Noisy inputs make neural networks play it safe, but latent inputs can restore the extremes.

Version March 26, 2026
Preprint typeset using LaTeX style openjournal v. 09/06/15

arXiv:2512.23138v2 [astro-ph.IM] 25 Mar 2026

WHY MACHINE LEARNING MODELS SYSTEMATICALLY UNDERESTIMATE EXTREME VALUES II: HOW TO FIX IT WITH LATENTNN

Yuan-Sen Ting
Department of Astronomy, The Ohio State University, Columbus, OH 43210, USA
Center for Cosmology and AstroParticle Physics (CCAPP), The Ohio State University, Columbus, OH 43210, USA and
Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany
Version March 26, 2026

The fix is conceptually old-fashioned: admit that measured inputs are not the true inputs.

ABSTRACT

Attenuation bias—the systematic underestimation of regression coefficients due to measurement errors in input variables—affects astronomical data-driven models. For linear regression, this problem was solved by treating the true input values as latent variables to be estimated alongside model parameters. In this paper, we show that neural networks suffer from the same attenuation bias and that the latent variable solution generalizes directly to neural networks. We introduce LatentNN, a method that jointly optimizes network parameters and latent input values by maximizing the joint likelihood of observing both inputs and outputs. We demonstrate the correction on one-dimensional regression, multivariate inputs with correlated features, and stellar spectroscopy applications. LatentNN reduces attenuation bias across a range of signal-to-noise ratios where standard neural networks show large bias. This provides a framework for improved neural network inference in the low signal-to-noise regime characteristic of astronomical data. This bias correction is most effective when measurement errors are less than roughly half the intrinsic data range, in the regime of very low signal-to-noise and few informative features. Code is available at `https://github.com/tingyuansen/LatentNN`.

Subject headings: methods: statistical — methods: data analysis — techniques: spectroscopic — stars: fundamental parameters

1. INTRODUCTION

Modern astronomy increasingly relies on neural networks for inference (Cranmer et al. 2020; Huertas-Company & Lanusse 2023; Ting 2025a). Data-driven approaches now routinely map observed quantities to physical parameters across a wide range of applications, from stellar parameters derived from spectra (e.g., Ness et al. 2015; Fabbro et al. 2018; Leung & Bovy 2019; Ting et al. 2019) to photometric redshifts from imaging surveys (e.g., Li et al. 2022a; Lin et al. 2022; Zhou et al. 2022) to stellar properties from time series light curves (Pasquet et al. 2019; Hon et al. 2021).

With the rise of neural network approaches, however, input uncertainties are often ignored or left unmodeled. This is largely a design choice rather than an oversight. Classical machine learning methods have long grappled with measurement errors in inputs, developing explicit frameworks for handling uncertain covariates. Neural network modeling, by contrast, arose in domains like computer vision (Simonyan & Zisserman 2014; Szegedy et al. 2014; He et al. 2016) and natural language processing (Devlin et al. 2018; Brown et al. 2020; OpenAI et al. 2023; Touvron et al. 2023) where input noise is negligible, and its standard architectures have little mechanism for directly incorporating input uncertainties. When applied to astronomical data, this simplification can lead to systematic errors that are difficult to diagnose (Kelly 2007; Ting 2025b).

Astronomy operates in a regime that differs from most machine learning applications. Astronomical observations often have signal-to-noise ratios (SNR) of order 10, yet we seek inference precision at the percent level or better. This tension between noisy inputs and stringent precision requirements creates an unique challenge for neural network inference.





The rare extremes are often the scientifically interesting objects, precisely where this bias hurts most.

This problem is known as attenuation bias (Spearman 1904; Fuller & Thompson 2002; Kelly 2007). In Paper I of this series (Ting 2025b), we demonstrated this effect in detail for linear regression and explored its implications for astronomical data analysis. Attenuation bias manifests as a compression of the predicted dynamic range: high values are consistently predicted too low, while low values are predicted too high. The effect depends on the ratio of measurement uncertainty to intrinsic signal range and, for finite samples, on sample size. As a rule of thumb, attenuation bias becomes important at signal-to-noise ratios below 10, especially at low-dimensional problems. The exact magnitude depends on the dimensionality and correlation structure of the inputs, but the bias is never zero when input uncertainties are present. We refer interested readers to Paper I for detailed derivations and further discussion of attenuation bias in the context of linear regression.

The consequences of attenuation bias are not merely increased scatter in predictions. When we train a regression model to predict some quantity from noisy inputs, treating the observed inputs as exact is a reasonable approximation only when measurement errors are small compared to the intrinsic variation in the data. When measurement errors constitute even a modest fraction of the signal, a systematic compression of the predicted dynamic range emerges—high values are underestimated

ting.74@osu.edu

More training data cannot average this bias away; the input noise changes the relationship being learned.

Measurement error stretches the input axis, so an ordinary fit becomes systematically too shallow.

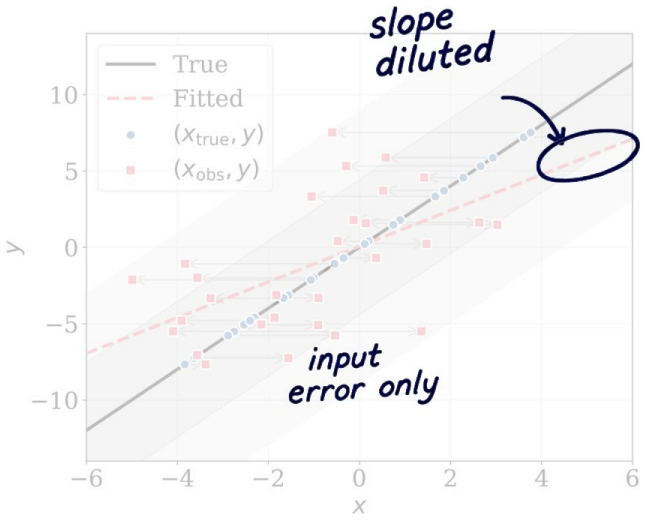


FIG. 1.— Schematic illustration of attenuation bias. Blue circles show true positions (x_{true}, y) following $y = 2x$ (black line); red squares show observed positions (x_{obs}, y) after adding measurement error with $\text{SNR}_x = 1$. Grey arrows indicate horizontal displacement from true to observed x positions. The tilted grey shaded bands show the $\pm 1\sigma_x$ (darker) and $\pm 2\sigma_x$ (lighter) uncertainty regions around the true relationship. A linear fit to the observed data (red dashed line) yields a shallower slope, demonstrating how input measurement errors systematically bias regression coefficients toward zero.

The red points move only sideways, yet that is enough to rotate the fitted line.

and this bias persists even with infinite data. It cannot be removed by collecting more samples or using more sophisticated models. It is a fundamental consequence of treating noisy measurements as if they were exact.

For linear models, the statistics literature has long recognized this problem and developed solutions such as Deming regression (Deming 1943; Fuller 1987; Carroll et al. 2006; Ting 2025c). The key insight is to treat the true input values as latent variables to be estimated alongside model parameters (Loredo 2004; Kelly 2007). Rather than taking the noisy observations at face value, the model simultaneously infers what the true inputs likely were and learns the relationship between those true inputs and the outputs. The present paper extends this idea to neural networks. We show that neural networks suffer from the same attenuation bias as linear regression, and that the latent variable approach provides a general solution applicable to neural networks.

We introduce LatentNN, a practical implementation using modern automatic differentiation frameworks, and validate it on one-dimensional regression, multivariate inputs, and spectral applications. Section 2 briefly reviews attenuation bias following Paper I notation. Section 3 demonstrates that neural networks exhibit the same bias. Section 4 presents the LatentNN formulation and results, first for the one-dimensional case, then extending to multivariate inputs. Section 5 applies the method to stellar spectra. Section 7 discusses limitations, extensions, and connections to hierarchical Bayesian modeling, followed by conclusions in Section 8.



2. ATTENUATION BIAS

To understand attenuation bias, consider a simple regression problem where we want to learn the relationship between an input variable x and an output variable y. Suppose the true relationship is $y_{\text{true}} = f(x_{\text{true}})$, which could be linear or nonlinear. In practice, we never observe the true values directly. Instead, we measure $x_{\text{obs}} = x_{\text{true}} + \delta_x$, where δ_x represents the measurement error. If these errors are unbiased with variance σ_x^2, then $\delta_x \sim \mathcal{N}(0, \sigma_x^2)$. The naive approach in machine learning is to train on the observed pairs $(x_{\text{obs}}, y_{\text{obs}})$, treating x_{obs} as if it were the true value. This seemingly innocuous simplification leads to systematic bias.

Measurement errors in x artificially inflate the observed variance. If the true values have variance $\sigma^2_{\text{range}} = \text{Var}(x_{\text{true}})$, representing the intrinsic spread or dynamic range of the data, then the observed variance is

$$\text{Var}(x_{\text{obs}}) = \sigma^2_{\text{range}} + \sigma_x^2. \tag{1}$$

This stretches the data horizontally without a corresponding vertical stretch, because the measurement errors in x do not affect y. The observed distribution appears wider in the x-direction than the true distribution, while the y-extent remains unchanged. Any regression method treating x_{obs} as exact will therefore fit a slope that is systematically shallower than the true relationship. This weakening of the apparent relationship is why attenuation bias is also known as regression dilution.

This effect can be visualized by plotting data points and their regression line. Figure 1 illustrates this with a simple example: true data points (blue) following a linear relationship $y = 2x$ are scattered horizontally by measurement error to their observed positions (red), with arrows indicating the displacement. Grey tilted shaded bands show the $\pm 1\sigma_x$ and $\pm 2\sigma_x$ uncertainty regions around the true relationship. A regression line fit to the observed data (red dashed) has a shallower slope than the true underlying relationship.

For linear regression with true relationship $y = \beta x$, where β is the true slope, we can quantify this effect precisely (see Appendix A for the full derivation). The ordinary least squares estimator $\hat{\beta}$ has expected value

$$\mathbb{E}[\hat{\beta}] = \beta \cdot \lambda_\beta, \quad \text{where} \quad \lambda_\beta = \frac{1}{1 + (\sigma_x/\sigma_{\text{range}})^2}. \tag{2}$$

This ratio—not the absolute noise level—is the quantity to watch.

The quantity λ_β is called the attenuation factor. It depends only on the ratio of measurement error to intrinsic signal range, $\sigma_x/\sigma_{\text{range}}$. When measurement errors are small compared to the signal range, λ_β is close to unity and the bias is negligible. When measurement errors are comparable to the signal range, λ_β can be much less than unity, indicating large bias. Since $\lambda_\beta < 1$, predictions are systematically pulled toward the mean: high values are predicted too low, while low values are predicted too high. In the extreme case where measurement errors dominate ($\sigma_x \gg \sigma_{\text{range}}$), λ_β approaches zero and the model predicts the mean for all inputs, obscuring the true relationship.

Consider some representative cases in the one-dimensional setting:

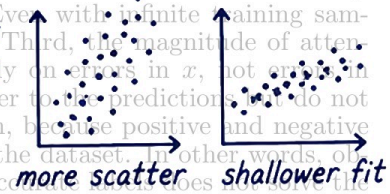
- At $\sigma_x/\sigma_{\text{range}} = 0.1$ (equivalently, $\text{SNR} = \sigma_{\text{range}}/\sigma_x = 10$), the attenuation factor is $\lambda_\beta \approx 0.99$, corresponding to 1% bias.
- At $\sigma_x/\sigma_{\text{range}} = 0.33$ ($\text{SNR} \approx 3$), the attenuation factor is $\lambda_\beta \approx 0.90$, corresponding to 10% bias.
- At $\sigma_x/\sigma_{\text{range}} = 1.0$ ($\text{SNR} = 1$), the attenuation factor is $\lambda_\beta = 0.50$, corresponding to 50% bias.

At SNR = 1, half the true dynamic range disappears from the fitted slope.

Attenuation is a data problem, not a shortage of samples, labels, or network capacity.

At SNR of 10 or below, a regime common in astronomical observations, the bias reaches percent level or more.

Several properties of this bias are worth emphasizing. First, the bias depends only on the ratio $\sigma_x/\sigma_{\rm range}$, not on the absolute values of either quantity. Second, it is independent of sample size n and cannot be resolved by collecting more data. Even with infinite training samples, the bias remains. Third, the magnitude of attenuation is entirely determined by errors in x, not errors in y. Errors in y add scatter to the predictions but do not systematically bias them, because positive and negative errors cancel out over the dataset. In other words, obtaining better or more accurate labels does not solve the problem.[1]



Output noise broadens the cloud; input noise tilts the learned relationship.

This asymmetry between x and y arises because they play fundamentally different roles in regression: we predict y given x, not the reverse. Errors in x change the question being asked—we are predicting from a corrupted version of the input—which systematically weakens the apparent relationship. However, as we will see in Section 4, while σ_y does not affect the bias itself, it does play a role in correcting the bias: the Deming regression solution (Deming 1943; Golub & Van Loan 1980; Ting 2025c) and its neural network generalization require specifying both σ_x and σ_y to properly weight the loss terms.

It is worth noting that this bias is sometimes mistakenly attributed to insufficient training data at the extremes or to noisy labels. Observing that models underpredict extreme values, one might assume the problem lies in having too few training examples at the extremes or in noisy labels. However, attenuation bias has nothing to do with either of these issues. The bias arises purely from measurement errors in the input, and it persists even with perfectly balanced training data and perfect labels.

3. ATTENUATION BIAS IN NEURAL NETWORKS

The previous section reviewed attenuation bias in the context of linear regression, where the effect can be derived analytically. Paper I further argued that the same bias extends to polynomial and nonlinear models. However, for nonlinear models or high-dimensional inputs, the attenuation factor can no longer be expressed as a simple function of σ_x and $\sigma_{\rm range}$. The relationship between measurement error and attenuation becomes more complex, depending on the specific functional form.

This complexity makes it difficult to predict and therefore analytically correct the bias in closed form for general machine learning models. Here we first demonstrate empirically that multi-layer perceptrons (MLPs) exhibit the same qualitative behavior. MLPs, also known as fully connected neural networks (Cybenko 1989; Hornik et al. 1989), are the simplest neural network architecture: each neuron in one layer is connected to every neuron in the next layer, with nonlinear activation functions between layers. Despite their simplicity, MLPs remain widely used in astronomy (e.g., Odewahn et al. 1992; Storrie-Lombardi et al. 1992; Collister & Lahav 2004; Guo et al. 2020; Wong et al. 2020; Wang et al. 2021; Li et al. 2022b; Vynatheya et al. 2022; Winecki & Kochanek 2024) and serve as the foundation for more complex architectures.

