

# Nonlinear reconciliation: Error reduction theorems

Anonymous authors

Paper under double-blind review

## Abstract

Forecast reconciliation, an ex-post technique applied to forecasts that must satisfy constraints, has been a prominent topic in the forecasting literature over the past two decades. Recently, several efforts have sought to extend reconciliation methods to the probabilistic settings. Nevertheless, formal theorems demonstrating error reduction in nonlinear contexts, analogous to those presented in Panagiotelis et al. (2021), are still lacking. This paper addresses that gap by establishing such theorems for various classes of nonlinear hypersurfaces and vector-valued functions. Specifically, we derive an exact analog of Theorem 3.1 from Panagiotelis et al. (2021) for hypersurfaces with constant-sign curvature. Additionally, we provide an error reduction theorem for the broader case of hypersurfaces with non-constant-sign curvature and for general manifolds with codimension  $> 1$ . To support reproducibility and practical adoption, we release a JAX-based Python package, *to be released upon publication*, implementing the presented theorems and reconciliation procedures.

## 1 Introduction

Forecast reconciliation is a post-processing technique used to minimally correct an original set of forecasts, ensuring they satisfy known structural constraints. A straightforward example is the case of additive signals, grouped into increasingly larger sets, such as the quantity of a given good sold at the regional, state, and national levels. In such cases, the forecasts must obey additive relationships, a setting typically referred to as hierarchical forecasting (HF) in the literature.

Forecast reconciliation has become a widely adopted technique because it can be applied *ex-post*, independently of the forecasting model. This makes it a model-agnostic method for enforcing known structural constraints on forecasted signals. Reconciliation is typically motivated by two main considerations. An *existential* reason, if forecasts not respecting the known constraints are unusable or of little value. In this case reconciling is not an option. The second is a *utilitarian* reason: knowing the signals lie on a submanifold induces a bias in the production of the forecasts; this can lower the forecasting error. In the case of linear constraints, Theorem 3.1 of Panagiotelis et al. (2021) shows that orthogonally projecting forecasts onto the consistent subspace reduces the overall root mean squared error (RMSE) across the hierarchy. This refers to the RMSE of a single forecast vector at a given time point, computed as the square root of the average of its squared components (as opposed to the RMSE over an entire test set). In the nonlinear case, such a reduction is not always guaranteed. In this paper, we present theorems that establish conditions under which this RMSE reduction holds. In other words, these conditions answer the question: Should I expect an RMSE reduction when applying nonlinear reconciliation to the current forecasted point?

### 1.1 Linear reconciliation

Originally studied in industrial management as family-based forecasting (Fliedner, 2001), HF remains an active research topic, as evidenced by its inclusion in the 2020 M5 competition, where the organizers concluded that new methods capable of minimizing errors across all levels could offer significant accuracy gains (Makridakis et al., 2022). The simplest reconciliation strategy, bottom-up (Orcutt et al., 1968; Dunn et al., 1976), forecasts only the most disaggregated series, aggregating them upward. However, bottom-level series often have low signal-to-noise ratios, limiting performance. Forecasting all series independently leads to incoherent

results unless reconciled. A major step forward was presented in (Hyndman et al., 2011), where the authors proposed an optimal reconciliation based on generalize least square (GLS). Later refinements (Athanasopoulos et al., 2009; Wickramasuriya et al., 2019) proposed the use of covariance of base forecast errors. Later modification of the method include dropping the unbiasedness assumption (Ben Taieb et al., 2017; Di Modica et al., 2021), extension to temporal hierarchies, where series are aggregated over time (Athanasopoulos et al., 2017; Yang et al., 2017) and learning dynamic hierarchies (Cini et al., 2020; 2024). While point forecast reconciliation is well established, probabilistic reconciliation remained at first under-explored due to its complexity. Closed-form solutions exist only under strong assumptions (e.g., Gaussianity) (Corani et al., 2021), or errors’ independence, when aggregation can be performed using convolutions. Other techniques based on sample reordering or copulas have been proposed (Jeon et al., 2019; Taieb et al., 2021), though they require knowledge of the joint cumulative density function (CDF) at each level or rely on empirical approximations. A recent framework reconciles samples from joint but incoherent distributions via constrained optimization based on scoring rules (Panagiotelis et al., 2023), outperforming copula methods but still assuming known joint CDF.

## 1.2 Non-linear reconciliation

The nonlinear reconciliation problem has been less studied in the context of forecasting. Since the nonlinear reconciliation problem doesn’t usually have an analytical, closed-form solution, its application is less attractive than its linear counterpart. A method that could be used to produce a set of coherent forecasts for nonlinear manifolds, even if the authors applied it only to linear hierarchies, has been proposed by Spiliotis et al. (2021), and later applied in Rombouts et al. (2025), where a nonlinear regressor is trained to produce forecasts for the independent signal starting from the independent forecasts of all the signals. However, this method needs to train a machine learning model, and one has no control on the kind of projection applied, since this is implicitly learned by the model. On the other side, in the context of state estimation, nonlinear reconciliation techniques are well known. At the intersection of forecasting and nonlinear state estimation, the forecast aided state estimation (FASE) method has been proposed in the context of power grid state estimation (Brown Do Coutto Filho & De Souza, 2009). In FASE the forecasting model is a linear state space, and the reconciliation step is equivalent to an iterated extended Kalman filter. In appendix A.1, we show that the FASE objective function represents a special case of the nonlinear reconciliation problem we introduce in the next section. That is, the theory presented in this paper is also valid for FASE and its application to power flow state estimation and forecasting.

## 1.3 Contributions

The main contributions are the following:

1. In section 2 we define the nonlinear reconciliation problem and an algorithm for its solution. In appendix A.1 we show that this algorithm is a generalization of the FASE algorithm.
2. In section 3 we present theorems providing point-wise RMSE reduction conditions for the nonlinear reconciliation problem, for different classes of manifolds  $M$  implicitly defined by level sets of a function  $f$ . We propose two theorems, 1 and 2 for constant sign curvature manifolds and a third one, 3, providing a probabilistic reduction condition, which can be applied in the general case of  $M$  with codimension  $>1$  and without assumptions on sub-level sets of  $f$ .
3. In section 4 we empirically assessed their correctness. The results confirm the soundness of the theorems. Furthermore, we explore the possibility of using 3 to craft a reconciliation strategy beating the naive ones (always reconcile, never reconcile), and shows that these strategies lead to RMSE reduction.
4. To ease the application of these theorems and reconciliation methods, we release a JAX-based python library *to be released upon acceptance*.

## 2 Nonlinear reconciliation problem

We would like to generate predictions for  $n$  signals,  $\hat{z} \in \mathbb{R}^n$  which are known to satisfy one or more nonlinear constraints encoded by the level set  $f(z) = 0$  of a vector-valued function  $f(z) : \mathbb{R}^n \rightarrow \mathbb{R}^m$  with  $m$  constraints ( $1 < m < n$ ). The nonlinear reconciliation problem minimally displacing the original forecasts  $\hat{z}$  onto the constraint manifold  $f(z) = 0$  can be written as:

$$\tilde{z} = s(\hat{z}) := \arg \min_{z \in M} \|z - \hat{z}\|_W, \quad (1)$$

where  $M := \{z \in \mathbb{R}^n \mid f(z) = 0\}$  is an  $(n - m)$ -dimensional differentiable manifold embedded in  $\mathbb{R}^n$ , and  $W \succ 0$  is a weight matrix. In what follows we set  $W = I_n$ , so the solution is the orthogonal projection of  $\hat{z}$  onto  $M$ . The mapping  $s : \mathbb{R}^n \rightarrow M$  is called the nonlinear reconciliation operator. The linear reconciliation problem is a subset of 1, where  $M := \{z \in \mathbb{R}^n \mid z - Az_b = 0\}$ ,  $A \in \mathbb{R}^{m \times (n-m)}$  and  $z_b \in \mathbb{R}^{n-m}$  is usually referred to as bottom time series. The hierarchical reconciliation problem is a further subset, in which  $A$  encodes a (usually unweighted) tree structure,  $a_{i,j} \in \{0, 1\}$ .

The nonlinear optimisation problem in equation 1 can be solved by forming the Lagrangian

$$\mathcal{L}(z, \lambda) = (z - \hat{z})^\top W (z - \hat{z}) + \lambda^\top f(z), \quad \lambda \in \mathbb{R}^m. \quad (2)$$

A stationary point is characterised by  $\nabla_z \mathcal{L} = 0$  and  $\nabla_\lambda \mathcal{L} = 0$ . Collecting these conditions gives the vector equation

$$F(z, \lambda) := \begin{bmatrix} \nabla_z \mathcal{L} \\ \nabla_\lambda \mathcal{L} \end{bmatrix} = \begin{bmatrix} 2W(z - \hat{z}) + J^\top \lambda \\ f(z) \end{bmatrix} = 0, \quad J := \nabla f(z) \in \mathbb{R}^{m \times n}. \quad (3)$$

Starting from  $(z_k, \lambda_k)$ , a Newton–Raphson step  $(\delta z, \delta \lambda)$  solves

$$\underbrace{\begin{bmatrix} 2W & J^\top \\ J & 0 \end{bmatrix}}_{J_F(z_k, \lambda_k)} \begin{bmatrix} \delta z \\ \delta \lambda \end{bmatrix} = -F(z_k, \lambda_k), \quad (4)$$

after which we update

$$(z_{k+1}, \lambda_{k+1}) = (z_k, \lambda_k) + (\delta z, \delta \lambda). \quad (5)$$

### 2.1 Coherent probability distributions on nonlinear manifolds

We recall the definition of reconciled probability distribution introduced in Panagiotelis et al. (2023) and Zambon et al. (2023). This concept underpins the probabilistic guarantees we derive for manifolds with non-constant-sign curvature.

#### Definition 1: Coherent probability distribution on a nonlinear manifold

Let  $\hat{\nu} \in \mathcal{P}(\mathbb{R}^n)$  be a predictive distribution on  $(\mathbb{R}^n, \mathcal{B}_{\mathbb{R}^n})$  and let  $s : \mathbb{R}^n \rightarrow M$  be a Borel-measurable reconciliation operator whose image equals the constraint manifold  $M := \{z \in \mathbb{R}^n : f(z) = 0\}$ . The coherent predictive distribution  $\tilde{\nu} \in \mathcal{P}(M)$  is the pushforward

$$\tilde{\nu} = s_{\#} \hat{\nu}, \quad \tilde{\nu}(F) = \hat{\nu}(s^{-1}(F)), \quad F \in \mathcal{B}_M,$$

where  $\mathcal{B}_M := \{F \subseteq M : s^{-1}(F) \in \mathcal{B}_{\mathbb{R}^n}\}$ .

