# Generalization by Composition: Primitives for Robot Control

INVITED TALK / DARTMOUTH AUTONOMY SEMINARS


Avik De

Adjunct Associate Professor, Electrical and Systems Engineering, University of Pennsylvania

Chief Science Officer, Ghost Robotics


Nov 2026

# Intelligence in Motor Skills

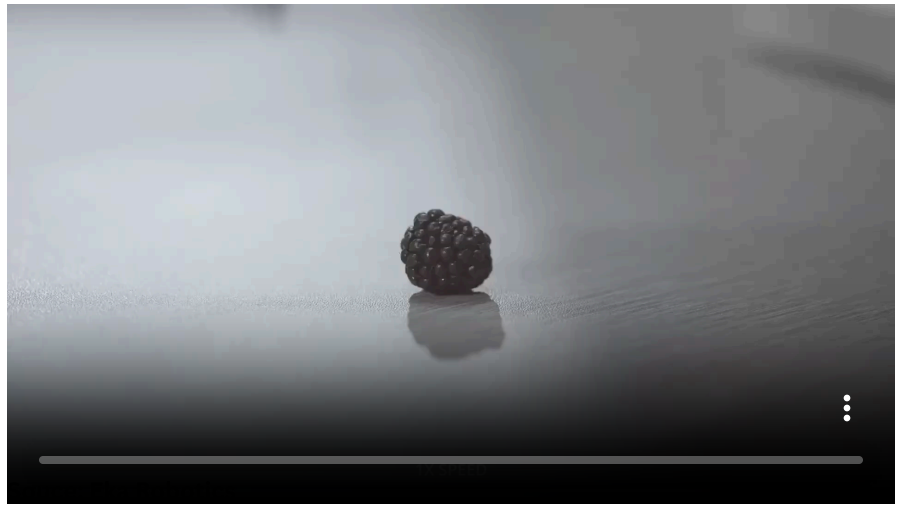
## What is intelligence?

- Compression (Shannon, 1951) (Buchanan et al., 2026)
- Prediction (Clark & others, 2013)
- Prediction error (Friston, 2010)
- Embodiment (Brooks, 1991)
- Generalization (Legg & Hutter, 2007)

## Applications of intelligence

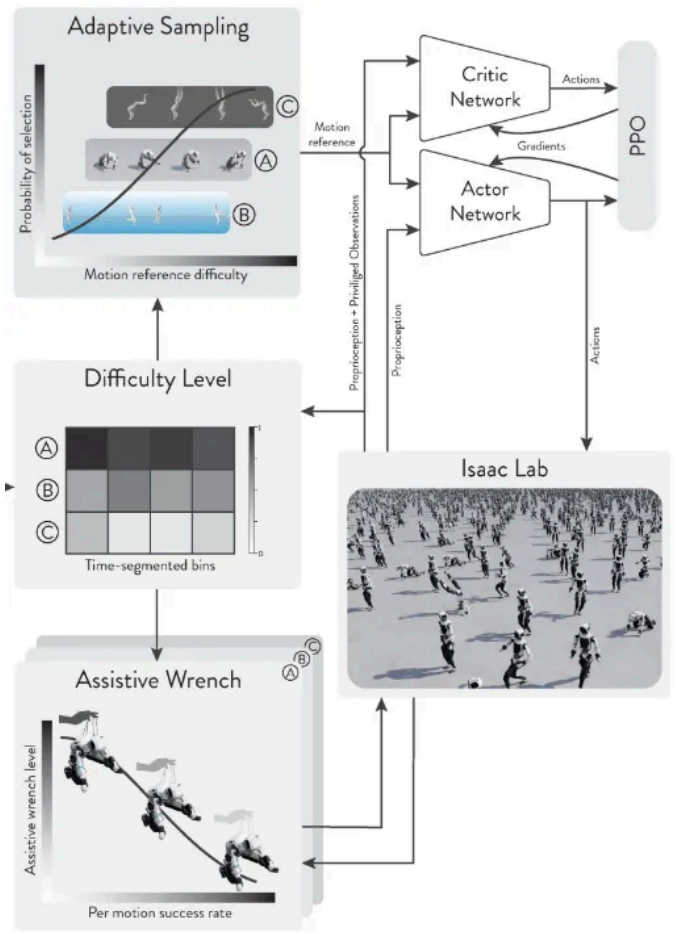
> *I like to talk about intelligence having two parts, an athletic part and a cognitive part.*
>
> *— Marc Raibert*

SoTA performance (EKA-Robotics, 2026) (Sleiman et al., 2026)





# RL With Fine-Grained Actions

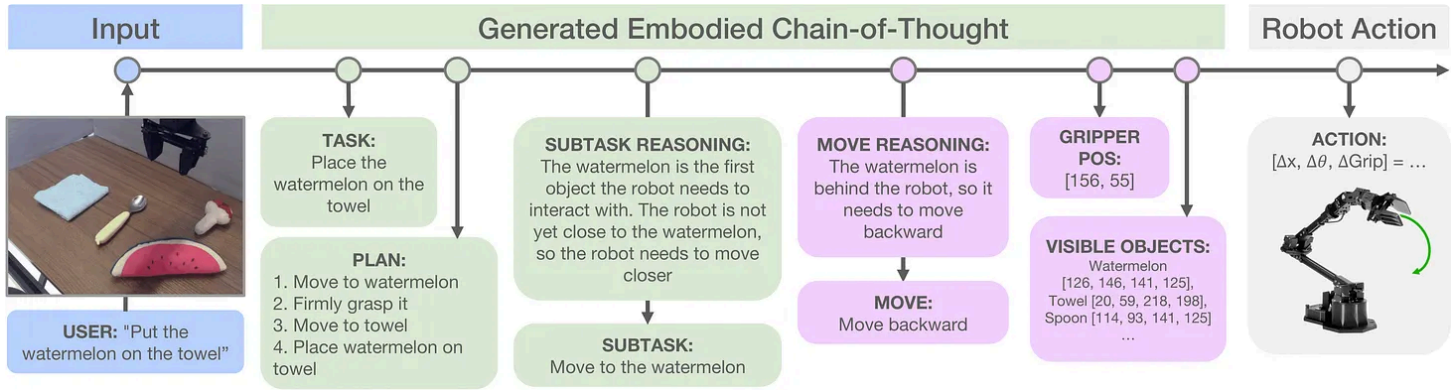



ZEST uses residual joint-space corrections to reference positions with a fixed impedance controller (Sleiman et al., 2026)

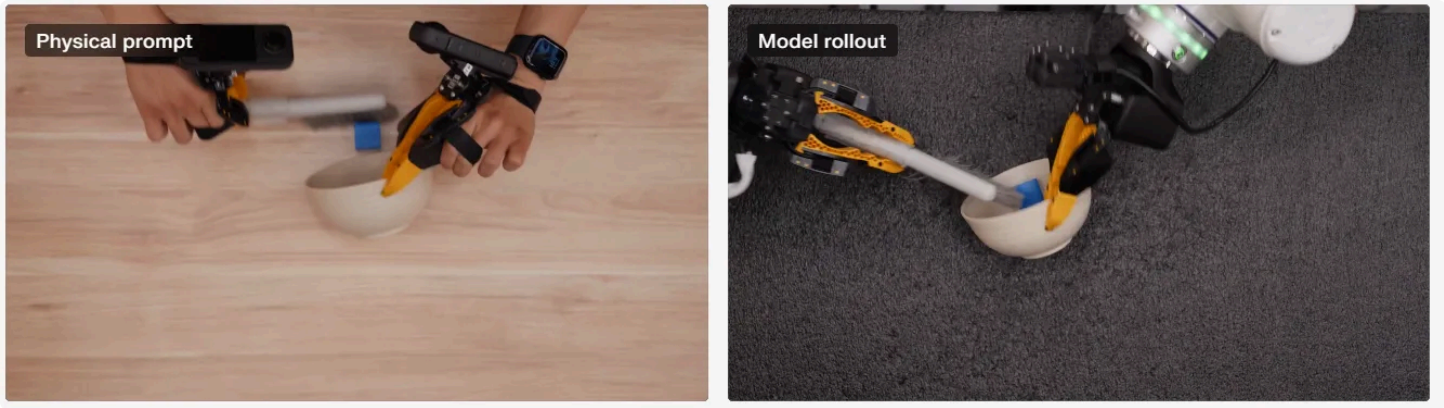
- Using joint actions for RL first showed success in animation (Peng & Van De Panne, 2017)
- Sim2real successes starting with (Haarnoja et al., 2018)
- Fits with the bitter lesson (Sutton, 2019)
- RL's main *challenge* is generalization across **task**, **environment**, and **embodiment** (Ibarz et al., 2021) (Kirk et al., 2023)

RL facilitates flexible, robust, embodied intelligence, but is weak at **generalization**

# Task, Embodiment Generalization in Foundation Models




pi0.7 uses a scratchpad for subtask reasoning (Intelligence et al., 2026)




GEN-1.5 in-context learning for generalization to new tasks

## Task generalization

Observed:

- Sequential task composition
- Text-based subtask generation
- In-context learning (ICL)

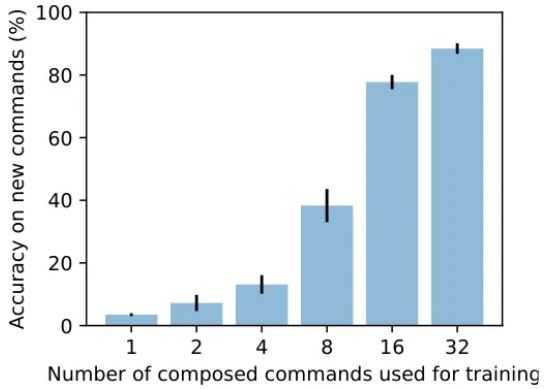
## Embodiment generalization

- Tool improvisation
- Large-scale cross-embodiment pretraining*

> *Missing: RL's robustness!*

# Extrapolation ("Jump") is Not Solved

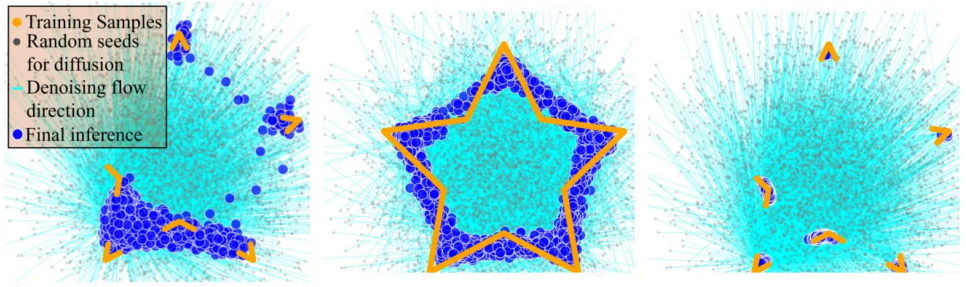
## No systematicity




Composed examples needed for combined commands (Lake & Baroni, 2018)

RNNs need training examples in context

## Test-time extrapolation




(a) low-capacity model, small data (SD)

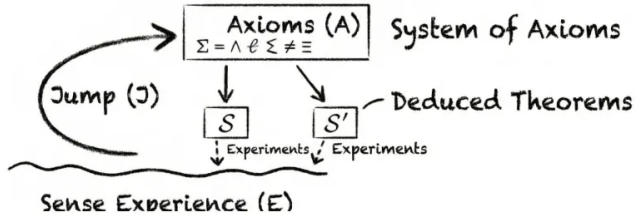
(b) low-capacity model, large data (LD)

(c) high-capacity model, SD (e.g. Diffusion Policy)

Test-time extrapolation of diffusion policies, ACT (He et al., 2026)

Diffusion policies, ACT behave like lookup tables or interpolators, with poor OOD results

## LLMs can't jump

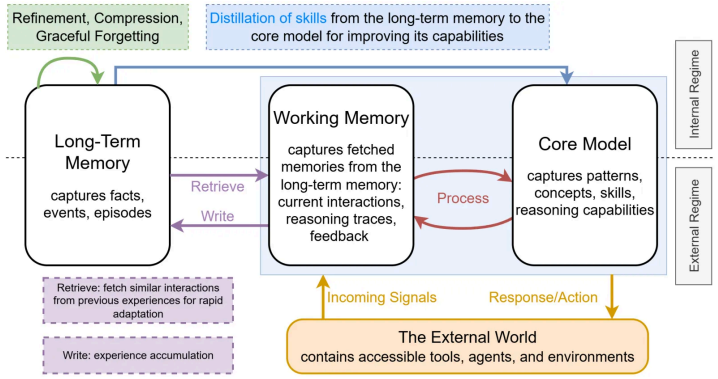



Einstein's Experience → Jump → Action Diagram (Zahavy, 2026)

LLMs cannot form novel hypotheses ("jumps") when there is no error signal gradient for hillclimbing

# Adaptation is Not Solved

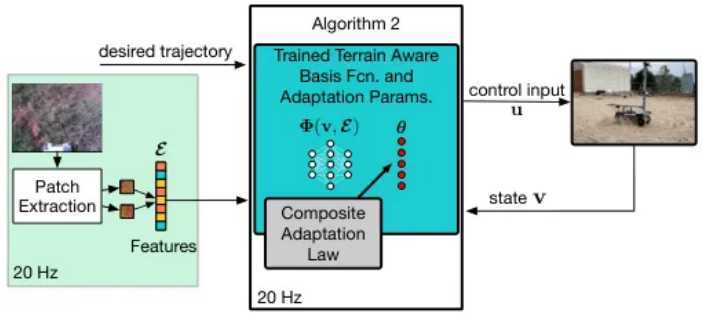
## Continual learning




Modular memory architecture (Dorovatas et al., 2026)

We cannot yet continuously acquire new skills without catastrophic forgetting (IWL) or blowing up inference compute (ICL)

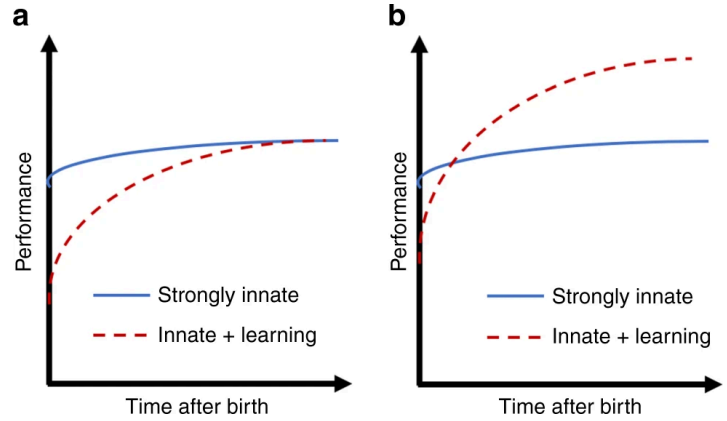
## Disturbance rejection




"Last layer" adaptation (Lupu et al., 2024)

- Terrain/environment robustness poor
- Loops and hierarchies acting at different time-scales are key to online adaptation

