A Toy Model of Predictive Classification in a Quasi-Periodic Environment

Mark Tracy

1 Introduction

We seek to explore the following question: Under what conditions can a predictive agent converge on stable internal representations of dynamical invariants from relational observations of a structured yet non-repeating process?

To begin, we construct a minimal dynamical environment together with an agent that attempts to model its evolution through prediction. The purpose of the model is to study how a system can infer latent structure from relational observations and iteratively refine its internal representation through prediction error.

Rather than directly supervising latent parameters of the environment, the agent predicts future relational observations. Deviations between predicted and observed evolution provide the learning signal.

This toy setting provides a controlled environment for studying how representation and classification may emerge from prediction in a relational dynamical system.

2 Quasi-Periodic Environment

Let the environment consist of three cyclic variables

$$\theta_1(t), \theta_2(t), \theta_3(t) \in S^1 \cong \mathbb{R}/2\pi\mathbb{Z}.$$

Their dynamics are

$$\theta_i(t+1) = \theta_i(t) + \omega_i \pmod{2\pi}, \qquad i = 1, 2, 3.$$

Equivalently,

$$\theta(t+1) = \theta(t) + \omega \pmod{2\pi}$$

where

$$\theta(t) = (\theta_1(t), \theta_2(t), \theta_3(t)), \qquad \omega = (\omega_1, \omega_2, \omega_3).$$

Definition 1 (Rational Independence)**.** *Real numbers*

$$\omega_1, \omega_2, \omega_3$$

are rationally independent *if the only integer solution to*

$$k_1\omega_1 + k_2\omega_2 + k_3\omega_3 = 0$$

with

$$k_1, k_2, k_3 \in \mathbb{Z}$$

is

$$k_1 = k_2 = k_3 = 0.$$

Definition 2 (Quasi-Periodic System)**.** *The dynamical system defined above is* quasi-periodic *if the frequencies*

$$\omega_1, \omega_2, \omega_3$$

are rationally independent.

In this case the trajectory is dense on the torus

$$T^3 = S^1 \times S^1 \times S^1.$$

3 Observation Model

The environment state is

$$x_t = \theta(t).$$

The agent observes only relational quantities:

$$o_t = h(x_t).$$

We define

$$h(x_t) = (\cos \Delta_{12}, \sin \Delta_{12}, \cos \Delta_{13}, \sin \Delta_{13}, \cos \Delta_{23}, \sin \Delta_{23})$$

where

$$\Delta_{ij}(t) = \theta_i(t) - \theta_j(t) \pmod{2\pi}.$$

Thus the agent observes only relational phase differences between the cyclic processes.

4 Predictive Agent

A predictive agent is defined by three functions

$$o_t = h(x_t)$$

$$s_{t+1} = u(s_t, o_t)$$

$$\hat{o}_{t+1} = g(s_{t+1})$$

where s_t is internal state and $\hat{o}_{t+1}$ is the predicted observation.

4.1 Neural Parameterization

In the simplest implementation the state update function is parameterized by a neural network

$$s_{t+1} = \mathrm{MLP}_u([s_t, o_t]).$$

The prediction head is taken to be a *linear readout* of the internal state

$$\hat{o}_{t+1} = W s_{t+1} + b$$

for matrix W and bias vector b.

Restricting the prediction head to be linear creates a representational bottleneck: the model cannot rely on a powerful nonlinear decoder to transform arbitrary latent features into accurate predictions. Instead, the internal state must organize information about the environment in a form that is directly readable through simple linear transformations.

5 Prediction Error

Prediction error is

$$e_{t+1} = \hat{o}_{t+1} - o_{t+1}.$$

Training minimizes

$$\mathcal{L}_{t+1} = \|\hat{o}_{t+1} - o_{t+1}\|_2^2.$$

6 Time-Rescaling Symmetry

Let $k > 0$ and define

$$\omega' = k\omega.$$

Then

$$\theta'(t) = \theta(0) + k\omega t$$

which is equivalent to

$$\theta'(t) = \theta(\tau), \quad \tau = kt.$$

Thus the orbit is unchanged.