To test whether neural networks suffer from attenuation bias, we conduct numerical experiments following the setup of Paper I. We consider the true function $y = 2x$, chosen for direct comparison with the linear theory, and sample $N_{\rm train} = 1000$ training values and $N_{\rm test} = 200$ test values of $\mathbf{x}_{\rm true}$ uniformly from $[-5, 5]$. We add Gaussian noise to obtain $\mathbf{x}_{\rm obs} = \mathbf{x}_{\rm true} + \delta_x$ where $\delta_x \sim \mathcal{N}(0, \sigma_x^2)$, and vary σ_x to test different SNR regimes. We define $\mathrm{SNR}_x \equiv \sigma_{\rm range}/\sigma_x$ as the signal-to-noise ratio in the input.

We train two models on each dataset. The first is ordinary least squares linear regression. The second is an MLP with 2 hidden layers of 64 units each and ReLU activations, trained for 20000 epochs using the Adam optimizer with learning rate 0.03.

For neural networks, we cannot directly measure a single coefficient β as we did for linear regression in Section 2. The MLP used here has over 4000 parameters distributed across multiple layers, rather than a single interpretable slope. Instead, we quantify attenuation through the predictions themselves. Here, predictions are always made using the observed (noisy) inputs, $y_{\rm pred} = f_{\boldsymbol{\theta}}(x_{\rm obs})$, not the unknown true inputs. We define λ_y as the slope of the regression of $y_{\rm pred}$ against $y_{\rm true}$. If a model suffers from attenuation bias, its predictions will be compressed toward the mean: when we plot $y_{\rm pred}$ versus $y_{\rm true}$, the slope will be less than unity. A value of $\lambda_y = 1$ indicates no attenuation, while $\lambda_y < 1$ indicates that predictions are systematically pulled toward the sample mean.

I measure compression through λ_y because a neural network has no single slope coefficient to inspect.

As discussed in Paper I, for linear regression in one dimension with true relationship $y = \beta x$, these two measures of attenuation are equivalent. The predicted value is $\hat{y} = \hat{\beta} x_{\rm obs}$, and when we regress predictions against true values, the attenuation in the slope estimate $\hat{\beta}$ translates directly to attenuation in the predictions. In this case, $\lambda_y = \lambda_\beta = 1/(1 + (\sigma_x/\sigma_{\rm range})^2)$. For nonlinear models or multivariate inputs, λ_y remains a useful empirical measure of how much the model compresses the dynamic range of predictions.

Figure 2 shows the results. The top row shows MLP predictions on held-out test data at three specific SNR_x values: at $\mathrm{SNR}_x = 0.3$ (severe noise), predictions are compressed to $\lambda_y \approx 0.1$; at $\mathrm{SNR}_x = 1$ (noise comparable to signal range), $\lambda_y \approx 0.5$; and at $\mathrm{SNR}_x = 3$, $\lambda_y \approx 0.9$. The bottom panel shows that both linear regression and the MLP follow the theoretical attenuation curve $\lambda_y = 1/(1 + (\sigma_x/\sigma_{\rm range})^2)$ across the full range of SNR_x. For the linear true function $y = 2x$, the MLP effectively learns a linear mapping, so it follows the same theoretical curve as linear regression. The important point is that our MLP here, despite having thousands of parameters and the flexibility to fit arbitrary nonlinear functions, exhibits the same systematic bias as the much simpler linear regression. As we saw in Paper I, model complexity does not protect against attenuation bias.

Regularization does not resolve the issue either. One might hope that L_2 regularization (weight decay), which penalizes large weights, could counteract the attenuation

Weight decay controls complexity, but it cannot reconstruct information lost before training.

[1] This study focuses on discriminative models that map observables to labels. When labels are more accurate than inputs, generative models that predict observables from labels can mitigate attenuation bias, as the roles of x and y are reversed. See Paper I for further discussion.

A large neural network follows the same attenuation curve as plain linear regression.

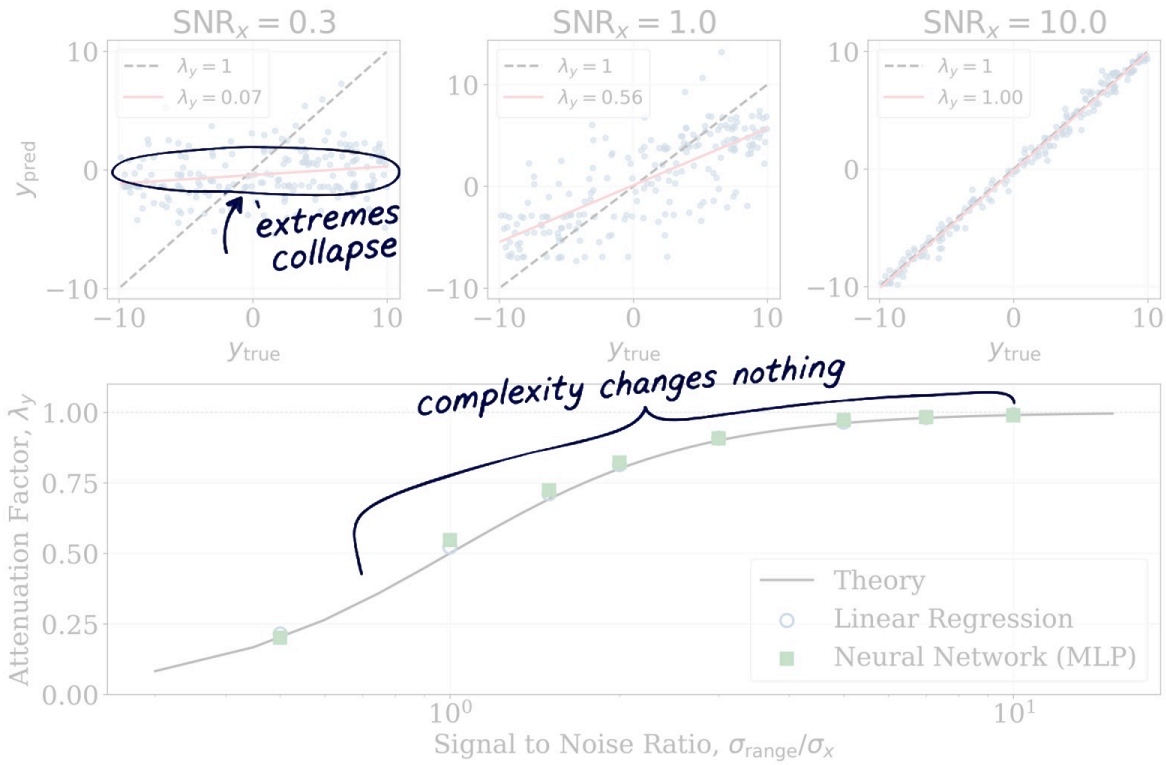


FIG. 2.— Attenuation bias in neural networks for the true function $y = 2x$. **Top:** Predicted versus true y on held-out test data for an MLP (2 hidden layers, 64 units) at three SNR_x values. At $\text{SNR}_x = 0.3$, predictions collapse to $\lambda_y \approx 0.1$; at $\text{SNR}_x = 1$, $\lambda_y \approx 0.5$; at $\text{SNR}_x = 3$, $\lambda_y \approx 0.9$. **Bottom:** Attenuation factor λ_y versus SNR_x for linear regression (blue open circles) and neural network (green filled squares). Both follow the theoretical curve $\lambda_y = 1/(1 + \text{SNR}_x^{-2})$ (black solid line), demonstrating that model complexity does not protect against attenuation bias.

Thousands of parameters did not help; the flexible model simply learned the biased conditional mean more accurately.

(Tikhonov 1963; Hoerl & Kennard 1970; Bishop 1995; Ting 2025c). However, regularization addresses overfitting by constraining model complexity, not systematic bias in the inputs. The bias arises from the mismatch between $\mathbf{x}_{\text{obs}}$ and $\mathbf{x}_{\text{true}}$, which regularization cannot correct. We verified this empirically by sweeping over a range of L_2 regularization strengths: the attenuation factor λ_y remains unchanged regardless of the regularization parameter.

4. LATENTNN: CORRECTING ATTENUATION BIAS WITH LATENT VARIABLES

The escape route is to optimize plausible true inputs rather than freezing noisy observations in place.

Having established that neural networks suffer from attenuation bias, we now develop a method to correct it. The approach is inspired by the statistics literature on errors-in-variables regression (Carroll et al. 1995; Carroll et al. 2006; Bishop 2006): rather than treating noisy observations as exact, we treat the unknown true values as latent variables to be estimated alongside the model parameters. For linear regression, this leads to Deming regression (Deming 1943). Here we show that the same principle extends naturally to neural networks, yielding a practical correction method we call LatentNN.

4.1. *The Latent Variable Solution in Linear Regression*

Deming regression is the linear ancestor of LatentNN.

Standard regression treats x_{obs} as exact, which leads to bias. An alternative approach recognizes that x_{true} is an unknown quantity for each data point. Rather than treating the observed values as truth, we can treat the true values as latent variables to be estimated alongside the model parameters.

Consider observations $(x_{\text{obs},i}, y_{\text{obs},i})$ for $i = 1, \ldots, N$. We assume that the observed input $x_{\text{obs},i}$ is a noisy measurement of the true value $x_{\text{true},i}$, and the observed output $y_{\text{obs},i}$ is a noisy measurement of the true output $y_{\text{true},i}$, which follows the linear relationship $y_{\text{true},i} = \beta x_{\text{true},i}$. Mathematically,

$$x_{\text{obs},i} = x_{\text{true},i} + \delta_{x,i}, \quad \delta_{x,i} \sim \mathcal{N}(0, \sigma_x^2), \tag{3}$$

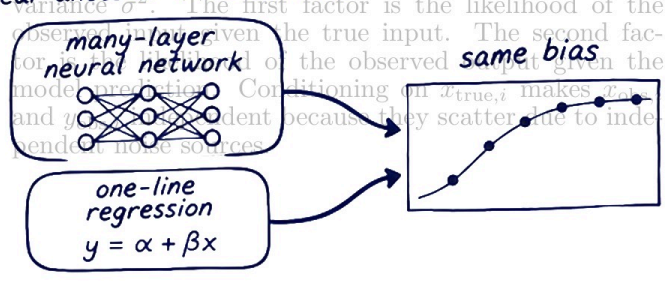
$$y_{\text{obs},i} = \beta x_{\text{true},i} + \delta_{y,i}, \quad \delta_{y,i} \sim \mathcal{N}(0, \sigma_y^2). \tag{4}$$

Here σ_y^2 is the known measurement uncertainty in the output. In principle, one could add a jitter term to account for intrinsic scatter about the true relationship, but we assume pure measurement noise for simplicity.

Given these assumptions, the joint probability of observing both $x_{\text{obs},i}$ and $y_{\text{obs},i}$ given the true value $x_{\text{true},i}$ and the model parameter β is

$$p(x_{\text{obs},i}, y_{\text{obs},i} | x_{\text{true},i}, \beta) = \mathcal{N}(x_{\text{obs},i} | x_{\text{true},i}, \sigma_x^2) \times \mathcal{N}(y_{\text{obs},i} | \beta x_{\text{true},i}, \sigma_y^2). \tag{5}$$

Here $\mathcal{N}(x|\mu, \sigma^2)$ denotes the probability density of observing x from a Gaussian distribution with mean μ and variance σ^2. The first factor is the likelihood of the observed input given the true input. The second factor is the likelihood of the observed output given the model prediction. Conditioning on $x_{\text{true},i}$ makes x_{obs} and y_{obs} independent because they scatter due to independent noise sources.



The correction comes from balancing fidelity to observations against consistency with one shared model.

Over all N observations, the full likelihood is

$$\mathcal{L}(\beta, \{x_{\text{true},i}\}) = \prod_{i=1}^{N} \mathcal{N}(x_{\text{obs},i} | x_{\text{true},i}, \sigma_x^2) \times \mathcal{N}(y_{\text{obs},i} | \beta x_{\text{true},i}, \sigma_y^2). \quad (6)$$

This formulation introduces $N + 1$ parameters—the slope β plus the true value $x_{\text{true},i}$ for each of the N data points—for $2N$ observed quantities.

The problem looks over-parameterized until the common relationship ties every latent input together.

At first glance this appears underdetermined. However, the problem is well-posed for two reasons. First, the linear relationship provides an inductive bias that constrains all $x_{\text{true},i}$ to be consistent with a single slope. The true values are not independent; they must all be consistent with the same underlying linear relationship. Second, the term involving x_{obs} acts as a regularizer, preventing the x_{true} values from drifting arbitrarily far from the observations. This balance between fitting the model and staying close to the data leads to the Deming regression solution, which has a closed form (see Appendix B).

This approach of treating unknown true values as parameters to be inferred is an example of hierarchical modeling (Loredo 2004; Kelly 2007; Gelman et al. 2013). In this hierarchical structure, the model parameter β sits at the top level, the true values $x_{\text{true},i}$ form the intermediate level, and the actual measurements $(x_{\text{obs}}, y_{\text{obs}})$ sit at the bottom level. The model constrains the intermediate level, while the intermediate level generates the observations.

The resulting Deming regression estimator (derived in Appendix B) recovers the true slope β without attenuation. When the error variances σ_x^2 and σ_y^2 are known, defining the error variance ratio $r = \sigma_y^2/\sigma_x^2$, the estimator is

$$\hat{\beta}_{\text{Deming}} = \frac{S_{yy} - rS_{xx} + \sqrt{(S_{yy} - rS_{xx})^2 + 4rS_{xy}^2}}{2S_{xy}}, \quad (7)$$

where $S_{xx} = \sum_i (x_{\text{obs},i} - \bar{x})^2$, $S_{yy} = \sum_i (y_{\text{obs},i} - \bar{y})^2$, and $S_{xy} = \sum_i (x_{\text{obs},i} - \bar{x})(y_{\text{obs},i} - \bar{y})$ are the sum of squared deviations and cross-products, with $\bar{x}$ and $\bar{y}$ denoting the sample means. In the limit $\sigma_x \to 0$ (equivalently $r \to \infty$), this reduces to the ordinary least squares estimator $\hat{\beta}_{\text{OLS}} = S_{xy}/S_{xx}$. When $\sigma_x > 0$, the Deming estimator yields a steeper slope than OLS, correcting for the attenuation bias (see also Ping 2025b).

