# From Primates to Robots:
# Emerging Oscillatory Latent-Space Dynamics for Sensorimotor Control

**Alexander Luis Mitchell (mitch@robots.ox.ac.uk)**
A2I & DRS, Oxford Robotics Institute, 23 Banbury Rd
Oxford, OX2 6NN, United Kingdom

**Oiwi Parker Jones (oiwi@robots.ox.ac.uk)**
A2I, Oxford Robotics Institute, 23 Banbury Rd
Oxford, OX2 6NN, United Kingdom

**Jun Yamada (jyamada@oxfordrobotics.institute)**
A2I, Oxford Robotics Institute, 23 Banbury Rd
Oxford, OX2 6NN, United Kingdom

**Wolfgang Merkt (wolfgang@robots.ox.ac.uk)**
DRS, Oxford Robotics Institute, 23 Banbury Rd
Oxford, OX2 6NN, United Kingdom

**Ioannis Havoutis (ioannis@oxfordrobotics.institute)**
DRS, Oxford Robotics Institute, 23 Banbury Rd
Oxford, OX2 6NN, United Kingdom

**Ingmar Posner (ingmar@robots.ox.ac.uk)**
A2I, Oxford Robotics Institute, 23 Banbury Rd
Oxford, OX2 6NN, United Kingdom

## Abstract

**Oscillatory signals drive reaching and locomotion for both robots and primates. Recent results in neuroscience have shown that periodic signals are present in the motor cortex of primates during rhythmic tasks such as locomotion as well as during linear movements notably reaching. Our works on learning latent motion representations in robotics revealed that oscillatory latent dynamics emerge automatically from training data for quadruped locomotion as well as manipulation tasks. Inspired by these works, we recreate the locomotion latent-spaces as well as create manipulation specific versions. With the latter we show that manipulation problems can be solved using periodic signals in a suitable latent-space. We see that these trajectories are reminiscent of those seen in the motor cortex. We artificially lesion the decoder of the locomotion model to understand how the correlations are captured. This results in deformation of both the oscillatory signals in the latent space and degradation in the locomotion trajectories.**

**Keywords:** Representation learning; machine learning; robotics; neuroscience; sensorimotor control

## Introduction

Understanding the intricacies of locomotion and manipulation is a key driver for both neuroscientists and roboticists. Indeed, recent work in neuroscience has revealed that locomotion in primates is controlled by oscillations in the monkey's motor cortex (Churchland et al., 2012). These oscillations are observed in low-dimensional projections of neural populations in the motor cortex, whilst the monkey is performing rhythmic movement such as swimming and walking. As a matter of fact, these oscillations are also found in the motor cortex when the primate is moving its limbs linearly during reaching.

In conjunction with the previous results, the utilisation of generative models in the field of robotics (Mitchell et al., 2022, 2022) has revealed that robot locomotion can be represented as oscillatory signals in low-dimensional representations. In our previous work (Mitchell et al., 2022, 2022) we create a structured latent-space for locomotion utilising a deep generative model in order to control a real quadruped robot. In particular, a variational auto-encoder (VAE) (Kingma & Welling, 2014; Rezende, Mohamed, & Wierstra, 2014) is utilised to create a low-dimensional representation of locomotion data (e.g. joint positions, torques, etc.) via an information bottleneck, which is the latent-space. It is found that the properties of locomotion such as cadence and footstep height are disentangled in the space and that cyclic signals injected within produce continuous locomotion.

The work in Churchland et al., 2012 broke new ground with their discovery that reaching is also periodic in the motor cortex. Similarly, the oscillations, which emerge in the latent space, are a consequence of their learnt structure. We thus posit that the cyclic signals in the motor cortex are also a result of order in the neural populations of the primate's motor cortex.

