frozen SBERT sentence embedding is passed through a small MLP that outputs a single predicted duration. (f) The predicted duration is matched to its nearest neighbor on the archive's duration axis; the matched batch's own intervention and trajectory, as in (b)–(c), are what a vision-language model scores against the prompt, and that score is the reward used to train P2I.

contains opposing goals – e.g. "stop" versus "accelerate" – and a single policy trained across such goals is known to suffer objective interference, where the gradient for one goal actively degrades another (Yu et al., 2020; Chen et al., 2018), so naively pooling all objectives into one unconstrained shared model is not a safe default either. Together, these rule out the two most direct approaches: brute-forcing the mapping with enough paired data, and training one policy across all objectives without accounting for how they interact.

We address both properties by treating language-to-intervention as an offline learning problem. Archives of intervention-outcome pairs already exist, collected for purposes other than language-conditioned learning; we treat such an archive as a fixed dataset over which a policy is trained without any new experiments. To supply a reward signal in the absence of labeled prompt-outcome alignment, we use a pretrained VLM as a zero-shot scorer (Rocamonde et al., 2023), judging how well an archived outcome matches a natural-language description, and check that judgment against independent, non-VLM ground truth rather than trust it by assumption. The policy uses one output head per behavioral category over a shared backbone.

Contributions. This work makes three key contributions, each tested directly in Section 4 rather than only asserted: (1) we demonstrate an archive-based learning paradigm for language-to-intervention translation, using a vision-language model's own judgment as the reward in place of new experimental feedback; (2) we validate this approach end-to-end, generalizing to entirely new instructions on archive data withheld from training at 80.0% ground-truth accuracy, against a 66.7% chance baseline; (3) we show this VLM-derived reward is not redundant even with only two behaviors currently distinguished: replacing it with a single fixed duration chosen directly from real velocity data trains to a comparable loss but collapses held-out generalization to 59.1%, below chance, while the VLM-fit reward reaches 80.0%.

2 Related Work

Large language models serve as general-purpose interfaces to complex domains, generating functional code (Chen et al., 2021), aiding molecular design (Jumper et al., 2021), and assisting sci-

entific discovery (Zheng et al., 2025), typically letting users specify goals in language rather than low-level parameters. Biological control is harder: language has no closed-form semantics for what an electrical stimulus should mean, and paired (prompt, intervention, outcome) data is expensive to collect. Synthetic biology has produced xenobots, living *Xenopus laevis* cell assemblies exhibiting coordinated motion (Kriegman et al., 2021; Blackiston et al., 2021a), where targeted electrical or optogenetic interventions elicit reproducible behavioral change (Zajdel et al., 2020; Levin, 2021; Emiliani et al., 2022), but control remains manual, tuned parameter-by-parameter without a language interface.

Vision-language models can serve as zero-shot reward models (Rocamonde et al., 2023; Radford et al., 2021; Liang et al., 2023; Wang et al., 2024), though rarely validated against ground truth independent of the VLM itself – a gap we address directly. Learning across opposing behavioral goals is a multi-task problem, for which task-specific heads over a shared backbone are a standard remedy (Yu et al., 2020; Chen et al., 2018; Ruder, 2017; Caruana, 1997). Training from a fixed archive with only terminal rewards places this work in the offline, contextual-bandit regime (Levine et al., 2020; May et al., 2012), for which derivative-free search such as CMA-ES (Hansen & Ostermeier, 2001) is standard; we show the reward is in fact differentiable, so both backpropagation and gradient-free search apply, and use their agreement to validate the optimum found.

Two prior studies explore the underlying idea of prompt-to-intervention learning as a proof-of-concept in simulated, artificial-life cells; we position this work explicitly against the specific gaps each left open. Le et al. (2025) introduced a P2I network mapping a prompt embedding to a spatial vector field steering simulated 2D cell collectives, scored by a VLM against a fixed one-word vocabulary (e.g. "clustering" vs. "scattering"); the reward was treated as non-differentiable by design, optimized only with $(1{+}1)$-ES and a genetic algorithm, and the paper flags gradient interference between opposing objectives as an open problem at multi-prompt scale. Le et al. (2026) extends this, replacing the fixed vocabulary with a continuous, free-form VLM score in $[0, 1]$ – closer in spirit to the reward used here – but trained on a single prompt ("form a cluster"), again assumed non-differentiability, and lists gradient-trainable optimization and real-cell validation as explicit future work. This paper answers both directly rather than by assumption: the VLM-derived reward is fit as a smooth, differentiable function and trained by direct backpropagation, cross-validated against four independent optimizers; we train and evaluate on 33 prompts rather than one; and every experiment here is on a real, living, non-neural organism rather than a physics simulation.

