hugomcp@di.ubi.pt, 2026/27

Machine Learning

Practical Sheet 11

Physics-Inspired Machine Learning

The goal of this exercise is to predict the cooling of a device from six noisy readings/sensors.

Compare two identical neural networks: one trained with observations and an initial condition, and one that also uses the cooling equation. Explain the resulting curves, including cases in which the assumed physics is wrong.

The physical model

Assume a spatially uniform object, constant ambient temperature T_a, constant cooling rate $k > 0$, and no internal heat source. Newton's law gives $dT/dt = -k(T - T_a)$.

With $y = (T - T_a)/(T_0 - T_a)$ and $\tau = kt$, the problem becomes:

$$\mathbf{dy/d\tau + y = 0, \qquad y(0) = 1.}$$

The analytic solution $y(\tau) = \exp(-\tau)$ is available for generating this synthetic example and scoring predictions. Do not use its values as extra training labels or in the physics loss.

Setup and first prediction

Open p_11_ml_26_27_Physics_starter.ipynb.

The notebook supplies the data, model, optimiser and plots.

Before running anything, sketch the cooling curve. What would change if the device remained powered or the ambient temperature varied?

hugomcp@di.ubi.pt, 2026/27

1. A Continuous Function from Finite Weights

The network receives a scalar τ and returns a scalar $y_\theta(\tau)$. Two hidden layers contain 32 tanh units each. All queries share the same weights θ.

1.1 Run the setup and model cells. Evaluate the untrained network at $\tau = 0.37$, 1.45 and 2.83. Why can it return values at locations absent from the observations?

1.2 Evaluate the same network on 25 and 250 query points. Plot both. Does the number of trainable parameters change? Explain what a denser plot changes.

1.3 Explain why an ordinary tanh regressor already defines a continuous function. Which extra ingredient will make it physics-informed?

2. Controlled Baseline

τ	Observed y
0.0	1.00000
0.2	0.83873
0.4	0.64032
0.6	0.55881
0.8	0.42933
1.0	0.38788

The notebook computes these values as $\exp(-\tau)$ plus fixed noise $[0, 0.02, -0.03, 0.01, -0.02, 0.02]$. The table is rounded for reading; train with the notebook tensors.

Use $L_{baseline} = L_{data} + 10\, L_{initial}$, where L_{data} is mean squared observation error and $L_{initial} = [y_\theta(0) - 1]^2$. The baseline receives the same initial condition as the PINN.

2.1 Combine the supplied loss terms and physics_weight. Train with physics_weight = 0, seed 7, Adam, learning rate 0.001 and 2,000 updates.

2.2 Plot predictions on $0 \leq \tau \leq 5$. Identify the observed-data region. Describe what the model does beyond $\tau = 1$ without deciding yet whether it is physically correct.

Save the initial model before training. Start every comparison from a deep copy of these weights. Reusing a trained model would confound the comparison.

3. The PINN

Enforce the equation at 64 equally spaced collocation points in $0 \leq \tau \leq 4$. These points have no measured y labels.

$r_\theta(\tau) = dy_\theta/d\tau + \beta\, y_\theta(\tau)$. For the correct law, $\beta = 1$.

$L_{physics}$ = mean $[r_\theta(\tau_j)^2]$. The PINN uses $L_{data} + 10\, L_{initial} + L_{physics}$. All three terms use dimensionless variables.

3.1 Complete TODO 2: use torch.autograd.grad to calculate dyθ/dτ. Complete TODO 3: return the residual. The notebook creates coordinate tensors with requires_grad=True.

```
dy = torch.autograd.grad(
    y, tau, grad_outputs=torch.ones_like(y),
    create_graph=True
)[0]
```

Here the MLP processes each row independently. create_graph=True keeps the derivative in the graph so training can differentiate the physics loss with respect to θ. Do not detach y and do not compute the residual inside torch.no_grad().

3.2 Check the residual function using $y = \exp(-\tau)$, $y = 0$ and $y = \exp(-2\tau)$. Predict each residual before running the provided checks. Which candidate still needs the initial condition?

3.3 Train a fresh copy of the same initial model with physics_weight = 1 and $\beta = 1$. Plot the baseline, PINN and analytic reference together. Use identical architecture and optimiser settings.

3.4 Inspect the separate data, initial and physics losses. Why is comparing only the total losses of the two models misleading?

4. Evaluation beyond the observations

Complete TODO 4: implement RMSE as the square root of the mean squared difference. Use 201 reference points per interval; the notebook uses these points only after training.

Region	Interval	Available training information
Data	$0 \leq \tau \leq 1$	Observed values and equation
Physics only	$1 < \tau \leq 4$	Equation without observed values
Outside training	$4 < \tau \leq 5$	Neither observations nor equation checks

The evaluation grid includes interval endpoints for convenience. Shared endpoints have negligible impact here; describe the regions conceptually as above.

4.1 Fill the results table for seed 7.

Report residual RMS on the supplied 200-point midpoint grid in [0,4], independent of the training collocation grid. Also report $|y_\theta(0) - 1|$.

Metric	Baseline	PINN	Wrong law
Data-region RMSE			
Physics-only RMSE			
Outside-training RMSE			
Residual RMS for $y' + y$			
Initial-condition error			

4.2 Repeat the baseline and PINN with seed 17. Preserve matched initial weights within each comparison. Does the size of the improvement remain identical?

4.3 Is prediction at $\tau = 3$ extrapolation beyond the observations? Is it also outside the physics-training domain? Answer both questions separately.

4.4 Explain why a small residual on the training grid is insufficient evidence of physical accuracy. Mention the equation assumptions and the initial condition.

Expected qualitative behaviour: the correct-law PINN often improves the region without observations in this small example. A different seed or optimisation setting can change the numbers. The exercise does not establish that PINNs always outperform other methods.

5. Wrong Physical Prior

Keep the observations unchanged and deliberately impose $y' + 2y = 0$. Train from the same initial weights with $\beta = 2$ and physics_weight = 1.

5.1 Add the wrong-law model to your plot and results table. Evaluate its residual using the true reference equation $y' + y = 0$. Separately inspect its residual under the imposed equation $y' + 2y = 0$.

5.2 Explain the conflict between the data and the imposed law. Why can satisfying the imposed law more closely make the prediction less accurate?

5.3 Discuss whether a smaller error in the far tail would, by itself, prove that the wrong-law model is better. Consider that both cooling curves approach zero.

5.4 State when you would use an analytic solution or ordinary differential-equation solver instead of training a PINN for this problem.