## What Biology does

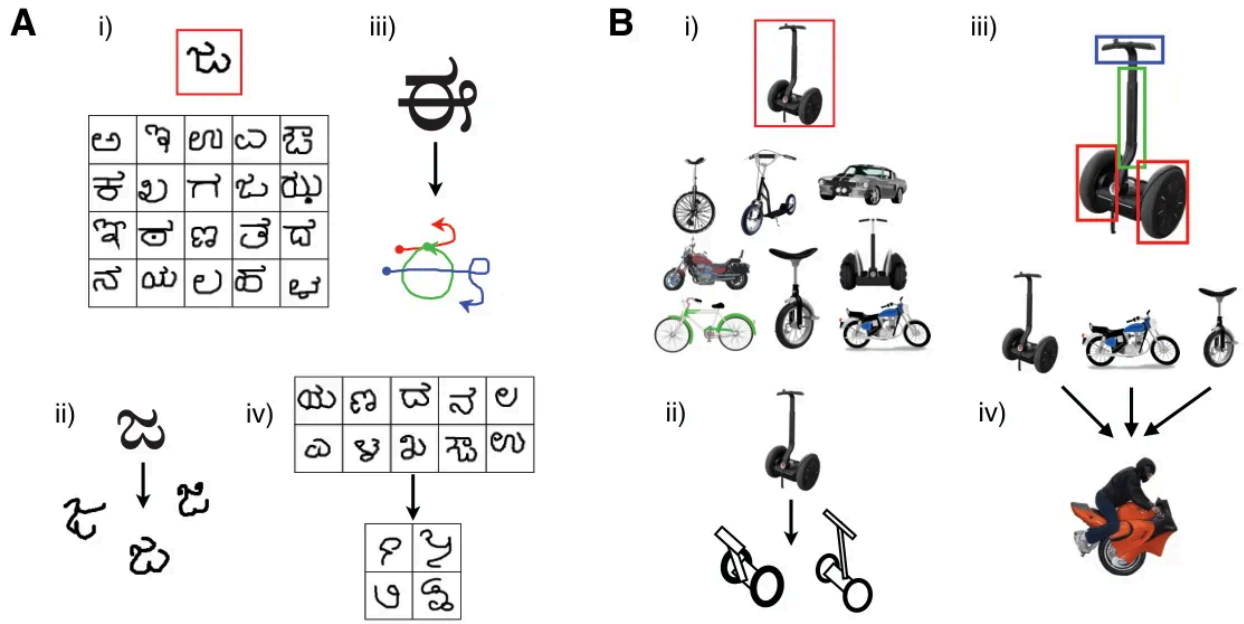



The species relying on a strongly innate strategy will outcompete the species employing a mixed strategy (Zador, 2019)

Animals are able to function well so soon after birth because they are born with highly **structured** brain connectivity

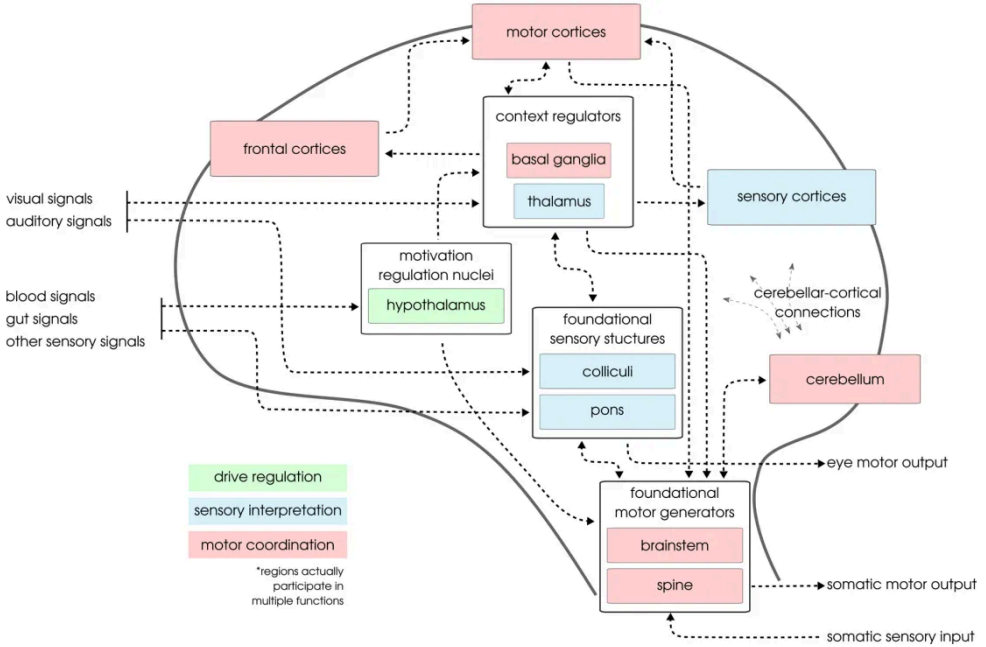
# Building Machines That Learn, Think, and Move Like People

## What kind of structure?




Pattern-recognition scaling vs. compositional, modular causal models (and physics/motor primitives) for sample-efficient generalization (Lake et al., 2017)

## Literally instantiated

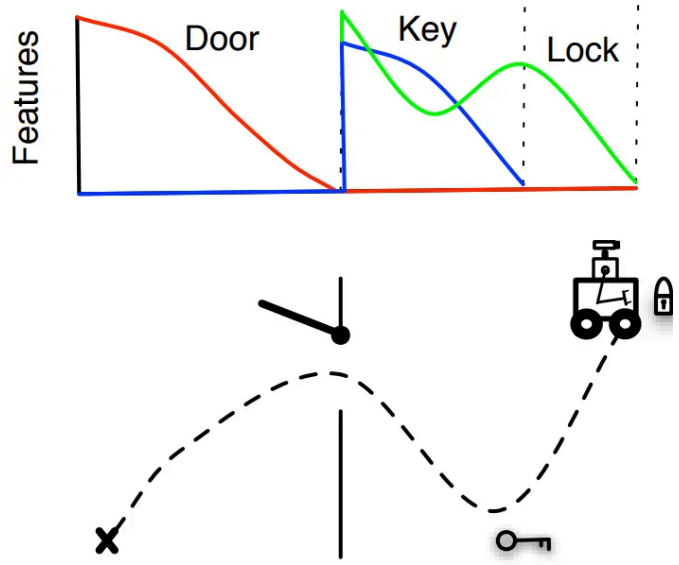



Modular, hierarchical organization of biological motor systems (spinal primitives, brainstem loops, cortical planning) (Merel et al., 2019)

> *Is it because Biology just doesn't have powerful enough computers?*

# Computational Challenges are Not Solved

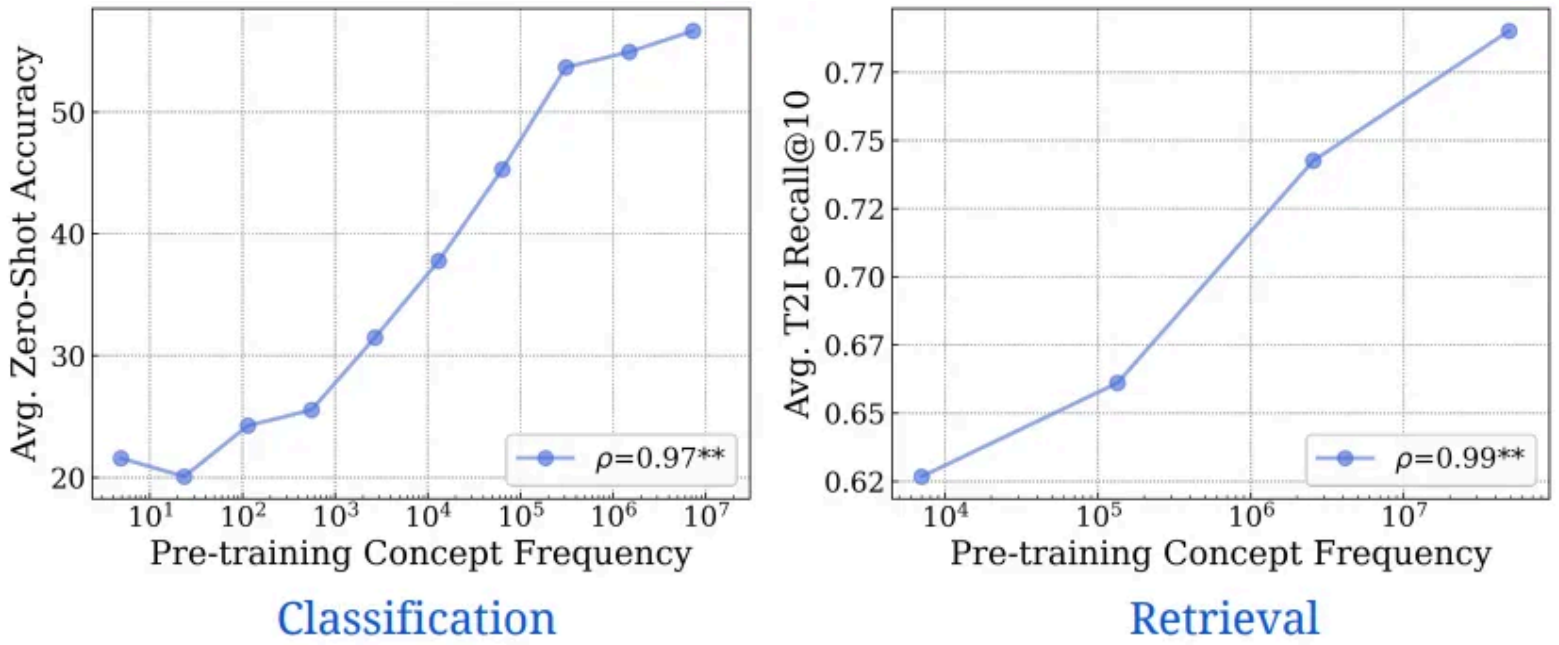
## Training complexity




Constructing Skill Trees (Konidaris et al., 2012)

$O(N_{\text{skills}})$ vs. $O(N_{\text{sequences}})$ demonstrations with explicit composition

## Exponential scaling




Linear improvements in downstream "zero-shot" capability require exponentially more pretraining data on those concepts (Udandarao et al., 2024)

# Outline

## Thesis

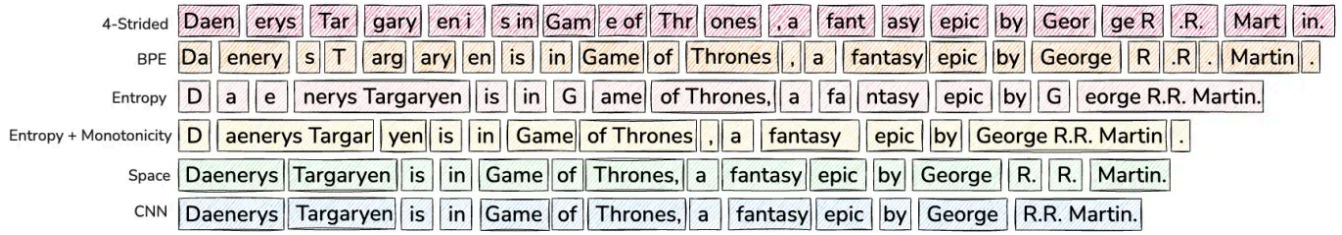
> *An intelligence architecture should be **generative, adaptive, efficient, and embodied** — able to understand and use the body's dynamics, synthesize novel behaviors, and efficiently adapt to new tasks and environments.*

## Organization

1. **Introduction**
2. The Case For Dynamical Action Primitives
3. Building Composable Action Primitives
4. Open Questions

# The Case For Dynamical Action Primitives

# An Analogy: Tokens in LLMs



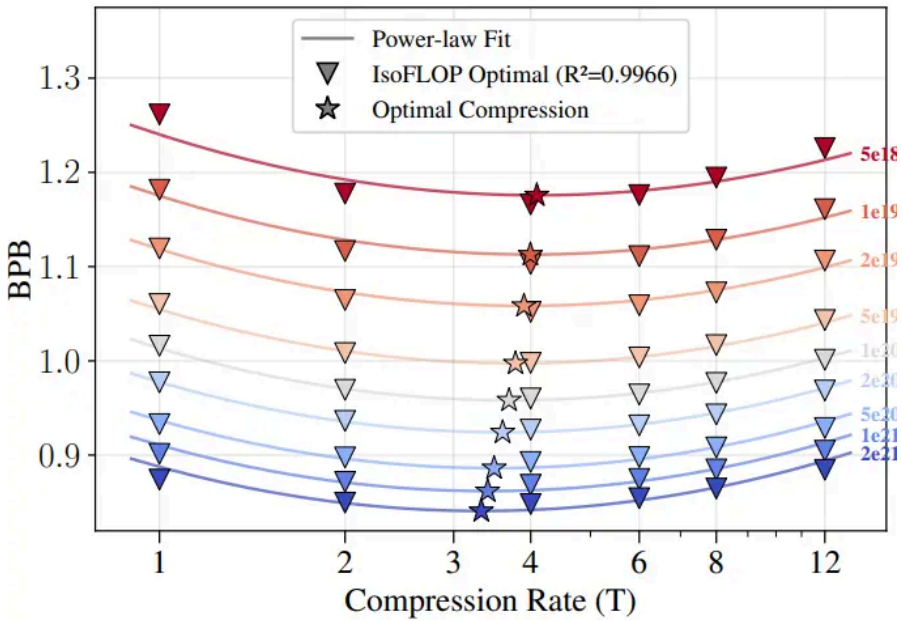

Strategies for grouping bytes into patches (Pagnoni et al., 2025)

## Vocabulary size / token size tradeoff

Letters (fine) vs. words (coarse)

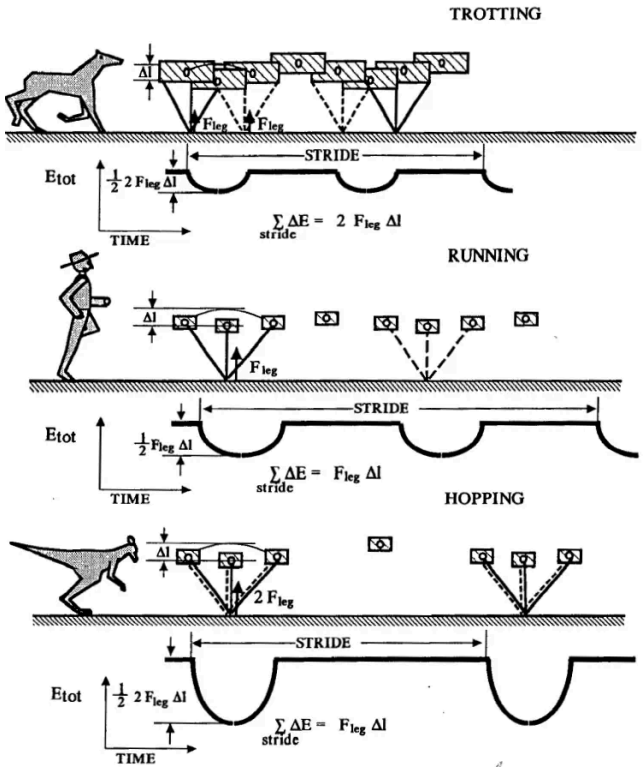
- Memory: higher for coarse tokens
- Computation: longer sequences for fine tokens
- Generation: sequences of words vs. letters
- Transfer: coarse are task-specific, but transfer better

E.g. medical journal, English-Italian dictionary

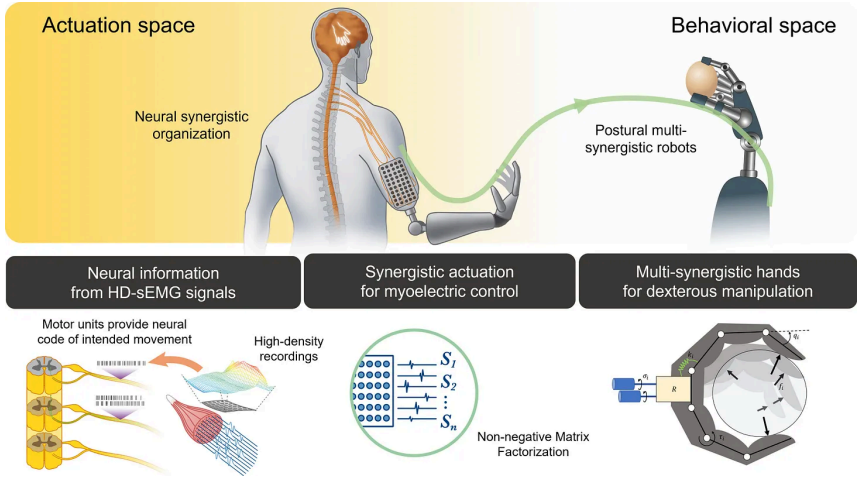



Loss (BPB) vs. bytes/token (compression rate) shows an optimum at each training budget (Limisiewicz et al., 2026)