3 METHODS

3.1 PROBLEM SETUP

We are given a fixed archive $\mathcal{A} = \{(d_i, x_i)\}_{i=1}^{N}$ of N previously-recorded interventions on a xenobot: $d_i \in \mathbb{R}_{>0}$ is the real electrical-stimulation duration applied in trial i, and x_i is that trial's real, tracked PRE/POST trajectory; $\mathcal{A}$ is fixed throughout training and evaluation, with no new intervention ever applied to collect it.

A user supplies a natural-language prompt $p \in \mathcal{P}$ describing a desired change in behavior (e.g. "get the bot to slow down"). We seek a policy $f_\theta : \mathcal{P} \to \mathbb{R}_{>0}$ that maps a prompt to a predicted duration $\hat{d} = f_\theta(p)$; at evaluation time, $\hat{d}$ is rounded to its nearest real, archived duration (Section 3.4), and judged against that intervention's actual outcome.

Because no ground-truth (p, d) pairs exist – no human ever paired a prompt with a duration – we cannot supervise f_θ directly. Instead a vision-language model acts as a judge $J(p, x) \in [0, 1]$, scoring how well a real trajectory x matches what p describes; J never sees d, only trajectory x_i and prompt text, so its judgment can depend on duration only through what the trajectory itself reveals. Composing J with the archive gives a per-prompt reward over real, already-observed outcomes, $r_i(p) = J(p, x_i)$ for each archive point i, computed once, offline, and never queried live during training. The remainder of this section gives, in order: the real data $\mathcal{A}$ is built from and the train/test archive split (Section 3.2), the architecture of f_θ (Section 3.3), how $\{r_i(p)\}$ becomes a training signal (Section 3.4), and the held-out evaluation protocol (Section 3.5).

3.2 DATASET AND TRAJECTORY EXTRACTION

The archive consists of videos of individually-tracked xenobots undergoing a three-phase protocol: a PRE baseline period, a DURING period (an electrical stimulus of controlled duration d, the sole intervention parameter here), and a POST period. We refer to each such recorded trial as a *batch* b throughout (a full tracked trial, not an ML minibatch).

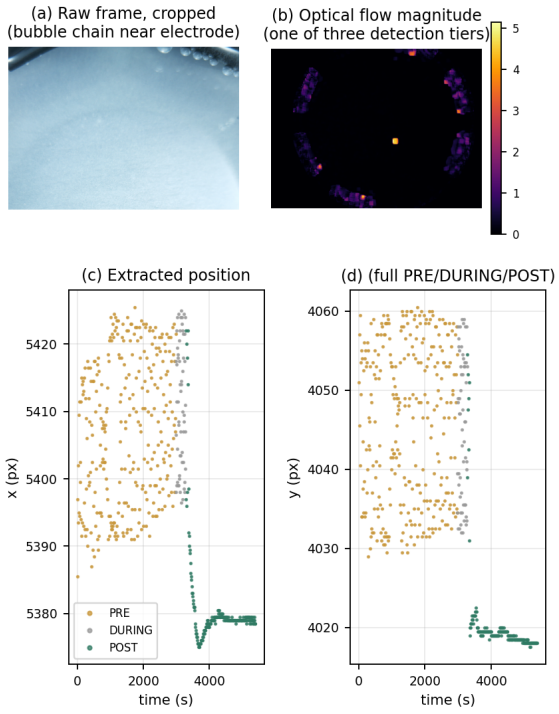

Figure 2: From raw video to trajectory, real batch. (a) A raw frame, cropped, with a stimulus-induced bubble chain near the dish edge. (b) Optical flow magnitude, one of three detection tiers used despite this noise – motion concentrates at the electrodes, with the true signal (bright point near center) far weaker. (c, d) Per-frame x/y position across the full PRE/DURING/POST protocol: a noisy PRE baseline settles into a distinct POST position – the behavior change the VLM judge assesses from the corresponding 2D trajectory.

From raw video to trajectory. Extracting the trajectory used throughout this paper (Figure 2) requires suppressing visual noise specific to this imaging setup – stimulus-induced air bubbles, specular reflection, and debris – that a naive brightness- or contour-based tracker confuses with the xenobot itself. We track position per-frame using the Segment Anything Model (SAM) (Kirillov et al., 2023) across three complementary detection strategies – bounding-box re-detection, optical-flow-guided prompting, and grid-search prompting – each cross-checked against dish geometry and the tracked xenobot's own appearance to reject false positives. This yields, for each batch, a PRE-phase and a POST-phase position trace, each a 2D trajectory in real pixel coordinates. This – not a precomputed velocity – is what the VLM judge J reads (Section 3.4): if J were instead shown a ready-made velocity number, judging a motion prompt could reduce to comparing two numbers, and a plain numerical rule could do that job just as well, defeating the point of testing whether a VLM reading raw trajectories is useful at all. Section 4.3 makes this comparison directly, training against a purely numerical, velocity-only target instead of J's judgment. Frames where all three detection tiers fail are excluded as missing data rather than treated as stillness.

