

Figure 6: HNN schema. The forward pass of an HNN is composed of a forward pass through a differentiable model as well as a backpropagation step through the model.

382 **Training details.** We selected hyperparameters using a coarse grid search over learning rates
 383 $\{10^{-1}, 10^{-2}, 10^{-3}\}$, layer widths $\{100, 200, 300\}$, activations $\{\tanh, \text{relu}\}$, and batch size where
 384 relevant $\{100, 200\}$. The main objective of this work was not to produce state-of-the-art results, so
 385 the settings we chose were aimed simply at producing models that gave good qualitative performance
 386 on the tasks at hand. We used weight decay of 10^{-4} on the first three tasks.

387 We trained all of these experiments on a desktop CPU.

388 **The large test loss on Task 3.** The test losses we report on Task 3 are significantly larger than the
 389 training losses. This discrepancy is a result of the way we partitioned the training and test sets. The
 390 dataset provided by [29] consisted of just a single trajectory from a real pendulum, as shown in the
 391 second panel of Figure 7(c). Also, we needed to evaluate our model’s performance over a series of
 392 adjacent time steps in order to measure the energy MSE metric. For this reason, we were forced
 393 to use the first 4/5 of this trajectory for training (black vectors in Figure 7(c)) and the last 1/5 for
 394 evaluation (red vectors).

395 The consequence of this train/test split is that our test set had a slightly different distribution from
 396 our training set, producing larger test losses compared to train losses. We found that the relative
 397 magnitudes of the test losses between the baseline and HNN models were informative, which is why
 398 we report them. We did not perform this ungainly train/test split on the other two tasks in this section.

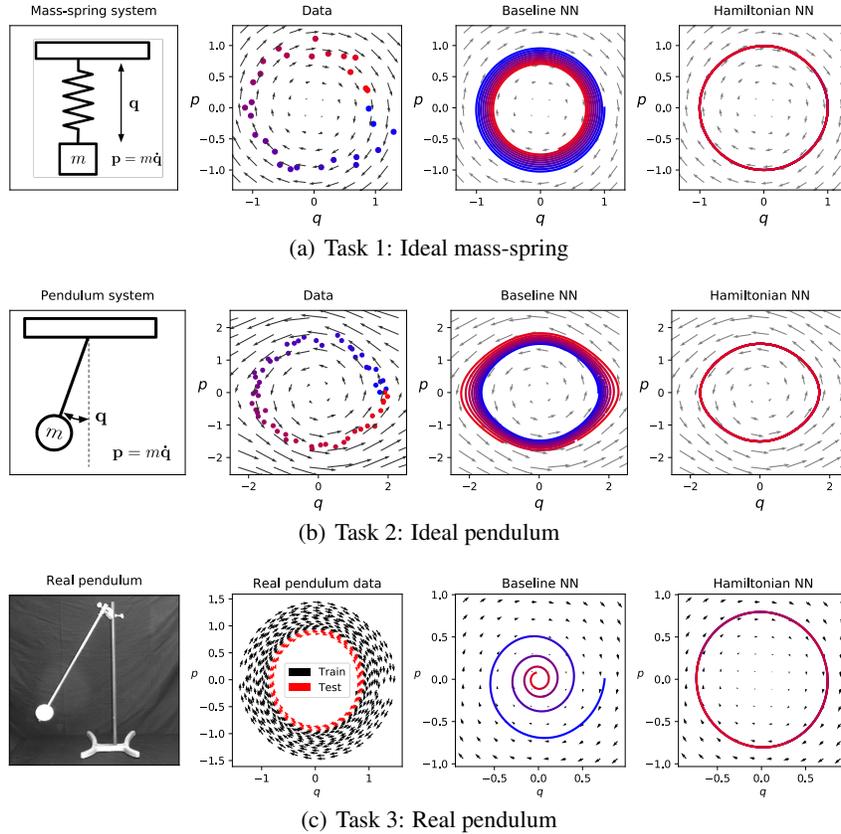


Figure 7: More qualitative results comparing the HNN to a baseline neural network on the first three physics tasks. From top to bottom: Task 1: Ideal mass-spring, Task 2: ideal pendulum, Task 3: Real pendulum.

399 **B Supplementary Information for Task 4: Two-body problem**

400 **Training details.** We selected hyperparameters with a grid search as described in the previous section.
 401 Again, the main objective of this work was not to produce state-of-the-art results, so the settings we
 402 chose were aimed simply at producing models that gave good qualitative performance on the tasks at
 403 hand. We did not use weight decay on this task, though when we tried a weight decay of 10^{-4} or
 404 results did not change significantly.