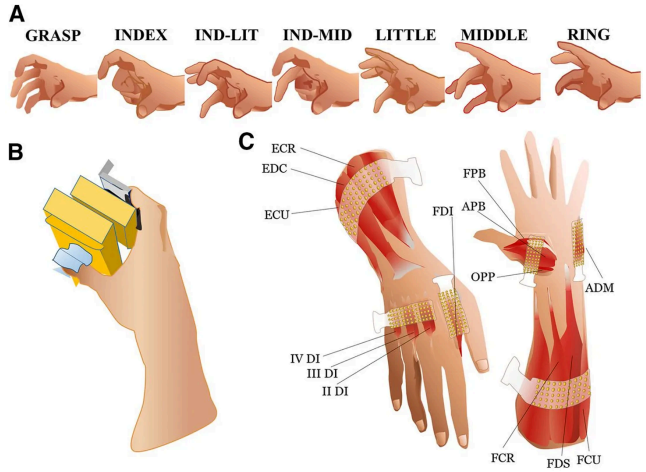
# Animals Exhibit Coarse Action Primitives




Bouncing like a monopode (Blickhan & Full, 1993)

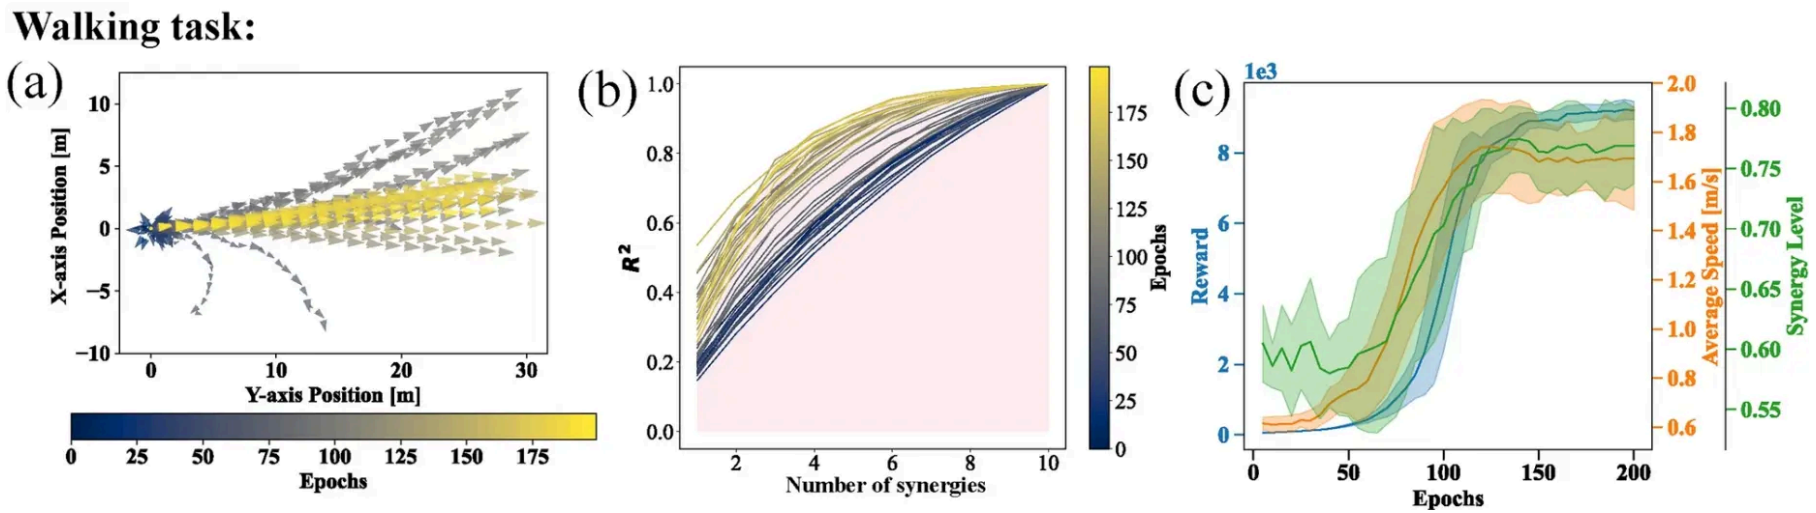



Neural decoding of synergistic behavior of spinal motoneurons (Capsi-Morales et al., 2025)

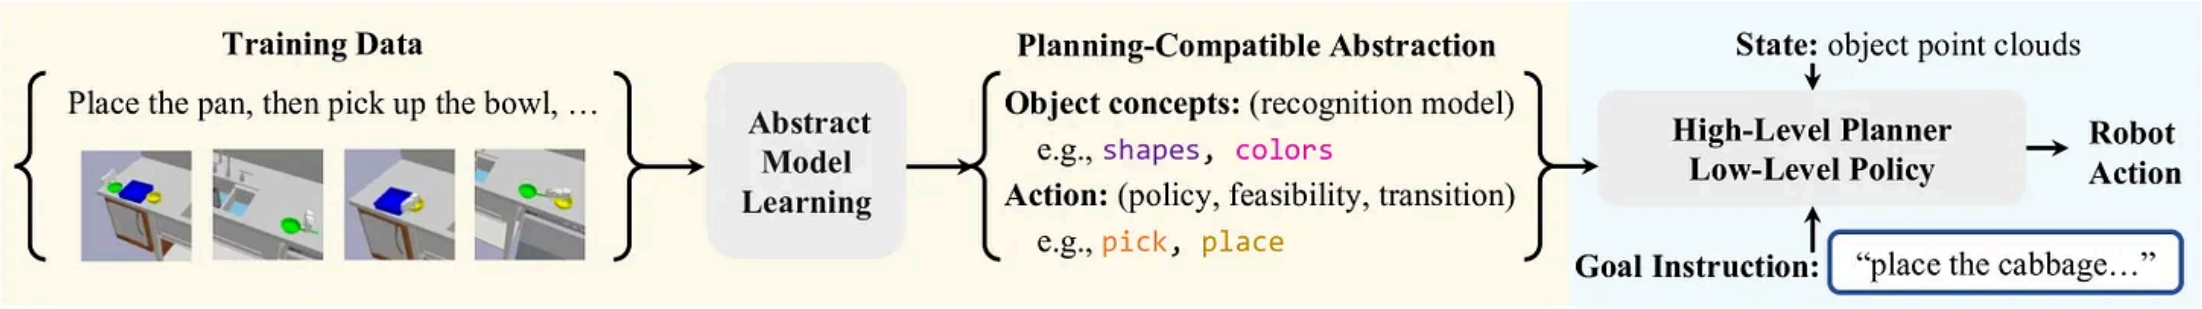



Motor neuron activities exhibit a synergistic organization in complex tasks; large dimensionality reduction in motor neuron outputs (Tanzarella et al., 2021)

# Learned Models Exhibit Coarse Action Primitives




With only energy and the motor performance in the reward, synergies (via PCA on joint torques) emerge (Li & Hayashibe, 2025)




Learning planning abstractions in sequential decision-making domains from language-annotated demonstrations (Liu et al., 2024)

# What We Want: Dynamical Primitives



Fish with no active control (Beal et al., 2006)

## Embrace dynamics

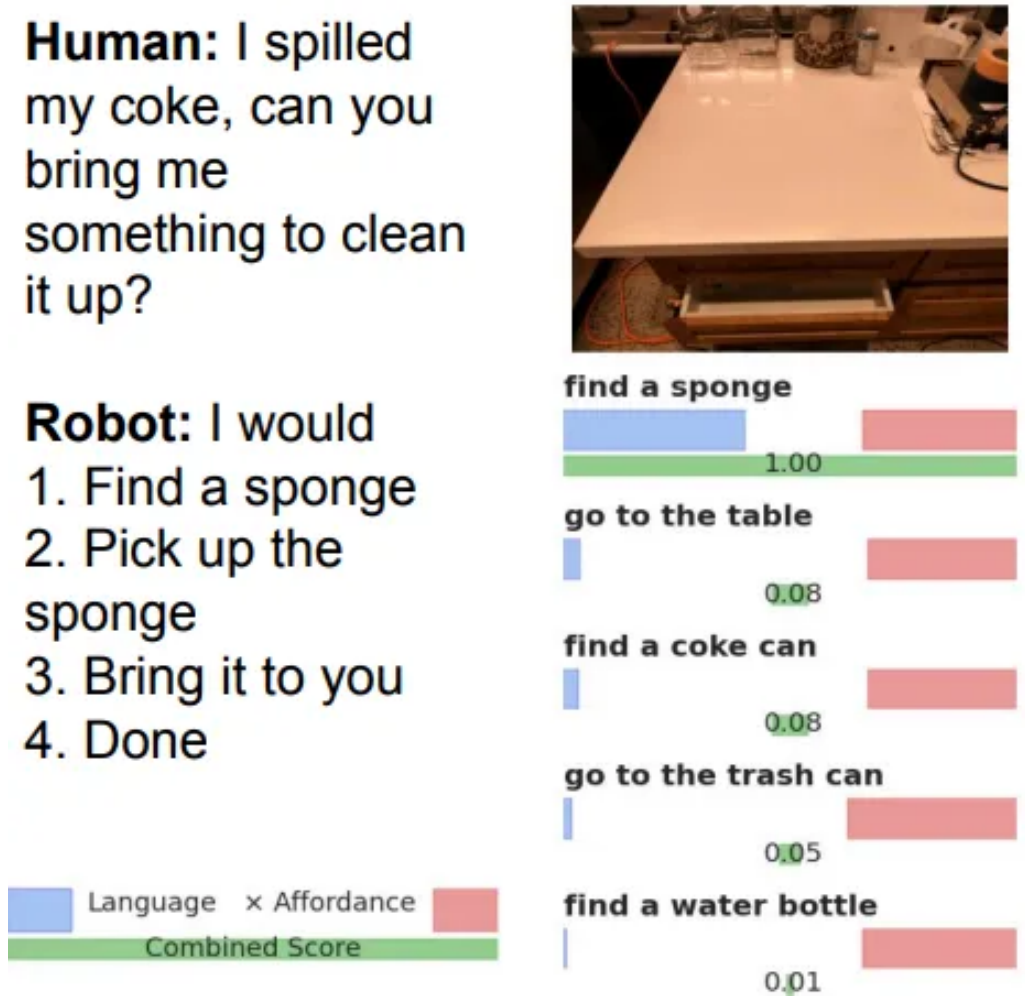
- Kangaroo hopping vs. human running need different strategies
- Unconventional robot designs
- Highly dynamic behaviors like gymnastics
- Mechanical intelligence is invisible when training from video data

## Dynamical stability and robustness

- Adapt to terrain disturbances (Lupu et al., 2024)
- The “ICL” equivalent for physical work

> *Cannot learn from imitation/pre-training!*

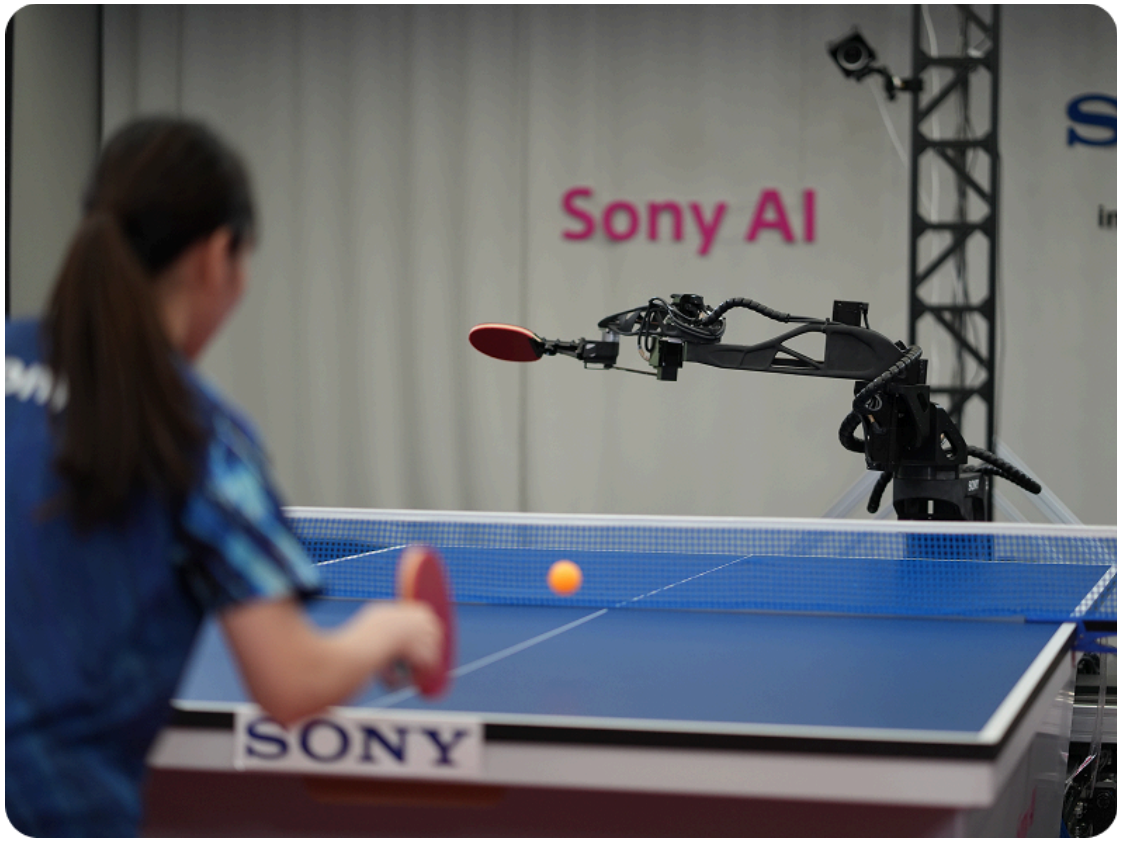
# What We Want: Explicit Sequential Composability




- "Emergent" language-driven sequencing → explicit, guaranteed, and dynamical sequential composition
- Dynamics pose correctness challenges: e.g. `walk` → `long jump` not feasible; `walk` → `run` → `long jump` is

SayCan (Ahn et al., 2022), similar to pi0.7 (Intelligence et al., 2026)

# What We Want: Explicit Parallel Composability




Sony Ace elite-level table tennis robot (Durr et al., 2026).