Assessment window. Classifying a batch requires a PRE baseline window and a POST assessment window, each a fixed number of frames – the last few PRE frames and the first few POST frames – which aggregate (mean or median) to a pre-window velocity v_b^{pre} and a post-window (assessment) velocity v_b^{assess} per batch. A window whose length instead scaled with d would confound duration with the amount of time available to observe behavior, since xenobots naturally decelerate after any stimulus, so we fix both window lengths to a constant, independent of d. The two window lengths and the aggregation statistic are themselves chosen by a grid search (7 pre-window sizes $\times$ 8 post-window sizes $\times$ 2 aggregation statistics, 112 combinations), maximizing how well $v_b^{\text{assess}} - v_b^{\text{pre}}$ separates real 'faster'- from 'slower'-labeled batches by duration (Mann–Whitney test). We validate the selected configuration with a permutation test: shuffling the (batch, duration) pairing

and rerunning the full 112-combination search 200 times against the real result. The full grid, the permutation protocol, and both results are in Section 4.2 and the Supplementary Material.

After quality filtering, 101 batches remain classified (26 'faster', 75 'slower') out of 149 tracked; 48 of 149 are excluded as unreliable rather than force-classified. Each of the 149 tracked batches corresponds to a distinct, individually stimulated xenobot construct – no construct is reused across batches, so batches constitute genuine biological replicates rather than repeated trials on the same organism. We use a 60/40 train/test split (61/40: 16/45 faster/slower train, 10/30 faster/slower test), stratified at the sub-class level so the rarest sub-class is represented in both splits.

3.3 Architecture

We embed prompts with SentenceBERT (Reimers & Gurevych, 2019), a frozen 384-dimensional sentence encoder. The P2I policy network (Figure 1e) takes this embedding as input and outputs a predicted intervention duration, through a fully-connected backbone: $384 \to 64$ (ReLU, dropout 0.2) $\to 32$ (ReLU, dropout 0.2), with one sigmoid output head per behavior category (26,786 parameters total for 2 heads) – one head per behavioral direction, so that opposing objectives (e.g. "stop" vs. "accelerate") never share an output. A single shared head is a real alternative, not ruled out by assumption: we test it directly, with a distributional (μ, σ) output and regularization tuned specifically for it, and it remains both worse and less stable on unseen language than 2 heads (Supplementary Material) – so we use 2 heads throughout.

3.4 Reward, baselines, and optimizer comparison

Given a prompt p, f_θ outputs a predicted duration $\hat{d} = f_\theta(p)$. Training requires an error signal computed from real data, without running any new experiment.

For each of the $|\mathcal{P}_{\text{train}}| = 33$ training prompts, every archive point is scored by the VLM judge: $r_i(p) = J(p, x_i)$ for $i = 1, \ldots, N$ (N=101), giving the full reward table $\{r_i(p)\}_{i,p}$, $|\mathcal{A}| \times |\mathcal{P}_{\text{train}}| = 101 \times 33 = 3{,}333$ real (archive point, prompt) pairs, computed once and reused for every subsequent optimization step. As an ablation, J is also queried from a precomputed velocity-vs-time plot instead of raw positions, to test whether withholding that signal and forcing J to infer pace from position and time alone produces a more reliable judge.

This table admits two different training signals for θ (Figure 3).

(a) Discrete route. Following prior work in this line of research (Le et al., 2025; 2026), the predicted duration is snapped to its nearest archive point,

$$i^*(\hat{d}) = \underset{i \in \{1,\ldots,N\}}{\arg\min} \left| d_i - \hat{d} \right|, \tag{1}$$

and its score $r_{i^*(\hat{d})}(p)$ is used directly as the reward, so θ is optimized to maximize $\mathbb{E}_{p \sim \mathcal{P}_{\text{train}}}\left[r_{i^*(f_\theta(p))}(p) \right]$. Because the $\arg\min$ in Equation 1 is piecewise-constant, this reward is non-differentiable with respect to $\hat{d}$ almost everywhere; we reproduce this route as an ablation, optimizing it with two methods that do not require a gradient through it – REINFORCE (policy gradient) and CMA-ES (Hansen & Ostermeier, 2001) (diagonal-covariance variant; Ros & Hansen, 2008) (full hyperparameters in the Supplementary Material).