Equation 8 is the intuition: combine the direct x measurement with the x implied by y.

The structure of the optimal $x_{\text{true},i}$ provides further intuition. For a given β, the optimal estimate is a weighted average

$$x_{\text{true},i} = w_x \cdot x_{\text{obs},i} + w_y \cdot \frac{y_{\text{obs},i}}{\beta}, \quad (8)$$

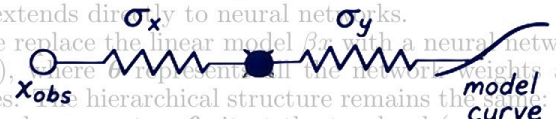
where the weights are (see Appendix B for derivation)

$$w_x = \frac{r}{r + \beta^2}, \quad w_y = \frac{\beta^2}{r + \beta^2}. \quad (9)$$

The weights sum to unity. When σ_x is small compared to σ_y/β (i.e., r is large), the weight w_x dominates and we trust the direct x measurement. When σ_x is large (i.e., r is small), the weight w_y dominates and we infer the true x from the y measurement and the model relationship. This optimal weighting is what allows the method to correct for attenuation bias. By appropriately balancing direct measurements against the constraints imposed by the model, we recover estimates of both the true parameter values and the true underlying data points.

4.2. *The LatentNN Formulation*

The linear form of Deming regression was not essential—what mattered was the hierarchical structure where latent values generate observations and the model constrains the latent values. The linear form permits an analytic solution, but the underlying principle of joint optimization over model parameters and latent true values extends directly to neural networks.



We replace the linear model βx with a neural network $f_{\boldsymbol{\theta}}(x)$, where $\boldsymbol{\theta}$ represents all the network weights and biases. The hierarchical structure remains the same: the network parameters $\boldsymbol{\theta}$ sit at the top level (analogous to β), the latent values $x_{\text{latent},i}$ form the intermediate level (analogous to $x_{\text{true},i}$), and the observations $(x_{\text{obs}}, y_{\text{obs}})$ sit at the bottom level. The likelihood has exactly the same form as in the linear case,

$$\mathcal{L}(\boldsymbol{\theta}, \{x_{\text{latent},i}\}) = \prod_{i=1}^{N} \mathcal{N}(x_{\text{obs},i} | x_{\text{latent},i}, \sigma_{x,i}^2) \times \mathcal{N}(y_{\text{obs},i} | f_{\boldsymbol{\theta}}(x_{\text{latent},i}), \sigma_{y,i}^2), \quad (10)$$

with $f_{\boldsymbol{\theta}}(x_{\text{latent},i})$ replacing $\beta x_{\text{true},i}$. Here $\sigma_{x,i}$ and $\sigma_{y,i}$ denote the per-sample uncertainties, which may vary across observations.

Taking the negative log-likelihood, rearranging, and absorbing constants, we obtain the loss function

$$\mathcal{J}(\boldsymbol{\theta}, \{x_{\text{latent},i}\}) = \sum_{i=1}^{N} \frac{(y_{\text{obs},i} - f_{\boldsymbol{\theta}}(x_{\text{latent},i}))^2}{\sigma_{y,i}^2} + \sum_{i=1}^{N} \frac{(x_{\text{obs},i} - x_{\text{latent},i})^2}{\sigma_{x,i}^2}. \quad (11)$$

The neural-network loss keeps exactly this two-way tug-of-war, without requiring a linear function.

Comparing to the Deming regression objective in Appendix B, we see the same structure: a prediction loss term weighted by $1/\sigma_{y,i}^2$ and a latent likelihood term weighted by $1/\sigma_{x,i}^2$. The balance between these terms is controlled by the uncertainties, which determines how much the latent values deviate from the observations to improve prediction. This is the same trade-off we saw in the weighted average formula for $x_{\text{true},i}$: trust the direct measurement when it is precise, but infer the true value from the model when the measurement is noisy. The uncertainties $\sigma_{x,i}$ and $\sigma_{y,i}$ must be specified for proper correction. In the experiments that follow, to facilitate our study and visualization of the results, we simplify to homogeneous uncertainties σ_x and σ_y across all samples.

The optimization problem is to find

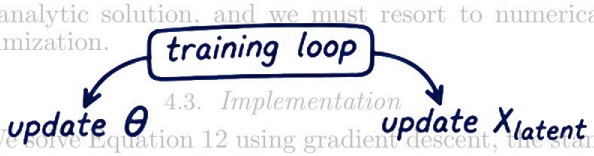
$$\hat{\boldsymbol{\theta}}, \{\hat{x}_{\text{latent},i}\} = \underset{\boldsymbol{\theta}, \{x_{\text{latent},i}\}}{\arg\min} \; \mathcal{J}(\boldsymbol{\theta}, \{x_{\text{latent},i}\}). \quad (12)$$

For a dataset with N samples and p input dimensions, this involves optimizing over the network parameters $\boldsymbol{\theta}$ plus $N \times p$ latent variables. The total number of parame-

Implementation is ordinary backpropagation with the training inputs promoted to learnable parameters.

ters is thus much larger than in standard neural network training, scaling with the training set size.

The key difference from the linear case is that Equation 12 has no closed-form solution. In Deming regression, the linear structure allows us to solve for both $\hat{\beta}$ and $\{\hat{x}_{\text{true},i}\}$ analytically (see Appendix B). For a neural network $f_{\boldsymbol{\theta}}$, the nonlinear dependence on $\boldsymbol{\theta}$ precludes an analytic solution, and we must resort to numerical optimization.



4.3. *Implementation*

We solve Equation 12 using gradient descent, the standard optimization algorithm for neural networks. Let $\Phi = (\boldsymbol{\theta}, \{x_{\text{latent},i}\})$ denote the full set of parameters to be optimized. At each iteration t, we update

$$\Phi^{(t+1)} = \Phi^{(t)} - \eta \nabla_{\Phi} \mathcal{J}, \tag{13}$$

where η is the learning rate and $\nabla_{\Phi}\mathcal{J}$ is the gradient of the loss with respect to all parameters. This is the same update rule used in standard neural network training, but with Φ now including both the network weights $\boldsymbol{\theta}$ and the latent inputs $\{x_{\text{latent},i}\}$. This represents a conceptual shift from typical neural network applications, where we optimize only over $\boldsymbol{\theta}$. Here, we treat the observed inputs as noisy measurements rather than ground truth, and optimize jointly over $\boldsymbol{\theta}$ and the latent inputs.

In modern deep learning frameworks such as PyTorch or JAX, this is straightforward to implement. The latent values $x_{\text{latent},i}$ are declared as additional learnable parameters, and the framework automatically computes gradients via backpropagation (Rumelhart et al. 1986; LeCun et al. 1998).

Several practical considerations affect performance. First, initialization matters. We initialize the latent values at the observed values:

$$x^{(0)}_{\text{latent},i} = x_{\text{obs},i}. \tag{14}$$

I initialize each latent input at its observation; starting elsewhere needlessly enlarges an already difficult search.

In practice, this means declaring a learnable parameter array of shape (N, p) initialized to the observed input matrix. This ensures that the initial loss is reasonable and provides a starting point consistent with our probabilistic model.

Second, additional L_2 regularization (weight decay) on the network parameters $\boldsymbol{\theta}$ is important. The loss function in Equation 11 has two terms: a prediction loss that encourages the network to fit the data, and a latent likelihood term that keeps the latent values close to the observations. With $N \times p$ latent variables plus thousands of network parameters, the optimization can find solutions where the latent values drift from the observations if the network can learn a systematic bias that creates a more favorable mapping. Weight decay on the network parameters prevents this by penalizing large weights, encouraging smoother functions that generalize better. Weight decay is applied only to $\boldsymbol{\theta}$, not to the latent values x_{latent}—the latent values are already constrained by the latent likelihood term.

Regularize the network weights, not the latent inputs—the likelihood already anchors those inputs.

5. NUMERICAL EXPERIMENTS

We now validate the LatentNN approach on test cases where the true relationship and measurement errors are known, allowing direct comparison between estimated and true attenuation factors. We begin with the one-dimensional case then extend to multivariate inputs with correlated features.

5.1. *One-Dimensional Case*

We use the same data setup and network architecture as in Section 3: the true relationship $y = 2x_{\text{true}}$ with 1000 training samples, and an MLP with 2 hidden layers of 64 units trained for 20000 epochs. For LatentNN, we add a learnable parameter $x_{\text{latent},i}$ for each training sample, initialized to $x_{\text{obs},i}$.

To ensure fair comparison with no data leakage, we use a three-fold data split: 1000 samples for training, 200 for validation (hyperparameter tuning), and 200 for final evaluation. The validation set is used exclusively for selecting hyperparameters—for LatentNN we maximize λ_y on the validation set, while for the standard MLP we minimize validation loss. Final performance metrics are computed on the held-out test set, which is never used during hyperparameter selection and the training process.

Validation matters because low-SNR solutions can be unbiased on average yet unstable across random starts.

Figure 3 shows a representative comparison at $\text{SNR}_x = 1$, using $\text{SNR}_y = 10$ as in Section 3. At $\text{SNR}_x = 1$, the measurement error equals the signal range ($\sigma_x = \sigma_{\text{range}}$)—an extreme regime where the noise is 100% of the data's dynamic range. The top panels show predicted versus true values on the held-out test set. The standard MLP shows $\lambda_y \approx 0.5$, demonstrating large attenuation bias. LatentNN recovers $\lambda_y \approx 1$, correcting the bias. The bottom-left panel shows the learned functions: the standard MLP learns an attenuated slope, while LatentNN recovers the true relationship $f(x) = 2x$.

The bottom-right panel shows the training dynamics for LatentNN. At initialization, $x_{\text{latent}} = x_{\text{obs}}$, so the x_{latent} likelihood term starts near zero. As training proceeds, the prediction loss decreases as the network learns the mapping, but the x_{latent} likelihood increases. This increase is expected and desirable: to achieve accurate predictions, the latent values must shift from the noisy observations x_{obs} toward the true values x_{true}, which increases the likelihood penalty. The total loss decreases overall because the improvement in prediction accuracy outweighs the likelihood cost. At convergence, the optimization has found the balance point where the latent values have moved sufficiently toward their true positions to enable unbiased prediction.

To systematically evaluate performance, we repeat each experiment 8 times with different random seeds, using the same three-fold data split (train/validation/test) for each run. Figure 4 shows λ_y versus SNR_x across a range of noise levels from $\text{SNR}_x = 1$ to $\text{SNR}_x = 10$, testing three SNR_y values (3, 10, and 30). As expected from the theory in Section 2, varying SNR_y has little effect on attenuation—the bias depends primarily on input noise. The standard MLP follows the theoretical attenuation curve $\lambda_y = 1/(1 + (\sigma_x/\sigma_{\text{range}})^2)$ regardless of SNR_y. LatentNN maintains $\lambda_y \approx 1$ across all tested SNR_x and SNR_y values. LatentNN provides robust correction even at $\text{SNR}_x = 1$, though with increased run-to-run scatter at lower SNR_x values as the problem becomes increasingly ill-conditioned.

The mechanism of correction differs between SNR regimes. At high SNR where $\sigma_x/\sigma_{\text{range}} \ll 1$, the x_{latent} likelihood term in Equation 11 dominates and the latent values remain close to the observations ($x_{\text{latent}} \approx x_{\text{obs}}$).

At SNR = 1, LatentNN recovers the missing slope by allowing training inputs to move.

The real test is the held-out set, where no test inputs were optimized.