- E.g. combine upper and lower-body skills
- Poses similar dynamic stability challenges

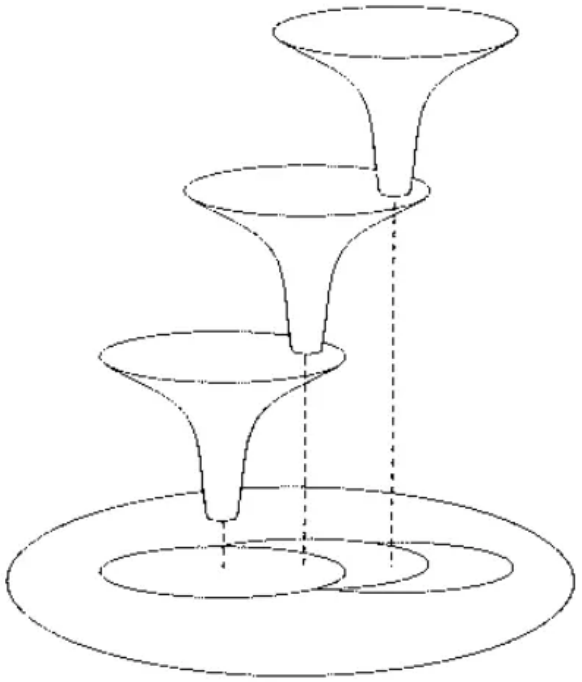
# Building Composable Action Primitives

# Formal Sequential Composition using Dynamical Models

"Buegler" robot palming, catching, juggling

(Burridge et al., 1999)





- Formal guarantee of composability
- Idea: each controller's goal state inside domain of attraction of next

# Empirical Dynamical Parallel Composition




"Control in 3 parts" (Raibert, 1986)

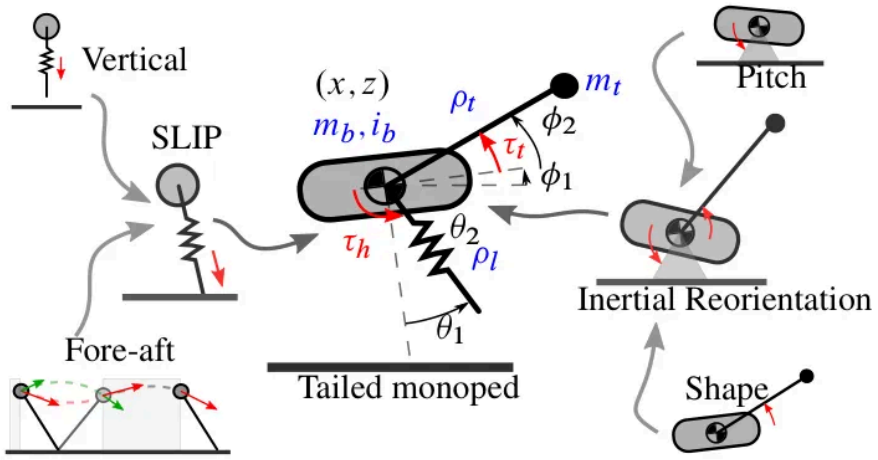
- Control of hopping height
- Control of forward speed
- Control body attitude

**Control of Running Decomposed into Three Parts**

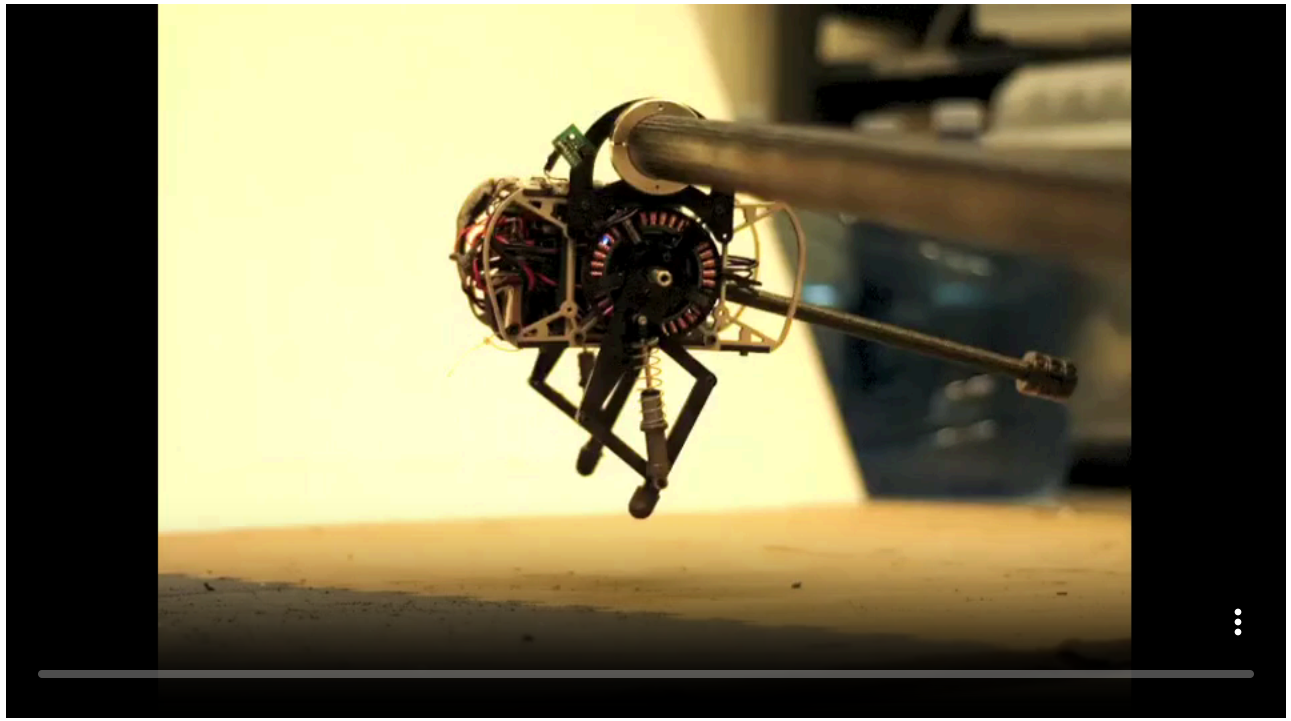
The control system we explored for the planar one-legged machine treats hopping, forward speed, and body attitude as three separate control problems. One part of the control excites the hopping motion and regulates

Specific to a particular type of design – would not work with a high-inertia leg.

# Formalized Parallel Composition for Tailed Hopping



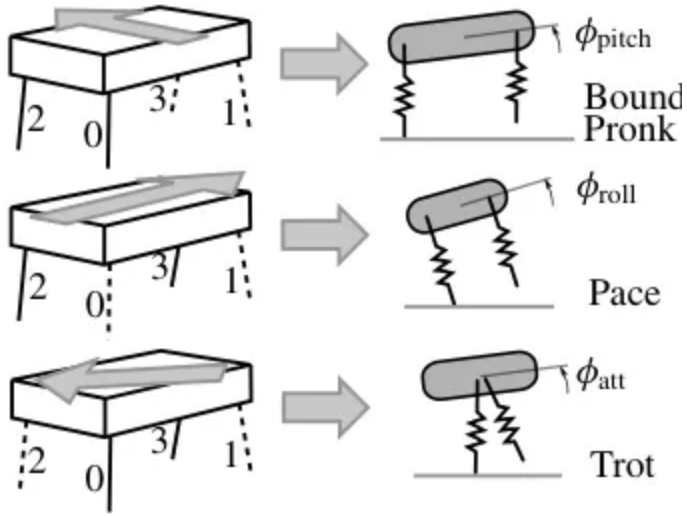

(De & Koditschek, 2015) (Shamsah et al., 2018)



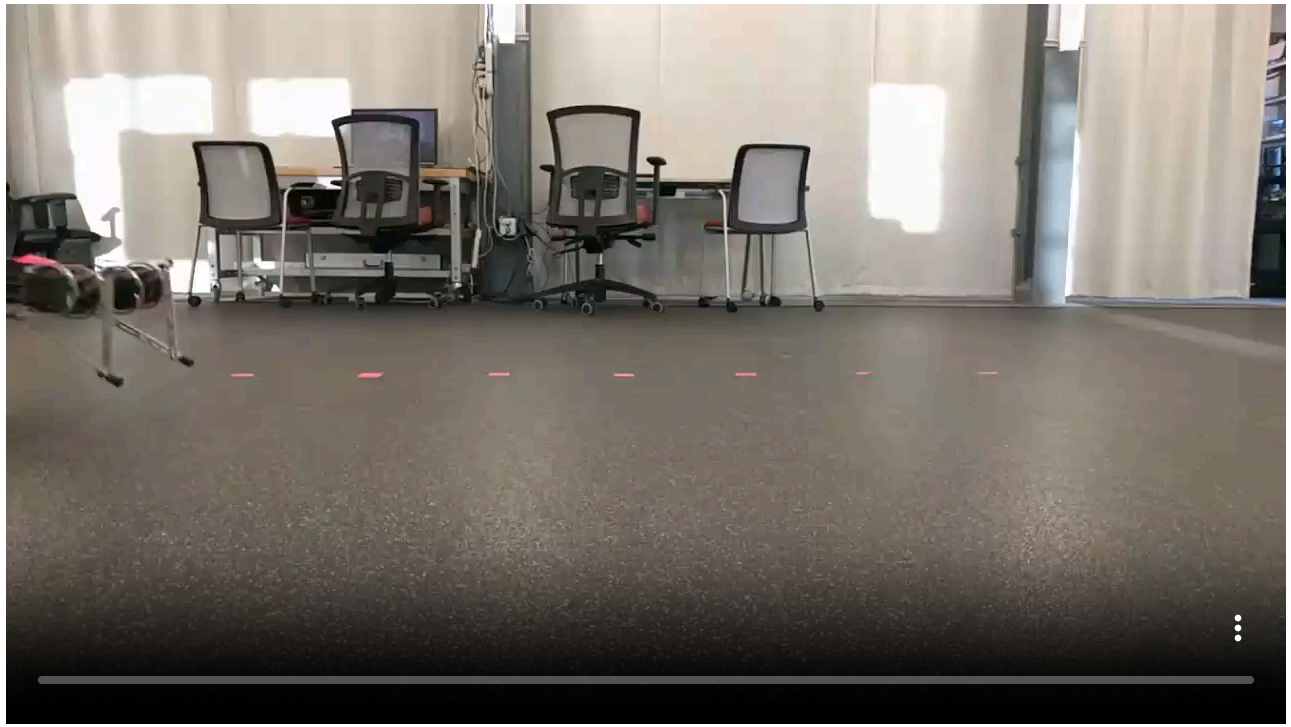
## Template-based control

- Projection: $p_T = \Pi(p)$
- Template dynamics: $\dot{x}_T = f(x_T, u_T; p_T)$
- Anchoring (controller embedding): $u = \Pi^{\dagger}(u_T)$

# Minitaur Vertical Hopping Compositions




Virtual biped permutations (De & Koditschek, 2018)



- Decoupled, decentralized control + mechanical coordination
- **Embodiment transfer** between Minitaur/Jerboa (different projection & embedding)
- **Task transfer** (generate different gaits)

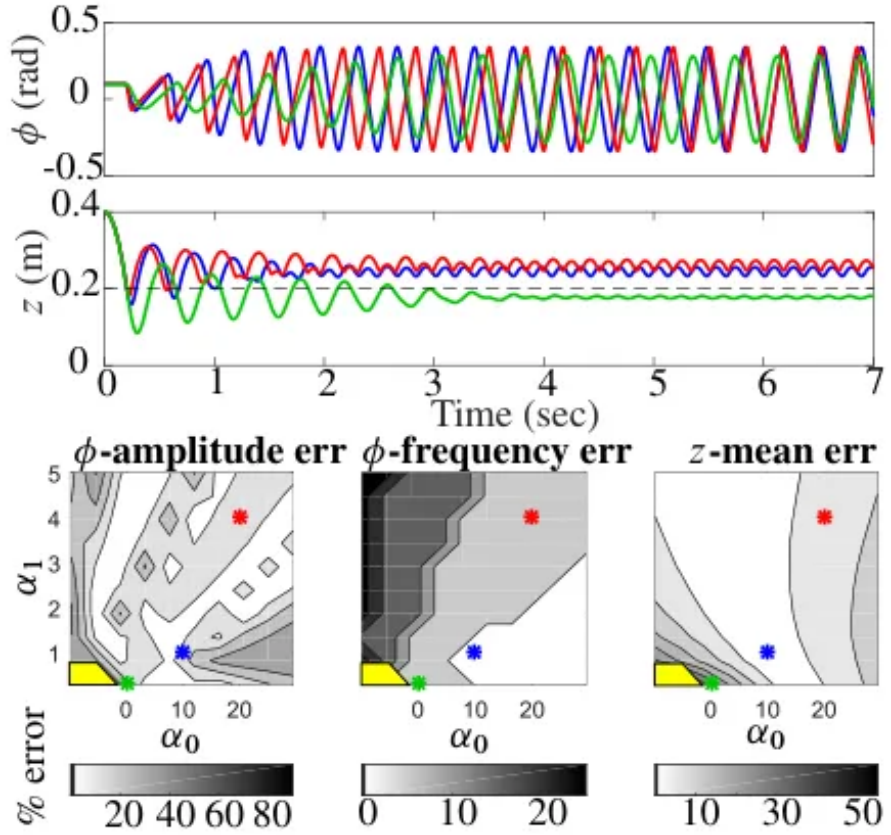
# Minitaur Vertical Hopping Model-Robustness

## Robust to modeling errors

- $\alpha$ non-dimensionalized parameters varied $\sim 10\times$
- Similar to RL with aggressive domain randomization
- **Weakness of model-based optimal control methods**

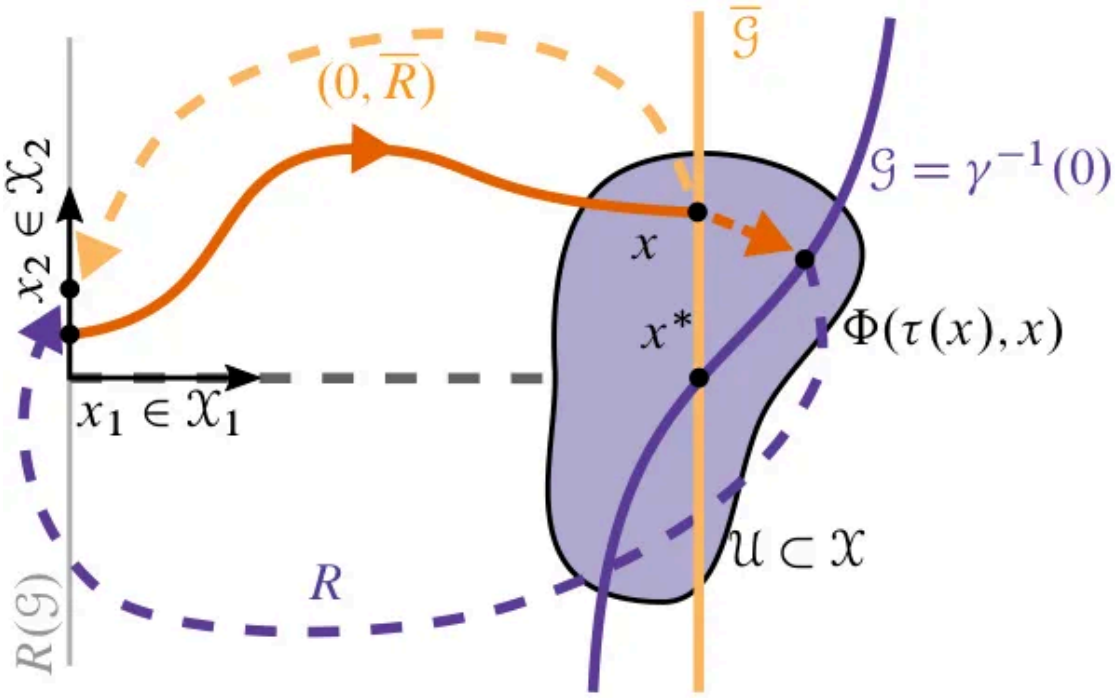
## Why does it work?