(b) Differentiable route. For each prompt p, we fit a logistic curve to its $\{(d_i, r_i(p))\}_{i=1}^N$ pairs,

$$\hat{r}_p(d) = \sigma\Big(w_p \, \tilde{d} + b_p \Big), \qquad \tilde{d} = \frac{d - d_{\min}}{d_{\max} - d_{\min}}, \tag{2}$$

with (w_p, b_p) fit per prompt by logistic regression and $d_{\min}, d_{\max}$ the archive's own minimum and maximum duration. We reduce $\hat{r}_p(\cdot)$ to a single fixed regression target $d^\dagger(p)$ – the duration at the midpoint of $\hat{r}_p$'s own achieved range, clipped into the nearest real, archive-verified interval where every point is correctly labeled for p's intended behavior if it would otherwise fall outside one (construction detail in the Supplementary Material) – and train

$$\min_\theta \; \mathbb{E}_{p \sim \mathcal{P}_{\text{train}}}\Big[\big(f_\theta(p) - d^\dagger(p) \big)^2 \Big] \tag{3}$$

by ordinary backpropagation (Adam). No VLM, archive lookup, or simulation is queried during a training step under this route – $d^\dagger(p)$ is computed once, offline, directly from the already-computed reward table. The same table independently validates both routes' outcomes against ground truth (Section 4.2).

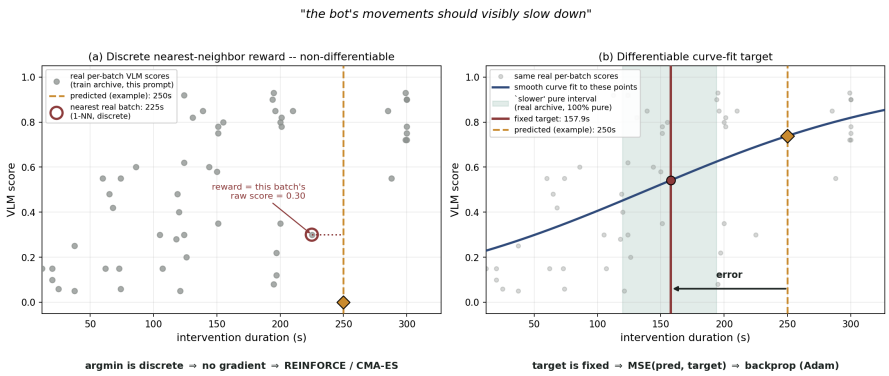

Figure 3: The two ways to turn the precomputed VLM-score table into a training signal, for the same real prompt. (a) Discrete nearest-neighbor reward: a predicted duration snaps to its nearest real archive batch, and that batch's raw score is read off directly as the reward – the argmin is discrete, so this route needs REINFORCE or CMA-ES. (b) Differentiable curve-fit target: the same real points, fit into one smooth curve and reduced to a single fixed target, so the error is an ordinary mean-squared error and the whole thing trains by backpropagation.

Agreement between all three resulting optimizers – one gradient-based (Equation 3), one policy-gradient, one derivative-free, all consuming the identical reward table – validates the optimum each finds rather than assuming one is necessary.

3.5 Evaluation protocol

We verify no held-out prompt duplicates a training prompt and bound cross-split similarity (full list in the Supplementary Material). A held-out prompt's output head is inferred from its embedding via nearest-centroid similarity to training-prompt centroids, never looked up from a known category.

Let $\mathcal{P} = \mathcal{P}_{\text{train}} \sqcup \mathcal{P}_{\text{test}}$ and $\mathcal{A} = \mathcal{A}_{\text{train}} \sqcup \mathcal{A}_{\text{test}}$ be the disjoint prompt and archive splits (Section 3.2). Four evaluation cells follow from the product of prompt side $\times$ archive side:

$$\text{Train} = \mathcal{P}_{\text{train}} \times \mathcal{A}_{\text{train}}, \quad \text{GT3} = \mathcal{P}_{\text{train}} \times \mathcal{A}_{\text{test}}, \quad \text{GT1} = \mathcal{P}_{\text{test}} \times \mathcal{A}_{\text{train}}, \quad \text{GT2} = \mathcal{P}_{\text{test}} \times \mathcal{A}_{\text{test}}. \quad (4)$$

For a cell $\mathcal{P}' \times \mathcal{A}'$, a prediction $f_\theta(p)$, $p \in \mathcal{P}'$, is scored by k-nearest-neighbor vote (k=3) among $\{d_i : i \in \mathcal{A}'\}$ against each neighbor's true, independently-tracked 'behavior_group' label, and is correct if the majority vote matches p's intended direction. GT2 is the only cell where both $p \notin \mathcal{P}_{\text{train}}$ and every $i \in \mathcal{A}'$ satisfies $i \notin \mathcal{A}_{\text{train}}$ – generalization along both axes at once.

Chance on this task is 66.7%, not 50%: prompts are drawn evenly from three buckets (motion_increase, motion_reduction, stop), but 'stop' and 'motion_reduction' are both scored against the same 'slower' archive direction, so two of every three prompts need 'slower' – a baseline that always answers 'slower' is right two-thirds of the time by construction, not by tracking any real signal. Two further checks reuse this setup without retraining: an $8\times$-paraphrase stress test (Section 4.1) and a post-hoc sub-label split of 'slower' into genuine stops vs. mild slowdowns (Section 4.4).