## 3 Error reduction theorems

In the linear case, it has been shown that reconciliation via orthogonal projection always reduces the overall point-wise prediction error. This result relies on the Pythagorean theorem and was demonstrated in Panagiotelis et al. (2021). In this section, we derive analogous theorems for the nonlinear case.

All proofs are based on a geometric interpretation of problem 1, as illustrated in Figure 1, which depicts the reconciliation process for two examples of predicted points, denoted by  $\hat{z}$ .

Let  $\hat{z}$  be a prediction of the true point  $z$ , and let  $\tilde{z}$  denote the orthogonal projection of  $\hat{z}$  onto the manifold  $M$ . We define three key vectors: the original prediction displacement  $\hat{\delta} = z - \hat{z}$ , the reconciled displacement after projection  $\tilde{\delta} = z - \tilde{z}$ , and the reconciliation adjustment  $\delta_\pi = \tilde{z} - \hat{z}$ .

The condition for RMSE reduction can then be stated as:

$$\|\hat{\delta}\|^2 \geq \|\tilde{\delta}\|^2 \quad (6)$$

Using the vectorial relation  $\hat{\delta} = \delta_\pi + \tilde{\delta}$  this becomes:

$$\|\delta_\pi\|^2 + 2\delta_\pi^T \tilde{\delta} \geq 0 \quad (7)$$

that can be rewritten as:

$$\delta_\pi^T \tilde{\delta} \geq -\frac{\|\delta_\pi\|^2}{2} \quad (8)$$

### 3.1 Constant sign curvature hypersurfaces

A hypersurface in  $\mathbb{R}^n$  is a smooth  $(n-1)$ -dimensional submanifold that can be locally defined as the level set  $M = \{z \in \mathbb{R}^n : f(z) = c\}$ , where  $f : \mathbb{R}^n \rightarrow \mathbb{R}$  is a smooth function and  $c$  is a regular value (i.e.,  $\nabla f(z) \neq 0$  for all  $z \in M$ ). As an example, Figure 1 shows the hypersurface given by  $M = \{z \in \mathbb{R}^2 : f(z) \triangleq y - x^2 = 0\}$  as a thick black line, while blue and violet represent regions of the negative and positive values of  $f(z)$ .




Figure 1: Conceptual plot for Theorem 1 ( $\hat{z}$  with negative  $f(\hat{z})$  in the bottom right) and Theorem 3 ( $\hat{z}$  with positive  $f(\hat{z})$ ), shown along atoms of its predictive distribution and their projection onto  $M$ ).

For hypersurfaces with a constant sign curvature, we can exploit the fact that they possess a supporting hyperplane at the point of projection. If the forecasted point  $\hat{z}$  is on the right side of the hypersurface, it will see the surface *bending away* from it, which is sufficient to guarantee a reduction of the error *independently* from where the true point  $z$  is located.

### Theorem 1: Constant sign curvature hypersurfaces

Let  $M = \{z : f(z) = 0\} \subset \mathbb{R}^n$  be a hypersurface defined by  $f : \mathbb{R}^n \rightarrow \mathbb{R} \in C^2$  with  $\nabla f(z) \neq 0$  on  $M$  and with convex sub or super level sets. Given a prediction  $\hat{z} \in \mathbb{R}^n$ , denote by  $\tilde{z} \in M$  its orthogonal projection onto  $M$ . Projecting  $\hat{z}$  into  $\tilde{z}$  will decrease RMSE if the following condition holds:

$$\text{sign}(f(\hat{z})) \cdot \lambda_{\min}(H_{\tan}(\tilde{z})) > 0 \quad (9)$$

where  $H_{\tan}$  denotes the Hessian  $H = D^2 f(z)$  restricted to the tangent space of  $M$ .

**Proof 1** A bounding condition for equation 7 to be positive is  $\delta_\pi^T \tilde{\delta} \geq 0 \ \forall \ \tilde{\delta}$ . Since  $\delta_\pi$  is an orthogonal projection onto  $M$ , it is parallel to the gradient at  $\tilde{z}$ , so we can rewrite

$$\mu \nabla f(\tilde{z})^T \tilde{\delta} \geq 0 \quad \forall \tilde{\delta} \quad (10)$$

The sign of  $\mu$  is given by the side  $\hat{z}$  is located with respect to the level set  $f(z) = 0$ , that is

$$\text{sign}(\mu) = -\text{sign}(f(\hat{z})) \quad (11)$$

since if  $f(\hat{z})$  is positive, the projecting vector  $\delta_\pi$  will oppose the gradient and vice versa. Since  $M$  is the boundary of a convex sub or super level set of  $f(z)$ , it has a supporting hyperplane in  $\tilde{z}$ , that implies the sign of  $\nabla f(\tilde{z})^T \tilde{\delta}$  is constant on  $M$ . This also implies:

$$\text{sign}(\nabla f(\tilde{z})^T \tilde{\delta}) = -\text{sign}(\lambda_{\min}(H_{\tan}(\tilde{z}))) \quad (12)$$

(see A.3 for the proof). Putting together equation 10 equation 11 and equation 12 we obtain the sufficient condition of Theorem 1.

### Corollary 1: Constant sign curvature hypersurfaces

Under the same assumptions of Theorem 1, defining  $\delta_\pi = \mu \nabla f(\tilde{z})$ , orthogonally projecting  $\hat{z}$  into  $\tilde{z}$  will decrease RMSE if the following condition holds:

$$\mu = \begin{cases} < 0 & \text{if } \{z : f(z) < 0\} \text{ convex} \\ > 0 & \text{if } \{z : f(z) > 0\} \text{ convex} \end{cases} \quad (13)$$

**Proof of corollary 1** The projection vector  $\delta_\pi$  at  $\tilde{z}$  is by definition parallel to  $\nabla f(\tilde{z})$  so that  $\delta_\pi = \mu \nabla f(\tilde{z})$  holds. Since sub-level set of  $f(z)$  is convex iff  $\lambda_{\min}(H_{\tan}) > 0$  (see A.2), considering equation 11 we can re-write condition 9 as equation 13.

Corollary 1 states that, knowing if the sub or super level sets of  $f(z)$  are convex, we can restate Theorem 1 using the gradient at  $\tilde{z}$  without computing the Hessian. We will use this fact to craft a theorem for vector valued functions in the next section.

## 3.2 Vector-valued functions

The next result extends Corollary 1 to manifolds defined by vector-valued constraints. The key idea is the same: if each defining function  $f_i$  has convex sub or super level set, and if the projection direction lies in the positive cone spanned by the gradients at projection point  $\{\nabla f_i(\tilde{z})\}_{i=1}^m$ , then this direction defines a supporting hyperplane to the intersection  $C = \cap_{i=1}^m C_i$  of the corresponding convex sets  $C_i = \{z : f_i(z) \leq 0\}$ . This yields a uniform decrease of squared error. Figure 2 illustrates the construction.




Figure 2: Concept for Theorem 2 with  $f = (f_1, f_2) = [x - y, x^2 - z]$ . The zero level sets  $f_1 = 0$  and  $f_2 = 0$  are shown in blue and violet. The intersection  $C = \{f_1 \leq 0\} \cap \{f_2 \leq 0\}$  is in green. If the projection direction is a positive combination of the normals  $\nabla f_1(\tilde{z}), \nabla f_2(\tilde{z})$ , it defines a supporting hyperplane for  $C$  (gray).

### Theorem 2: Manifolds of codimension $m > 1$ with convex/concave level geometry

Let  $M = \{z \in \mathbb{R}^n : f(z) = 0\}$  with  $f = (f_1, \dots, f_m) : \mathbb{R}^n \rightarrow \mathbb{R}^m$ ,  $f \in C^2$ , and assume  $\nabla f_i(z) \neq 0$  and  $\text{rank } J(z) = m$  on  $M$ , where  $J(z) = \nabla f(z)$ . Suppose that for each  $i$ , the sublevel set  $C_i = \{z : f_i(z) \leq 0\}$  is convex. For any  $\hat{z} \in \mathbb{R}^n$  let  $\tilde{z} \in M$  be its orthogonal projection and define

$$\delta_\pi := \tilde{z} - \hat{z}, \quad \tilde{\delta}(z) := z - \tilde{z}.$$

If there exists  $\mu \in \mathbb{R}_+^m$  such that

$$\delta_\pi = - \sum_{i=1}^m \mu_i \nabla f_i(\tilde{z}), \quad (14)$$

then for every  $z \in C := \cap_{i=1}^m C_i$ ,  $\|z - \hat{z}\|^2 \geq \|z - \tilde{z}\|^2$ ; in particular, the inequality holds for all  $z \in M$ .

If instead each  $f_i$  is concave (equivalently, the superlevel sets  $\{f_i \geq 0\}$  are convex), the same conclusion holds with the sign in equation 14 reversed:  $\delta_\pi = \sum_{i=1}^m \mu_i \nabla f_i(\tilde{z})$ .

**Proof 2** Fix  $\tilde{z}$  and let  $z \in C := \cap_{i=1}^m C_i$ . Since  $C_i$  is convex and  $f_i(\tilde{z}) = 0$ , the hyperplane with normal  $\nabla f_i(\tilde{z})$  supports  $C_i$  at  $\tilde{z}$ :

$$\nabla f_i(\tilde{z})^\top (z - \tilde{z}) \leq 0 \quad \text{for all } z \in C_i,$$

hence also for all  $z \in C$ . If equation 14 holds with  $\mu_i \geq 0$ , then

$$\delta_\pi^\top \tilde{\delta}(z) = - \sum_{i=1}^m \mu_i \nabla f_i(\tilde{z})^\top (z - \tilde{z}) \geq 0.$$

Using  $\hat{\delta} = \delta_\pi + \tilde{\delta}(z)$ ,

$$\|\hat{\delta}\|^2 - \|\tilde{\delta}(z)\|^2 = \|\delta_\pi\|^2 + 2\delta_\pi^\top \tilde{\delta}(z) \geq \|\delta_\pi\|^2 \geq 0,$$

which proves the claim. In the concave case, the supporting inequality reverses sign for the convex superlevel sets, yielding the stated flip of equation 14.

### 3.3 Applicability of Theorem 1 and 2

Theorems 1 and 2 provide *sufficient* conditions for a decrease of RMSE, conditioned on whether the forecast  $\hat{z}$  lies in the relevant sublevel or superlevel region of the function(s) defining  $M$ . How often those conditions hold depends on the geometry of  $M$ .