- Simplicity of $f$ and the (lack of) sensitivity of the $\Pi$ function




Parametric robustness (De & Koditschek, 2018)

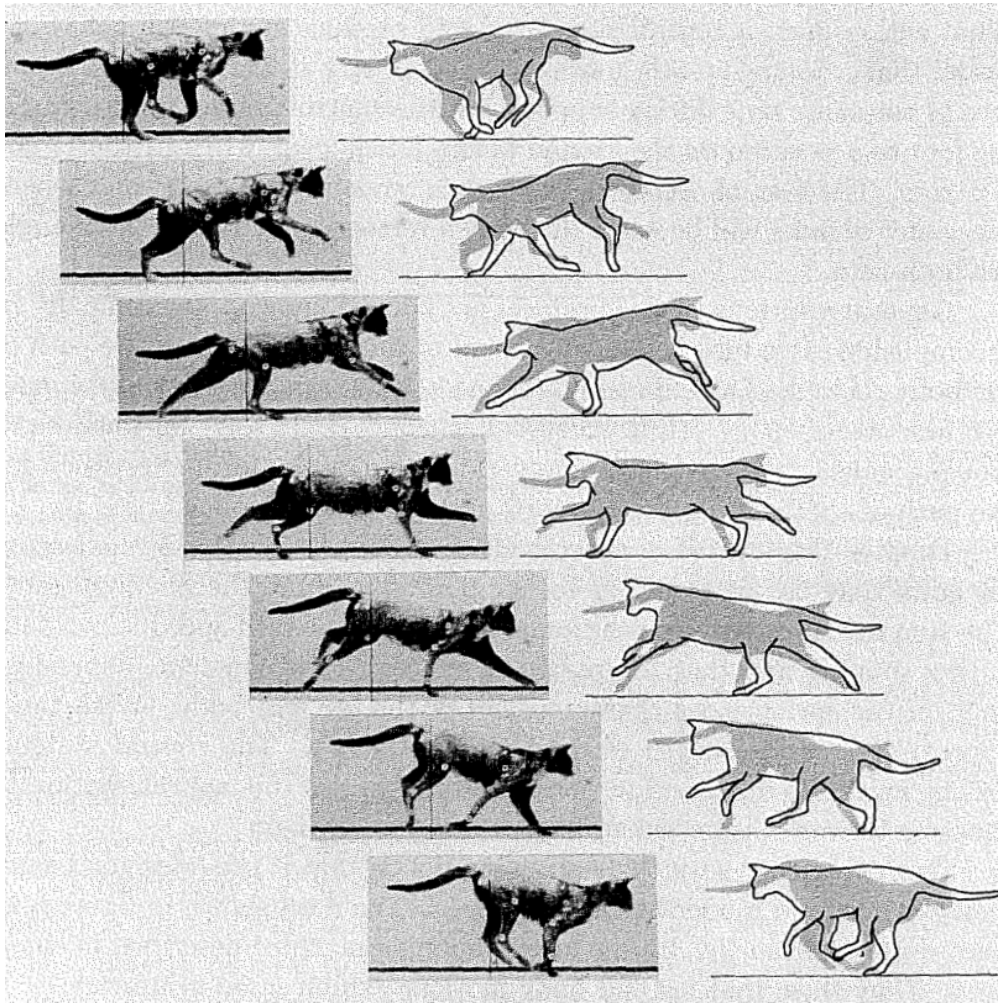
# “Correct By Design”: Hybrid Averaging




Hybrid averaging theory to find asymptotic behavior of composed cyclic systems
(De et al., 2018)

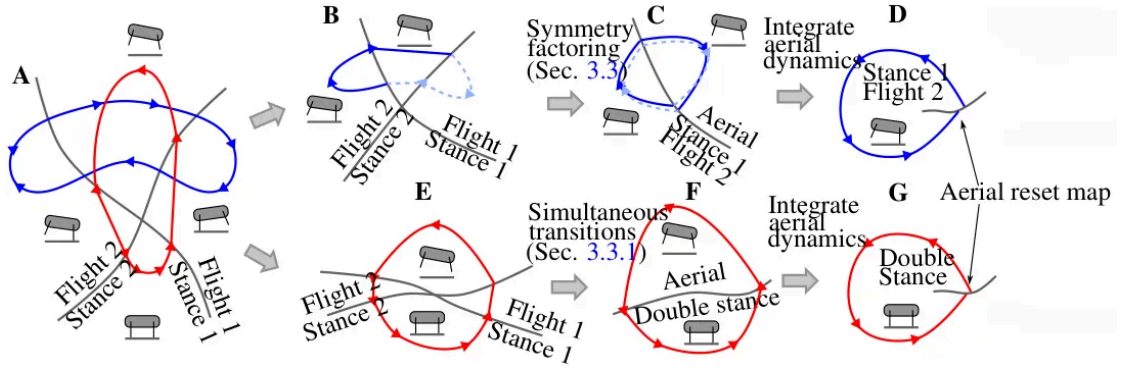
- **Analytically:** Reduce complex behaviors to simple ones
- **Synthetically:** Allows empirical composition to work

# “Correct By Design”: Time-Reversal Symmetry



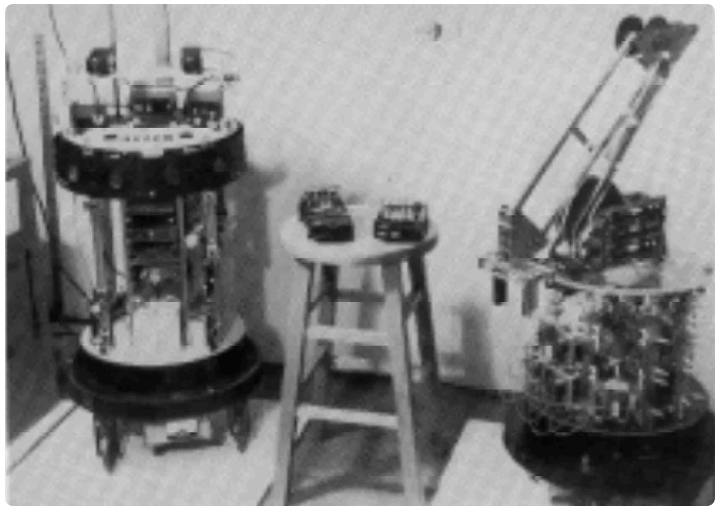
Gallopping cat image from (Raibert, 1986)



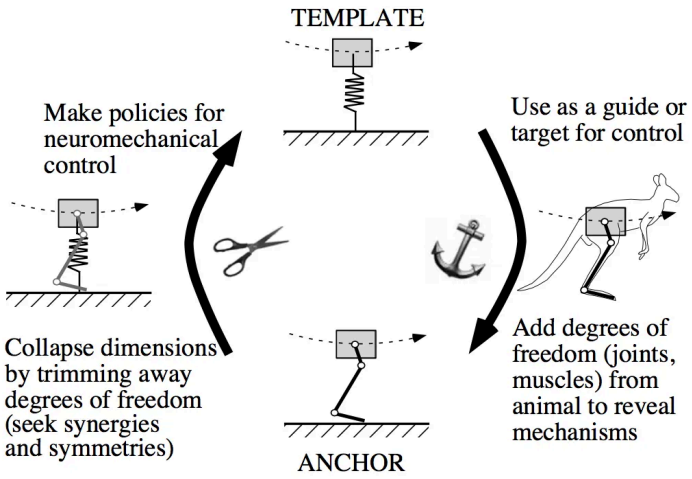



Dimension reduction using time-reversal symmetry complements averaging.

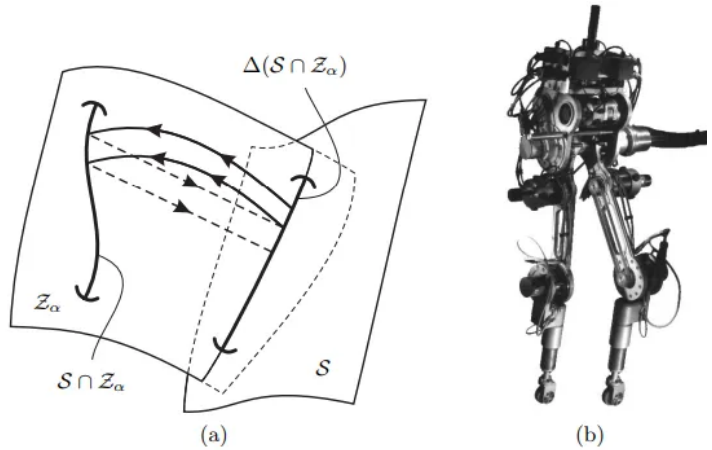
# Hierarchical Control Modularity



Subsumption architecture (Brooks, 1991)



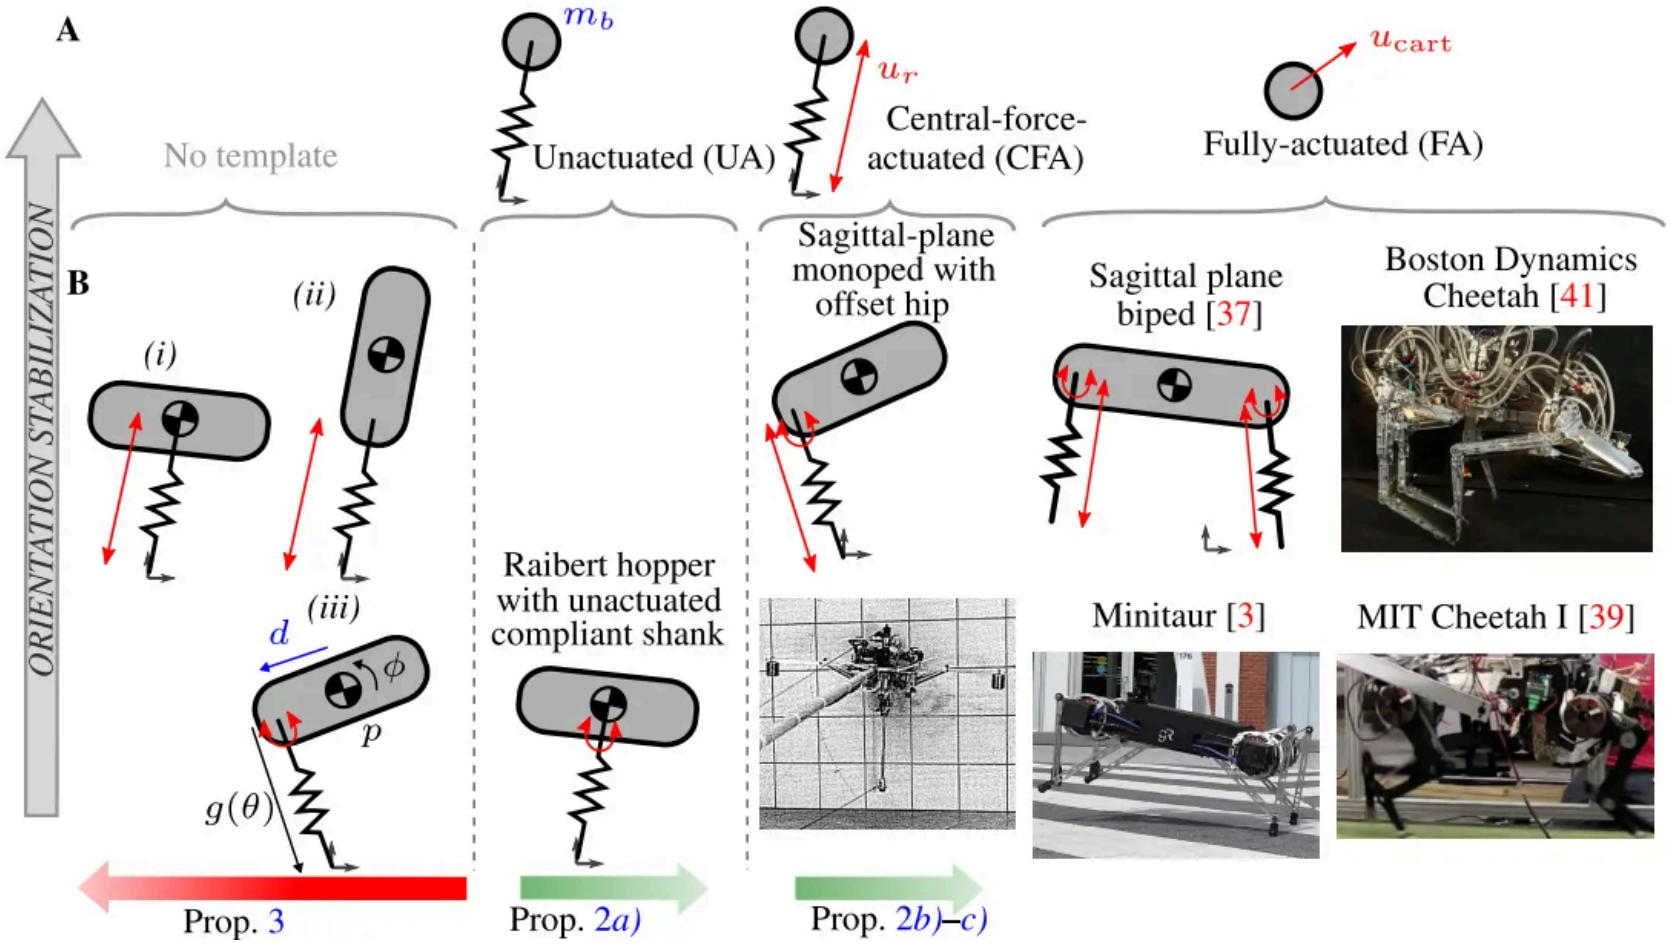

Template and anchors to describe the behavior of animals and robots (Full & Koditschek, 1999)




Hybrid Zero Dynamics formalism of dimension reduction with posture control (Westervelt et al., 2003).

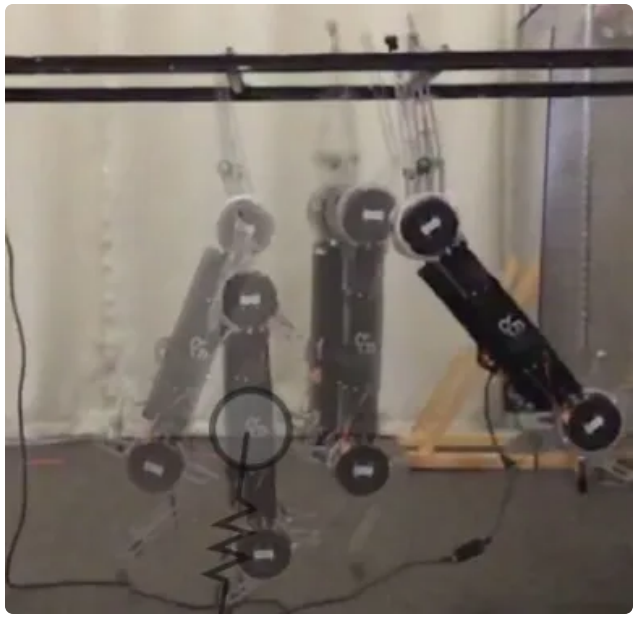
- Alternative to a `Sense -> Plan -> Act`, or a linear VLA architecture
- Explicit hierarchies facilitate multi-timescale control, embodiment transfer