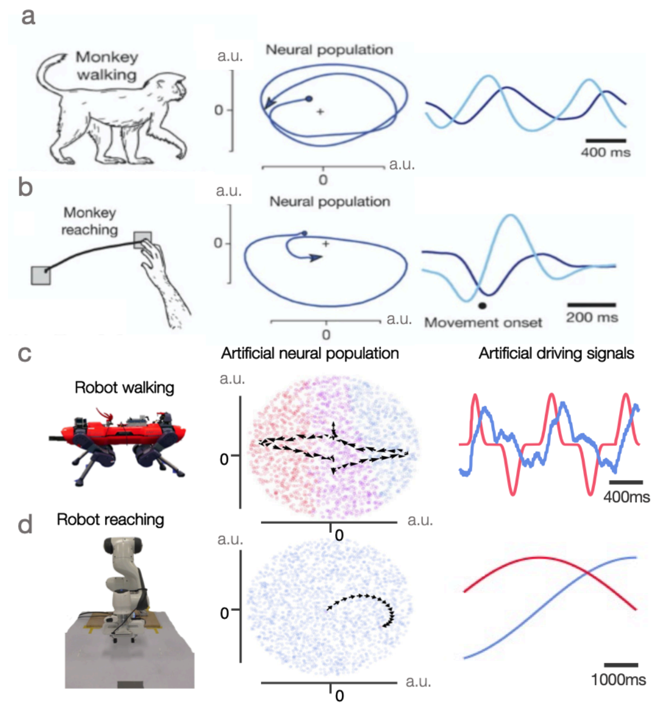


Figure 1: Oscillatory signals are found in neural populations of primate's motor cortex in both locomotion (panel (a)) and reaching (panel (b)). These two panels are reproduced with permission from Churchland et al., 2012. Quadruped locomotion can be solved via planning in a structured latent-space (panel (c)). The trajectories utilised are also cyclic. A suitably trained latent space for robot manipulation also facilitates solving reaching tasks (panel (d)).

To investigate this, we present the following paper, which is a summary of our previous work Parker Jones et al., 2022. Here, we explore the structure of the learnt latent-space for locomotion and show that key gait-specific properties emerge automatically. As a result, we are able to generate new movements unseen during training. In order to investigate if robotic manipulation is also cyclic in a learnt latent-space, we create another similar space and show that a low-dimensional representation with interpretable properties emerges. With this understanding, we inject oscillatory signals similar to those used during locomotion to solve the reaching tasks. When the latent-space signals for locomotion and manipulation are plotted and compared to those found in the motor cortex (Churchland et al., 2012) as seen in Figure 1, there is a visual similarity. Finally, we take the analogy between the latent space and the motor cortex further and artificially lesion the VAE's decoder and visualise the results. We find that the locomotion trajectories degrade in a predictable and repeatable way and that encoding the resulting locomotion reveals deformation of the cyclic signals in the latent space.

## Methods

We posit that the oscillations which exist in latent space and in the motor cortex are a result of their structures. Since we are unable to artificially construct a motor cortex in the lab, we created a learnt latent-space. Therefore, we describe how the latent-space is created; how we inspect and discover the structure; and finally, how the models are lesioned in order to inspect how correlations are captured by the models.

A VAE is trained for the quadruped and for the manipulator. The inputs to the VAEs' encoders consist of a history of robot states (e.g. joint positions, torques, contact forces) and the decoder output are the predicted next states for a preview horizon. The VAEs are trained using short trajectories of both the manipulator and the quadruped operating in their environments. For the manipulator, the ELBO loss is minimised. The quadruped continually makes and breaks contact as it takes steps resulting in discontinuously changing dynamics between swing versus stance. Therefore, the latent space is constrained by learnt multi-layer perceptrons (MLP) called performance predictors (PP) that predict which feet are in contact. This adds an additional binary cross-entropy loss term to the locomotion training loss during optimisation.

The latent-space structure is investigated in two ways: Firstly, trajectories from the test dataset are encoded and the resulting trajectories inspected. Secondly, a short oscillatory trajectory is injected into latent space and decoded. The decoded trajectory is sent to a tracking controller and the robots' movements are visualised.

Finally, we artificially impair the locomotion VAE in order to investigate how the latent-space trajectories map to the resulting robot trajectories. This is achieved by applying a cascading dropout to the decoder layers. This filter either passes the output of neurons in the decoder through or zeros them. The probability of zeroing the neuron output increases with time and once the filter zeros a neuron, it continues to do so – akin to neural degradation.

## Results

The learnt latent-spaces are straight-forward to inspect compared to the motor cortex. Therefore, we probe our latent-space structures to investigate the structure and compare the resulting trajectories to those observed in motor cortex. Additionally, we modulate the signals injected into the latent-space to discover how these affect the robot's movement. Finally, we impair the locomotion model and visualise how this causes the locomotion trajectories to degrade.