We estimate an *a priori* likelihood as follows. For several hypersurfaces ( $m = 1$ ) and manifolds of codimension  $m > 1$ , we sample points  $z_i \in M$ . For each  $z_i$  we draw an isotropic error  $\varepsilon_i \sim \mathcal{N}(0, \Sigma)$  with  $\Sigma = \sigma_I^2 I_n$  at multiple noise levels  $\sigma_I$ , form the forecast  $\hat{z}_i = z_i + \varepsilon_i$ , and compute its projection  $\tilde{z}_i$  by solving equation 1. We then: (i) check the corresponding theorem's condition at  $\tilde{z}_i$  and record the indicator of it holding, and (ii) record whether reconciliation reduces squared error, i.e.  $\mathbf{1}\{\|z_i - \tilde{z}_i\|^2 < \|z_i - \hat{z}_i\|^2\}$ . Averaging these indicators over  $i$  yields the empirical probabilities shown in Figure 3 (left: theorem's condition holds; right: squared-error reduction).

Across the tested manifolds, the sufficient conditions identify only a subset of cases that actually benefit from reconciliation. Moreover, real forecasting errors are rarely isotropic around a nonlinear  $M$ . The next example shows that even with optimal mean-squared prediction, forecasts can systematically fall into the region where Theorems 1 and 2 do *not* trigger.

#### Example 1: Optimal prediction on a parabola

Let  $M = \{z : f(z) = 0\}$  with  $f(z) = z_1^2 - z_2$ . Given data  $\mathcal{D} = \{z_i\}_{i=1}^n$ , the independent least-squares predictors are the empirical means

$$\hat{z}_1^* = \frac{1}{n} \sum_i z_{1,i}, \quad \hat{z}_2^* = \frac{1}{n} \sum_i z_{2,i} = \frac{1}{n} \sum_i z_{1,i}^2 \geq \left( \frac{1}{n} \sum_i z_{1,i} \right)^2 = \hat{z}_1^{*2},$$

where the inequality is Jensen's. Hence  $f(\hat{z}^*) = \hat{z}_1^{*2} - \hat{z}_2^* \leq 0$ , i.e., the optimal independent prediction lies in the sublevel region  $\{z_2 \geq z_1^2\}$ .

Therefore, the rates in Figure 3 are optimistic regarding the frequency with which the sufficient conditions will hold in practice: in many realistic settings,  $\hat{z}$  is likely to lie in the region where a reduction in error is not guaranteed by Theorems 1 and 2. This motivates the probabilistic guarantees for general nonconvex manifolds of codimension  $m > 1$  developed in the next section.

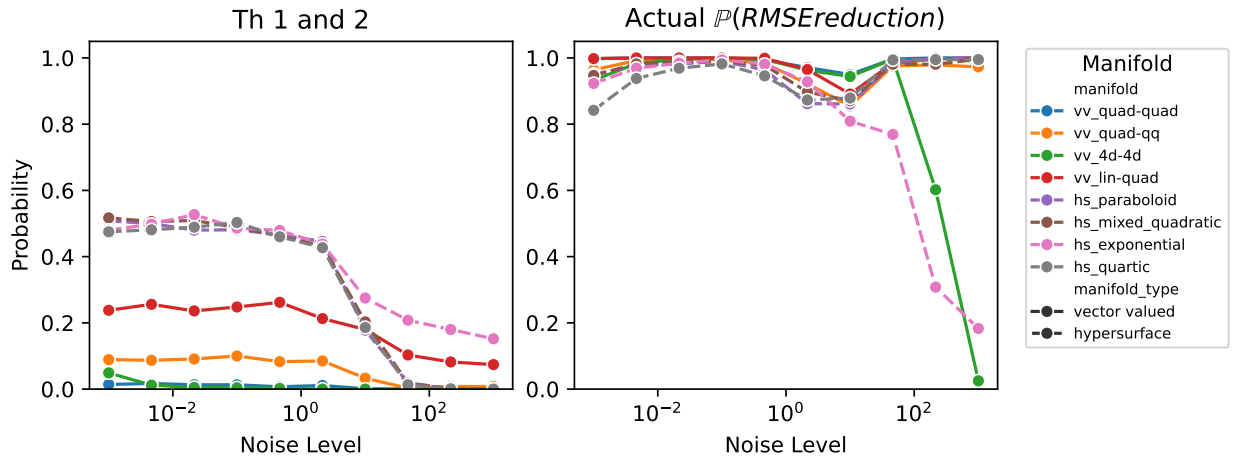


Figure 3: Left: probability of observing a positive condition from Theorem 1 and 2 under isotropic Gaussian perturbation of points on the manifold. Right: empirical probability of reducing RMSE by reconciling. The x axis shows the noise level  $\sigma_I$ .

### 3.4 Non-constant sign curvature manifolds

In Theorem 1 we used the orientation condition in equation 11 together with equation 12 to obtain a sufficient error-reduction result for constant-sign curvature hypersurfaces. Even for that class of manifolds, if the prediction  $\hat{z}$  lies on the side where  $M$  bends *towards*  $\hat{z}$ , the theorem does not apply. For those cases, and for manifolds without constant-sign curvature, we can still leverage the general inequality described in equation 8, which says that the post-projection error vector  $\tilde{\delta}$  may “climb” along the projection direction  $\delta_\pi$  by at most half its length. Geometrically, this defines a half-space of candidate true points  $z$  for which reconciliation reduces the RMSE. Figure 1 illustrates, in two dimensions, the corresponding half-plane for the left prediction point. For surfaces that can be sampled, this half-space intersects  $M$  in a region where we will observe a reduction.

While informative, this condition does not directly quantify the likelihood of error reduction. We therefore propose a more practical probabilistic test to answer: *What is the probability that projecting the prediction onto  $M$  reduces RMSE?* Since the only unknown in equation 8 is the post-projection error  $\tilde{\delta}$ , which is fully determined by  $\tilde{z}$  and  $z$ , this question reduces to estimating the probability that the true point lies in the favorable half-space. Equivalently, this requires access to a reconciled predictive distribution  $\tilde{\nu}$  supported on  $M$ .

#### Theorem 3: Probabilistic RMSE reduction

Let  $\tilde{\nu}$  be a predictive distribution supported on the constraint manifold  $M$  and let  $[\tilde{z}_i]_{i=1}^S$  be sampled atoms from  $\tilde{\nu}$  with associated normalized weights  $\tilde{\pi} = [\tilde{\pi}_i]_{i=1}^S$ , where  $\sum_i \tilde{\pi}_i = 1$ . Then, if  $\tilde{\pi}_i = \frac{1}{k} \forall i$ , a consistent estimator under the predictive law  $\tilde{\nu}$  for  $\mathbb{P}(\|\hat{\delta}\| \geq \|\tilde{\delta}\|)$  is:

$$e((\tilde{z}_i, \tilde{\pi}_i)_{i=1}^S) = \sum_{i=1}^S \tilde{\pi}_i \mathbf{1}_{\{\delta_\pi^T \tilde{\delta}_i \geq -\frac{1}{2} \|\delta_\pi\|^2\}} \quad (15)$$

**Proof 3** The proof is a straightforward application of the third definition of error reduction in equation 8. Since in this equation the only unknown variable is  $\tilde{\delta}$ , which is completely defined by the known  $\tilde{z}$  and  $z$ , which is unknown at prediction time, we can write the estimated probability for an improvement as an integral over the true distribution of  $z$ ,  $\nu$ :

$$\mathbb{P}(\|\hat{\delta}\| \geq \|\tilde{\delta}\|) = \int_M \mathbf{1}_{\{\delta_\pi^T \tilde{\delta}(z) \geq -\frac{1}{2} \|\delta_\pi\|^2\}} \nu(z) dz \quad (16)$$

$$\simeq \int_M \mathbf{1}_{\{\delta_\pi^T \tilde{\delta}(z) \geq -\frac{1}{2} \|\delta_\pi\|^2\}} \tilde{\nu}(z) dz \quad (17)$$

$$\simeq \sum_{i=1}^S \tilde{\pi}_i \mathbf{1}_{\{\delta_\pi^T \tilde{\delta}_i \geq -\frac{1}{2} \|\delta_\pi\|^2\}} \quad (18)$$

In equation 17 we approximate the true distribution with the forecasted  $\tilde{\nu}$ , while in the second approximation equation 18 we further approximate it with the sample-based representation  $(\tilde{z}_i, \tilde{\pi}_i)_{i=1}^k$ , where  $\tilde{z}_i$  enters the expression implicitly since  $\tilde{\delta}_i = \tilde{z}_i - s(\hat{z})$ . With equal weights,  $\mathbb{E}[e_k] = \frac{1}{k} \sum_{i=1}^k \mathbb{E}[\mathbf{1}_A(\tilde{z}_i)] = \mathbb{P}_{\tilde{\nu}}(A) = P_{\tilde{\nu}}(A)$ , and strong consistency follows from the strong law of large numbers (SLLN).

Since equation 15 estimation is based on a binary test on samples from  $\tilde{\nu}$ , we can use a binomial proportion confidence interval to consider the uncertainty introduced by the sample based approximation, like the Clopper–Pearson intervals (Clopper & Pearson, 1934), as we did in section 4.

One possible practical procedure to obtain  $e(\tilde{\pi})$  is the following:

1. Obtain the reconciled point forecast via the nonlinear projection operator:  $\tilde{z} = s(\hat{z})$
2. Starting from atoms of a non-reconciled predictive distribution  $\hat{z}_i$ , obtain the consistent samples always via projection  $\tilde{z}_i = [s(\hat{z}_i)]_{i=1}^S$



3. Count how many reconciled samples fall in the right half-space via equation 15, using  $\tilde{\delta}_i = \tilde{z}_i - s(\hat{z})$

We emphasize that Theorem 3 does not require the use of an orthogonal projection, in contrast to Theorem 1 and Theorem 2. This represents a distinct advantage, since, at least in the linear case, non-orthogonal projections have been shown to yield even greater RMSE reductions in many settings.

## 4 Numerical tests

To facilitate the application of the proposed theorems, we provide an open-source JAX-based Python library, which supports just-in-time compilation and GPU acceleration. We use this library to conduct numerical tests for the presented theorems. The tests cover a collection of manifolds with codimension 1 and 2, including cases with constant-sign curvature and convex sub-level sets. The applicability of each theorem to the different classes of manifolds is summarized in Table 1.

Table 1: Applicability of the presented theorems.

	constant sign curvature	non-constant sign curvature
codimension 1	1 2 3	3
codimension >1	23	3

Furthermore, as stated before, Theorem 3 also applies when using non-orthogonal projections.