4 Results

Each subsection below tests one of the three contributions claimed in the introduction directly, rather than asserting it: Section 4.1 tests contribution (2), the headline generalization number, and is itself the direct evidence for contribution (1) – that training works at all without any new experiment; Section 4.2 checks the precondition contribution (1) depends on, that J's judgment and the archive it draws from are themselves trustworthy; Section 4.3 tests contribution (3), that this reward is not redundant with a simpler numerical alternative; Section 4.4 is further robustness analysis not tied to a specific claimed contribution.

4.1 Main result: generalization to held-out prompts and archives

One of the contributions we state in the introduction is that this offline paradigm generalizes to entirely new instructions at well above chance; this section runs that experiment directly. Table 1 shows Backprop outperforming the two cross-validating optimizers, CMA-ES and REINFORCE,

Table 1: Ground-truth accuracy, all four evaluation cells. 30 seeds for all three optimizers.

	Train	GT3	GT1	GT2
Backprop (reference model)	$\mathbf{89.1 \pm 1.0\%}$	$\mathbf{88.9 \pm 0.0\%}$	$\mathbf{81.1 \pm 5.6\%}$	$\mathbf{80.0 \pm 0.0\%}$
CMA-ES	$66.7 \pm 0.0\%$	$66.5 \pm 1.0\%$	$62.4 \pm 7.1\%$	$62.9 \pm 6.7\%$
REINFORCE	$69.4 \pm 10.2\%$	$69.4 \pm 10.2\%$	$66.9 \pm 9.0\%$	$66.9 \pm 8.7\%$

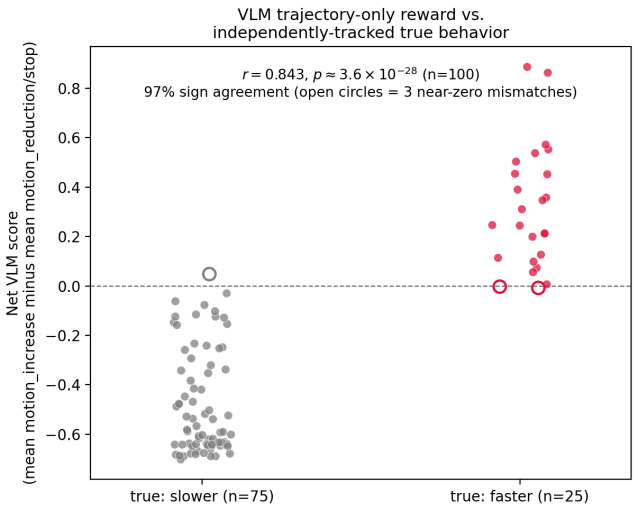

Figure 4: Net VLM score $s(b)$ per batch (jittered for visibility) against true, independently-tracked behavior label. Open circles: the 3 sign-disagreement batches.

on every evaluation cell, each averaged over independent seeds (full analysis in the Supplementary Material). We use Backprop as the reference model for the rest of this paper.

These four cells test generalization along two independent axes: unseen *prompts* and unseen *archive* data. Accuracy degrades gracefully across all three held-out cells relative to Train, and GT2 – stressing both axes at once – stays well above the 66.7% chance baseline. This is neither a leakage nor a small-sample artifact: no held-out prompt duplicates a training prompt (max 0.60 Jaccard overlap), and stress-testing the single saved, already-trained model at $8\times$ scale (120 new paraphrases, no retraining) holds accuracy at 85.0%, close to the original 80.0%, with 'stop' the most reliable bucket (full detail in the Supplementary Material).

4.2 Validating the VLM judge and the archive

We first ask whether J's judgment tracks a batch's real, independently-tracked behavior at all. For batch b, define its net score $s(b)$ as the mean of $r_b(p)$ across motion_increase prompts minus the mean across motion_reduction/stop prompts, so $s(b) > 0$ means J judged b as more consistent with speeding up than slowing down. Figure 4 plots $s(b)$ for every batch against that batch's true, independently-tracked behavior label. The two agree at $r = 0.843$ with 97% sign agreement; every disagreement sits within ± 0.02px/s of a true zero velocity change – genuinely ambiguous cases, not J misreading a clear signal. Correlating $s(b)$ against the raw velocity delta itself, instead of the binary label, gives a markedly weaker $r = 0.49$: J reliably separates the two behavior categories but does not linearly track fine-grained speed within one, consistent with how the reward is used throughout this paper (a per-prompt target duration, not a speed regression). Querying J from a precomputed velocity-vs-time plot instead of the raw trajectory reaches only $r = 0.49$–0.66 per category – withholding the precomputed answer and forcing genuine inference produces a more reliable J, not merely a harder-to-game one.