405 We trained this experiment on a desktop CPU.

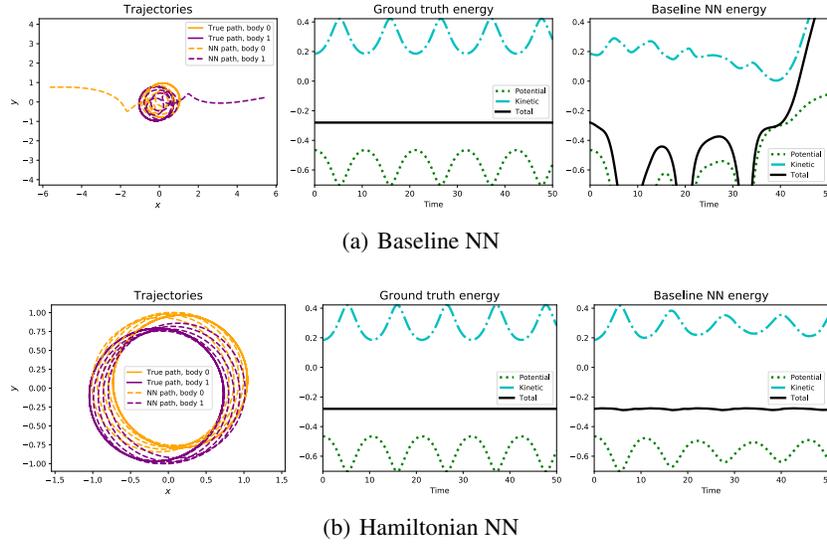


Figure 8: More qualitative results for the orbit task. Numerical errors accumulate in the baseline model until the bodies end up traveling in opposite directions. The total energy diverges towards infinity as well. In comparison, the HNN’s trajectory diverges from the ground truth but continues to roughly conserve the total energy of the system.

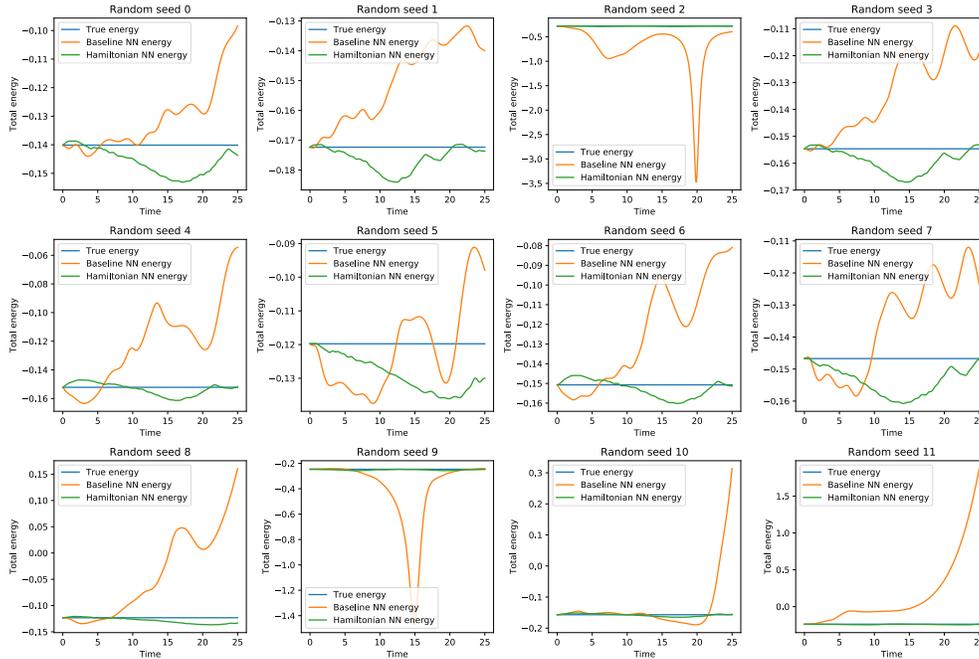


Figure 9: Comparison of how well the HNN conserves total energy compared to the baseline its baseline on the two-body task.

406 **Three body problem.** As mentioned briefly in the body of the paper, we also trained our models
 407 on the three body problem. The results we report here show a relative advantage to using the HNN
 408 over the baseline model. However, both models struggled to accurately model the dynamics of the
 409 three-body problem, which is why we relegated these results to the Appendix. Going forward, we
 410 hope to improve these results to the point where they can play a more substantial role in Section 4.