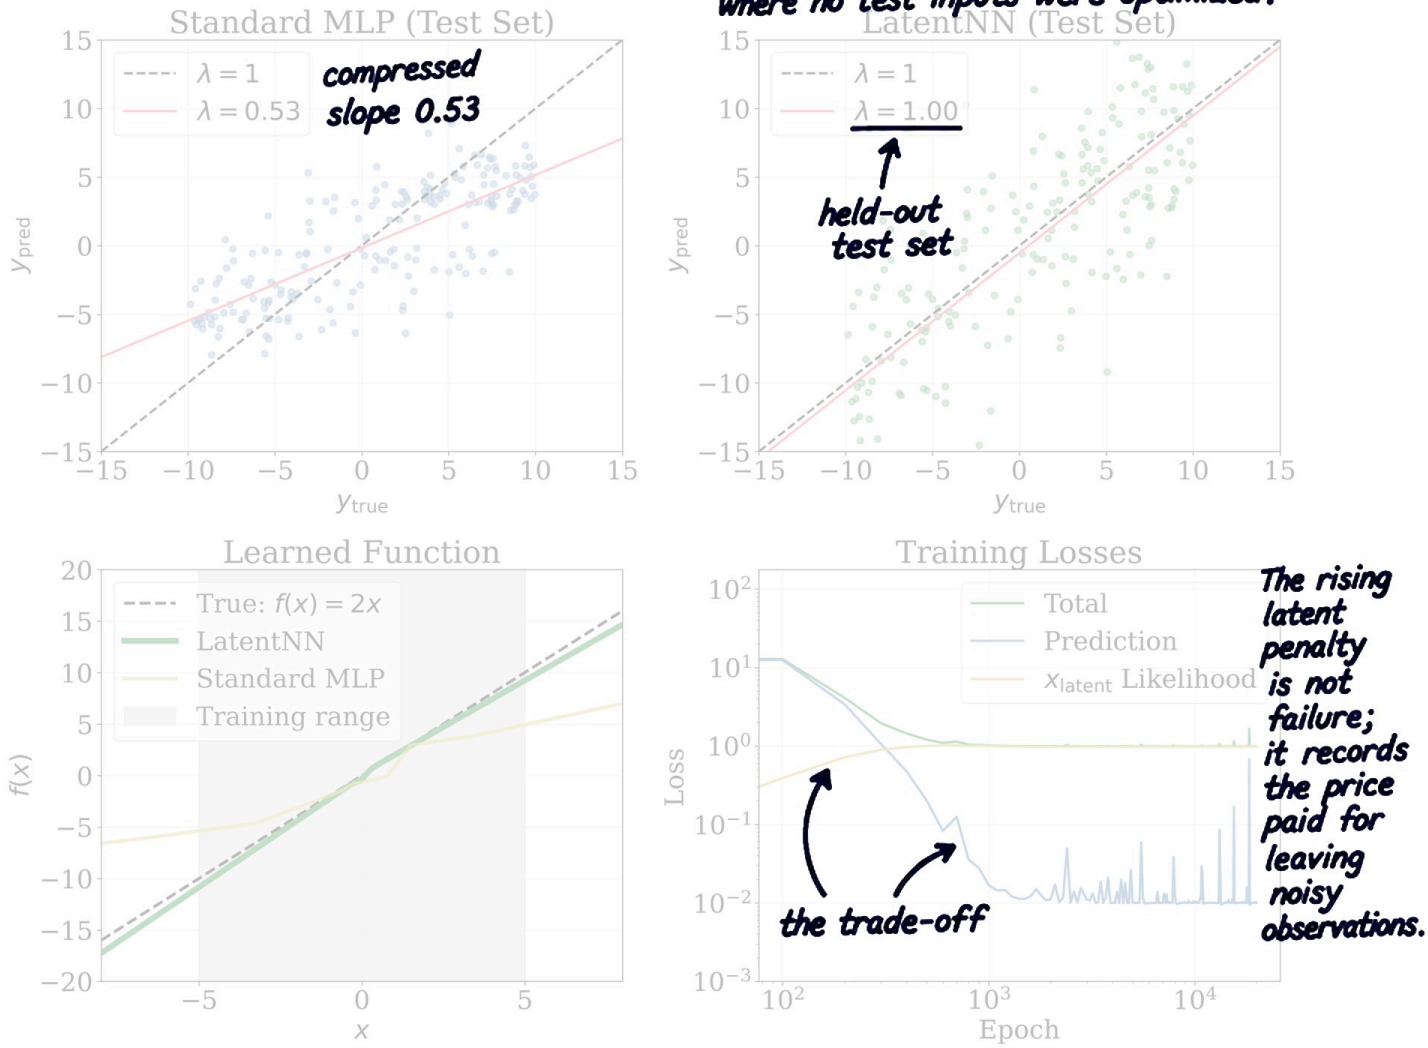


FIG. 3.— LatentNN demonstration at $\mathrm{SNR}_x = 1$. **Top:** Predicted versus true y on the test set for standard neural networks with the MLP arhitecture (left, $\lambda_y \approx 0.5$) and LatentNN (right, $\lambda_y \approx 1$). LatentNN corrects the attenuation bias. **Bottom left:** Learned functions. The standard MLP (orange) learns an attenuated slope, while LatentNN (green) recovers the true function $f(x) = 2x$ (black dashed). The gray shaded region indicates the noiseless training x-value range. **Bottom right:** LatentNN training losses. The prediction loss decreases while the x_{latent} likelihood increases, reflecting the trade-off as latent values shift from noisy observations toward true values.

At moderate SNR ($\mathrm{SNR}_x \approx 2$–3), the two loss terms compete more evenly, and weight decay tuning becomes important for LatentNN. The standard MLP, by contrast, fails regardless of weight decay—it lacks the mechanism to correct attenuation bias, and adjusting regularization only trades off between overfitting and underfitting without addressing the systematic bias.

LatentNN generalizes well to held-out test data. Weight decay regularization applied only to the network parameters $\boldsymbol{\theta}$ prevents the network from overfitting to the specific denoised distribution of the training latent values. The test set performance closely tracks the training set performance across all SNR values, indicating that the learned function $f_{\boldsymbol{\theta}}$ captures the true underlying relationship rather than artifacts of the training data.

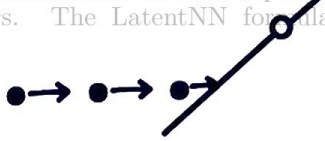
5.2. *Multivariate Inputs*

Real astronomical applications involve multivariate inputs—for example, stellar spectra contain many flux measurements that serve as inputs to predict stellar parameters. The LatentNN formulation generalizes directly: the latent variables $\mathbf{X}_{\mathrm{latent}}$ form an $N \times p$ matrix, where p is the number of input features. The loss function from Equation 11 extends naturally to

$$\mathcal{J} = \sum_{i=1}^{N} \frac{(y_i - f_{\boldsymbol{\theta}}(\mathbf{X}_{\mathrm{latent},i}))^2}{\sigma_{y,i}^2} + \sum_{i=1}^{N} \sum_{j=1}^{p} \frac{(X_{\mathrm{obs},ij} - X_{\mathrm{latent},ij})^2}{\sigma_{x,ij}^2}, \tag{15}$$

where $\sigma_{y,i}$ is the output uncertainty for sample i and $\sigma_{x,ij}$ is the input uncertainty for sample i and dimension j. For simplicity, we assume homogeneous noise σ_x and σ_y across all samples and dimensions, which allows us to visualize results as a function of a single SNR_x.

As shown in Paper I, for uncorrelated features each dimension attenuates independently, recovering the one-dimensional result. We therefore focus on features that are intrinsically correlated in the signal, which better represent astronomical data like stellar spectra where different pixels respond coherently to changes in stellar param-



The learned function, rather than perfect recovery of every latent point, is the object we need.

LatentNN stays near unbiased across noise levels, while correlated feature redundancy helps only after enough dimensions accumulate.

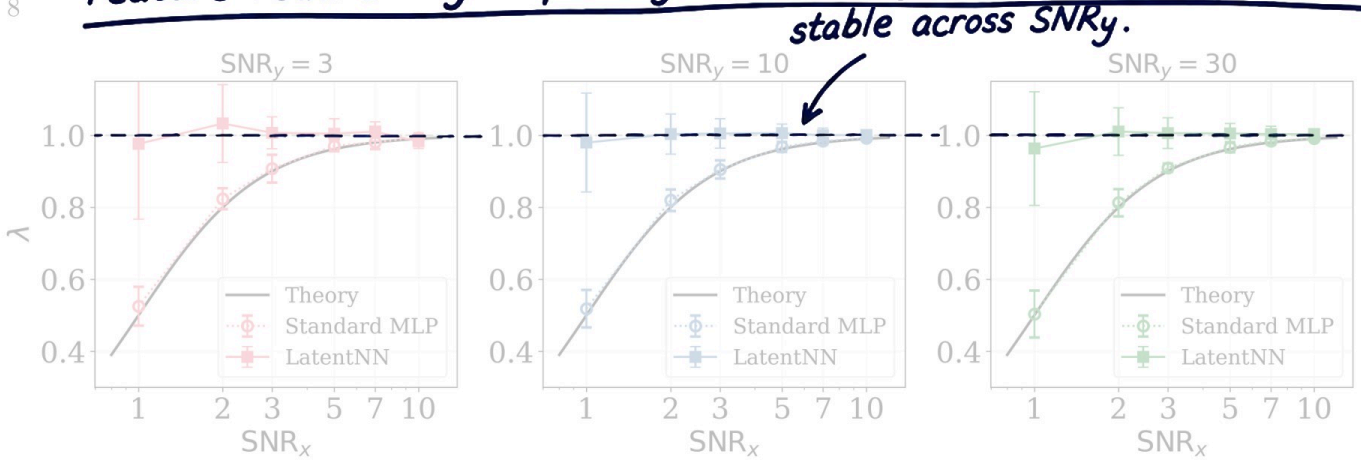


FIG. 4.— Attenuation factor λ_y versus SNR_x for three SNR_y levels (3, 10, 30; different colors). Error bars show standard deviation over 8 runs with different random seeds, using a three-fold data split (train/validation/test) to avoid data leakage in hyperparameter selection. Standard MLP (open circles, dotted) follows the theoretical attenuation curve (black solid) regardless of SNR_y. LatentNN (filled squares, solid) maintains mean $\lambda_y \approx 1$ across all tested SNR values. Even at $\mathrm{SNR}_x = 1$, LatentNN achieves unbiased predictions on average, though with increased run-to-run scatter at lower SNR_x as the problem becomes increasingly ill-conditioned.

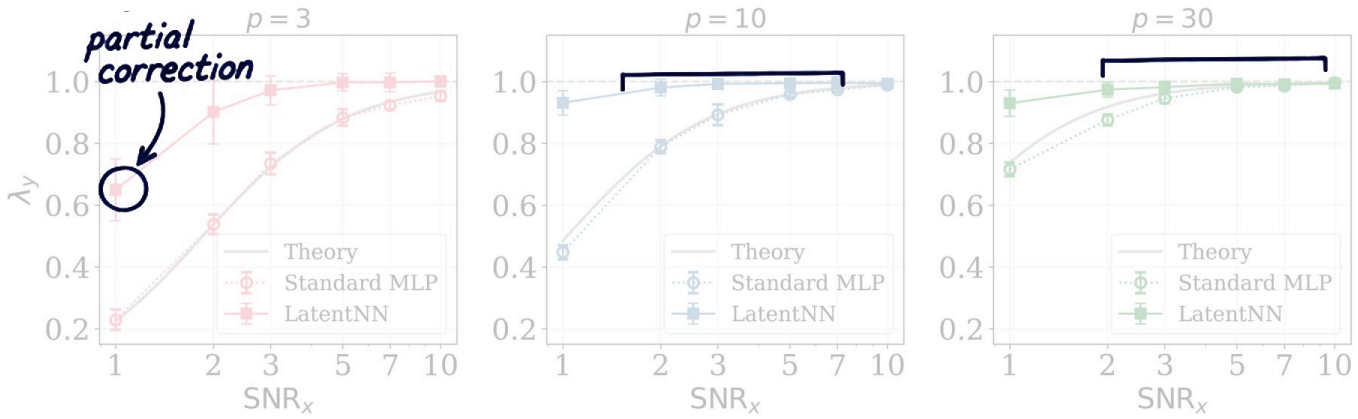


FIG. 5.— Attenuation factor λ_y versus SNR_x for correlated multivariate inputs with $p = 3, 10, 30$ dimensions (panels). Error bars show standard deviation over 8 runs. Black dashed line marks $\lambda_y = 1$ (no attenuation). Gray solid curve shows the analytic prediction for linear regression from Paper I; standard MLP (open circles, dotted) closely follows this theoretical curve, confirming that model complexity does not mitigate the bias. With increasing p, attenuation is mitigated because independent noise in different dimensions must conspire coherently to displace observations along the signal direction. LatentNN (filled squares, solid) maintains $\lambda_y \approx 1$ for $p \geq 10$ at $\mathrm{SNR}_x \geq 2$, with mild degradation ($\lambda_y \approx 0.9$) at $\mathrm{SNR}_x = 1$. For $p = 3$, LatentNN correction degrades at low SNR ($\lambda_y \approx 0.7$ at $\mathrm{SNR}_x = 1$), but still outperforms the standard MLP ($\lambda_y \approx 0.2$).

Changing label SNR barely moves the standard model; input noise remains the controlling variable.

eters. By "correlated features" we mean that the true input values $x_{\mathrm{true},j}$ are correlated with each other—not that the measurement errors are correlated. The measurement errors remain independent across features; it is the underlying signal that exhibits correlation structure.

To create such correlated features, we generate $N_{\mathrm{train}} = 1000$ training samples where all features derive from a single latent variable: $x_j = a_j z$ where $z \sim \mathrm{Uniform}(-0.5, 0.5)$ and a_j varies linearly from 0.2 to 0.4 across dimensions. This construction ensures that the p input features are perfectly correlated (all driven by the same underlying z), mimicking how spectral pixels respond coherently to changes in a stellar parameter. The true output is $y = \mathbf{x}^\top \mathbf{a}$, centered at zero. We fix $\mathrm{SNR}_y = 10$ as in the single-variate experiments. We use the same [illegible] as Section 3, with the same three-fold split (600 train, 200 validation, 200 test) and weight decay grid search strategy. We test $p = 3, 10, 30$ dimensions at SNR_x values from 1 to 10, repeating each experiment 8 times.

Figure 5 shows the results. Paper I derived an analytic expression for the attenuation factor in correlated linear regression: $\lambda_\beta = \sum_j a_j^2/(1/\mathrm{SNR}_x^2 + \sum_j a_j^2)$. The gray curves show this theoretical prediction. Just as in the single-variate case, the standard MLP closely follows this [illegible] that additional model complexity does not mitigate the fundamental attenuation bias. With correlated features, attenuation is weaker than the $p = 1$ case because independent noise in different dimensions must conspire coherently to displace observations along the signal direction—a coincidence that becomes increasingly unlikely as p grows. With increasing p, the standard MLP attenuation factor is mitigated compared to smaller p, but remains significant at lower SNR.

Ten or thirty coherent features can cross-check one another's noisy measurements.

For LatentNN, the results reveal a non-monotonic relationship between dimensionality and correction difficulty. For $p \geq 10$, LatentNN maintains $\lambda_y \approx 1$ for $\mathrm{SNR}_x \geq 2$, with only mild degradation ($\lambda_y \approx 0.9$) at $\mathrm{SNR}_x = 1$. For $p = 3$, however, [illegible] degrades at low SNR: $\lambda_y \approx 0.7$ at $\mathrm{SNR}_x = 1$ and $\lambda_y \approx 0.9$ at $\mathrm{SNR}_x = 2$. Comparing to the single-variate case (Figure 4), where

Three correlated features are an awkward middle ground: more latent freedom, but not yet much redundancy.

three features (disagree) → thirty features (consensus)

Spectra are a natural stress test because many pixels vary only weakly compared with their measurement noise.

$p = 1$ achieves $\lambda_y \approx 1$ even at $\mathrm{SNR}_x = 1$, we see that intermediate dimensionality is harder than both low and high p.

This pattern reflects competing effects. Increasing p generally helps because correlated features can lean on each other—the network learns their covariance structure and uses redundancy for more robust inference. However, this benefit only manifests when p is large enough. For $p = 1$, the problem is simple with a tractable one-dimensional posterior. For $p = 3$, the posterior over latent values becomes more complex due to the increased dimensionality, yet there is insufficient redundancy among the three features to leverage. For $p \geq 10$, the higher-dimensional correlated structure provides enough constraints that features can collaboratively inform each other's latent values, enabling robust correction for $\mathrm{SNR}_x \geq 2$. Nevertheless, even in the challenging $p = 3$ regime, LatentNN substantially outperforms the standard MLP: at $\mathrm{SNR}_x = 1$, LatentNN achieves $\lambda_y \approx 0.7$ compared to the standard MLP's $\lambda_y \approx 0.2$, more than tripling the recovered dynamic range.