The archive itself is equally load-bearing. The selected window configuration (30 pre-frames, 50 post-frames, median aggregation) separates $\{d_b\}$ by behavior at Mann–Whitney $p = 1.57 \times 10^{-4}$, against a best-achievable $p = 1.6 \times 10^{-5}$ elsewhere in the same 112-point grid – trading a little separability for more readable per-batch plots for J. But this configuration was *chosen* by searching all 112 combinations rather than fixed in advance, so the achieved p-value alone is not trustworthy: searching many configurations and reporting whichever looks best is a form of p-hacking, where a

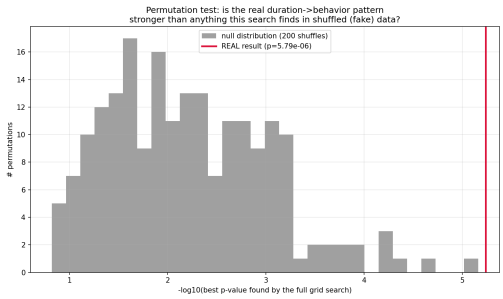

Figure 5: Null distribution of the best p-value achievable by the full 112-way grid search across 200 random duration-label shuffles, against the real result.

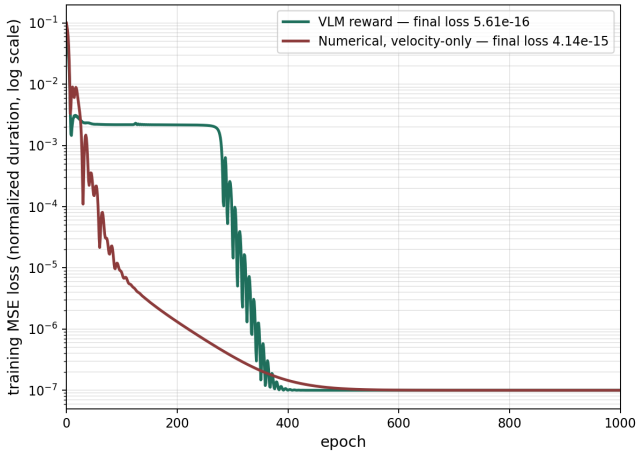

Figure 6: Training loss (seed 42, log scale) for the VLM-fit target $d^\dagger(p)$ vs. the purely numerical, velocity-only target – same architecture, optimizer, and epochs. Both converge to a comparable near-zero loss; the numerical target is not harder to fit.

small p-value can reflect the size of the search rather than a real effect. We check for exactly this with a permutation test: shuffle the (b, d_b) pairing, rerun the full 112-way search on the shuffled data, and keep only its best-achievable p-value; repeating this 200 times gives a null distribution of what the search alone could produce by chance, with no real effect present. The real configuration's $p = 1.57 \times 10^{-4}$ is more extreme than all 200 shuffled best-case results (Figure 5), so the separation reflects a real effect and not an artifact of having searched over many configurations.

4.3 VLM REWARD VS. A NUMERICAL TARGET

The purpose of this section is to check whether the VLM-derived target $d^\dagger(p)$ (Equation 3) is doing real work, or whether a much simpler numerical target – with no VLM involved at all – would train f_θ just as well. With only two archive behaviors currently distinguished, this is a real possibility worth ruling out directly rather than assuming.

We replace $d^\dagger(p)$ with a single fixed duration per behavior direction, chosen purely from real velocity data: the training-archive batch whose measured velocity change ($v_b^{\text{assess}} - v_b^{\text{pre}}$) is strongest in that direction, no VLM or per-prompt differentiation involved. Every 'faster' prompt is given the same target duration, and every 'slower' prompt the same (different) target duration. f_θ is trained against these fixed targets with the identical objective as Equation 3 (mean squared error, Adam, 30 seeds), so training loss and held-out accuracy are directly comparable to the VLM-fit target.

Figure 6 shows both targets converge to a comparable near-zero training loss, so the numerical target is not simply easier or harder to fit than $d^\dagger(p)$, and optimization difficulty cannot explain a generalization gap between them if one exists. Table 2 shows one does.

Train accuracy is nearly identical (88.9% vs. 89.1%), but GT2 falls to $59.1 \pm 2.3\%$ – *below* the 66.7% chance baseline. Per-prompt inspection shows why: the single chosen 'faster' duration happens to

Table 2: Same architecture and optimizer (Backprop), 30 seeds; only the training target differs.

	Train	GT3	GT1	GT2
VLM reward (Table 1)	$\mathbf{89.1 \pm 1.0}\%$	$\mathbf{88.9 \pm 0.0}\%$	$\mathbf{81.1 \pm 5.6}\%$	$\mathbf{80.0 \pm 0.0}\%$
Numerical, velocity-only	$88.9 \pm 0.0\%$	$66.7 \pm 0.0\%$	$79.6 \pm 1.7\%$	$59.1 \pm 2.3\%$

land in a region of the held-out archive's duration axis dominated by 'slower'-labeled real batches, so every held-out 'faster' prompt is misclassified as 'slower' – not a training failure but a placement failure, invisible from training accuracy alone. A curve fit to the full VLM-scored distribution avoids this; a single extremal real data point does not. Even with only two behaviors currently defined, the VLM reward is not redundant.