411 Table 2 gives a summary of quantitative results and Figure 10 shows a qualitative analysis of the
 412 models we trained on this task.

Table 2: Quantitative results for the three-body problem.

	Train loss		Test loss		Energy MSE	
	Baseline	HNN	Baseline	HNN	Baseline	HNN
Task 4b: 3-body problem	0.096	0.080	0.380	0.488	103.9	0.039

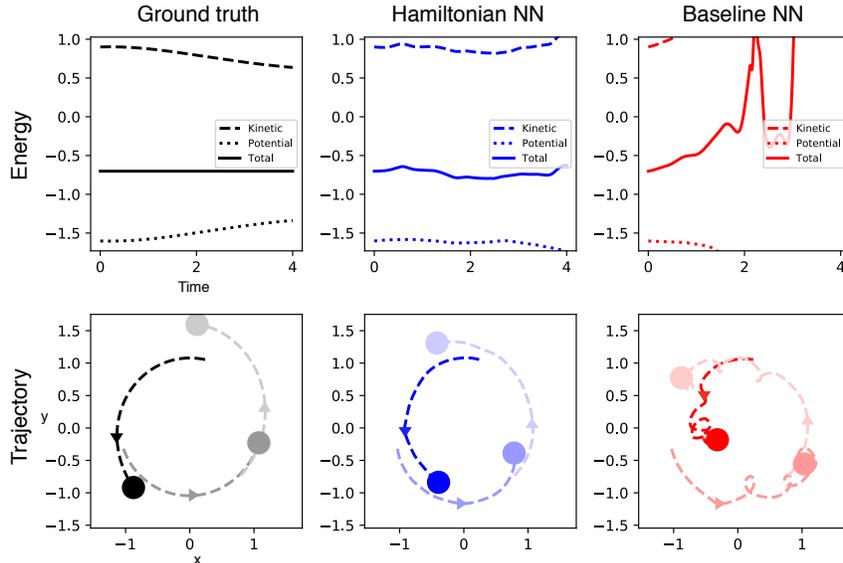


Figure 10: Analysis of an example three-body trajectory. The baseline model does not conserve total energy and quickly diverges from ground truth. The HNN, meanwhile, roughly conserves total energy and its trajectories resemble the ground truth.

413 C Supplementary Information for Task 5: Pixel Pendulum

414 **Training details.** We selected hyperparameters with a grid search as described in the previous section.
 415 We used a weight decay of 10^{-5} on this experiment. We found that, unlike previous experiments,
 416 weight decay had a significant impact on results. We suspect that this is because the scale of the
 417 gradients on the weights of the HNN portion of the model were different from the scale of the
 418 gradients of the weights of the autoencoder portion of the model.

419 We trained this experiment on a desktop CPU.

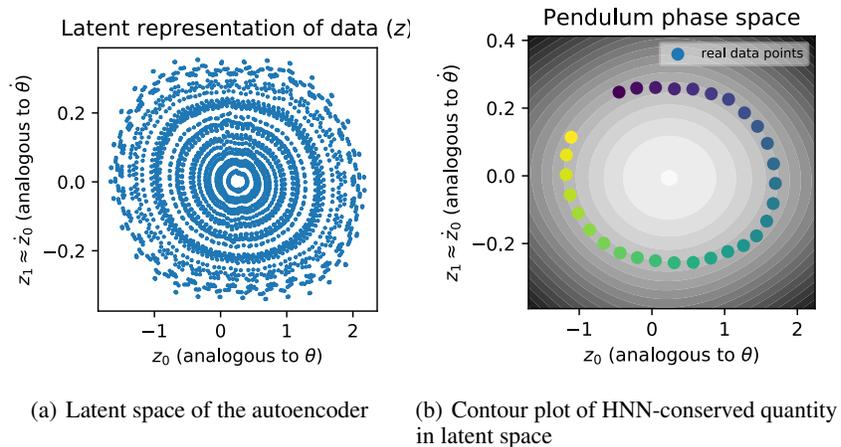


Figure 11: Latent space plots from the Pixel Pendulum model. Note that the learned latent space bears a strong resemblance to the true phase space of a pendulum. In particular, there is a faint diamond shape to the outer contour lines of Figure 11(b). This pattern is reminiscent of the nonlinear dynamics we observed in the ideal pendulum phase space plot of Figure 2 (row 2, column 1)