6. APPLICATION TO STELLAR SPECTRA

Having validated LatentNN on synthetic examples, we now apply it to a more realistic astronomical setting.

Stellar spectroscopy provides a natural and important application motivating this work. To clarify the notation: in this spectroscopic context, the input vector $\mathbf{x}$ consists of flux measurements at different wavelength pixels, while the output y is the stellar parameter we wish to infer (e.g., metallicity [M/H]). This follows our previous examples where $\mathbf{x}$ denotes the observed quantities with measurement error and y is the label to be predicted. Data-driven models have been applied extensively to spectroscopic surveys (Blanco-Cuaresma et al. 2014; Ness et al. 2015; Ting et al. 2017; Fabbro et al. 2018; Leung & Bovy 2019; Straumit et al. 2022; Xiang et al. 2022), particularly for low-resolution and low-SNR observations where theoretical spectral models often fall short. These applications operate in the regime where attenuation bias is large. Even at high resolution ($R \sim 20000$), the intrinsic variation in continuum-normalized flux typically spans only ∼5–20% ($\sigma_{\mathrm{range}} \approx 0.05$–$0.2$).

conventional SNR ≠ feature SNR

For the weaker lines of $\sigma_{\mathrm{range}} \approx 0.05$, a spectrum with conventional SNR = 20 (defined relative to the normalized flux) has per-pixel uncertainty $\sigma_x = 0.05$, giving $\mathrm{SNR}_{\mathrm{range}} = \sigma_{\mathrm{range}}/\sigma_x = 1$. Thus, conventional SNR = 20 corresponds to $\mathrm{SNR}_{\mathrm{range}} = 1$. We emphasize that what matters for attenuation bias is the ratio of signal range to measurement uncertainty, not the conventional spectroscopic SNR. Since the ratio $\sigma_{\mathrm{range}}/\sigma_x$ is often of order 10 or less for many spectral pixels in typical spectroscopic applications, they fall squarely in the regime where attenuation bias becomes important—especially for elements with weak or limited spectral features.

Spectroscopic SNR can sound healthy while the relevant feature-variation SNR is close to one.

Real applications of LatentNN to survey data are underway, but this paper focuses on introducing the technique and establishing its theoretical foundation as an extension of Deming regression to neural networks. Here we focus on inferring overall metallicity [M/H] for clarity, but the formalism generalizes directly to multiple labels such as individual elemental abundances.

We use the pretrained synthetic spectral emulator with The Payne (Ting et al. 2019) in The Payne github repository[2] to generate spectra at resolution $R \simeq 22500$, following Paper I. The underlying spectral model uses atomic and molecular line data from Kurucz (1970, 1993, 2005, 2017).[3] We focus on a metallicity-sensitive region around 16200 Å.

Concretely, the experimental setup is as follows: we generate synthetic spectra for stars with varying [M/H] values, then train a neural network to predict [M/H] from the observed (noisy) flux values. The input $\mathbf{x}_i$ for sample i is a vector of continuum-normalized flux measurements at p wavelength pixels, and the output y_i is the metallicity [M/H]. We add Gaussian noise to the flux values to simulate observational uncertainty, giving $\mathbf{x}_{\mathrm{obs},i} = \mathbf{x}_{\mathrm{true},i} + \boldsymbol{\delta}_x$ where $\boldsymbol{\delta}_x \sim \mathcal{N}(0, \sigma_x^2 \mathbf{I})$. The attenuation factor λ_y is defined as before: the slope of the regression of predicted [M/H] against true [M/H] on a held-out test set.

To explore how the number of informative features affects attenuation and correction, we test three pixel configurations: 3, 10, and 30 contiguous pixels centered at 16200 Å. These choices represent the range of spectral information available for different labels. While major elements like Fe have many strong lines, attenuation bias is most problematic for elements with only a few spectral features (e.g., K, V; see fig. 7 and 8 of Xiang et al. 2019). In practice, even surveys with thousands of pixels often have only a small number of pixels that contain useful information for a given label.

I vary the number of pixels because information for rare elements may live in only two or three lines.

We generate samples spanning [M/H] from −2.0 to 0.0 dex, with fixed stellar parameters ($T_{\mathrm{eff}} = 4750$ K, $\log g = 2.5$, typical of red clump stars). All abundances scale with [M/H] for simplicity. We use the same threefold data split as in the synthetic experiments (1000 train, 200 validation, 200 test) with hyperparameters selected on the validation set and final evaluation on the held-out test set. We add Gaussian noise with standard deviation σ_x to each flux pixel and $\sigma_{[\mathrm{M/H}]} = 0.05$ dex to the metallicity labels, representing conservative uncertainties in training labels from surveys like APOGEE (Holtzman et al. 2015; García Pérez et al. 2016). We test flux error values $\sigma_x \in \{1\%, 2\%, 3\%, 5\%, 10\%\}$ of the normalized flux. We use the same network architecture and training settings as Section 3, with the same weight decay grid search strategy from Section 5. We repeat each experiment 8 times to estimate uncertainties.

Figure 6 shows the results. The top panel illustrates how the spectral flux varies with [M/H] across the selected wavelength region, with color indicating metallicity from −2.0 (blue) to 0.0 dex (red). Feature depths range from 5–20%, but even features with 10–20% total depth do not give $\mathrm{SNR}_x = 2$–4 at 5% flux error as one might naively expect. This is because spectral features vary approximately quadratically with metallicity (Rix et al. 2016). In Paper I, we showed that the effec-

[2] https://github.com/tingyuansen/The_Payne

[3] This emulator from the github repository of The Payne is trained on stellar models spanning [Fe/H] down to −1.5 dex. Our experiment extends to −2.0 dex, so we are extrapolating slightly. However, as seen in the top panel of Figure 6, the spectral variation remains smooth and physically reasonable. The exact accuracy of the spectral model is not critical to our conclusions about attenuation bias—the goal is to generate realistic correlated features mimicking the physical coupling between spectral pixels.

The synthetic spectra isolate attenuation cleanly before we confront survey systematics.

In mock stellar spectra, ten informative pixels are usually enough for LatentNN to preserve the metallicity range.

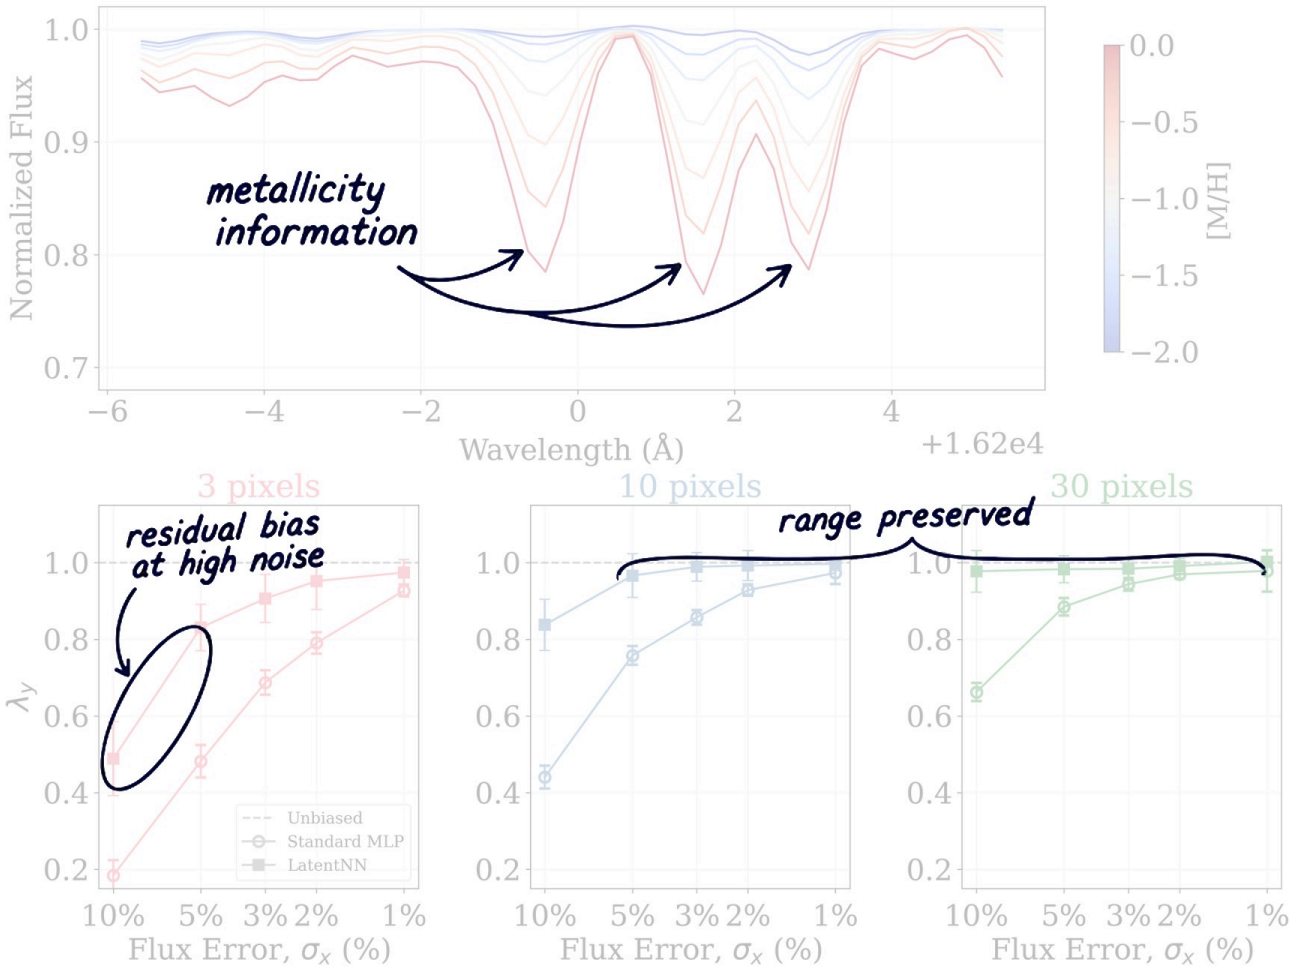


FIG. 6.— Application to stellar spectra at $R \simeq 22500$ for inferring [M/H]. **Top:** Continuum-normalized flux variation with [M/H] from −2.0 dex (blue) to 0.0 dex (red) in a metallicity-sensitive region around 16200 Å. **Bottom:** Attenuation factor λ_y versus flux error (%) for 3, 10, and 30 contiguous pixels (panels). Error bars show standard deviation over 8 runs. The standard MLP (open circles) exhibits attenuation that worsens with fewer pixels and higher noise. LatentNN (filled squares) maintains $\lambda_y \gtrsim 0.95$ for 10 and 30 pixels across all noise levels. Even for 3 pixels where correction is more challenging, LatentNN also ameliorate the attenuation bias compared to the standard MLP (e.g., $\lambda_y \approx 0.5$ vs. 0.2 at 10% flux error).

Quadratic spectral responses make the effective SNR worse than line depth divided by flux error suggests.

tive SNR is reduced by at least a factor of n for polynomial features of degree n. For quadratic features, even $\sigma_{\rm range} = 0.1$ with $\sigma_x = 0.05$ yields an effective $\mathrm{SNR}_x \approx 1$ rather than 2.

The bottom panels show λ_y versus flux error for each pixel configuration. The behavior mirrors the synthetic multivariate experiments (Section 5.2): the standard MLP exhibits attenuation that worsens with fewer pixels and higher noise, while LatentNN provides robust correction when sufficient pixels are available. For 10 and 30 pixels, LatentNN maintains $\lambda_y \gtrsim 0.95$ across all tested noise levels. For 3 pixels, however, correction degrades at high noise: $\lambda_y \approx 0.5$ at 10% flux error and $\lambda_y \approx 0.8$ at 5% flux error. This matches the intermediate-dimensionality challenge discussed in Section 5.2—with few pixels, the posterior over latent values is complex yet there is insufficient redundancy to leverage. Even so, LatentNN outperforms the standard MLP across all configurations: at 10% flux error with 3 pixels, LatentNN achieves $\lambda_y \approx 0.5$ compared to the standard MLP's $\lambda_y \approx 0.2$.

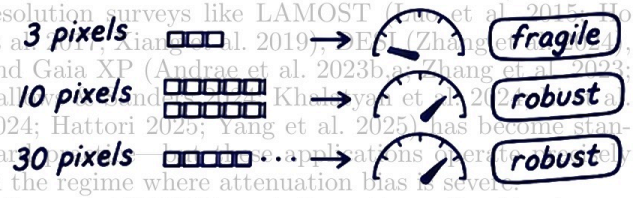
With only three pixels and heavy noise, the correction helps substantially but cannot perform miracles.

For spectral features with $\sigma_{\rm range} \approx 0.10$–$0.20$ (strong lines), flux errors of $\sigma_x \lesssim 3$–5% yield effective $\mathrm{SNR}_x \gtrsim 2$, naively accounting for the factor of ∼2 reduction from quadratic features—consistent with our findings throughout that LatentNN works well for $\mathrm{SNR}_x \geq 2$. For weaker features with $\sigma_{\rm range} \approx 0.05$, correspondingly lower flux errors ($\sigma_x \lesssim 1$–2%) are needed. At higher noise levels with few informative pixels, careful hyperparameter tuning becomes essential, and users should expect some residual bias.

7. DISCUSSION

7.1. *Implications for Spectroscopic Surveys*

The approach of transferring stellar parameters and abundances from high-resolution surveys like APOGEE (Holtzman et al. 2015; Majewski et al. 2017) to lower-resolution surveys like LAMOST (Luo et al. 2015; Ho et al. 2017; Xiang et al. 2019), SDSS (Zhang et al. 2023), and Gaia XP (Andrae et al. 2023b,a; Zhang et al. 2023; Fallows & Sanders 2024; Khalatyan et al. 2024; Li et al. 2024; Hattori 2025; Yang et al. 2025) has become standard practice—but these applications operate precisely in the regime where attenuation bias is severe.