4.4 Revisiting 'stop' vs. 'slower'

This section is further robustness analysis, not a test of a claimed contribution. 'Stop' and 'slower' prompts are merged throughout this paper into one 'motion_reduction' class because the data cannot support a clean 3-way split: a direct stop/mild-slower/faster classifier reaches only 0.477 balanced accuracy (barely above chance), against 0.713 once merged.

Is that gap because stop and mild slowdown are genuinely indistinguishable, or simply data-starved? Retrospectively splitting 'slower' by a sub-label rule (velocity within 10% of baseline at any point) finds a real, significant separation between genuine stops (n=48) and mild slowdowns (n=14; $p = 7.5 \times 10^{-6}$, non-overlapping interquartile ranges) – the distinction exists in the data. But three separate attempts to learn it from language all collapsed to the same majority answer regardless of prompt, consistent with data scarcity (only 14 mild-slower examples) rather than a model limitation (detail in the Supplementary Material). This does not affect the paper's main result, which only relies on the coarser merge, and is left for future work.

A downstream consequence. Checking where each held-out 'stop'/'motion_reduction' prediction actually lands (90 held-out prompts; Table 3): 'stop' is highly reliable (97.8% land on a true stop); 'motion_reduction' over-shoots to full-stop behavior in 33.3% of cases, traced to the same class imbalance within held-out 'slower' data (19 true-stop vs. 6 true-mild-slower) that motivated the merge in the first place.

Table 3: Where each held-out prompt's predicted duration lands, against the archive's true sub-label (n=45 each; full per-source breakdown in the Supplementary Material).

prompt intent	lands on true STOP	lands on true MILD-SLOWER
'stop' prompt	$\mathbf{97.8}\%$ (44/45)	2.2% (1/45)
'motion_reduction' prompt	33.3% (15/45)	$\mathbf{66.7}\%$ (30/45)

5 Discussion and Future Work

In this paper, we have demonstrated that a natural-language interface for a living organism can be learned entirely offline, from an existing archive and a vision-language model's own judgment, without any new experiments. The resulting P2I mapping generalizes to entirely new instructions at 80.0% ground-truth accuracy against a 66.7% chance baseline, holding across three independent optimizers sharing the same VLM reward; this reward is also not redundant with a simpler numerical alternative, which collapses held-out generalization to 59.1%, below chance (full analysis in the Supplementary Material).

This paper's scope has real limits: one intervention parameter, a small archive (101 of 149 tracked batches), predefined behavior categories, and no main-text result from a new physical experiment. The system retrieves rather than generates – every prediction is the archive's own closest already-observed match, with compositional recombination beyond it left untested (Supplementary Material). The most direct next step is prospective: does a predicted intervention, applied to a newly-stimulated xenobot never seen in training, actually reproduce its intended behavior? Beyond xenobots, the same offline, archive-plus-VLM-judgment paradigm applies wherever an intervention-outcome archive already exists – other synthetic living machines, cell and tissue cultures under electrical or optogenetic control, and, further out, clinical bioelectric therapy (Ebrahimkhani & Levin, 2021; Zajdel et al., 2020; Blackiston et al., 2021b; Levin, 2021) – with this offline half a tractable place to start.

AI USE STATEMENT

A vision-language model (Claude) is a core methodological component of this work: it supplies the training reward for P2I by judging real archive trajectories against natural-language prompts (Section 3.4), and its reliability as a judge is itself directly evaluated against independent, non-VLM ground truth (Section 4.2) rather than assumed. Separately, in preparing this manuscript, a generative AI tool (Claude) was used to assist with paraphrasing. All AI-assisted text was reviewed by the authors. We take responsibility for the final content of this work, including text, claims, and artifacts produced with the aid of generative AI.

REFERENCES

Emmanuel Bengio, Moksh Jain, Maksym Korablyov, Doina Precup, and Yoshua Bengio. Flow network based generative models for non-iterative diverse candidate generation. In *Advances in Neural Information Processing Systems*, volume 34, 2021.

Douglas Blackiston, Emma Lederer, Sam Kriegman, Simon Garnier, Joshua Bongard, and Michael Levin. A cellular platform for the development of synthetic living machines. *Science Robotics*, 6 (52):eabf1571, 2021a.

Douglas J Blackiston, Alexis K Trinh, Bianca Tran, and Michael Levin. Stimulating neuroplasticity to restore vision. *Eneuro*, 8(6), 2021b.

Rich Caruana. Multitask learning. *Machine learning*, 28(1):41–75, 1997.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.