# Center-of-Mass (CoM) Primitives Guaranteed to Appear

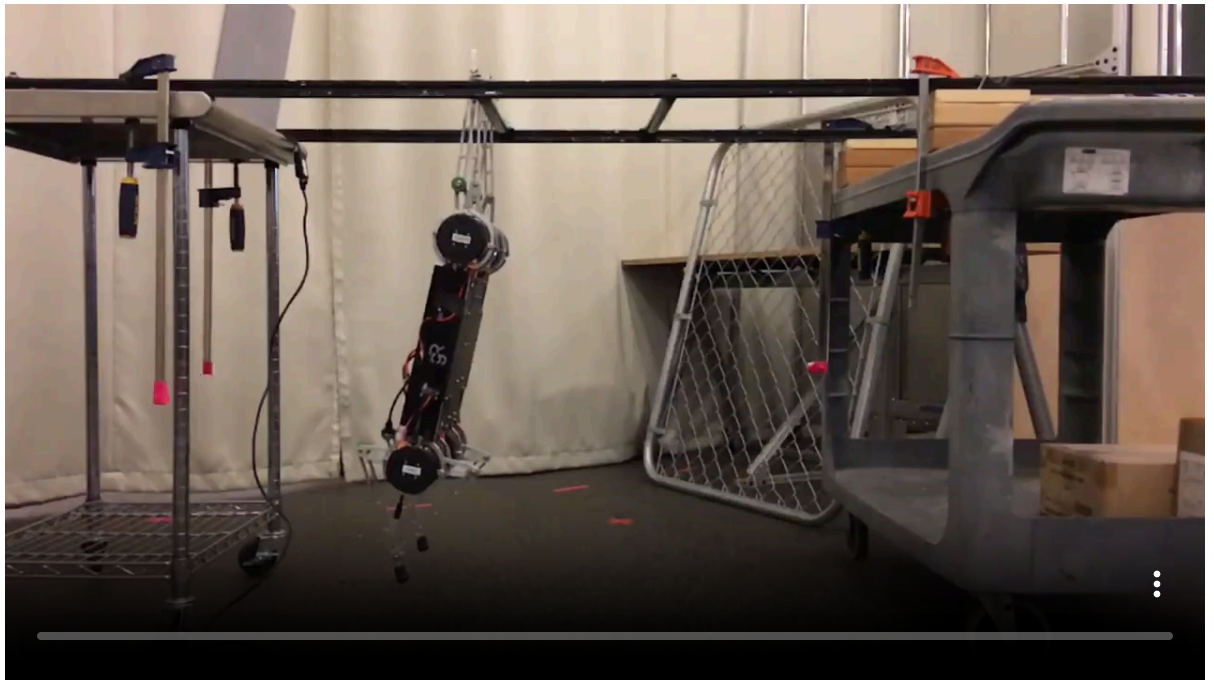



In low-dimensional sagittal-plane models, posture control results in necessary dimension reduction to a specific form of CoM dynamics (De et al., 2022)

# Embed CoM Primitive Task Controller



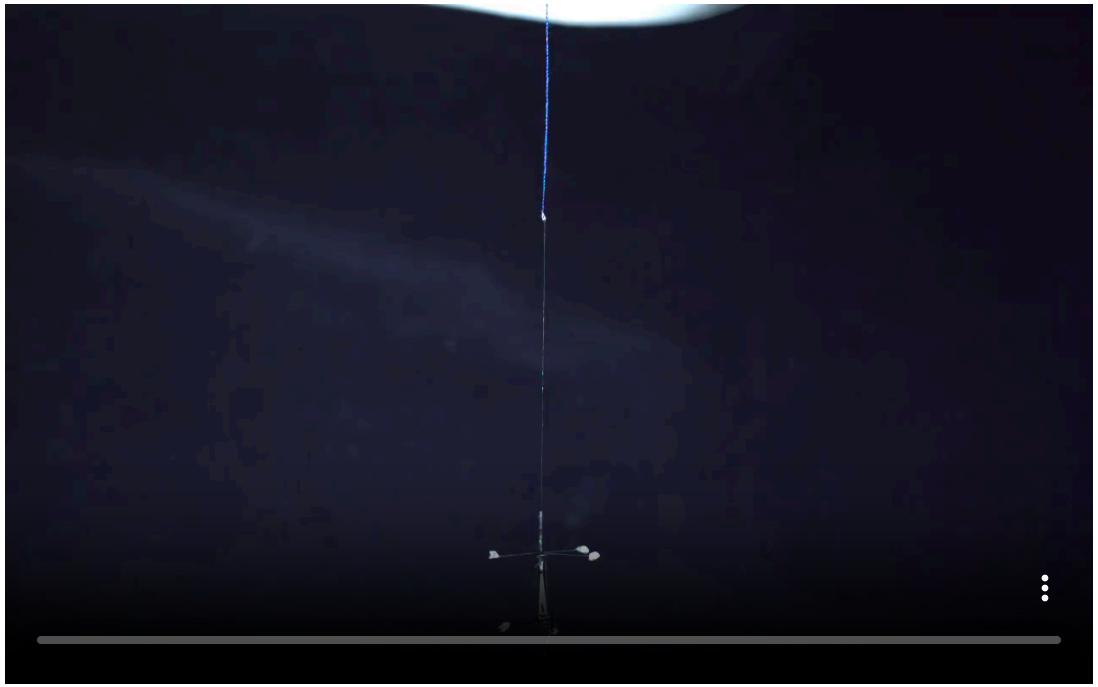
Mode-reactive templates (De et al., 2022)

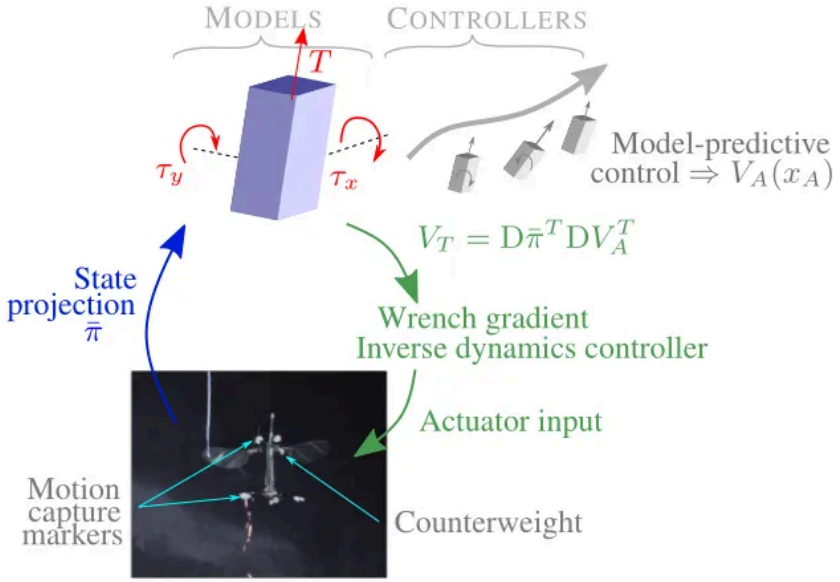


```
Analysis:  Full model        ->    postural dimension reduction    ->       CoM dynamics
Synthesis: Robot control     <-           posture control          <-   CoM task control
```

> *"Automated"* $\Pi$ *function in certain settings*

# Inverse Dynamics Action Primitive: MPC for RoboBee

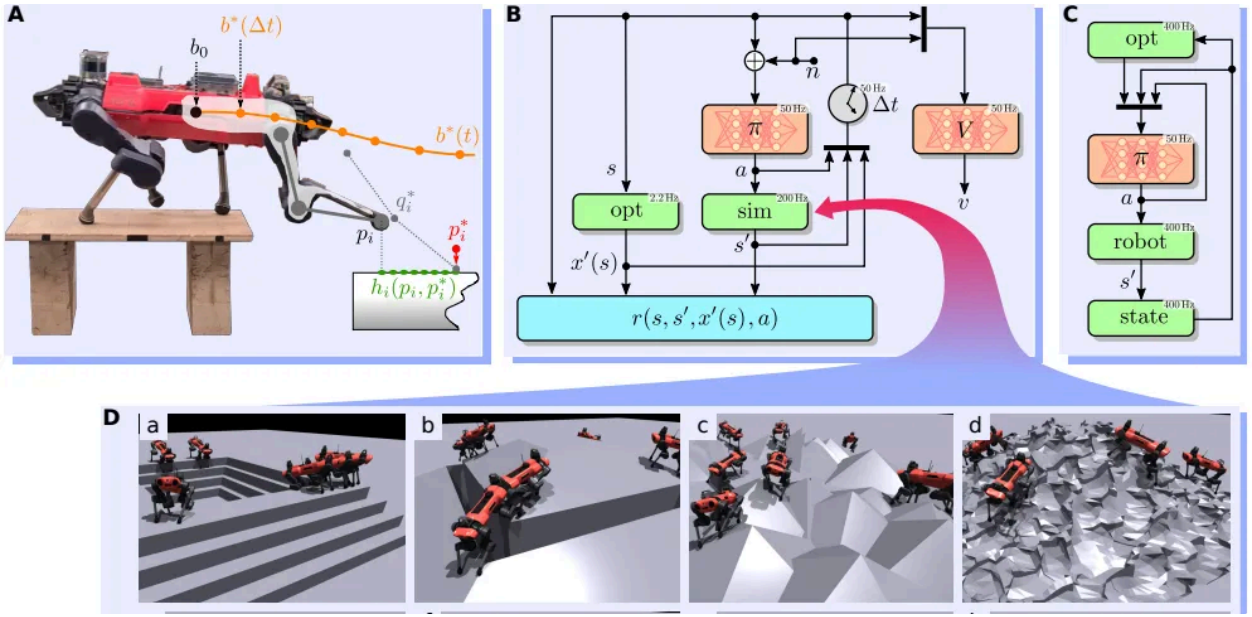




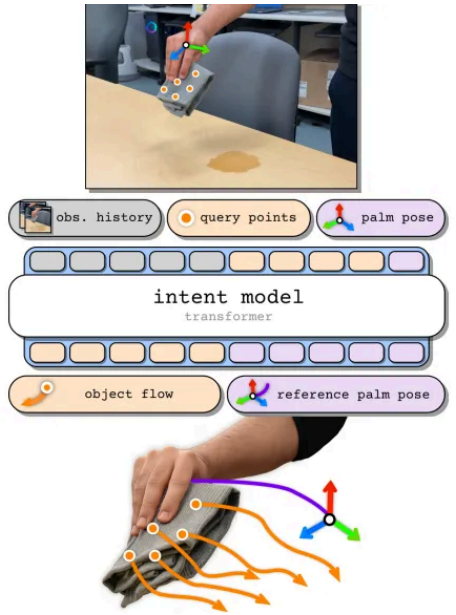

Hierarchical MPC for RoboBee (De et al., 2022)

- MPC (higher-level): online generalization
- Inverse Dynamics (lower-level): model-robustness, embodiment transfer

# RL Action Primitive (Fixed Interface)




Traj opt → RL tracking – Deep Tracking control (DTC) traj opt → RL tracking (Jenelten et al., 2024)

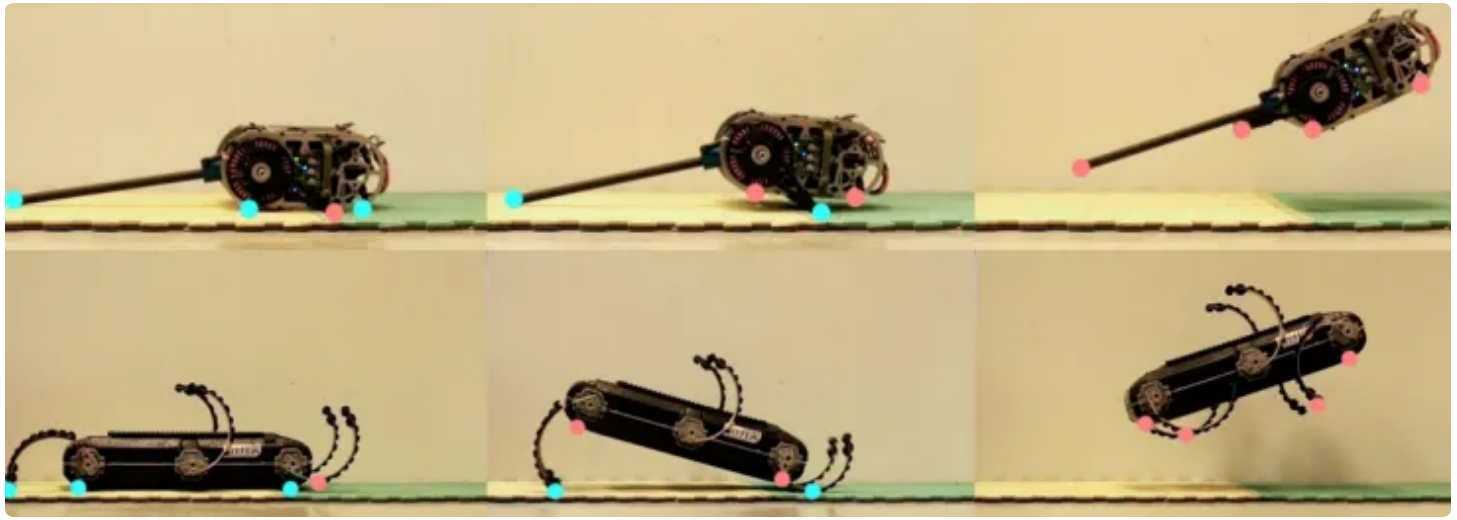



Intent model → RL tracking – LUCID (Gupta et al., 2026)

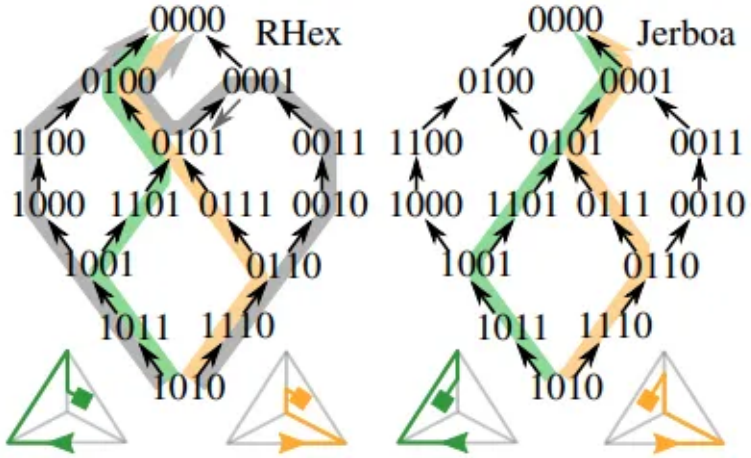
- RL can allow for faster synthesis of robust primitives when sufficiently specified
- Dynamical, embodied, robust, though not guaranteed to be composable

# Symbolic Planning for Leaping

- Contact-heavy planning with torque control ⟶ difficult online (Posa et al., 2014)
- Characterize “letters” ⟶ cheap online planning



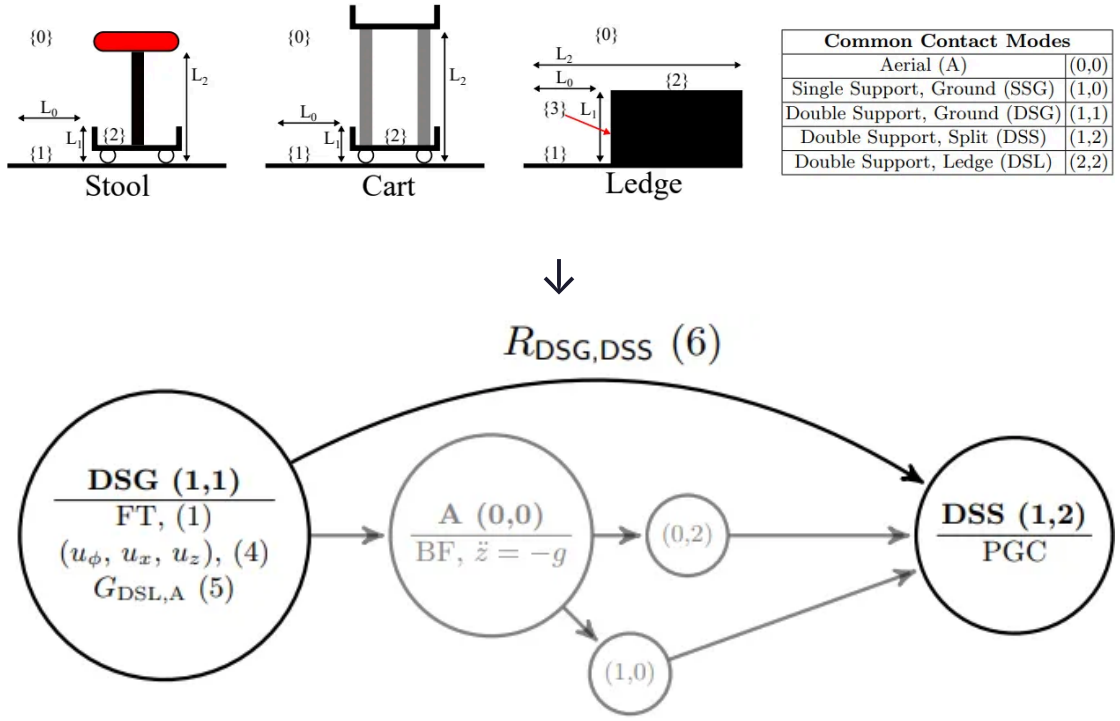
Leaping behaviors built by composing sequences of hardcoded transitions (Brill et al., 2015)