### 4.1 Experiments setup and evaluation metrics

All the tested manifolds are defined by one or two graphs  $g : \mathbb{R}^2 \rightarrow \mathbb{R}$ , so that  $M : \{z : f(z) = 0\}$ ,  $f = [g(z_1, z_2)]$  for manifolds with codimension 1 and  $f = [g_1(z_1, z_2), g_2(z_1, z_2)]$  for manifolds with codimension 2. This allows us to sample the manifolds by defining a two-dimensional data generation process spanning the surfaces:

$$\begin{aligned} z_1^{t+1} &= \theta_1 z_1^t + w_1 \\ z_2^{t+1} &= \theta_2 z_2^t + w_2 \end{aligned} \tag{19}$$

where  $w_1, w_2 \sim \mathcal{N}(0, \sigma^2)$  are two independent random variables with normal distribution and  $\theta_1, \theta_2 \in [0, 1]$  are two arbitrary parameters defining the autoregressive process. We then use the graphs  $g$  to obtain coherent tuples,  $z^t$ . We built a dataset of  $T=1\text{E}4$  samples  $\mathcal{D} = \{([z_i^t]_{i=1}^n)_{t=1}^T\}$ , applied a random reshuffle and trained a LightGBM regressor to predict  $z^{t+1}$  from  $z^t$  on 10% of  $\mathcal{D}$ . We obtained a predictive distribution in terms of samples  $\{(\tilde{z}_i, \tilde{\pi}_i)_{i=1}^S\}$  with  $S = 200$  via bootstrap on a calibration set with 40% of the  $\mathcal{D}$ , and we used the remaining 50% to test the theorems. Figure 4 shows a hypersurface (paraboloid) and a manifold with codimension >1 (quadratic-exponential) as examples of  $M$ . On both panels, three examples of predictions in terms of deterministic points  $\hat{z}$ , associated bootstrapped samples, their projection onto  $M$ , and true points  $z$  are shown.

For Theorems 1 and 2, we retrieved the number of predicted RMSE reduction cases and the false positives (cases in which the theorems predict a decrease of RMSE while observing an increase). We anticipate that we observed a total of 0 false positives across all the functions.

Since Theorem 3 provides a probabilistic test, we assess it with a randomized study. We generate 25 studies with associated datasets  $\mathcal{D}_r$  of  $10^4$  points from  $M$  using the data generation process equation 19, with random parameters  $\theta_1$  and  $\theta_2$ . We follow the previously described procedure to obtain predictive distributions via bootstrap on each test set. For each point on the test set we then evaluate 3,  $e((\tilde{z}_i, \tilde{\pi}_i)_{i=1}^S)$ . Since Theorem 3 provides a probability for a binary event, a standard way to test calibration is to group observations in bins based on the predicted probability, and observe if the grouped points have a ratio of success close to the one

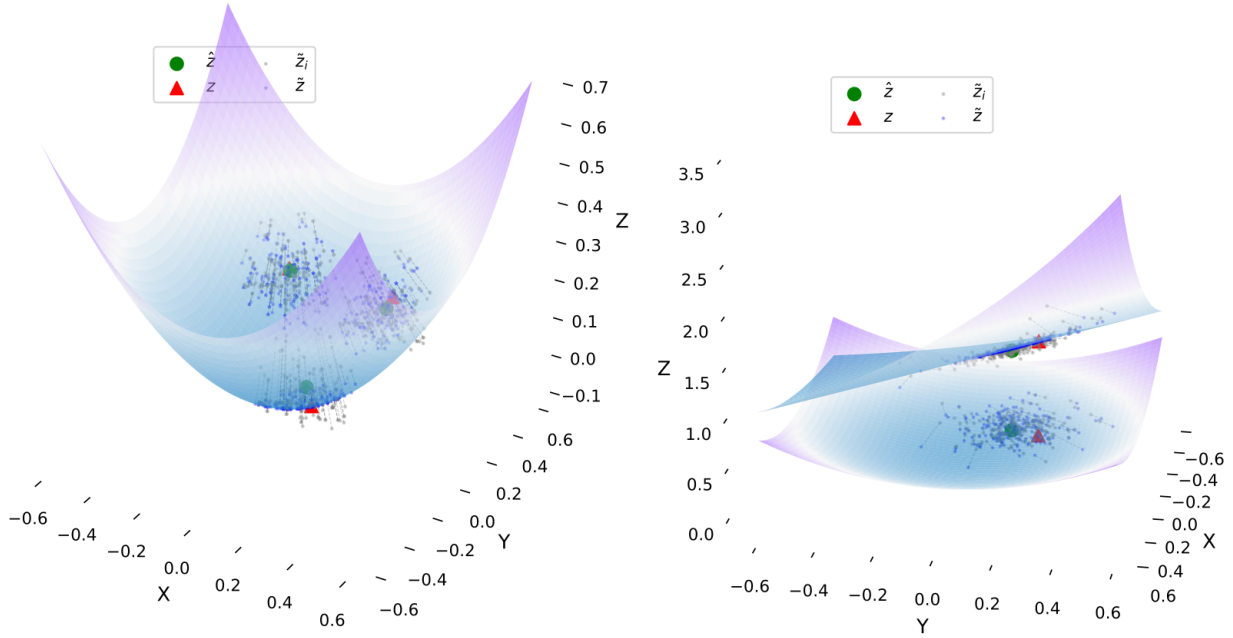


Figure 4: Examples of test cases in terms of tuples of predictions (green dots), ground truth (red triangles), and samples from predictive and reconciled distributions (gray and blue dots). Left: paraboloid case, Right: vector-valued function. In this case, just one tuple is plotted, but two points are visible since  $f : \mathbb{R}^4 \rightarrow \mathbb{R}^2$ .

predicted by the test in that bin. Instead, we compute a moving average calibration curve defined as:

$$r(t) = \frac{1}{w} \sum_{\max(t-w/2, 0)}^{\min(t+w/2, T)} \mathbf{1}_{\{\|\tilde{\delta}_{p(t)}\| < \|\hat{\delta}_{p(t)}\|\}} \quad (20)$$

where  $w$  is the smoothing window width, and the max and min operators handle boundary effects. The indicator function evaluates the ex-post condition (i.e., whether reconciliation led to a lower RMSE after the uncertainty resolved), and  $p(t)$  denotes a permutation that orders the samples based on the ex-ante predictions from Theorem 3.

Since we equip the probability estimation with upper and lower boundaries via Clopper-Pearson intervals, we also test for the binary coverage of the test, defined as:

$$c = \sum_{i \in \mathcal{D}_{te}} c_i \quad c_i = \begin{cases} e_l((\tilde{z}_i, \tilde{\pi}_i)_{i=1}^S) > 0.5 & \text{if } \|\tilde{\delta}\| < \|\hat{\delta}\| \\ e_u((\tilde{z}_i, \tilde{\pi}_i)_{i=1}^S) < 0.5 & \text{otherwise} \end{cases} \quad (21)$$

where  $e_u, e_l$  are the upper and lower bounds.

Finally, considering the case in which the forecasting task doesn't explicitly require  $\hat{z} \in M$ , we investigate if 3 can be used to craft an effective reconciliation strategy, reconciling just the points for which the estimated probability of a RMSE reduction is higher than a given threshold  $\theta$ :

$$z_i^\theta = \begin{cases} \tilde{z}_i := s(\hat{z}_i) & \text{if } e_i > \theta \\ \hat{z}_i & \text{otherwise} \end{cases} \quad (22)$$

The results are shown as the reduction in RMSE from the baseline RMSE  $\|\hat{\delta}\|$  (never reconcile), normalized for the RMSE of the optimal strategy, where the forecast  $\hat{z}$  is reconciled if this guarantees a reduction in RMSE:

$$z_i^* = \begin{cases} \tilde{z}_i := s(\hat{z}_i) & \text{if } \|\hat{\delta}\| > \|\tilde{\delta}\| \\ \hat{z}_i & \text{otherwise} \end{cases} \quad (23)$$

That is,  $z_i^*$  represents the best possible ex-post strategy, or equivalently, the optimal reconciliation strategy with perfect knowledge of the future. The final normalized score is:

$$\Delta_{rel,opt}^{RMSE} = \frac{\|\hat{\delta}\| - \|z_i^\theta - z_i\|}{\|\hat{\delta}\| - \|z_i^* - z_i\|} \quad (24)$$

## 4.2 Hypersurfaces

Plots of the tested hypersurfaces are shown in 11 in the appendix. Table 2 shows the ratio of cases with an observed RMSE reduction (first column), with a predicted reduction from Theorem 1 (second column), and of false positives, for different levels of noise affecting the data generation process described in equation 19. Across all tested hypersurfaces and noise levels, we observed zero false positives, providing strong empirical support for Theorem 1.

		$\ \tilde{\delta}\  < \ \hat{\delta}\ $	Th. 1 reduction	Th. 1 FP
$\sigma = 0.1$	Abs	0.59	0.13	0.00
	Exponential	0.87	0.46	0.00
	MixedQuadratic	0.37	0.15	0.00
	Paraboloid	0.34	0.10	0.00
	Quartic	0.34	0.11	0.00
$\sigma = 0.3$	Abs	0.59	0.13	0.00
	Exponential	0.68	0.36	0.00
	MixedQuadratic	0.38	0.15	0.00
	Paraboloid	0.33	0.10	0.00
	Quartic	0.33	0.11	0.00
$\sigma = 0.5$	Abs	0.59	0.13	0.00
	Exponential	0.60	0.36	0.00
	MixedQuadratic	0.38	0.15	0.00
	Paraboloid	0.33	0.10	0.00
	Quartic	0.30	0.12	0.00
$\sigma = 0.7$	Abs	0.59	0.13	0.00
	Exponential	0.55	0.33	0.00
	MixedQuadratic	0.38	0.15	0.00
	Paraboloid	0.33	0.10	0.00
	Quartic	0.29	0.09	0.00

Table 2: Predictions and false positives for Theorem 1 for different hypersurfaces with constant sign curvature, increasing levels of noise  $\sigma$ .

In Figure 5, two examples of calibration plots are shown for a constant sign curvature (paraboloid, left) and a non-constant sign curvature (Himmelblau, right). Blue points show the collected  $1.25 \times 10^5$  (25 studies on the 50% test sets of the  $10^4$  datasets) RMSE differences  $\|\hat{\delta}\| - \|\tilde{\delta}\|$  before and after reconciliation. Lines show the moving averages  $r$  from equation 20. The top and bottom panels refer to the upper and lower bounds obtained from Clopper-Pearson intervals, respectively, while the middle one shows  $r$  for the nominal response of Theorem 3. While a higher inter-study variance is expected in the central region of the plot, we can also observe a high variability for low density regions: in the Himmelblau plot very low probabilities assigned by Theorem 3 has a very high variance, indicating that this manifold presents a low number of instances where only few atoms of the predictive distribution are projected onto the half-space defined by equation 8. This is also likely due to the forecast  $\tilde{\nu}$  being unreliable in these conditions.