Even at APOGEE resolution, attenuation bias is non-negligible at moderate SNR. For lower-resolution surveys where individual spectral features are blended and per-

This is the practical boundary: enough correlated information must survive in the inputs.

Compressed predictions can erase the metal-poor tail and distort the Galactic history inferred from it.

pixel $\sigma_{\rm range}$ is smaller, the ratio $\sigma_x/\sigma_{\rm range}$ is larger and attenuation bias is more severe. Crucially, $\sigma_{\rm range}$ scales approximately proportionally with spectral resolution R: a survey like LAMOST ($R \sim 1800$) has $\sigma_{\rm range}$ roughly 10 times smaller than APOGEE ($R \sim 22500$). This means that even at high conventional SNR ~ 100 ($\sigma_x \sim 1\%$), the effective SNR_x remains in the regime where attenuation bias is significant.

Our spectroscopic experiments (Figure 6) quantify the severity. For 10 informative pixels at 10% flux error (conventional SNR ≈ 10), the standard MLP yields $\lambda_y \approx 0.5$. Assuming training data centered around [M/H] $= -1$, a halo star with true [M/H] $= -2$ would be predicted as [M/H] ≈ -1.5, potentially misclassified as a thick-disk star. Even at 5% flux error (conventional SNR ≈ 20), we find $\lambda_y \approx 0.7$ for 10 pixels—a star at [M/H] $= -2$ would be predicted as [M/H] ≈ -1.7. While this might seem acceptable for elements with many features, it is problematic for elements with fewer than 10 informative spectral lines. For 3 pixels at 10% error, λ_y drops to ≈ 0.2, meaning a [M/H] $= -2$ star would be predicted as [M/H] ≈ -1.2.

A halo star can be moved toward the thick-disk regime without any obvious failure flag.

This has practical consequences—metal-poor stars are systematically predicted to be more metal-rich, making the rarest objects harder to identify. This bias directly affects the inferred metallicity distribution function (MDF), particularly its metal-poor tail, which is crucial for constraining the formation history of the proto-Milky Way (Rix et al. 2022; Chen et al. 2024, 2025). If attenuation bias compresses the metallicity scale, the true slope of the metal-poor MDF tail would be steeper than observed, with implications for inferred gas masses, star formation efficiencies, and inflow histories in galactic chemical evolution models. Stellar ages suffer from compression at both ends, with the oldest stars underestimated and the youngest overestimated. Distances are similarly affected, leading to systematic errors in the inferred geometry and dynamics of the Milky Way. Obtaining better training labels does not solve the problem because the bias arises from noisy inputs, not labels.

When information is concentrated in sparse spectral features such as individual absorption lines, attenuation is most severe. Consider an element like K or V with only 3 usable lines, compared to Fe with dozens of strong features: K and V abundances are effectively in the low-dimensional regime where $\sigma_x/\sigma_{\rm range}$ is large, while Fe benefits from the high-dimensional correlated regime. This pattern—elements with fewer lines showing larger systematic errors at extreme values—has been observed in practice (Xiang et al. 2019; Zhang et al. 2024) but is often attributed to model limitations or data quality issues. Attenuation bias is a fundamental contributor, and the LatentNN framework provides a path toward unbiased parameter estimation for data-driven deep learning application to spectroscopic data.

Better labels do not repair this; the damaging uncertainty sits in the spectrum supplied as input.

7.2. *Connection to Hierarchical Bayesian Modeling*

The errors-in-variables problem has a long history in statistics (Deming 1943; Fuller 1987). Deming regression and orthogonal regression (Golub & Van Loan 1980; Van Huffel & Vandewalle 1991) treat the true input values as latent variables to be estimated alongside the regression coefficients—but this approach has seen limited application in astronomy (but see Kelly 2007), in part because astronomical relationships are often nonlinear and proper feature engineering can be difficult. Standard neural network frameworks do not include errors-in-variables handling by default because they arose in domains like computer vision where input noise is negligible. LatentNN extends the classical latent variable approach to neural networks, providing a principled framework for astronomical applications where input uncertainties are non-negligible.

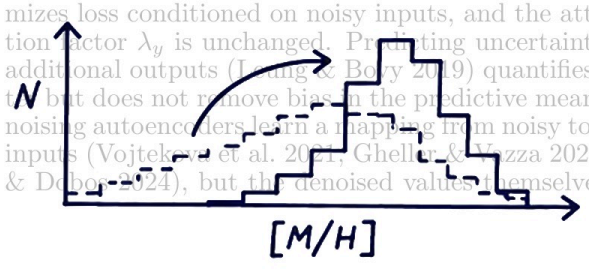
Our LatentNN formulation can be understood as the maximum a posteriori (MAP) limit of a hierarchical Bayesian model. In the full Bayesian formulation, the true values $\mathbf{x}_{\rm true}$ are latent variables with prior distributions centered on the observed values $\mathbf{x}_{\rm obs}$, and we would infer the posterior distribution over both the model parameters $\boldsymbol{\theta}$ and all the latent values. The LatentNN optimization finds the mode of this posterior. Neural networks provide a flexible function approximator that avoids the need for explicit feature engineering, while the latent variable framework provides the statistical foundation for handling input uncertainties.

This formulation provides the foundation for understanding why correlated features can lean on each other, as discussed in Section 5.2. When the model learns the correlated structure of the inputs, it can use this structure to constrain the latent values: each feature's latent value is informed not only by its own observation but also by what the learned model expects given the other features.

It is important to note that the latent values $\mathbf{x}_{\rm latent}$ should not be interpreted strictly as "denoised" inputs. In high-dimensional settings with large p, the posterior distribution over $\mathbf{x}_{\rm latent}$ can be broad and complex—many combinations of latent values may be consistent with the observed data and the learned model. What matters for our purposes is that the network parameters $\boldsymbol{\theta}$ remain well-constrained: even when individual latent values have substantial uncertainty, the function f_θ that maps inputs to outputs is well-defined. The optimization finds a configuration where the network generalizes correctly to new data, achieving unbiased predictions ($\lambda_y \approx 1$) even if the individual latent values do not perfectly recover the true inputs. This is the sense in which LatentNN corrects attenuation bias—through the learned function rather than through perfect recovery of latent values. At low SNR or intermediate dimensionality where the posterior is particularly complex (the transition regime in Figure 5), case-by-case hyperparameter tuning may be needed, but a systematic exploration of these regimes is beyond the scope of this paper.

Latent values are scaffolding—the scientifically useful product is an unbiased mapping that generalizes.

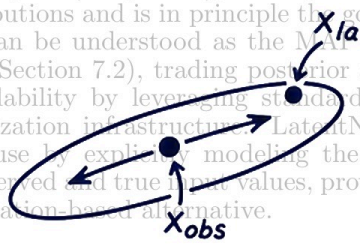
It is worth comparing LatentNN with alternative strategies that might address input uncertainties. Training with noise-augmented data (Fabbro et al. 2018), resampling input noise at each epoch—improves robustness but does not correct the systematic bias. Each augmented sample is still drawn from $\mathcal{N}(x_{\rm obs}, \sigma_x^2)$, so the expected input remains $x_{\rm obs}$; the optimization still minimizes loss conditioned on noisy inputs, and the attenuation factor λ_y is unchanged. Predicting uncertainties as additional outputs (Lueng & Bovy 2019) quantifies scatter but does not remove bias in the predictive mean. Denoising autoencoders learn a mapping from noisy to clean inputs (Vojtekova et al. 2021; Ghellar & Vazza 2022; Pál & Dobos 2024), but the denoised values themselves can



The latent-input likelihood generalizes cleanly, but its uncertainty model must match the instrument.

exhibit attenuation bias when the autoencoder is trained on noisy data. Bayesian neural networks place distributions over the network *weights* to capture epistemic uncertainty, but do not model uncertainty in the *inputs*. The posterior predictive mean $\mathbb{E}_{\boldsymbol{\theta}}[f_{\boldsymbol{\theta}}(\mathbf{x}_{\rm obs})]$ is still conditioned on the noisy $\mathbf{x}_{\rm obs}$ rather than the unknown $\mathbf{x}_{\rm true}$; so attenuation bias in the mean prediction is unchanged by construction. Marginalizing over weight uncertainty does not correct a bias that arises from the inputs. Variational autoencoders (Kingma & Welling 2013) learn latent representations, but their latent space is a lower-dimensional learned embedding rather than a corrected reconstruction of the physical input space with known noise properties. Gaussian process regression can in principle incorporate input noise, but standard GP scaling ($\mathcal{O}(N^3)$ in time) makes this impractical for the dataset sizes typical in astronomy. The approach most closely related to LatentNN is full hierarchical Bayesian inference with latent true inputs (Kelly 2007), which provides full posterior distributions and is in principle the gold standard. LatentNN can be understood as the MAP limit of this framework (Section 7.2), trading posterior characterization for scalability by leveraging standard neural network optimization infrastructure. LatentNN addresses the root cause by explicitly modeling the relationship between observed and true input values, providing a scalable optimization-based alternative.

Noise augmentation teaches robustness around x_{obs}; it does not move the expected input back toward x_{true}.



7.3. *Generalizations*

The core idea of maximizing a likelihood with latent variables extends to many settings beyond the homoscedastic Gaussian case demonstrated here. Known weights in Equation 15, heteroscedastic uncertainties are handled by per-sample weights. If measurement uncertainties are correlated across pixels, as is common in spectroscopy due to sky subtraction and continuum normalization, the latent likelihood term incorporates the noise covariance matrix $\boldsymbol{\Sigma}_{x,i}$:

Correlated pixel noise belongs in a covariance matrix, not in independent per-pixel error bars.

$$\mathcal{J} = \sum_{i=1}^{N} \frac{1}{\sigma_{y,i}^2}(y_i - f_{\boldsymbol{\theta}}(\mathbf{x}_{{\rm latent},i}))^2 + \sum_{i=1}^{N} (\mathbf{x}_{{\rm obs},i} - \mathbf{x}_{{\rm latent},i})^\top \boldsymbol{\Sigma}_{x,i}^{-1} (\mathbf{x}_{{\rm obs},i} - \mathbf{x}_{{\rm latent},i}). \tag{16}$$

For Poisson noise, outliers, or other non-Gaussian errors, the Gaussian likelihood can be replaced with the appropriate distribution, giving $\mathcal{J} = -\log p(\mathbf{x}_{\rm obs}|\mathbf{x}_{\rm latent}) - \log p(\mathbf{y}_{\rm obs}|f_{\boldsymbol{\theta}}(\mathbf{x}_{\rm latent}))$ for whatever form p takes.

For classification tasks with noisy inputs, the same latent variable approach applies. The prediction term becomes cross-entropy loss instead of mean squared error, while the latent likelihood term remains unchanged:

$$\mathcal{J} = \text{CrossEntropy}(f_{\boldsymbol{\theta}}(\mathbf{x}_{\rm latent}), y) + \sum_{i=1}^{N} (\mathbf{x}_{{\rm obs},i} - \mathbf{x}_{{\rm latent},i})^\top \boldsymbol{\Sigma}_{x,i}^{-1} (\mathbf{x}_{{\rm obs},i} - \mathbf{x}_{{\rm latent},i}). \tag{17}$$

Extension to other neural network architectures requires only that the network be differentiable with respect to its inputs. Convolutional neural networks for imaging data, transformers for sequential data, or any other architecture can be used by simply treating the input as learnable parameters with an appropriate latent likelihood term.

7.4. *Limitations*

The method requires specifying both input uncertainties σ_x (or $\sigma_{x,i}$ in the heteroscedastic case) and output uncertainties σ_y. In practice, σ_x may need to be estimated based on the noise model of the instrument, while σ_y reflects label uncertainties or, when labels are precise, a nominal value that sets the relative weighting in the loss function. Misspecified σ_x or σ_y leads to suboptimal correction: underestimating σ_x causes undercorrection, while overestimating it can cause overcorrection. A possible remedy is adding jitter terms to be optimized along with the other parameters.

LatentNN performs optimization rather than sampling. We obtain point estimates of $x_{{\rm latent},i}$ and $\boldsymbol{\theta}$, but characterizing the full posterior over the high-dimensional space $(\boldsymbol{\theta}, \mathbf{x}_{\rm latent})$ is challenging. Full Bayesian treatment of neural networks is already difficult, and adding the latent variables makes it more so. Approximate Bayesian approaches such as dropout (Hinton et al. 2012; Srivastava et al. 2014) or variational inference (Kingma & Welling 2013) may provide practical uncertainty estimates.

The increased number of parameters also leads to a more complex loss landscape. The optimization may converge to different solutions from different initializations, particularly at low SNR where the problem is ill-conditioned. Careful hyperparameter tuning, especially of the weight decay parameter as discussed in Section 5, is often necessary.

In practice, however, the results are not highly sensitive to the weight decay value. In the one-dimensional experiments at $\text{SNR}_x = 2$, sweeping weight decay over three orders of magnitude (10^{-4} to 10^{-1}) yields λ_y values between 0.94 and 0.98—a variation of only $\sim$4%. For the multivariate case with $p = 10$ correlated features, the variation is even smaller ($\sim$2%). The dominant source of variability is not hyperparameter choice but rather run-to-run scatter from random initialization, which increases at low SNR: the standard deviation across 8 random seeds grows from $\sim$0.01 at $\text{SNR}_x = 10$ to $\sim$0.09 at $\text{SNR}_x = 1$. We recommend searching over a logarithmically spaced grid (e.g., $10^{-4}, 3\times10^{-4}, 10^{-3}, \ldots, 10^{-1}$) and selecting the value that yields λ_y closest to unity on a held-out validation set. When uncertainties are available for computing λ_y directly, no search is needed. An alternative is to monitor the balance between the prediction loss and latent likelihood during training, and a well-tuned weight decay produces a model where neither term dominates excessively at convergence.

Misstated uncertainties directly misstate how far latent inputs are allowed to move.