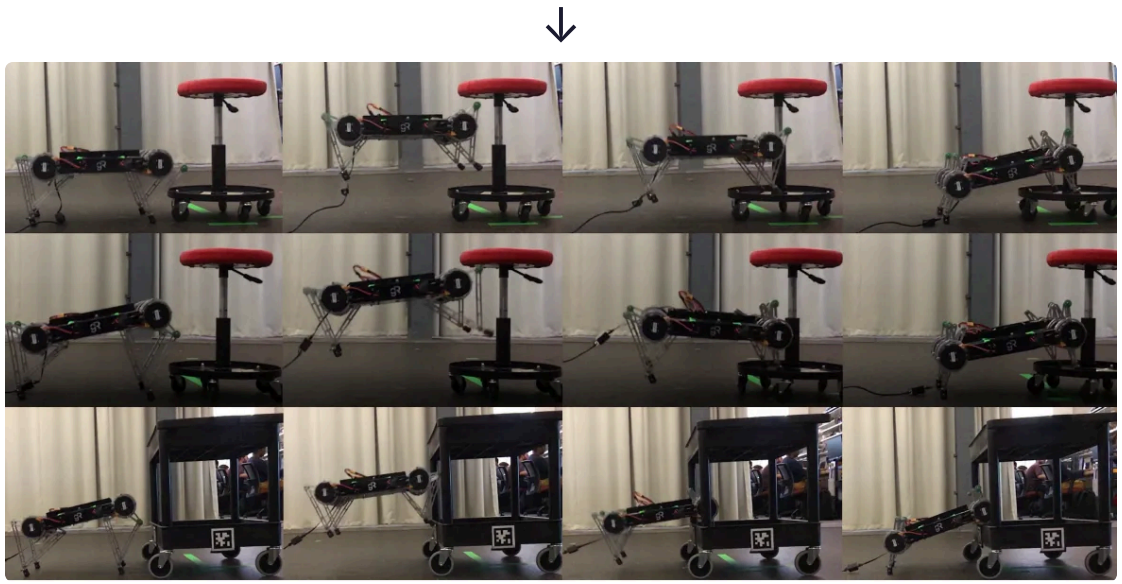

Ground reaction complex (GRC)

# Symbolic Planning for Manipulation

- Contact-heavy planning with torque control ⟶ difficult online (Mordatch et al., 2012)
- Characterize “letters” ⟶ cheap online planning




| Common Contact Modes | |
|---|---|
| Aerial (A) | (0,0) |
| Single Support, Ground (SSG) | (1,0) |
| Double Support, Ground (DSG) | (1,1) |
| Double Support, Split (DSS) | (1,2) |
| Double Support, Ledge (DSL) | (2,2) |



(Topping et al., 2019)

# What Have We Built So Far?

Action primitives that are...

**Dynamic**

Hopping, leaping, running, etc.; utilizing passive dynamics

**Generative**

Online recombination sequentially, parallelly

**Efficient**

Traj opt → cheap online discrete symbol-space planning

**Transferable**

Robust to modeling errors; swappable body params

We (still) need an architecture capable of synthesizing, as well as autonomously tasking and composing primitives

# Open Questions

# Best of Both Worlds: Semantic FM + Dynamical Primitives

**Foundation model (high-level)**

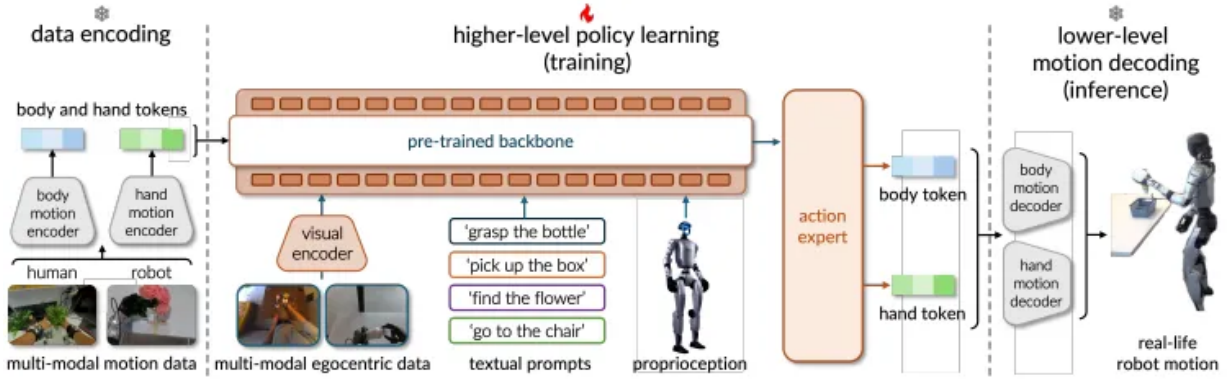
Semantic “common sense”, task decomposition, perception

↓

**Coarse action head (low-level)**

- Online generation of novel behaviors via primitive composition
- Few-shot learning via symbolic planning or online RL over primitives
- Dynamic stability, robustness, cross-embodiment transfer

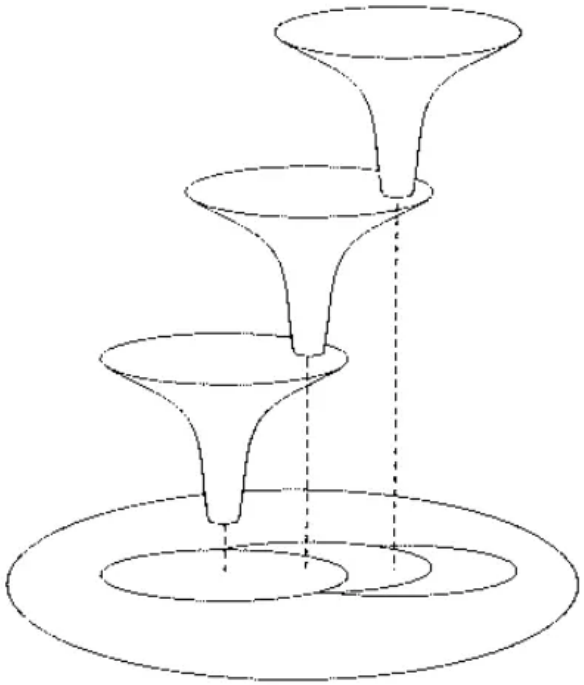
## Joint actions to latent interfaces



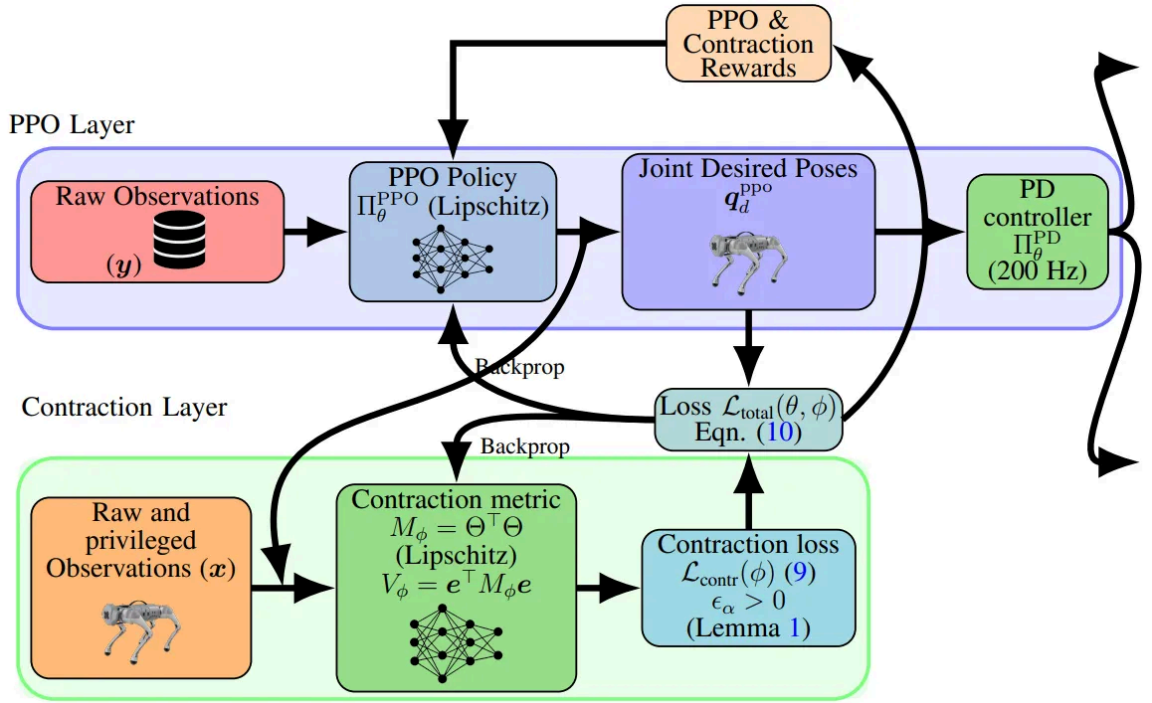

VioLA hierarchical connection of a “frozen encoder” with a latent (designed) interface (Albaba et al., 2026). Zero-shot locomotion success 100%, versus 16.7% for GR00T N1.7.

**Open challenge:** Aligning discrete/latent foundation model interfaces with formal composability when not trained end-to-end

# Formal Composability of RL Primitives



Sequential composition formalism provides a preimage condition for composing controllers (Burridge et al., 1999)



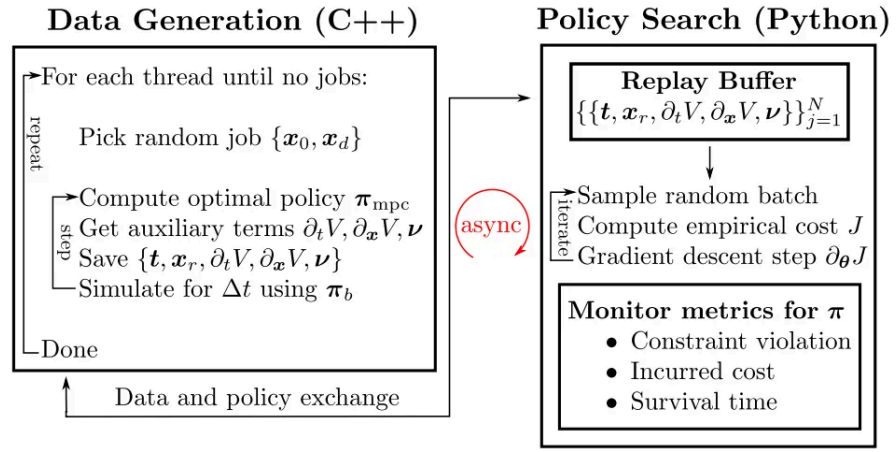

ContractionPPO adds a contraction loss during RL training for robustness (Zinage et al., 2026)

**Idea:** Add composability criteria to training loss for guaranteed online recombination

# Generalization's Computational Footprint

| | **Optimization-based** $\min J(x,u)\ \text{s.t. } g(x)=0$ | **RL** $a=\pi_\theta(s)$ |
|---|---|---|
| Online generalization | Strong | Weak |
| Computational requirement | High | Low |

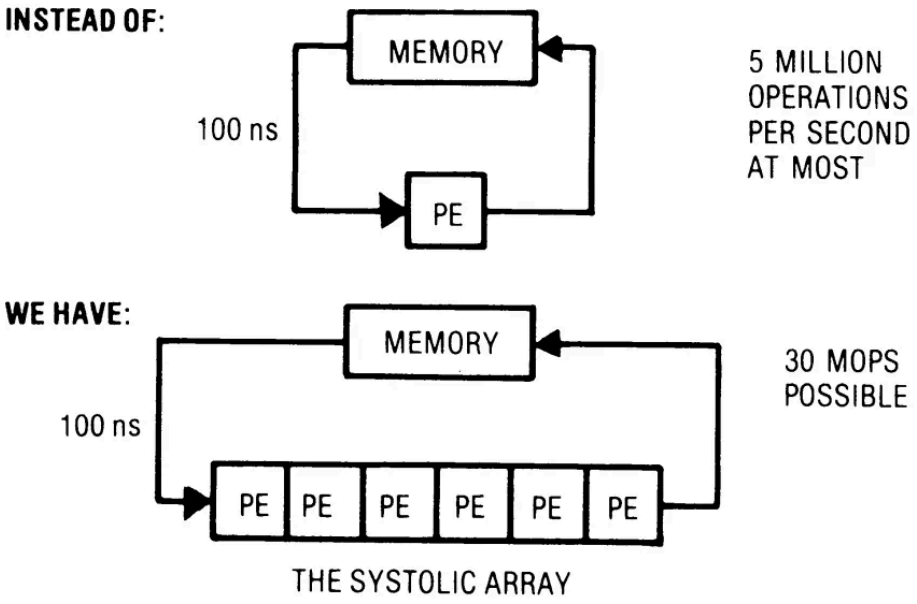
- Online MPC → NN distillation can be 10–70× (Khadiv et al., 2023) faster
- NN fails to generalize when OOD or into unseen contact modes where online MPC would re-solve (Reske et al., 2021)
- Model-based traj-opt inside RL training loop can be the primary bottleneck (Jenelten et al., 2024) (though not always (Bravo-Palacios et al., 2026)!)