Theorem 1. *The frequency vector ω is identifiable only up to multiplication by a positive scalar when inferred from relational phase observations alone.*

7 Observable Invariants

Observable invariants are frequency ratios

$$\frac{\omega_1}{\omega_3}, \qquad \frac{\omega_2}{\omega_3}.$$

More generally

$$[\omega_1 : \omega_2 : \omega_3].$$

Any representative of this projective class can be chosen by normalization.

8 Normalization of the Frequency Vector

Normalize via

$$\omega_1 + \omega_2 + \omega_3 = 1.$$

Then environments correspond to points in the interior of a 2-simplex.

9 Environment Distribution

To generate a family of environments, frequency vectors are sampled from the normalized simplex

$$\Delta^2 = \left\{\omega \in \mathbb{R}^3 : \omega_i > 0,\ \omega_1 + \omega_2 + \omega_3 = 1\right\}.$$

Sampling uniformly from this simplex produces a distribution over quasi-periodic dynamical systems parameterized by the ratios of the underlying frequencies.

Each sampled vector defines a distinct environment whose dynamics follow

$$\theta(t+1) = \theta(t) + \omega \pmod{2\pi}.$$

10 Relational Dynamics

$$\Delta_{ij}(t+1) = \Delta_{ij}(t) + (\omega_i - \omega_j) \pmod{2\pi}.$$

11 Linear Readout and Koopman Representation

Many nonlinear dynamical systems admit representations in which their evolution becomes linear in a suitable set of observables. This perspective is formalized in Koopman operator theory, where the dynamics of observable functions evolve linearly even when the underlying state dynamics are nonlinear.

In the present setting the relational observables

$$(\cos \Delta_{ij}, \sin \Delta_{ij})$$

are periodic functions on the torus T^3. Their evolution corresponds to phase shifts determined by the frequency differences $(\omega_i - \omega_j)$. Writing the complex observable

$$z_{ij}(t) = e^{i\Delta_{ij}(t)},$$

the relational dynamics imply

$$z_{ij}(t+1) = e^{i(\omega_i - \omega_j)} z_{ij}(t).$$

Thus the observable evolves through multiplication by a constant complex phase factor determined by the frequency difference. This is precisely a linear evolution in observable space, analogous to a Koopman eigenfunction.

The linear prediction head therefore tests whether the agent has learned an internal representation approximating this observable coordinate system. If prediction can be achieved through a linear transformation of the internal state, this suggests that the learned representation captures dynamical invariants analogous to Koopman eigenfunctions.

12 Interpretation

Successful prediction implies internal representation of the underlying dynamical structure. Because the prediction head is constrained to be linear, accurate prediction requires that the internal state encode the relevant dynamical information in a linearly accessible form.

13 Questions Testable in this Environment

13.1 Learning Latent Structure via Linear Probing

We test whether internal state s_T encodes the latent structure of the environment.

Running the predictive agent on multiple sampled environments produces final states

$$s_T^{(k)}.$$

The probing dataset consists of pairs

$$(s_T^{(k)}, y^{(k)})_{k=1}^{N}$$

with targets

$$y^{(k)} = \begin{pmatrix} \omega_1^{(k)}/\omega_3^{(k)} \\ \omega_2^{(k)}/\omega_3^{(k)} \end{pmatrix}.$$

A linear probe predicts

$$\hat{y} = Ws + b.$$

The probe is trained with mean squared error

$$\mathcal{L}_{probe} = \left\| W s_T^{(k)} + b - y^{(k)} \right\|_2^2.$$

13.2 Generalization Across Environments

Training the probe on some environments and evaluating on unseen ones tests whether the representation captures general structure.

13.3 Robustness to Nonstationarity

The environment can be extended to weakly nonstationary dynamics by allowing slow drift in the frequencies:

$$\omega(t+1) = \omega(t) + \epsilon_t.$$

14 Conclusion

This toy environment provides a controlled setting in which an agent predicts relational observations generated by quasi-periodic dynamics. Prediction error drives learning, while probing tests whether latent dynamical structure is represented internally.