The computational cost scales as $\mathcal{O}(N \times p)$, introducing N additional parameters per input dimension. For full spectra with thousands of pixels and training sets of tens of thousands of stars, this becomes a consideration. To give concrete estimates: for an APOGEE-like training set with $N = 50{,}000$ and $p = 100$ informative pixels, LatentNN introduces 5×10^6 additional latent parameters ($\sim$20 MB in float32), well within the capacity of a single modern GPU. Using the full spectral vector ($p \sim 7000$) increases this to $\sim 3.5\times10^8$ parameters ($\sim$1.4 GB), which

remains feasible on GPUs with ≥16 GB memory. For larger training sets ($N \sim 10^6$) at full spectral dimensionality, the latent parameters alone would require ∼28 GB, pushing the limits of single-GPU memory but accessible with high-memory GPUs or multi-GPU configurations. In practice, restricting to informative pixels ($p \sim 30$–100) keeps the problem comfortably tractable even for large N. Scalability requires care, particularly for high-dimensional inference problems. Possible strategies include using subsets of informative pixels. Dimensionality reduction before LatentNN is less straightforward because the noise properties in the reduced space may not be well characterized. In summary, LatentNN is most directly applicable to problems with moderate training set sizes ($N \sim 10^3$–10^5) and a manageable number of informative input features ($p \sim 10$–100), in the signal-to-noise regime where attenuation bias is significant ($\mathrm{SNR}_x \lesssim 10$) but correction remains reliable ($\mathrm{SNR}_x \gtrsim 2$). In astronomy, this encompasses many spectroscopic applications, particularly for elements with limited spectral features, as well as photometric and time-series problems where input uncertainties are non-negligible. Extension to larger datasets or higher-dimensional inputs is feasible with appropriate feature selection, though full-dimensional application to very large training sets remains a computational challenge.

Storing latent inputs is manageable after feature selection; full spectra for enormous surveys are another matter.

My practical target is roughly SNR_x between 2 and 10 with tens of informative features.

Despite these limitations, the method provides a practical solution for a fundamental problem for which existing remedies—such as full Bayesian inference over latent variables—have not yet been demonstrated in the neural network context and face substantial computational challenges when scaling to large datasets.

8. CONCLUSIONS

Neural networks have become essential tools for astronomical inference, from deriving stellar parameters from spectra to estimating photometric redshifts from imaging data. These applications often involve low-SNR observations where measurement uncertainties are non-negligible. When inputs are noisy, attenuation bias becomes a fundamental concern.

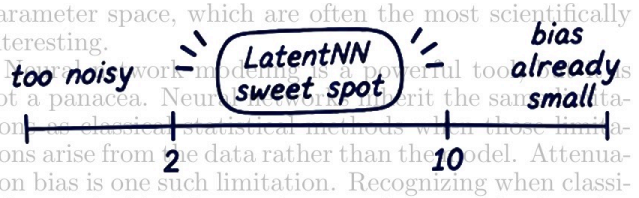
The broader lesson is humbling: modern networks still inherit classical errors-in-variables failures.

This effect, well known in the statistics literature for linear regression (Fuller 1987), extends to nonlinear networks. Neural networks trained on noisy inputs suffer from the same bias as linear regression, with coefficient slopes biased toward zero by a factor that scales with $\sigma_x/\sigma_{\mathrm{range}}$. Increasing model complexity does not mitigate the effect. This issue has received little attention in the machine learning literature because most applications involve high-SNR inputs where the effect is negligible. Astronomy is different, and as neural networks become increasingly central to astronomical inference, classical statistical considerations such as errors-in-variables become important.

Attenuation bias in linear regression is well understood, and solutions such as Deming regression have existed for decades. These methods treat the true input values as latent variables to be estimated alongside the model parameters. We show that this approach generalizes from linear to nonlinear models. By simultaneously inferring what the true inputs were, the model learns the correct input-output relationship rather than one diluted by measurement noise.

We propose LatentNN as a practical implementation of this idea. The method adds a learnable latent value for each training input, constrained by a latent likelihood term to stay close to the observed value, and jointly optimizes these latent values with the network parameters. This is a simple modification that can be applied to any differentiable architecture. We validated LatentNN on one-dimensional regression, where it matches Deming regression for the linear case, on multivariate inputs with varying correlation structure, and on realistic spectral applications using mock data with varying numbers of informative pixels. LatentNN effectively mitigates attenuation bias, particularly for $\mathrm{SNR}_x \gtrsim 2$, with the degree of correction depending on input dimensionality and correlation structure.

For astronomy, this work provides a framework for improved inference with neural networks in the low-SNR regime that characterizes much of astronomical data. The method is relevant for spectroscopic surveys, photometric applications, and any data-driven model where inputs carry non-negligible measurement uncertainty. The biases we address affect measurements at the extremes of parameter space, which are often the most scientifically interesting.



Neural network modeling is a powerful tool, but it is not a panacea. Neural networks inherit the same limitations as classical statistical methods when those limitations arise from the data rather than the model. Attenuation bias is one such limitation. Recognizing when classical statistical insights apply to modern machine learning methods, and adapting solutions accordingly, is essential for reliable scientific inference in astronomical research.

CODE AVAILABILITY

We provide a reference implementation of LatentNN at `https://github.com/tingyuansen/LatentNN`. The code demonstrates the method for simple fully-connected architectures as presented in this paper. The approach is general and can be adapted to any differentiable architecture by treating the input as learnable parameters with an appropriate latent likelihood term.

I thank David Weinberg and Hans-Walter Rix for useful discussion and comments on the manuscript. The original idea about errors-in-variables stemmed from my teaching of Astron 5550 (Advanced Data Analysis), during which I decided to discuss attenuation bias as explored in Paper I. While preparing for that lecture, I was inspired to read more into the errors-in-variables literature, also suggested by an anonymous referee during Paper I. Further discussion with Jeffrey Newman and Brett Andrews during a colloquium at UPitt inspired the extension of this work to deep learning, which ultimately led to this paper.

This research was supported by NSF Grant AST-2406729 and a Humboldt Research Award from the Alexander von Humboldt Foundation.

Claude 4.5 was used for assistance with code development and manuscript copy-editing. The scientific content, analysis, and conclusions are the author's, and all LLM-assisted changes were verified by the author.

LatentNN is most useful in the middle regime where bias matters but the signal remains recoverable.

The algebra shows exactly where attenuation enters and why Deming regression removes it.

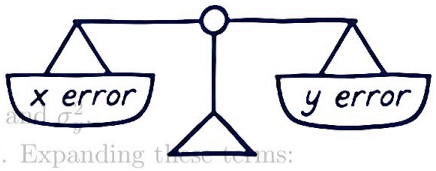
APPENDIX

DERIVATION OF ATTENUATION BIAS

Consider a true linear relationship $y_{\text{true}} = \beta x_{\text{true}}$. Our observations are

$$x_{\text{obs}} = x_{\text{true}} + \delta_x, \tag{A1}$$

$$y_{\text{obs}} = \beta x_{\text{true}} + \delta_y, \tag{A2}$$



where δ_x and δ_y are independent measurement errors with variances σ_x^2 and σ_y^2.

The ordinary least squares estimator is $\hat{\beta} = \text{Cov}(x_{\text{obs}}, y_{\text{obs}})/\text{Var}(x_{\text{obs}})$. Expanding these terms:

OLS divides by a variance inflated by σ_x^2, while the covariance numerator receives no matching inflation.

$$\text{Cov}(x_{\text{obs}}, y_{\text{obs}}) = \text{Cov}(x_{\text{true}} + \delta_x, \beta x_{\text{true}} + \delta_y) \tag{A3}$$

$$= \beta \,\text{Var}(x_{\text{true}}) = \beta \sigma_{\text{range}}^2, \tag{A4}$$

since measurement errors are independent of true values and of each other. For the denominator:

$$\text{Var}(x_{\text{obs}}) = \text{Var}(x_{\text{true}}) + \text{Var}(\delta_x) = \sigma_{\text{range}}^2 + \sigma_x^2. \tag{A5}$$

The expected value of the standard estimator is therefore

$$\mathbb{E}[\hat{\beta}] = \beta \frac{\sigma_{\text{range}}^2}{\sigma_{\text{range}}^2 + \sigma_x^2} = \beta \cdot \lambda_\beta, \tag{A6}$$

where $\lambda_\beta = 1/(1 + (\sigma_x/\sigma_{\text{range}})^2) < 1$ is the attenuation factor. The estimated slope is systematically biased toward zero.

DERIVATION OF DEMING REGRESSION

To correct attenuation bias, we treat the true values $\{x_{\text{true},i}\}$ as latent variables to be estimated alongside β. The joint likelihood, conditioned on the true values, is

$$\mathcal{L} = \prod_{i=1}^{N} \mathcal{N}(x_{\text{obs},i} | x_{\text{true},i}, \sigma_x^2) \cdot \mathcal{N}(y_{\text{obs},i} | \beta x_{\text{true},i}, \sigma_y^2). \tag{B1}$$

Expanding the normal distributions and taking the negative log-likelihood (dropping constants):

$$E(\{x_{\text{true},i}\}, \beta) = \sum_{i=1}^{N} \left[\frac{(x_{\text{obs},i} - x_{\text{true},i})^2}{2\sigma_x^2} + \frac{(y_{\text{obs},i} - \beta x_{\text{true},i})^2}{2\sigma_y^2} \right]. \tag{B2}$$

For a fixed β, each $x_{\text{true},i}$ appears only in the ith term, so we can optimize them independently. Taking the derivative and setting it to zero:

$$\frac{\partial E}{\partial x_{\text{true},i}} = -\frac{x_{\text{obs},i} - x_{\text{true},i}}{\sigma_x^2} - \frac{\beta(y_{\text{obs},i} - \beta x_{\text{true},i})}{\sigma_y^2} = 0. \tag{B3}$$

Rearranging and letting $r = \sigma_y^2/\sigma_x^2$ denote the error variance ratio:

$$x_{\text{true},i} = \frac{r\, x_{\text{obs},i} + \beta y_{\text{obs},i}}{r + \beta^2}. \tag{B4}$$

This can be rewritten as a weighted average:

$$x_{\text{true},i} = w_x \cdot x_{\text{obs},i} + w_y \cdot \frac{y_{\text{obs},i}}{\beta}, \tag{B5}$$

where $w_x = r/(r + \beta^2)$ and $w_y = \beta^2/(r + \beta^2)$ sum to unity. When $r \gg 1$ (precise x measurements), $w_x \approx 1$ and we trust x_{obs}. When $r \ll 1$ (precise y measurements), $w_y \approx 1$ and we infer x from y/β.

Taking the derivative of E with respect to β, substituting the expression for $x_{\text{true},i}$, and using the summary statistics $S_{xx} = \sum_i (x_{\text{obs},i} - \bar{x})^2$, $S_{yy} = \sum_i (y_{\text{obs},i} - \bar{y})^2$, and $S_{xy} = \sum_i (x_{\text{obs},i} - \bar{x})(y_{\text{obs},i} - \bar{y})$, we obtain a quadratic equation:

$$S_{xy}\beta^2 - (S_{yy} - rS_{xx})\beta - rS_{xy} = 0. \tag{B6}$$

Applying the quadratic formula and taking the positive root:

$$\hat{\beta} = \frac{S_{yy} - rS_{xx} + \sqrt{(S_{yy} - rS_{xx})^2 + 4rS_{xy}^2}}{2S_{xy}}. \tag{B7}$$

In the limit $\sigma_x \to 0$ (equivalently $r \to \infty$), this reduces to $\hat{\beta} \to S_{xy}/S_{xx}$, the ordinary least squares estimator.

Optimizing each latent x produces a weighted compromise; solving the remaining quadratic gives Deming's slope.

The same latent-input idea restores a nonlinear cubic that an ordinary MLP visibly flattens.

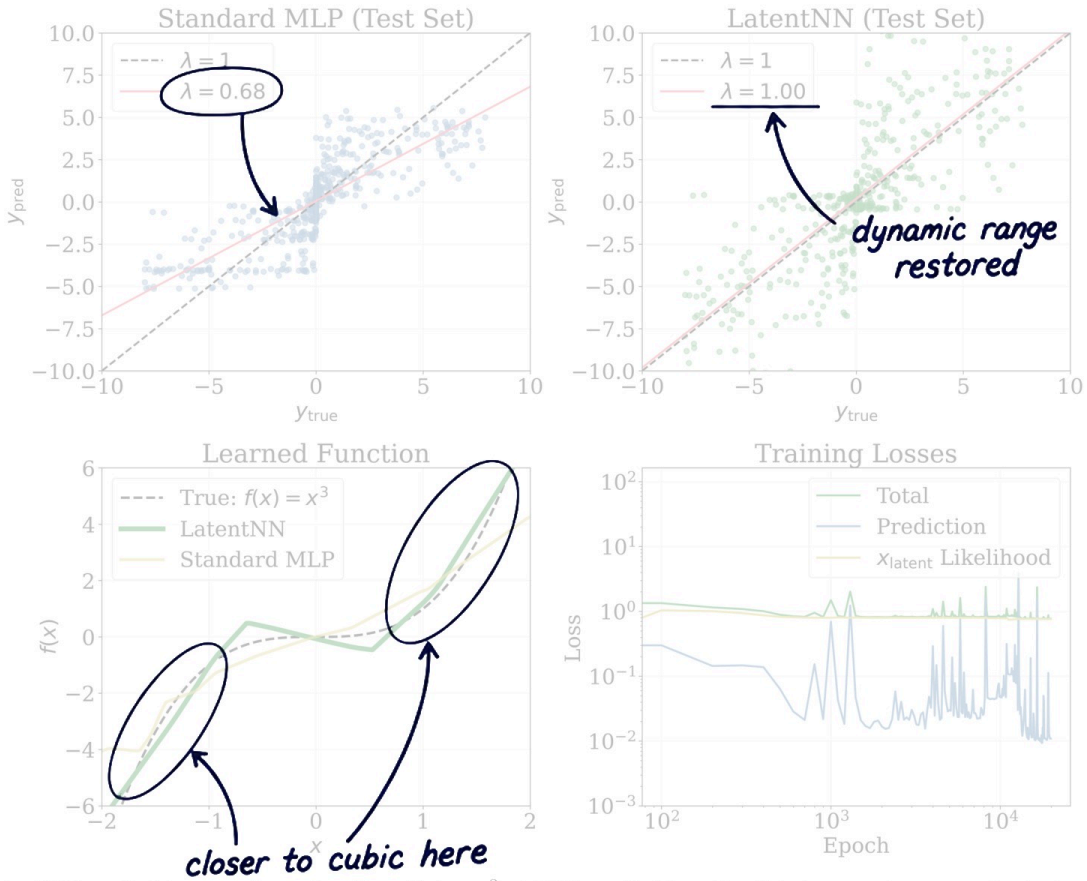


FIG. 7.— LatentNN applied to a nonlinear function $f(x) = x^3$ at $\text{SNR}_x = 2$. **Top:** Predicted versus true y on the test set for standard MLP (left, $\lambda_y \approx 0.7$) and LatentNN (right, $\lambda_y \approx 1$). **Bottom left:** Learned functions compared to the true cubic. The standard MLP learns a flattened curve, while LatentNN recovers the correct shape. **Bottom right:** LatentNN training losses showing the trade-off between prediction accuracy and latent variable likelihood.