MPC-Net "compiles" an MPC controller (Reske et al., 2021)

# Regularity vs. Co-Design Potential

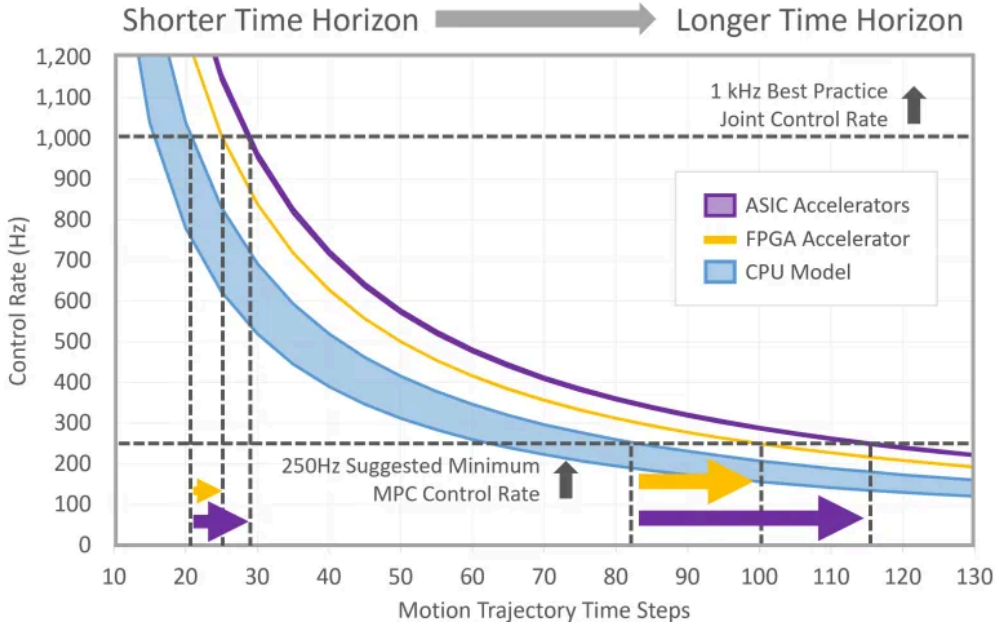
## DNN regularity allows codesign




Systolic arrays make DNN computation faster and more efficient (Kung et al., 1979)

- TPU (Jouppi et al., 2017) – 30–80× perf/W
- Serial Riccati recursion resists acceleration (Plancher & Kuindersma, 2018)

## ASICs: performant, but inflexible




Robotics codesign successes: Navion (Suleiman et al., 2019), Robomorphic ASIC (Neuman et al., 2021)

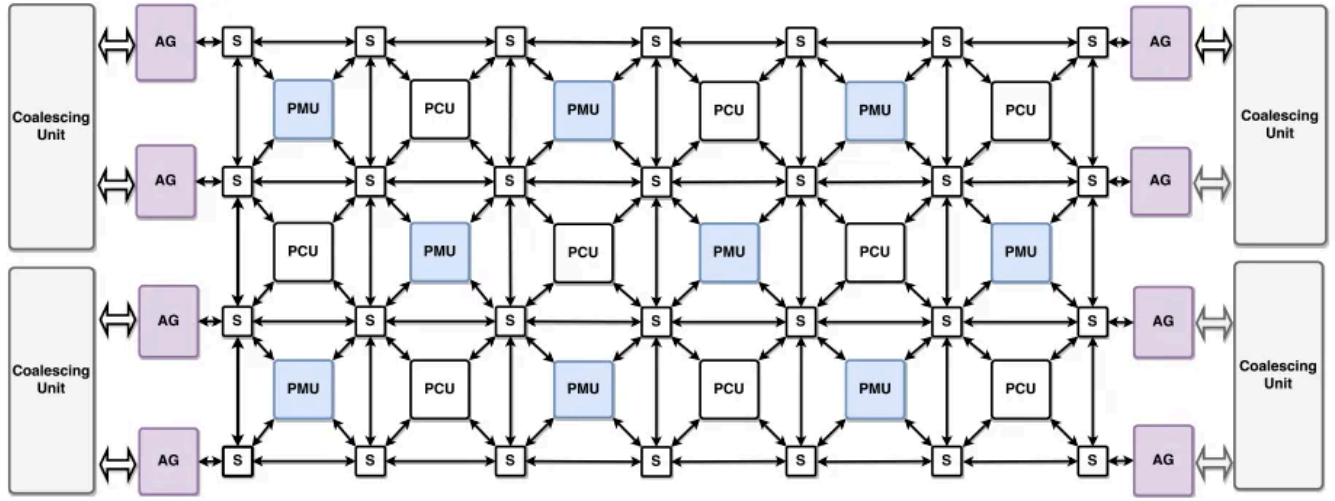
Huge perf/eff gains (e.g. 2 mW VIO; 40× dynamics grad), but inflexible

# Computer Architecture Affects Throughput and Energy

## Von Neumann vs. Dataflow

- **Stored-program bottleneck:** Data movement costs **10× (SRAM) to >150× (DRAM)** more energy than arithmetic (Horowitz, 2014)
- **MPC has low arithmetic intensity:** Dominated by memory movement (e.g., TinyMPC (Nguyen et al., 2024))
- **Data-flow and domain flow:** Organize computation around data movement (Omtzigt, 1990)

## Flexible dataflow hardware is emerging




Plasticine (Prabhakar et al., 2017), SambaNova RDU compile whole compute graphs into spatial pipelines on-chip (up to 77× perf/W)

### The opportunity for robotics:

- Systolic arrays (Kung et al., 1979) unlocked DNNs
- Domain-flow could unlock optimization-based control

# Conclusion

### What we want

- Robots that exhibit dynamical stability and robustness
- Can generate new behaviors online
- Generalize to new tasks and environments

### What we have so far

- Ways to hand-craft composable dynamical primitives from templates
- Ways to build RL primitives that may not be composable

### What we need to build

- Architectures that can flexibly and easily synthesize composable primitives
- Architectures that can compose them online to generate new behaviors

# Thank you!

More information: avikde.github.io

# References (1/9)


Ahn, M., Brohan, A., Brown, N., Chebotar, Y., Cortes, O., et al. "*Do as i can, not as i say: Grounding language in robotic affordances*." *arXiv preprint arXiv:2204.01691*, 2022.

Albaba, M., Bei\sswenger, J., Manasyan, A., Marta, D., Black, M., et al. "*VioLA: Learning Generalist Humanoid Control Policies from Human Data*." *arXiv preprint arXiv:2610.12435*, 2026.

Beal, D., Hover, F., Triantafyllou, M., Liao, J., and Lauder, G. "*Passive propulsion in vortex wakes*." *Journal of fluid mechanics*, 2006.

Blickhan, R., and Full, R. "*Similarity in multilegged locomotion: bouncing like a monopode*." *Journal of Comparative Physiology A*, 1993.

Bravo-Palacios, G., Zhang, J., Pestrikov, Z., Plancher, B., and Lew, T. "*TurboMPC: Fast, Scalable, and Differentiable Model Predictive Control on the GPU*." *arXiv preprint arXiv:2606.24039*, 2026.

Brill, A., De, A., Johnson, A., and Koditschek, D. "*Tail-assisted rigid and compliant legged leaping*." *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015.

Brooks, R. "*Intelligence without representation*." *Artificial intelligence*, 1991.

Buchanan, S., Pai, D., Wang, P., and Ma, Y. "*Principles and Practice of Deep Representation Learning: or a Mathematical Theory of Memory*." *arXiv preprint arXiv:2606.06624*, 2026.

# References (2/9)


Burridge, R., Rizzi, A., and Koditschek, D. "*Sequential composition of dynamically dexterous robot behaviors*." *The International Journal of Robotics Research*, 1999.

Capsi-Morales, P., Barsakcioglu, D., Catalano, M., Grioli, G., Bicchi, A., et al. "*Merging motoneuron and postural synergies in prosthetic hand design for natural bionic interfacing*." *Science Robotics*, 2025.

Clark, A., and others. "*Whatever next? Predictive brains, situated agents, and the future of cognitive science*." *Behavioral and brain sciences*, 2013.

De, A., and Koditschek, D. "*Parallel composition of templates for tail-energized planar hopping*." *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015.

De, A., and Koditschek, D. "*Vertical hopper compositions for preflexive and feedback-stabilized quadrupedal bounding, pacing, pronking, and trotting*." *The International Journal of Robotics Research*, 2018.

De, A., Burden, S., and Koditschek, D. "*A hybrid dynamical extension of averaging and its application to the analysis of legged gait stability*." *The International Journal of Robotics Research*, 2018.

De, A., McGill, R., and Wood, R. "*An efficient, modular controller for flapping flight composing model-based and model-free components*." *The International Journal of Robotics Research*, 2022.

De, A., Topping, T., Caporale, J., and Koditschek, D. "*Mode-reactive template-based control in planar legged robots*." *IEEE Access*, 2022.

# References (3/9)


Dorovatas, V., Schwerin, M., Bagdanov, A., Caccia, L., Carta, A., et al. "*Position: Modular Memory is the Key to Continual Learning Agents*." *arXiv preprint arXiv:2603.01761*, 2026.

Durr, P., El Gheche, M., Maeda, G., Mukai, N., Takahashi, N., et al. "*Outplaying elite table tennis players with an autonomous robot*." *Nature*, 2026.

EKA-Robotics. "*EKA Robotics*." 2026. [link]

Friston, K. "*The free-energy principle: a unified brain theory?*." *Nature reviews neuroscience*, 2010.

Full, R., and Koditschek, D. "*Templates and anchors: neuromechanical hypotheses of legged locomotion on land*." *Journal of experimental biology*, 1999.

Gupta, H., Shi, G., and Yuan, W. "*LUCID: Learning Embodiment-Agnostic Intent Models from Unstructured Human Videos for Scalable Dexterous Robot Skill Acquisition*." 2026. [link]

Haarnoja, T., Ha, S., Zhou, A., Tan, J., Tucker, G., et al. "*Learning to walk via deep reinforcement learning*." *arXiv preprint arXiv:1812.11103*, 2018.

He, C., Liu, X., Camps, G., Bruno, J., Sartoretti, G., et al. "*Demystifying robot diffusion policies: Action memorization and a simple lookup table alternative*." *The Fourteenth International Conference on Learning Representations*, 2026.

# References (4/9)


Horowitz, M. "*Computing's energy problem (and what we can do about it).*" *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, 2014.

Ibarz, J., Tan, J., Finn, C., Kalakrishnan, M., Pastor, P., et al. "*How to train your robot with deep reinforcement learning: lessons we have learned.*" *The International Journal of Robotics Research*, 2021.

Intelligence, P., Ai, B., Amin, A., Aniceto, R., Balakrishna, A., et al. "*pi-0.7: a Steerable Generalist Robotic Foundation Model with Emergent Capabilities.*" *arXiv preprint arXiv:2604.15483*, 2026.

Jenelten, F., He, J., Farshidian, F., and Hutter, M. "*DTC: Deep tracking control.*" *Science Robotics*, 2024.

Jouppi, N., Young, C., Patil, N., Patterson, D., Agrawal, G., et al. "*In-datacenter performance analysis of a tensor processing unit.*" *Proceedings of the 44th Annual International Symposium on Computer Architecture (ISCA)*, 2017.

Khadiv, M., Meduri, A., Zhu, H., Righetti, L., and Sch\""olkopf, B. "*Learning Locomotion Skills from MPC in Sensor Space.*" *Proceedings of The 5th Annual Learning for Dynamics and Control Conference (L4DC)*, 2023.

Kirk, R., Zhang, A., Grefenstette, E., and Rockt\""aschel, T. "*A survey of zero-shot generalisation in deep reinforcement learning.*" *Journal of Artificial Intelligence Research*, 2023.

Konidaris, G., Kuindersma, S., Grupen, R., and Barto, A. "*Robot learning from demonstration by constructing skill trees.*" *The International Journal of Robotics Research*, 2012.

# References (5/9)


Kung, H., Leiserson, C., and others. "*Systolic arrays (for VLSI).*" *Sparse Matrix Proceedings 1978*, 1979.

Lake, B., and Baroni, M. "*Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks.*" *International conference on machine learning*, 2018.

Lake, B., Ullman, T., Tenenbaum, J., and Gershman, S. "*Building machines that learn and think like people.*" *Behavioral and Brain Sciences*, 2017.

Legg, S., and Hutter, M. "*Universal intelligence: A definition of machine intelligence.*" *Minds and machines*, 2007.

Li, G., and Hayashibe, M. "*Motor synergy and energy efficiency emerge in whole-body locomotion learning.*" *Scientific Reports*, 2025.

Limisiewicz, T., Pagnoni, A., Iyer, S., Lewis, M., Mehta, S., et al. "*Compute optimal tokenization.*" *arXiv preprint arXiv:2605.01188*, 2026.

Liu, W., Chen, G., Hsu, J., Mao, J., and Wu, J. "*Learning planning abstractions from language.*" *International Conference on Learning Representations*, 2024.

Lupu, E., Xie, F., Preiss, J., Alindogan, J., Anderson, M., et al. "*MAGIC VFM-meta-learning adaptation for ground interaction control with visual foundation models.*" *IEEE Transactions on Robotics*, 2024.

# References (6/9)


Merel, J., Botvinick, M., and Wayne, G. "*Hierarchical motor control in mammals and machines*." *Nature communications*, 2019.

Mordatch, I., Popovi\''c, Z., and Todorov, E. "*Contact-invariant optimization for hand manipulation*." *Proceedings of the ACM SIGGRAPH/Eurographics symposium on computer animation*, 2012.

Neuman, S., Plancher, B., Bourgeat, T., Tambe, T., Devadas, S., et al. "*Robomorphic computing: a design methodology for domain-specific accelerators parameterized by robot morphology*." *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2021.

Nguyen, K., Schoedel, S., Alavilli, A., Plancher, B., and Manchester, Z. "*TinyMPC: Model-Predictive Control on Resource-Constrained Microcontrollers*." *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024.

Omtzigt, E. "*Domain flow and streaming architectures*." *Proceedings of the International Conference on Application Specific Array Processors*, 1990.

Pagnoni, A., Pasunuru, R., Rodriguez, P., Nguyen, J., Muller, B., et al. "*Byte latent transformer: Patches scale better than tokens*." *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2025.

Peng, X., and Van De Panne, M. "*Learning locomotion skills using deeprl: Does the choice of action space matter?*." *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, 2017.

# References (7/9)


Plancher, B., and Kuindersma, S. "*A performance analysis of parallel differential dynamic programming on a GPU.*" *International Workshop on the Algorithmic Foundations of Robotics*, 2018.

Posa, M., Cantu, C., and Tedrake, R. "*A direct method for trajectory optimization of rigid bodies through contact.*" *The International Journal of Robotics Research*, 2014.

Prabhakar, R., Zhang, Y., Koeplinger, D., Feldman, M., Zhao, T., et al. "*Plasticine: A reconfigurable architecture for parallel patterns.*" *Proceedings of the 44th Annual International Symposium on Computer Architecture (ISCA)*, 2017.

Raibert, M. "*Legged robots that balance.*" *MIT press*, 1986.

Reske, A., Carius, J., Ma, Y., Farshidian, F., and Hutter, M. "*Imitation Learning from MPC for Quadrupedal Multi-Gait Control.*" *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021.

Shamsah, A., De, A., and Koditschek, D. "*Analytically-guided design of a tailed bipedal hopping robot.*" *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018.

Shannon, C. "*Prediction and entropy of printed English.*" *Bell system technical journal*, 1951.

Sleiman, J., Li, H., Adu-Bredu, A., Deits, R., Kumar, A., et al. "*ZEST: Zero-shot embodied skill transfer for athletic robot control.*" *Science Robotics*, 2026.

# References (8/9)


Suleiman, A., Zhang, Z., Carlone, L., Karaman, S., and Sze, V. "*Navion: A 2-mw fully integrated real-time visual-inertial odometry accelerator for autonomous navigation of nano drones*." *IEEE Journal of Solid-State Circuits*, 2019.

Sutton, R. "*The bitter lesson*." *Incomplete Ideas (blog)*, 2019.

Tanzarella, S., Muceli, S., Santello, M., and Farina, D. "*Synergistic organization of neural inputs from spinal motor neurons to extrinsic and intrinsic hand muscles*." *The Journal of Neuroscience*, 2021.

Topping, T., Vasilopoulos, V., De, A., and Koditschek, D. "*Composition of templates for transitional pedipulation behaviors*." *The International Symposium of Robotics Research*, 2019.

Udandarao, V., Prabhu, A., Ghosh, A., Sharma, Y., Torr, P., et al. "*No ``Zero-Shot'' Without Exponential Data: Pretraining Concept Frequency Determines Multimodal Model Performance*." *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.

Westervelt, E., Grizzle, J., and Koditschek, D. "*Hybrid zero dynamics of planar biped walkers*." *IEEE transactions on automatic control*, 2003.

Zador, A. "*A critique of pure learning and what artificial neural networks can learn from animal brains*." *Nature communications*, 2019.

Zahavy, T. "*Position: LLMs can't jump*." *Forty-third International Conference on Machine Learning Position Paper Track*, 2026.

# References (9/9)

Zinage, V., Harutyunyan, N., Verheyden, E., Hadaegh, F., and Chung, S. "*ContractionPPO: Certified Reinforcement Learning via Differentiable Contraction Layers*." *IEEE Robotics and Automation Letters*, 2026.