As seen in Figure 1, the locomotion trajectories in latent-space form a periodic cycle. We inject these trajectories into the latent space and find that the time period of the red signal in Figure 1 controls the cadence of the robot and the amplitude governs the footstep height. The blue signal is inferred from the robot states and controls the footstep length of the robot, and is $\pi/2$ out of phase with the red signal. The two signals plotted against each other form the limit cycle in Figure 1.

We repeat this experiment for manipulation and find that

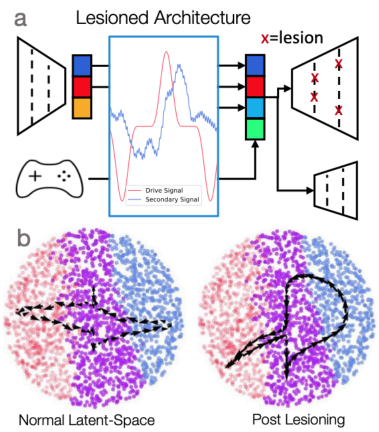


Figure 2: The VAE's decoder is lesioned and the trajectories decoded. These output trajectories are re-encoded and visualised. The latent-space trajectory has deformed after lesioning. The symmetrical structure has degraded.

two sine wave oscillations where the latter is also $\pi/2$ out of phase with the first can solve reaching tasks. This resembles what Churchland et al., 2012 have found in Figure 1 (b), where we see a cyclic pattern constructed from two slightly out of phase periodic signals for both reaching and locomotion. We emphasise that this is a curious resemblance and it is not possible for us to show that the similarity is more fundamental.

Lastly, we extend the analogy between the motor cortex and the latent space by artificially impairing the locomotion decoder. This is something that should not be done to the motor cortex on both practical and ethical grounds, but is possible for our *in-silico* models. We utilise the same oscillatory drive signals as in Figure 1 (c) and decode the results. We find that two pairs of legs are dragged along the ground and no longer take swing steps. The same sets of legs fail together since the locomotion is conditioned on the trot gait where the left front and right hind move together. We also encode the resulting motion and notice how the oscillations in latent space change shape and are no longer symmetrical, see Figure 2.

## Conclusion

There are visual similarities between the oscillations found in the primate's motor cortex and the robot's latent spaces. Unlike the motor cortex, we are able to interrogate the structure of the latent space. In doing so, we find that emergent properties of locomotion and the reaching workspace become embedded in the space. In the case of locomotion, cadence and footstep height are disentangled into separate orthogonal dimensions of the latent space. This facilitates the blending between dynamic trot gaits. Alternatively, for manipulation the workspace of the robot is embedded in the latent space. In essence the table top surface is mapped into the space. Finally, we are able to lesion our VAE models to observe how the latent-space trajectories degrade.

## Acknowledgements

This work was supported by a UKRI/EPSRC Programme Grant [EP/V000748/1], the EPSRC grant 'Robust Legged Locomotion' [EP/S002383/1], the EPSRC CDT [EP/L015897/1], the UKRI/EPSRC RAIN [EP/R026084/1] and ORCA [EP/R026173/1] Hubs and the EU H2020 Project MEMMO (780684). It was conducted as part of ANYmal Research, a community to advance legged robotics. The authors would like to acknowledge the use of the SCAN facility.

## References

Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Foster, J. D., Nuyujukian, P., Ryu, S. I., & Shenoy, K. V. (2012). Neural population dynamics during reaching. *Nature*, *487*.

Kingma, D., & Welling, M. (2014). Auto-encoding variational bayes. In *Int. conf. on learn. repr. (iclr).*

Mitchell, A. L., Merkt, W., Geisert, M., Gangapurwala, S., Engelcke, M., Jones, O. P., ... Posner, I. (2022). VAE-Loco: Versatile quadruped locomotion by learning a disentangled gait representation. *arXiv preprint arXiv:2205.01179*.

Mitchell, A. L., Merkt, W., Geisert, M., Gangapurwala, S., Engelcke, M., Parker Jones, O., ... Posner, I. (2022). Next steps: Learning a disentangled gait representation for versatile quadruped locomotion. *IEEE International Conference on Robotics and Automation (ICRA)*.

Parker Jones, O., Mitchell, A. L., Jun, Y., Merkt, W., Geisert, M., Havoutis, I., & Posner, I. (2022). Brain-like latent dynamics emerge in robot systems during walking and reaching. *Under review*.

Rezende, D. J., Mohamed, S., & Wierstra, D. (2014). Stochastic backpropagation and approximate inference in deep generative models. In *Int. conf. on mach. learn. (icml).*