EXTENSION TO NONLINEAR FUNCTIONS

Nonlinearity can amplify input errors most strongly where the derivative is steep.

The main text focuses on linear functions because they permit direct comparison with the theoretical attenuation factor λ_y (Paper I). However, the errors-in-variables problem and the LatentNN correction are more general and apply to arbitrary nonlinear functions.

To demonstrate this, we test LatentNN on a cubic polynomial $f(x) = x^3$. We use a larger network capacity (2 layers $\times$ 256 hidden units) to ensure the network can represent the nonlinear function accurately. Figure 7 shows results at $\text{SNR}_x = 2$, a challenging regime where noise equals 50% of the signal range. The same weight decay hyperparameter search procedure from Section 5 is applied.

The standard MLP exhibits severe attenuation, with $\lambda_y \approx 0.65$—even worse than the linear case at the same SNR as derived in Paper I. This occurs because measurement errors in x are amplified by the nonlinear function: an error δ_x produces a prediction error of approximately $f'(x)\delta_x$, which varies across the input domain. Near the edges of the training range where $|f'(x)| = 3x^2$ is largest, the effective noise is amplified.

I increased network width here so representation error could not masquerade as attenuation.

LatentNN corrects this bias. The bottom-left panel shows the learned functions: the standard MLP learns a flattened version of the cubic, while LatentNN recovers the true function shape. This demonstrates that the latent variable framework extends naturally beyond linear regression to nonlinear function approximation.

REFERENCES

Andrae, R., Rix, H.-W., & Chandra, V. 2023a, ApJS, 267, 8, doi: 10.3847/1538-4365/acd53e

Andrae, R., Fouesneau, M., Sordo, R., et al. 2023b, A&A, 674, A27, doi: 10.1051/0004-6361/202243462

Bishop, C. M. 2006, Pattern Recognition and Machine Learning (Information Science and Statistics) (Berlin, Heidelberg: Springer-Verlag)

Blanco-Cuaresma, S., Soubiran, C., Heiter, U., & Jofré, P. 2014, A&A, 569, A111, doi: 10.1051/0004-6361/201423945

Brown, T. B., Mann, B., Ryder, N., et al. 2020, arXiv e-prints, arXiv:2005.14165, doi: 10.48550/arXiv.2005.14165

Carroll, R. J., Ruppert, D., & Stefanski, L. A. 1995, Measurement Error in Non-linear Models (New York: Chapman and Hall)

Carroll, R. J., Ruppert, D., Stefanski, L. A., & Crainiceanu, C. M. 2006, Measurement error in nonlinear models: A modern perspective, Second edition (Chapman and Hall)

Chen, B., Orkney, M. D. A., Ting, Y.-S., & Hayden, M. R. 2025, arXiv e-prints, arXiv:2511.18901, doi: 10.48550/arXiv.2511.18901

Chen, B., Ting, Y.-S., & Hayden, M. 2024, PASA, 41, e063, doi: 10.1017/pasa.2024.56

Collister, A. A., & Lahav, O. 2004, PASP, 116, 345, doi: 10.1086/383254

Cranmer, K., Brehmer, J., & Louppe, G. 2020, Proceedings of the National Academy of Science, 117, 30055, doi: 10.1073/pnas.1912789117

Recovering the curve's shape matters more than assigning a unique denoised position to every point.

Cybenko, G. 1989, Mathematics of Control, Signals, and Systems, 2, 303, doi: 10.1007/BF02551274
Deming, W. E. 1943, Statistical adjustment of data (Wiley)
Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. 2018, arXiv e-prints, arXiv:1810.04805, doi: 10.48550/arXiv.1810.04805
Fabbro, S., Venn, K. A., O'Briain, T., et al. 2018, MNRAS, 475, 2978, doi: 10.1093/mnras/stx3298
Fallows, C. P., & Sanders, J. L. 2024, MNRAS, 531, 2126, doi: 10.1093/mnras/stae1303
Frost, C., & Thompson, S. G. 2002, Journal of the Royal Statistical Society Series A: Statistics in Society, 163, 173, doi: 10.1111/1467-985X.00164
Fuller, W. A. 1987, Measurement Error Models (New York: Wiley)
García Pérez, A. E., Allende Prieto, C., Holtzman, J. A., et al. 2016, AJ, 151, 144, doi: 10.3847/0004-6256/151/6/144
Gelman, A., Carlin, J. B., Stern, H. S., et al. 2013, Bayesian data analysis (3rd ed.) (Chapman and Hall)
Gheller, C., & Vazza, F. 2022, MNRAS, 509, 990, doi: 10.1093/mnras/stab3044
Golub, G. H., & Van Loan, C. F. 1980, SIAM Journal on Numerical Analysis, 17, 883
Guo, J., Bai, X., Deng, Y., et al. 2020, Sol. Phys., 295, 5, doi: 10.1007/s11207-019-1573-9
Hattori, K. 2025, ApJ, 980, 90, doi: 10.3847/1538-4357/ad9686
He, K., Zhang, X., Ren, S., & Sun, J. 2016, in 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 1, doi: 10.1109/CVPR.2016.90
Hinton, G. E., Srivastava, N., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. R. 2012, arXiv e-prints, arXiv:1207.0580, doi: 10.48550/arXiv.1207.0580
Ho, A. Y. Q., Ness, M. K., Hogg, D. W., et al. 2017, ApJ, 836, 5, doi: 10.3847/1538-4357/836/1/5
Hoerl, A. E., & Kennard, R. W. 1970, Technometrics, 12, 55
Holtzman, J. A., Shetrone, M., Johnson, J. A., et al. 2015, AJ, 150, 148, doi: 10.1088/0004-6256/150/5/148
Hon, M., Huber, D., Kuszlewicz, J. S., et al. 2021, ApJ, 919, 131, doi: 10.3847/1538-4357/ac14b1
Hornik, K., Stinchcombe, M., & White, H. 1989, Neural Networks, 2, 359
Hoyle, B. 2016, Astronomy and Computing, 16, 34, doi: 10.1016/j.ascom.2016.03.006
Huertas-Company, M., & Lanusse, F. 2023, PASA, 40, e001, doi: 10.1017/pasa.2022.55
Kelly, B. C. 2007, ApJ, 665, 1489, doi: 10.1086/519947
Khalatyan, A., Anders, F., Chiappini, C., et al. 2024, A&A, 691, A98, doi: 10.1051/0004-6361/202451427
Kingma, D. P., & Welling, M. 2013, arXiv e-prints, arXiv:1312.6114, doi: 10.48550/arXiv.1312.6114
Kurucz, R. L. 1970, SAO Special Report, 309
—. 1993, SYNTHE spectrum synthesis programs and line data
—. 2005, Memorie della Societa Astronomica Italiana Supplementi, 8, 14
—. 2017, ATLAS9: Model atmosphere program with opacity distribution functions, Astrophysics Source Code Library. http://ascl.net/1710.017
LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. 1998, Proceedings of the IEEE, 86, 2278
Leung, H. W., & Bovy, J. 2019, MNRAS, 483, 3255, doi: 10.1093/mnras/sty3217
Li, J., Wong, K. W. K., Hogg, D. W., Rix, H.-W., & Chandra, V. 2024, ApJS, 272, 2, doi: 10.3847/1538-4365/ad2b4d
Li, R., Napolitano, N. R., Feng, H., et al. 2022a, A&A, 666, A85, doi: 10.1051/0004-6361/202244081
Li, X., Wang, Z., Zeng, S., et al. 2022b, Research in Astronomy and Astrophysics, 22, 065018, doi: 10.1088/1674-4527/ac65e7
Lin, Q., Fouchez, D., Pasquet, J., et al. 2022, A&A, 662, A36, doi: 10.1051/0004-6361/202142751
Loredo, T. J. 2004, in American Institute of Physics Conference Series, Vol. 735, Bayesian Inference and Maximum Entropy Methods in Science and Engineering: 24th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering, ed. R. Fischer, R. Preuss, & U. V. Toussaint (AIP), 195–206, doi: 10.1063/1.1835214
Luo, A. L., Zhao, Y.-H., Zhao, G., et al. 2015, Research in Astronomy and Astrophysics, 15, 1095, doi: 10.1088/1674-4527/15/8/002
Majewski, S. R., Schiavon, R. P., Frinchaboy, P. M., et al. 2017, AJ, 154, 94, doi: 10.3847/1538-3881/aa784d
Ness, M., Hogg, D. W., Rix, H.-W., Ho, A. Y. Q., & Zasowski, G. 2015, ApJ, 808, 16, doi: 10.1088/0004-637X/808/1/16
Odewahn, S. C., Stockwell, E. B., Pennington, R. L., Humphreys, R. M., & Zumach, W. A. 1992, AJ, 103, 318, doi: 10.1086/116063
OpenAI, Achiam, J., Adler, S., et al. 2023, arXiv e-prints, arXiv:2303.08774, doi: 10.48550/arXiv.2303.08774
Pál, B., & Dobos, L. 2024, Astronomische Nachrichten, 345, e20240049, doi: 10.1002/asna.20240049
Pasquet, J., Pasquet, J., Chaumont, M., & Fouchez, D. 2019, A&A, 627, A21, doi: 10.1051/0004-6361/201834473
Rix, H.-W., Ting, Y.-S., Conroy, C., & Hogg, D. W. 2016, ApJ, 826, L25, doi: 10.3847/2041-8205/826/2/L25
Rix, H.-W., Chandra, V., Andrae, R., et al. 2022, ApJ, 941, 45, doi: 10.3847/1538-4357/ac9e01
Rumelhart, D. E., Hinton, G. E., & Williams, R. J. 1986, Nature, 323, 533
Simonyan, K., & Zisserman, A. 2014, arXiv e-prints, arXiv:1409.1556, doi: 10.48550/arXiv.1409.1556
Spearman, C. 1904, The American Journal of Psychology, 15, 72, doi: 10.2307/1412159
Srivastava, N., Hinton, G. E., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. R. 2014, Journal of Machine Learning Research, 15, 1929
Storrie-Lombardi, M. C., Lahav, O., Sodre, Jr., L., & Storrie-Lombardi, L. J. 1992, MNRAS, 259, 8P, doi: 10.1093/mnras/259.1.8P
Straumit, I., Tkachenko, A., Gebruers, S., et al. 2022, AJ, 163, 236, doi: 10.3847/1538-3881/ac5f49
Szegedy, C., Liu, W., Jia, Y., et al. 2014, arXiv e-prints, arXiv:1409.4842, doi: 10.48550/arXiv.1409.4842
Tikhonov, A. N. 1963, Soviet Mathematics Doklady, 4, 1035
Ting, Y.-S. 2025a, arXiv e-prints, arXiv:2510.10713, doi: 10.48550/arXiv.2510.10713
—. 2025b, The Open Journal of Astrophysics, 8, 95, doi: 10.33232/001c.142224
—. 2025c, arXiv e-prints, arXiv:2506.12230, doi: 10.48550/arXiv.2506.12230
Ting, Y.-S., Conroy, C., Rix, H.-W., & Cargile, P. 2019, ApJ, 879, 69, doi: 10.3847/1538-4357/ab2331
Ting, Y.-S., Rix, H.-W., Conroy, C., Ho, A. Y. Q., & Lin, J. 2017, ApJ, 849, L9, doi: 10.3847/2041-8213/aa921c
Touvron, H., Lavril, T., Izacard, G., et al. 2023, arXiv e-prints, arXiv:2302.13971, doi: 10.48550/arXiv.2302.13971
Van Huffel, S., & Vandewalle, J. 1991, The total least squares problem: Computational aspects and analysis (Society for Industrial and Applied Mathematics)
Vojtekova, A., Lieu, M., Valtchanov, I., et al. 2021, MNRAS, 503, 3204, doi: 10.1093/mnras/staa3567
Vynatheya, P., Hamers, A. S., Mardling, R. A., & Bellinger, E. P. 2022, MNRAS, 516, 4146, doi: 10.1093/mnras/stac2540
Wang, E. X., Nordlander, T., Asplund, M., et al. 2021, MNRAS, 500, 2159, doi: 10.1093/mnras/staa3381
Winecki, D., & Kochanek, C. S. 2024, ApJ, 971, 61, doi: 10.3847/1538-4357/ad5a0b
Wong, K. W. K., Ng, K. K. Y., & Berti, E. 2020, arXiv e-prints, arXiv:2007.10350, doi: 10.48550/arXiv.2007.10350
Xiang, M., Ting, Y.-S., Rix, H.-W., et al. 2019, ApJS, 245, 34, doi: 10.3847/1538-4365/ab5364
Xiang, M., Rix, H.-W., Ting, Y.-S., et al. 2022, A&A, 662, A66, doi: 10.1051/0004-6361/202141570
Yang, L., Yuan, H., Huang, B., et al. 2025, ApJS, 279, 7, doi: 10.3847/1538-4365/add5e3
Zhang, M., Xiang, M., Ting, Y.-S., et al. 2024, ApJS, 273, 19, doi: 10.3847/1538-4365/ad51dd
Zhang, X., Green, G. M., & Rix, H.-W. 2023, MNRAS, 524, 1855, doi: 10.1093/mnras/stad1941
Zhou, X., Gong, Y., Meng, X.-M., et al. 2022, MNRAS, 512, 4593, doi: 10.1093/mnras/stac786

This paper was built using the Open Journal of Astrophysics LaTeX template. The OJA is a journal which

provides fast and easy peer review for new papers in the astro-ph section of the arXiv, making the reviewing process simpler for authors and referees alike. Learn more at http://astro.theoj.org.