Table 3 reports the binary coverage  $c$  from 21. We see that the coverage tends to decrease with the noise  $\sigma$ , possibly due to the deterioration of the forecasts. If the predicted distribution  $\tilde{\nu}$  is not representative of the true distribution  $\nu$ , we expect the test to be unreliable.

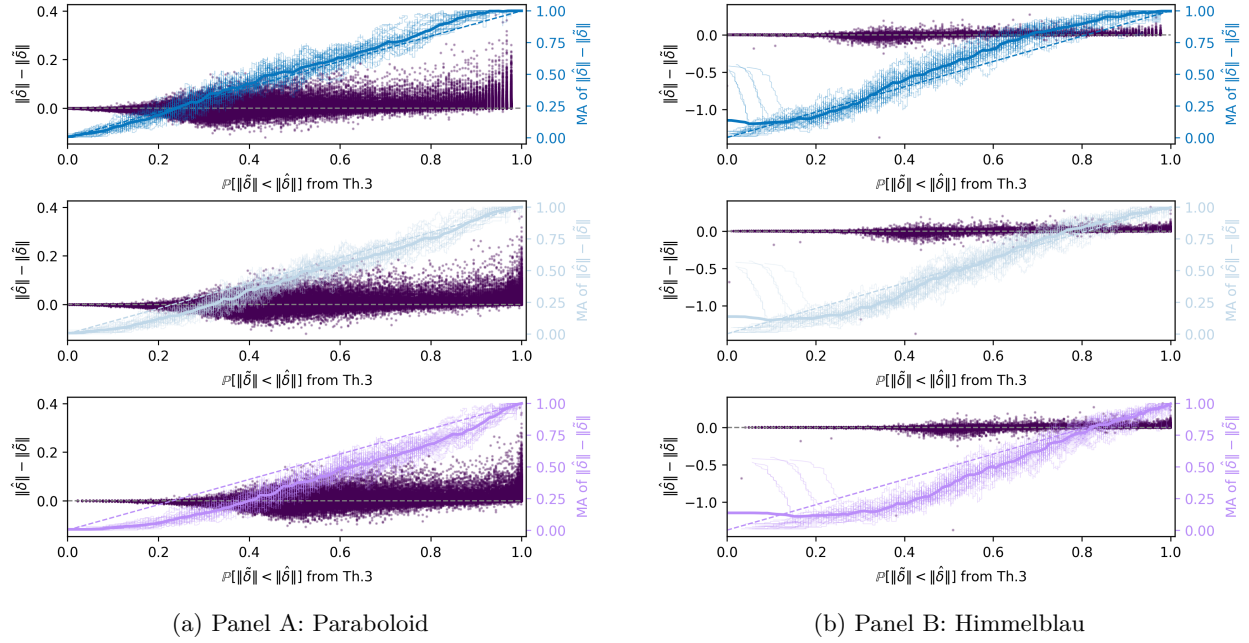


Figure 5: Theorem 3 calibration plots for two examples of constant sign curvature (left) and non-constant sign curvature (right) hypersurfaces, for noise level  $\sigma = 0.3$ .

	$\sigma = 0.1$	$\sigma = 0.3$	$\sigma = 0.5$	$\sigma = 0.7$
Ackley	0.56	0.41	0.25	0.15
Rastring	0.56	0.15	0.09	0.06
Rosenbrock	0.61	0.56	0.49	0.44
Quartic	0.72	0.67	0.66	0.65
Paraboloid	0.73	0.70	0.68	0.67
MixedQuadratic	0.74	0.71	0.70	0.68
Abs	0.76	0.76	0.76	0.75
Himmelblau	0.80	0.72	0.67	0.60
Exponential	0.90	0.78	0.72	0.71

Table 3: Binary coverage for Theorem 3 for different hypersurfaces, increasing levels of noise  $\sigma$ .

Figure 6 shows boxplots in terms of  $\Delta_{rel,opt}^{RMSE}$  defined in equation 24. Each boxplot contains  $\Delta_{rel,opt}^{RMSE}$  from the 25 randomized studies. From the definition of  $\Delta_{rel,opt}^{RMSE}$ , positive values indicate a reduction in RMSE compared to the unreconciled case, while a score of 1 represents the best binary reconciliation strategy.

The first boxplot represents the strategy in which the points are always reconciled, while the others represent the strategy of equation 22 with increasing values of  $\theta$ . For the exponential manifold, the best strategy, close to the optimal, is always reconciling. For the Rastring manifold, reconciling always worsens the RMSE, and choosing  $\theta = 0.6$  only slightly improves the RMSE over the unreconciled case. For the other 7 tested hypersurfaces, reconciling improves the RMSE, and the best reconciliation strategy based on Theorem 3 uses a threshold of  $\theta = 5$  or  $\theta = 6$ .

### 4.3 Manifolds with codimension $> 1$

We repeat the same tests for manifolds with codimension  $> 1$ , using Theorem 2 for manifolds with convex sub-level sets. Figure 12 in the appendix shows plots for the tested manifolds. Table 4 shows the ratio of cases with an observed RMSE reduction (first column), with a predicted reduction from Theorem 2 (second

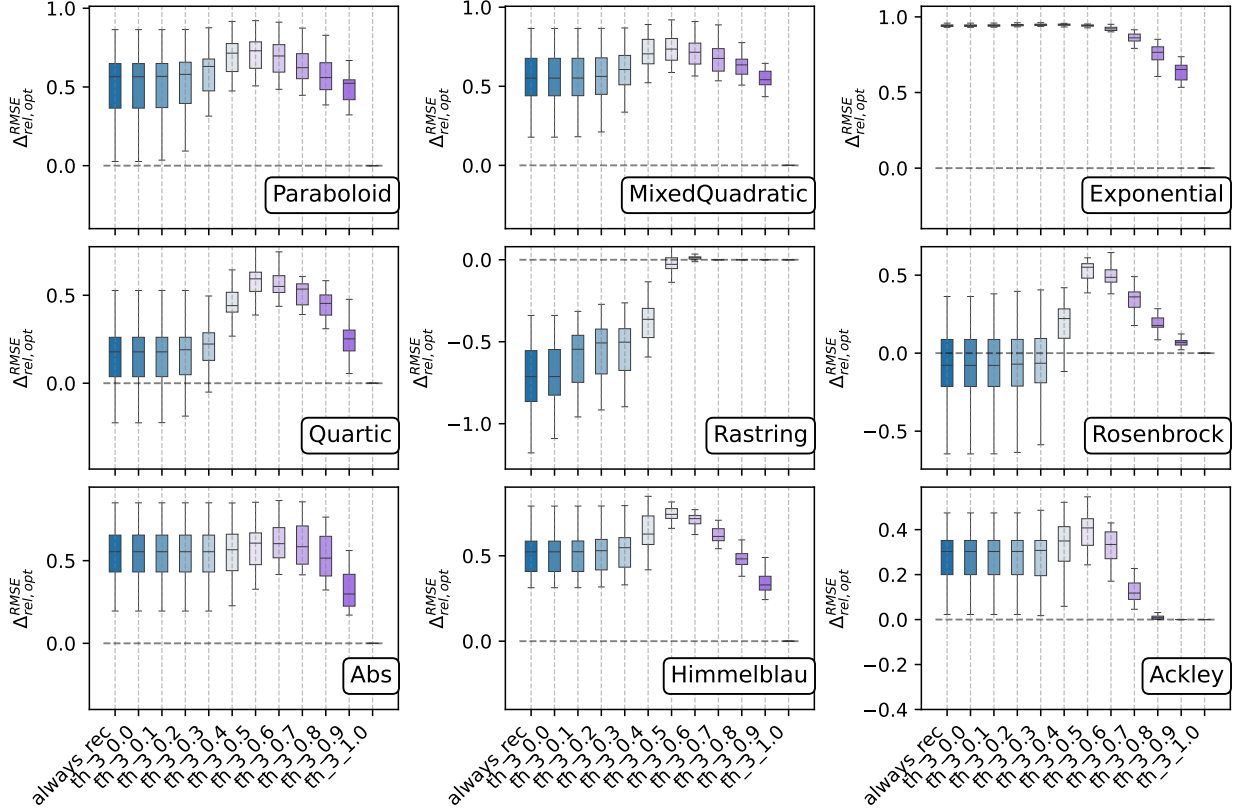


Figure 6: Boxplots of normalized improvements in RMSE for different reconciliation strategies for hypersurfaces.

column), and of false positives for manifolds with convex sub-level sets. Also in this case, no false positives are present. This empirical finding is fully consistent with Theorem 2 prediction that false positives should never occur under its assumptions.

Figure 7 shows the calibration plots for a quadratic-quartic and for a bowl-sin manifold. Also in this case, for the manifold shown in the first panel, we can observe some underestimation of the probability of an improvement due to reconciliation in low-density zones of the plot (low probabilities). As before, this is likely due to the forecast  $\tilde{\nu}$  being unreliable in these conditions.

The binary coverage is shown in Table 5, where the coverage is lower under high noise  $\sigma$  for the ring-trig and bowl-sin manifolds, which are the most irregular ones. Also in this case, this is likely the effect of the forecast  $\tilde{\nu}$  becoming unreliable.

Finally, Figure 8 shows boxplots in terms of  $\Delta_{rel,opt}^{RMSE}$  for different manifolds. Also in this case, for the saddle-poly, exp-cosh, and Rosenbrock style manifolds, always reconciling the sample is close to the optimal strategy. For the other manifolds, as for the hypersurface tests, the best reconciliation strategy based on Theorem 3 uses a threshold of  $\theta = 5$  or  $\theta = 6$ .

## 5 Note on geodesics and optimal deterministic predictions

This paper presented RMSE reduction conditions for nonlinear manifolds. However, for some forecasting tasks, other metrics could be more relevant. If we have a coherent forecast distribution  $\tilde{\nu}$  supported on  $M$ , finding the best point forecast is not as straightforward as in the linear case. In general, given a sample-based

		$\ \tilde{\delta}\  < \ \hat{\delta}\ $	Th. 2 reduction	Th. 2 FP
$\sigma = 0.1$	4d-4d	0.36	0.07	0.00
	bowl-sin	0.33	0.10	0.00
	e2-s2	0.34	0.09	0.00
	quad-qq	0.39	0.04	0.00
	quad-quad	0.33	0.10	0.00
$\sigma = 0.3$	4d-4d	0.35	0.08	0.00
	bowl-sin	0.36	0.07	0.00
	e2-s2	0.32	0.09	0.00
	quad-qq	0.39	0.02	0.00
	quad-quad	0.33	0.10	0.00
$\sigma = 0.5$	4d-4d	0.32	0.09	0.00
	e2-s2	0.29	0.08	0.00
	quad-qq	0.42	0.01	0.00
	quad-quad	0.33	0.10	0.00
$\sigma = 0.7$	4d-4d	0.28	0.09	0.00
	e2-s2	0.28	0.07	0.00
	quad-qq	0.48	0.01	0.00
	quad-quad	0.33	0.10	0.00

Table 4: Predictions and false positives for Theorem 2 for different  $M$  with  $\text{codimension} > 1$  and convex sub-level sets, increasing levels of noise  $\sigma$ .