Zhao Chen, Vijay Badrinarayanan, Chen-Yu Lee, and Andrew Rabinovich. Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In *International conference on machine learning*, pp. 794–803. PMLR, 2018.

Mo R Ebrahimkhani and Michael Levin. Synthetic living machines: A new window on life. *Iscience*, 24(5), 2021.

Valentina Emiliani, Emilia Entcheva, Rainer Hedrich, Peter Hegemann, Kai R Konrad, Christian Lüscher, Mathias Mahn, Zhuo-Hua Pan, Ruth R Sims, Johannes Vierock, et al. Optogenetics for light control of biological systems. *Nature Reviews Methods Primers*, 2(1):55, 2022.

Daniel Fried, Armen Aghajanyan, Jessy Lin, Sida Wang, Eric Wallace, Freda Shi, Ruiqi Zhong, Scott Yih, Luke Zettlemoyer, and Mike Lewis. Incoder: A generative model for code infilling and synthesis. In *International Conference on Learning Representations*, 2023.

Nikolaus Hansen and Andreas Ostermeier. Completely derandomized self-adaptation in evolution strategies. *Evolutionary computation*, 9(2):159–195, 2001.

Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. In *International Conference on Machine Learning*, pp. 9118–9147. PMLR, 2022.

John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *nature*, 596(7873):583–589, 2021.

Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. In *International Conference on Computer Vision*, pp. 4015–4026, 2023.

Sam Kriegman, Douglas Blackiston, Michael Levin, and Josh Bongard. Kinematic self-replication requires transient breaking of symmetry and multibody coordination. *Proceedings of the National Academy of Sciences*, 118(9):e2011466118, 2021.

Nam Le, Patrick Erickson, Yanbo Zhang, Michael Levin, and Josh Bongard. Giving simulated cells a voice: Evolving prompt-to-intervention models for cellular control. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pp. 2327–2335, 2025.

Nam H. Le, Patrick Erickson, Yanbo Zhang, Michael Levin, and Josh Bongard. Zapgpt: Free-form language prompting for simulated cellular control. In *Proceedings of the 2025 8th International Conference on Computational Intelligence and Intelligent Systems*, CIIS '25, pp. 119–126, New York, NY, USA, 2026. Association for Computing Machinery. ISBN 9798400719332. doi: 10.1145/3787256.3787274. URL https://doi.org/10.1145/3787256.3787274.

Michael Levin. Bioelectric signaling: Reprogrammable circuits underlying embryogenesis, regeneration, and cancer. *Cell*, 184(8):1971–1989, 2021.

Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.

Jacky Liang, Wenlong Huang, Fei Xia, Peng Xu, Karol Hausman, Brian Ichter, Pete Florence, and Andy Zeng. Code as policies: Language model programs for embodied control. In *2023 IEEE International conference on robotics and automation (ICRA)*, pp. 9493–9500. IEEE, 2023.

Benedict C May, Nathan Korda, Anthony Lee, David S Leslie, and Nicolo Cesa-Bianchi. Optimistic bayesian sampling in contextual-bandit problems. *Journal of Machine Learning Research*, 13(6), 2012.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PmLR, 2021.

Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2019.

Juan Rocamonde, Victoriano Montesinos, Elvis Nava, Ethan Perez, and David Lindner. Vision-language models are zero-shot reward models for reinforcement learning. *arXiv preprint arXiv:2310.12921*, 2023.

Raymond Ros and Nikolaus Hansen. A simple modification in cma-es achieving linear time and space complexity. In *International Conference on Parallel Problem Solving from Nature*, pp. 296–305. Springer, 2008.

Sebastian Ruder. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*, 2017.

Yufei Wang, Zhanyi Sun, Jesse Zhang, Zhou Xian, Erdem Biyik, David Held, and Zackory Erickson. Rl-vlm-f: Reinforcement learning from vision language foundation model feedback. *arXiv preprint arXiv:2402.03681*, 2024.

Tianhe Yu, Saurabh Kumar, Abhinav Gupta, Sergey Levine, Chelsea Finn, and Karol Hausman. Gradient surgery for multi-task learning. *Advances in Neural Information Processing Systems*, 2020.

Tom J Zajdel, Gawoon Shim, Linus Wang, Alejandro Rossello-Martinez, and Daniel J Cohen. Scheepdog: programming electric cues to dynamically herd large-scale cell migration. *Cell systems*, 10(6):506–514, 2020.

Yizhen Zheng, Huan Yee Koh, Jiaxin Ju, Anh TN Nguyen, Lauren T May, Geoffrey I Webb, and Shirui Pan. Large language models for scientific discovery in molecular property prediction. *Nature Machine Intelligence*, 7(3):437–447, 2025.

A APPENDIX

Full methodology, prompt lists, VLM query template, held-frame investigation, generalization matrices, and additional ablations are provided in the Supplementary Material accompanying this submission.