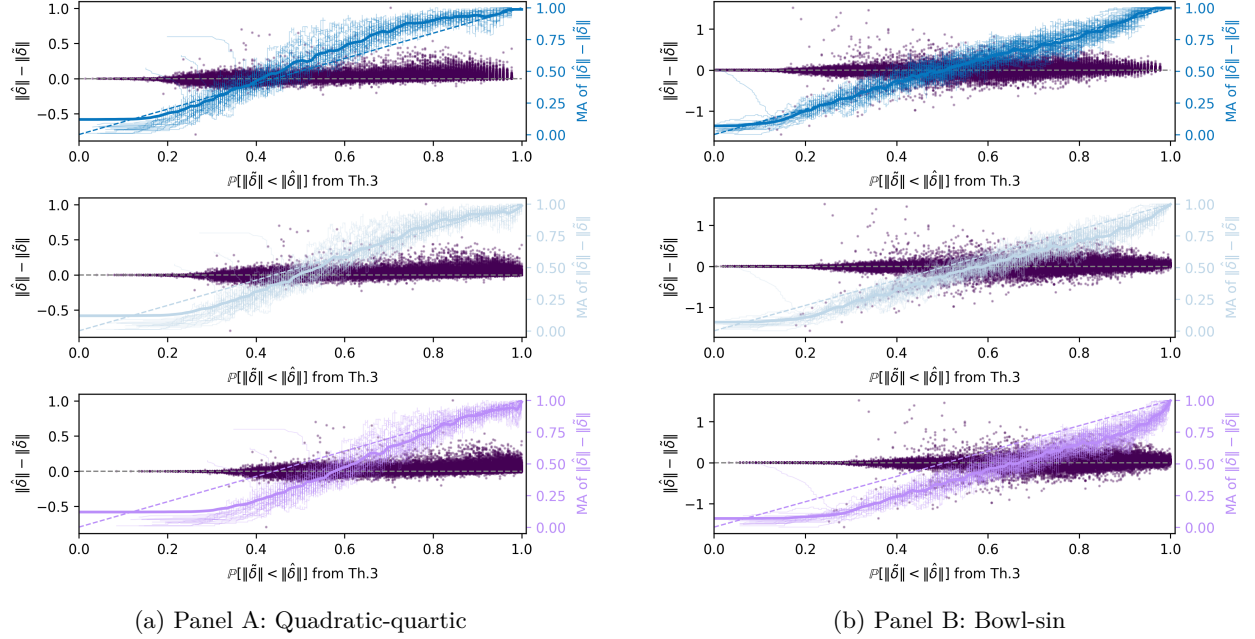
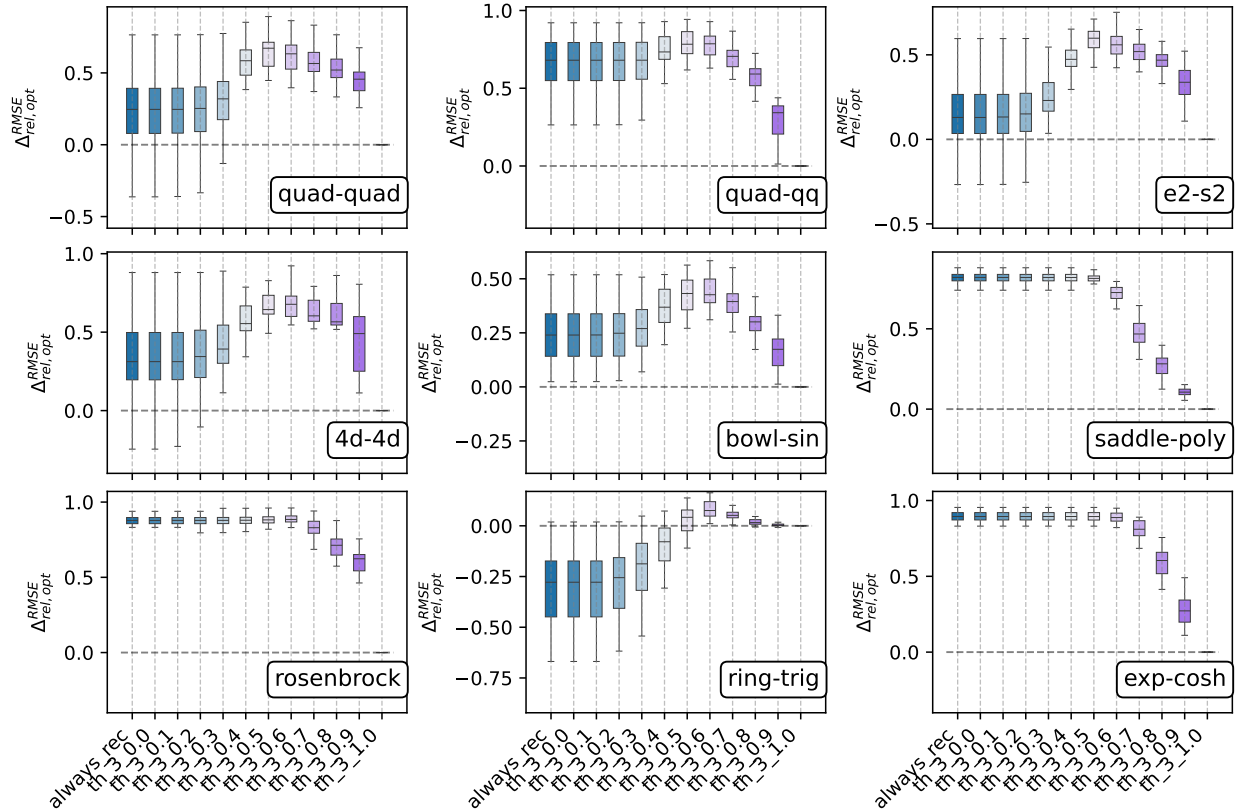


Figure 7: Theorem 3 calibration plots for two examples of  $M$  with  $\text{codimension} > 1$  and convex (left) and non-convex (right) sub-level sets, for noise level  $\sigma = 0.3$ .

representation of  $\tilde{\nu}$ ,  $(\tilde{z}_i, \tilde{\pi}_i)_{i=1}^S$ , the best deterministic point can be found solving:

$$\tilde{z}^* = \arg \min_{z \in M} \sum_{i=1}^S \tilde{\pi}_i l(z, \tilde{z}_i)^2 \quad (25)$$

	$\sigma = 0.1$	$\sigma = 0.3$	$\sigma = 0.5$	$\sigma = 0.7$
exp-cosh	0.66	0.64	0.58	0.52
ring-trig	0.67	0.33	0.35	0.30
quad-qq	0.69	0.66	0.60	0.52
saddle-poly	0.69	0.61	0.55	0.54
quad-quad	0.70	0.70	0.68	0.69
bowl-sin	0.71	0.56	0.46	0.37
e2-s2	0.71	0.69	0.69	0.70
4d-4d	0.72	0.68	0.69	0.68
rosenbrock	0.86	0.76	0.70	0.66

Table 5: Binary coverage for Theorem 3 for different  $M$  with codimension $>1$ , increasing levels of noise  $\sigma$ .Figure 8: Boxplots of normalized improvements in RMSE for different reconciliation strategies for manifolds of codimension $>1$ .

In the linear case  $M = \mathbb{R}^n$  and under RMSE loss  $l$ , the minimizer is the weighted average of the samples, while when  $M$  is a nonlinear manifold, this must be solved numerically. For general Manifolds, problem 25 is also known as Fréchet mean (Bačák, 2014; Lou et al., 2020), and is in general non-trivial to solve. For some forecasting tasks, it may be more meaningful to consider the geodesic distance  $g$  as the loss  $l$ , rather than the RMSE. Defined  $\gamma : [a, b] \rightarrow M$  a curve on the manifold, the geodesic between  $a$  and  $b$  is the distance of the shortest path on  $M$  between these two points:

$$g(a, b) = \inf \int_a^b \|\gamma'(t)\|_p dt \quad (26)$$



where the Riemannian metric  $\rho = (\rho_x)_{x \in M}$  is a smooth collection of inner products  $\rho_x : T_x M \times T_x M \rightarrow \mathbb{R}$ . The following example illustrates the importance of considering geodesic distance for some forecasting tasks.




Figure 9: Conceptual plot for the importance of using geodesics as scoring rules in forecasting. Red triangle: true point  $z$ ; green dots: examples of reconciled forecasts  $\tilde{z}_1, \tilde{z}_2$ ; blue lines: euclidean distances; red and green lines: shortest paths on  $M$  from predictions to  $z$ .

#### Example 2: Optimal forecasts under geodesic distance

*A heavy snowfall is forecast for tomorrow, with spatial probability distribution  $\nu(z = [x, y])$ , over the area surrounding a large lake. The event is expected to significantly disrupt road traffic. To reduce the economic impact, road maintenance crews can be preemptively deployed to a specific location along the road network, enabling a rapid response. Figure 9 illustrates two candidate deployment locations,  $\tilde{z}_1$  and  $\tilde{z}_2$ , marked as green dots and selected by the decision maker. While  $\tilde{z}_2$  is closer to the true disruption point  $z$  (red triangle) in terms of Euclidean distance, deploying crews there would incur much higher operational costs than choosing  $\tilde{z}_1$ . This is due to the road network constraints, as shown by the red and green geodesic paths, which represent the cost of relocating crews from each candidate point to  $z$  along the road infrastructure.*

## 6 Conclusions

This paper has advanced the theory of nonlinear forecast reconciliation through a geometric lens, establishing both deterministic guarantees and a practical probabilistic framework. Our first key contribution lies in Theorems 1 and 2, which provide sufficient conditions for strict RMSE reduction when reconciling forecasts via orthogonal projection onto constraint manifolds  $M = \{z : f(z) = 0\}$ . These deterministic results hold for codimension 1 and higher when all component functions  $f_i$  exhibit convex sub-level or super-level sets.

However, our simulation study under isotropic-Gaussian errors revealed significant practical limitations: these geometric conditions apply only to subsets of prediction points, particularly on highly curved manifolds. This observation motivated our second major contribution—Theorem 3—which addresses the fundamental question of when reconciliation is likely to help. By introducing a consistent Monte Carlo estimator for  $\Pr(\|z - \hat{z}\| \geq \|z - \tilde{z}\|)$ , this probabilistic framework accommodates arbitrary manifolds without curvature assumptions or codimension restrictions, while providing confidence intervals through exact or bootstrap methods.

Empirical validation confirmed the soundness of all theoretical contributions. Crucially, we observed zero false positives across all tests of Theorems 1 and 2, consistent with their theoretical guarantees. The estimator from Theorem 3 demonstrated excellent calibration, with empirical coverage of confidence bands aligning closely with nominal levels. Most significantly, when using this probability estimate to decide whether to reconcile—rather than naïve always/never approaches—we achieved measurable RMSE reduction across diverse manifolds.

Collectively, these results bridge geometric theory with forecasting practice: while Theorems 1 and 2 establish fundamental error-reduction principles under idealized conditions, Theorem 3 delivers a versatile decision-



making tool for real-world applications. This dual perspective enables more reliable forecast integration where theoretical guarantees meet practical uncertainty.

## 6.1 Future research directions

This paper lays the foundations for future research directions in the field of nonlinear reconciliation. An interesting research direction is to craft analogous theorems using the geodesic distance instead of the RMSE. An interesting research topic is to find a function of the geodesic that is a proper scoring rule for some classes of nonlinear manifolds. For example, the energy score is a proper scoring rule and it's a function of the Euclidean distance. Another interesting research direction is to find robust rules to set the threshold parameter  $\theta$  that maximizes the accuracy of the reconciliation strategy. Lastly, since the nonlinear projection problem is equivalent to an iterated extended Kalman filter, the RMSE reduction theorems and associated reconciliation strategies could be applied to nonlinear state estimation problems.

## References

- George Athanasopoulos, Roman A. Ahmed, and Rob J. Hyndman. Hierarchical forecasts for Australian domestic tourism. *International Journal of Forecasting*, 25(1):146–166, January 2009. ISSN 0169-2070. doi: 10.1016/j.ijforecast.2008.07.004. URL <https://www.sciencedirect.com/science/article/pii/S0169207008000691>.
- George Athanasopoulos, Rob J. Hyndman, Nikolaos Kourentzes, and Fotios Petropoulos. Forecasting with temporal hierarchies. *European Journal of Operational Research*, 262(1):60–74, October 2017. ISSN 0377-2217. doi: 10.1016/j.ejor.2017.02.046. URL <https://www.sciencedirect.com/science/article/pii/S0377221717301911>.
- Miroslav Bačák. Computing Medians and Means in Hadamard Spaces. *SIAM Journal on Optimization*, 24(3):1542–1566, January 2014. ISSN 1052-6234. doi: 10.1137/140953393. URL <https://epubs.siam.org/doi/10.1137/140953393>. Publisher: Society for Industrial and Applied Mathematics.
- Souhaib Ben Taieb, Jiafan Yu, Mateus Barreto, and Ram Rajagopal. Regularization in Hierarchical Time Series Forecasting with Application to Electricity Smart Meter Data. *Proceedings of the AAAI Conference on Artificial Intelligence*, 31(1), February 2017. ISSN 2374-3468, 2159-5399. doi: 10.1609/aaai.v31i1.11167. URL <https://ojs.aaai.org/index.php/AAAI/article/view/11167>.
- M. Brown Do Coutto Filho and J.C.S. De Souza. Forecasting-Aided State Estimation—Part I: Panorama. *IEEE Transactions on Power Systems*, 24(4):1667–1677, November 2009. ISSN 0885-8950, 1558-0679. doi: 10.1109/TPWRS.2009.2030295. URL <http://ieeexplore.ieee.org/document/5233865/>.
- Andrea Cini, Slobodan Lukovic, and Cesare Alippi. Cluster-based Aggregate Load Forecasting with Deep Neural Networks. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8, July 2020. doi: 10.1109/IJCNN48605.2020.9207503. URL <https://ieeexplore.ieee.org/document/9207503>. ISSN: 2161-4407.
- Andrea Cini, Danilo Mandic, and Cesare Alippi. Graph-based time series clustering for end-to-end hierarchical forecasting. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *ICML’24*, pp. 8985–8999, Vienna, Austria, July 2024. JMLR.org.
- C. J. Clopper and E. S. Pearson. The Use Of Confidence Or Fiducial Limits Illustrated In The Case Of The Binomial. *Biometrika*, 26(4):404–413, December 1934. ISSN 0006-3444. doi: 10.1093/biomet/26.4.404. URL <https://doi.org/10.1093/biomet/26.4.404>.
- Giorgio Corani, Dario Azzimonti, João P. S. C. Augusto, and Marco Zaffalon. Probabilistic Reconciliation of Hierarchical Forecast via Bayes’ Rule. In Frank Hutter, Kristian Kersting, Jefrey Lijffijt, and Isabel Valera (eds.), *Machine Learning and Knowledge Discovery in Databases*, Lecture Notes in Computer Science, pp. 211–226, Cham, 2021. Springer International Publishing. ISBN 978-3-030-67664-3. doi: 10.1007/978-3-030-67664-3\_13.

- Chiara Di Modica, Pierre Pinson, and Souhaib Ben Taieb. Online forecast reconciliation in wind power prediction. *Electric Power Systems Research*, 190:106637, January 2021. ISSN 03787796. doi: 10.1016/j.epsr.2020.106637. URL <https://linkinghub.elsevier.com/retrieve/pii/S0378779620304405>.
- D. M. Dunn, W. H. Williams, and T. L. Dechaine. Aggregate versus Subaggregate Models in Local Area Forecasting. *Journal of the American Statistical Association*, 71(353):68–71, March 1976. ISSN 0162-1459. doi: 10.1080/01621459.1976.10481478. URL <https://www.tandfonline.com/doi/abs/10.1080/01621459.1976.10481478>. Publisher: ASA Website \_eprint: <https://www.tandfonline.com/doi/pdf/10.1080/01621459.1976.10481478>.
- Gene Fliedner. Hierarchical forecasting: issues and use guidelines. *Industrial Management & Data Systems*, 101(1):5–12, February 2001. ISSN 0263-5577. doi: 10.1108/02635570110365952. URL <https://www.emerald.com/insight/content/doi/10.1108/02635570110365952/full/html>.
- Rob J. Hyndman, Roman A. Ahmed, George Athanasopoulos, and Han Lin Shang. Optimal combination forecasts for hierarchical time series. *Computational Statistics and Data Analysis*, 2011. ISSN 01679473. doi: 10.1016/j.csda.2011.03.006. ISBN: 0167-9473.
- Jooyoung Jeon, Anastasios Panagiotelis, and Fotios Petropoulos. Probabilistic forecast reconciliation with applications to wind power and electric load. *European Journal of Operational Research*, 279(2):364–379, December 2019. ISSN 03772217. doi: 10.1016/j.ejor.2019.05.020. URL <https://linkinghub.elsevier.com/retrieve/pii/S0377221719304242>.
- Aaron Lou, Isay Katsman, Qingxuan Jiang, Serge Belongie, Ser-Nam Lim, and Christopher De Sa. Differentiating through the Fréchet Mean. In *Proceedings of the 37th International Conference on Machine Learning*, pp. 6393–6403. PMLR, November 2020. URL <https://proceedings.mlr.press/v119/lou20a.html>. ISSN: 2640-3498.
- Spyros Makridakis, Evangelos Spiliotis, and Vassilios Assimakopoulos. M5 accuracy competition: Results, findings, and conclusions. *International Journal of Forecasting*, 38(4):1346–1364, October 2022. ISSN 01692070. doi: 10.1016/j.ijforecast.2021.11.013. URL <https://linkinghub.elsevier.com/retrieve/pii/S0169207021001874>.
- Guy H. Orcutt, Harold W. Watts, and John B. Edwards. Data Aggregation and Information Loss. *The American Economic Review*, 58(4):773–787, 1968. ISSN 0002-8282. URL <https://www.jstor.org/stable/1815532>. Publisher: American Economic Association.
- Anastasios Panagiotelis, George Athanasopoulos, Puwasala Gamakumara, and Rob J. Hyndman. Forecast reconciliation: A geometric view with new insights on bias correction. *International Journal of Forecasting*, 37(1):343–359, January 2021. ISSN 0169-2070. doi: 10.1016/j.ijforecast.2020.06.004. URL <https://www.sciencedirect.com/science/article/pii/S0169207020300911>.
- Anastasios Panagiotelis, Puwasala Gamakumara, George Athanasopoulos, and Rob J. Hyndman. Probabilistic forecast reconciliation: Properties, evaluation and score optimisation. *European Journal of Operational Research*, 306(2):693–706, April 2023. ISSN 0377-2217. doi: 10.1016/j.ejor.2022.07.040. URL <https://www.sciencedirect.com/science/article/pii/S0377221722006087>.
- Jeroen Rombouts, Marie Ternes, and Ines Wilms. Cross-temporal forecast reconciliation at digital platforms with machine learning. *International Journal of Forecasting*, 41(1):321–344, January 2025. ISSN 0169-2070. doi: 10.1016/j.ijforecast.2024.05.008. URL <https://www.sciencedirect.com/science/article/pii/S0169207024000475>.
- Evangelos Spiliotis, Mahdi Abolghasemi, Rob J. Hyndman, Fotios Petropoulos, and Vassilios Assimakopoulos. Hierarchical forecast reconciliation with machine learning. *Applied Soft Computing*, 112:107756, November 2021. ISSN 1568-4946. doi: 10.1016/j.asoc.2021.107756. URL <https://www.sciencedirect.com/science/article/pii/S1568494621006773>.

- Souhaib Ben Taieb, James W. Taylor, and Rob J. Hyndman. Hierarchical Probabilistic Forecasting of Electricity Demand With Smart Meter Data. *Journal of the American Statistical Association*, 116(533): 27–43, January 2021. ISSN 0162-1459, 1537-274X. doi: 10.1080/01621459.2020.1736081. URL <https://www.tandfonline.com/doi/full/10.1080/01621459.2020.1736081>.
- J. A. Thorpe. Convex Surfaces. In J. A. Thorpe (ed.), *Elementary Topics in Differential Geometry*, pp. 95–107. Springer, New York, NY, 1979. ISBN 978-1-4612-6153-7. doi: 10.1007/978-1-4612-6153-7\_13. URL [https://doi.org/10.1007/978-1-4612-6153-7\\_13](https://doi.org/10.1007/978-1-4612-6153-7_13).
- Shanika L. Wickramasuriya, George Athanasopoulos, and Rob J. Hyndman. Optimal Forecast Reconciliation for Hierarchical and Grouped Time Series Through Trace Minimization. *Journal of the American Statistical Association*, 114(526):804–819, April 2019. ISSN 0162-1459. doi: 10.1080/01621459.2018.1448825. URL <https://doi.org/10.1080/01621459.2018.1448825>. Publisher: ASA Website \_eprint: <https://doi.org/10.1080/01621459.2018.1448825>.
- Dazhi Yang, Hao Quan, Vahid R. Disfani, and Carlos D. Rodríguez-Gallegos. Reconciling solar forecasts: Temporal hierarchy. *Solar Energy*, 158:332–346, December 2017. ISSN 0038-092X. doi: 10.1016/j.solener.2017.09.055. URL <https://www.sciencedirect.com/science/article/pii/S0038092X17308423>.
- Lorenzo Zambon, Dario Azzimonti, and Giorgio Corani. Efficient probabilistic reconciliation of forecasts for real-valued and count time series. *Statistics and Computing*, 34(1):21, November 2023. ISSN 1573-1375. doi: 10.1007/s11222-023-10343-y. URL <https://doi.org/10.1007/s11222-023-10343-y>.

## A Appendix

### A.1 Equivalence of problem equation 1 and FASE method

The FASE method (Brown Do Coutto Filho & De Souza, 2009) formulates the state estimation procedure using a discrete time-variant space-state representation:

$$x_{k+1} = F_k x_k + g_k + w_k \quad (27)$$

$$y_k = h_k(x_k) + v_k \quad (28)$$

Usually, in FASE, the state  $x$  are the magnitudes and angles of the voltages, while  $y$  represents the active and reactive powers. Equation 28 represents an a-priori forecast of the state  $x$  at time  $t + 1$  given the state estimation at time  $t$ . Nonlinear forecasting models could also be used, leading to variations of the method, but linear models have been found to be difficult to beat.

The objective of the FASE method is formulated as:

$$\min_x \|\hat{y} - h(x)\|_R^2 + \|x - \hat{x}\|_M^2 \quad (29)$$

where  $R$  and  $M$  are positive definite weighting matrices for the forecast errors in  $y$  and  $x$ , respectively.

To turn the nonlinear projection problem of 1 into the FASE problem of 29, we define  $z = [x^T, y^T]^T$ ,  $M = \{f(z) = 0\} = \{y = h(x)\}$ . Then 1 becomes:

$$\begin{aligned} \min_z \quad & \|z - \hat{z}\|_W^2 \\ \text{s.t.} \quad & y = h(x) \end{aligned} \quad (30)$$

with weight matrix

$$W = \begin{bmatrix} M & 0 \\ 0 & R \end{bmatrix}$$

The constraint  $y = h(x)$  implies that any feasible point  $z$  must satisfy

$$z = \begin{bmatrix} x \\ h(x) \end{bmatrix}$$

We can therefore rewrite the projection problem as an unconstrained optimization over  $x$ :

$$\min_x \left\| \begin{bmatrix} x \\ h(x) \end{bmatrix} - \begin{bmatrix} \hat{x} \\ \hat{y} \end{bmatrix} \right\|_W^2$$

Expanding the norm with block-diagonal weights:

$$\left\| \begin{bmatrix} x - \hat{x} \\ h(x) - \hat{y} \end{bmatrix} \right\|_W^2 = \|x - \hat{x}\|_M^2 + \|h(x) - \hat{y}\|_R^2$$

This is exactly the objective function of the FASE formulation. This shows the two objectives are equivalent when considering a graph  $y = h(x)$  for the definition of the manifold and when the weight matrix  $W$  is block-diagonal, partitioned according to the dimensionality of  $x$  and  $y$ .

## A.2 Second fundamental form

Given  $f : \mathbb{R}^n \rightarrow \mathbb{R} \in C^2$ ,  $M = \{z : f(z) = c\} \subset \mathbb{R}^n$  is a level set of  $f$  and a regular hypersurface of  $f$  if  $\nabla f(z) \neq 0$  on  $M$ .  $H = D^2 f(z)$  is the Hessian of the function at  $z$ ,  $E_z \in \mathbb{R}^{n \times (n-1)}$  an orthonormal basis of the tangent space  $T_x M$  at  $x$ . Then the *second fundamental form*  $\mathbb{I}_p \in \mathbb{R}$  at a point  $p \in M$  is a symmetric bilinear form on the tangent space  $T_p M$ , defined by:

$$\mathbb{I}_p(z, w) = \langle D_z \nu, w \rangle, \quad z, w \in T_p M \quad (31)$$

$$\nu(z) = \frac{\nabla f(z)}{\|\nabla f(z)\|} \quad (32)$$

where  $\nu(z)$  is the unit normal vector field, measures the rate of change of the normal vector in the direction  $z$ , i.e., the normal curvature. In coordinates, this reduces to:

$$\mathbb{I}_p(z, z) = \frac{z^\top H(p) z}{\|\nabla f(p)\|}, \quad z \in T_p M \quad (33)$$

which describes the change of the normal acceleration of the surface in the direction of  $x$  belonging to the tangent space. A sub-level set is locally convex iff the second fundamental form is positive semidefinite (theorem 1 and 3 in Thorpe (1979)):

$$z^\top H(p) z \geq 0 \quad \text{for all } z \in T_z M \quad (34)$$

We can test this condition by inspecting if the Hessian projected on the tangent space is positive definite, or equivalently, by checking if the smallest eigenvalue of the restricted Hessian is greater than 0:

$$H_{\tan}(p) = E^\top H(p) E \quad (35)$$

$$\lambda_{\min}(H_{\tan}(p)) > 0 \Leftrightarrow H_{\tan}(p) \succ 0 \quad (36)$$

This shows that  $\lambda_{\min}(H_{\tan}(p))$  is a natural choice for inspecting the curvature of the levelset at  $p$ . Figure 10 shows different level sets for  $f(z) = z_1^4 - 3z_1^2 + z_2^2$  (left axis) and a specific level set for  $f(z) = z_1^4 - z_1^2 + z_2^2 + z_3^2 - 2z_1 z_3$  (right axis), both colored by  $\lambda_{\min}(H_{\tan})$ . The second fundamental form can be used to write the following second-order approximation for  $M$ :

$$z = \tilde{z} + t + \underbrace{\frac{1}{2} \nu(\tilde{z}) \mathbb{I}_{\tilde{z}}(t, t)}_{\text{scalar}} + o(\|t\|^2) \quad (37)$$

where  $t \in T_{\tilde{z}} M$  is in the tangent direction, while the correction introduced by the second fundamental form lies in the direction of the unit normal  $\nu(\tilde{z})$ ,  $\text{Im}(\nabla f) \subset \mathbb{R}^n$ .

### A.2.1 Vector valued functions

In the case of manifold with codimension higher than 1, that is defined by a vector valued constraint  $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ ,  $M = \{z \in \mathbb{R}^n : f(z) = 0\}$ ,  $m < n$ , the normal space  $\text{Im}(J_f(z))$  spans the Jacobian rows. In this case  $\mathbb{I}_{\tilde{z}}(t, t) \in \mathbb{R}^m$  is a vector and the analogous expression to equation 37 for second order approximation can be written as:

$$z = \tilde{z} + t + \frac{1}{2} J_f(\tilde{z})^\top \cdot \mathbb{I}_{\tilde{z}}(t, t) + o(\|t\|^2) \quad (38)$$

For the vector-valued case, an orthonormal basis  $E$  for the tangent space  $T_z M = \ker J(z)$  can be found by QR decomposition of the transposed Jacobian.




Figure 10: Two examples of hypersurfaces colored by  $\lambda_{\min}(H_{tan})$ . These examples illustrate how  $\lambda_{\min}(H_{tan})$  can be thought of as the degree with which the surface *bends away* or *bends towards* a point in the ambient space.

### A.3 Proof of equation 12

In the following, we prove equation 12 for hypersurface  $M = \{z : f(z) = 0\} \subset \mathbb{R}^n$  with  $f$  having a convex sub or super level set. Consider two points  $z$  and  $\tilde{z}$ , both belonging to  $M$  and minimally displaced  $\varepsilon = z - \tilde{z}$ . We can write the second-order Taylor expansion at  $\tilde{z}$ :

$$0 = f(z) = f(\tilde{z} + \varepsilon) = f(\tilde{z}) + \nabla f(\tilde{z})^T \varepsilon + \frac{1}{2} \varepsilon^T H(\tilde{z}) \varepsilon + o(\|\varepsilon\|^2) \quad (39)$$

Since  $f(\tilde{z}) = 0$ , and since  $\varepsilon$  becomes tangent to  $M$  in the limit, neglecting higher order terms, we can write:

$$\nabla f(\tilde{z})^T \varepsilon = -\frac{1}{2} \varepsilon^T H(\tilde{z}) \varepsilon \quad \forall \varepsilon \in T_{\tilde{z}} M \quad (40)$$

where  $T_{\tilde{z}} M$  is the tangent space at  $\tilde{z}$ . Equivalently, we can write:

$$\nabla f(\tilde{z})^T \varepsilon = -\frac{1}{2} E_{\tilde{z}}^T H(\tilde{z}) E_{\tilde{z}} \quad (41)$$

where  $E_{\tilde{z}} \in \mathbb{R}^{n \times (n-1)}$  is an orthonormal basis of the tangent space  $T_{\tilde{z}} M$  at  $\tilde{z}$ . Then:

$$\text{sign}(\nabla f(\tilde{z})^T \varepsilon) = -\text{sign}(E_{\tilde{z}}^T H(\tilde{z}) E_{\tilde{z}}) = -\text{sign}(\lambda_{\min}(H_{tan}(\tilde{z}))) \quad (42)$$

where the last equality comes from the definiteness of the restricted Hessian for convex / concave curves equation 34-equation 36. We proved equation 42 in the limit  $\varepsilon \rightarrow 0$ , but since  $M$  is the boundary of a convex set, it has a supporting hyperplane in  $\tilde{z}$ , which means that  $\text{sign}(\nabla f(\tilde{z})^T \varepsilon)$  is constant on  $M$ , and thus equation 42 holds also for  $\varepsilon = \tilde{\delta}$ .

#### A.4 Plots of tested manifolds




Figure 11: Plots of tested hypersurfaces.




Figure 12: Plots of tested manifolds with codimension  $> 1$ . Since it's hard to visualize level sets of  $f : \mathbb{R}^4 \rightarrow \mathbb{R}^2$ , we plotted level sets of  $f_1$  and  $f_2$  separately.