Les Cahiers du GERAD ISSN: 0711–2440

Temporally-constrained spatial patterns regression for reconstruction of high-resolution hydrological forcing

V. Houssou, J. Carreau

G–2026–46

August 2026

La publication de ces rapports de recherche est rendue possible grâce au soutien de HEC Montréal, Polytechnique Montréal, Université McGill, Université du Québec à Montréal, ainsi que du Fonds de recherche du Québec – Nature et technologies.

Dépôt légal – Bibliothèque et Archives nationales du Québec, 2026
– Bibliothèque et Archives Canada, 2026

The publication of these research reports is made possible thanks to the support of HEC Montréal, Polytechnique Montréal, McGill University, Université du Québec à Montréal, as well as the Fonds de recherche du Québec – Nature et technologies.

Legal deposit – Bibliothèque et Archives nationales du Québec, 2026
– Library and Archives Canada, 2026

GERAD HEC Montréal
3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec) Canada H3T 2A7

Tél. : 514 340-6053
Téléc. : 514 340-5665
info@gerad.ca
www.gerad.ca

Temporally-constrained spatial patterns regression for reconstruction of high-resolution hydrological forcing

Vihotogbé Houssou [a, b]

Julie Carreau [a, b]

[a] *Département de mathématiques et de génie industriel, Polytechnique Montréal (Qc), Canada*

[b] *GERAD, Montréal (Qc), Canada, H3T 1J4*

vihotogbe-2.houssou@polymtl.ca
julie.carreau@polymtl.ca

August 2026
Les Cahiers du GERAD
G–2026–46

Abstract : Accurate reconstruction of high-resolution hydrological forcing fields from sparse station observations remains challenging, particularly in poorly instrumented regions. In addition, most existing interpolation methods ignore temporal consistency, potentially introducing artificial jitter that propagates as significant errors when fed into hydrological models. To address this, we introduce Temporally-constrained Spatial Pattern Regression (T-SPR), an extension of the Spatial Pattern Regression (SPR) framework that combines anomaly-based interpolation with temporal regularization of the reconstructed forcing fields. An intermediate anomaly-based variant (A-SPR) is also considered to isolate the contribution of the climatological decomposition. These three variants (SPR, A-SPR and T-SPR) are evaluated using high-resolution regional climate model simulations and real station observations for daily precipitation, minimum temperature, and maximum temperature over southern Québec, Canada. Results show that anomaly-based interpolation substantially improves reconstruction accuracy under sparse observational conditions, particularly for temperature forcings. Temporal regularization further enhances robustness and temporal consistency by reducing spurious day-to-day variability and sensitivity to station configuration. The greatest benefits are observed at extremely low station densities, where information is most limited. These findings demonstrate that T-SPR, which combines climatological information and temporal coherence, provides an effective and computationally efficient approach for applications to hydrological forcing reconstruction in data-sparse catchments.

Keywords: Spatiotemporal interpolation; meteorological forcing; poorly gauged basins; vector autoregression; anomaly-based interpolation; hydrological modeling

Résumé : La reconstruction précise de champs de forçage hydrologique à haute résolution à partir d'observations éparses de stations demeure un défi, particulièrement dans les régions peu instrumentées. De plus, la plupart des méthodes d'interpolation existantes ignorent la cohérence temporelle, ce qui peut introduire un bruit artificiel se propageant sous forme d'erreurs significatives lorsqu'il est intégré dans les modèles hydrologiques. Pour répondre à ce problème, nous introduisons la régression de patrons spatiaux à contrainte temporelle (*Temporally-constrained Spatial Pattern Regression*, T-SPR), une extension du cadre de régression de patrons spatiaux (*Spatial Pattern Regression*, SPR) qui combine l'interpolation basée sur les anomalies avec une régularisation temporelle des champs de forçage reconstruits. Une variante intermédiaire basée sur les anomalies (A-SPR) est également considérée afin d'isoler la contribution de la décomposition climatologique. Ces trois variantes (SPR, A-SPR et T-SPR) sont évaluées à l'aide de simulations de modèle climatique régional à haute résolution et d'observations réelles de stations, pour les précipitations journalières, la température minimale et la température maximale sur le sud du Québec, Canada. Les résultats montrent que l'interpolation basée sur les anomalies améliore substantiellement la précision de la reconstruction en conditions d'observation éparses, en particulier pour les forçages de température. La régularisation temporelle renforce davantage la robustesse et la cohérence temporelle en réduisant la variabilité artificielle jour-à-jour et la sensibilité à la configuration des stations. Les bénéfices les plus importants sont observés aux densités de stations extrêmement faibles, là où l'information est la plus limitée. Ces résultats démontrent que la T-SPR, qui combine l'information climatologique et la cohérence temporelle, constitue une approche efficace et peu coûteuse en calcul pour la reconstruction de forçages hydrologiques dans des bassins versants pauvres en données.

Mots clés : Interpolation spatio-temporelle; forçage météorologique; bassins versants peu jaugés; autorégression vectorielle; interpolation basée sur les anomalies; modélisation hydrologique

---

Acknowledgements: The authors would like to acknowledge funding by the Natural Sciences and Engineering Research Council of Canada (NSERC), by Fonds de Recherche du Québec Nature et Technologies (FRQNT) and by IVADO.

1 Introduction

High-resolution gridded meteorological data are a fundamental input to a wide range of applications in climate science and hydrology, including distributed hydrological modelling, flood forecasting, drought monitoring, and climate impact assessment (Warren et al., 2022; Lucas-Picher et al., 2020; Livneh et al., 2015; Leduc et al., 2019). Precipitation, minimum temperature, and maximum temperature are among the most widely required variables (Tarek et al., 2020; Gasset et al., 2021; Faghih and Brissette, 2023), as they jointly drive the water and energy cycles at the land surface (Maurer et al., 2002; Spangler et al., 2019; Behnke et al., 2016). In regions where station networks are dense, these variables can be reconstructed with reasonable accuracy using classical spatial interpolation methods (Li and Heap, 2014; Ly et al., 2013; Hwang et al., 2012). However, a large fraction of the Earth's land surface — particularly in mountainous, arctic, and tropical regions — is instrumented by a very limited number of stations, making reliable reconstruction of daily gridded fields a persistent challenge (Kidd et al., 2021; Segond et al., 2007; Looper and Vieux, 2012; Hwang et al., 2012).

Classical spatial interpolation methods such as Inverse Distance Weighting (IDW), ordinary kriging, and kriging with external drift (KED) have been widely applied to reconstruct meteorological fields from station observations (Li and Heap, 2014; Ly et al., 2013; Li and Heap, 2008; Bokke, 2017; Hartkamp et al., 1999). These methods estimate values at ungauged locations as weighted combinations of neighboring observations and can optionally incorporate auxiliary information — such as topography or climatological means derived from regional climate models (RCMs) — as covariates or drift terms (Livneh et al., 2015; Werner et al., 2019; Varentsov et al., 2020; Hengl et al., 2003). While they may be effective under moderate network densities, these approaches exploit auxiliary gridded information in a limited way: they use it to constrain the spatial mean of the field, but do not leverage the rich spatial structure encoded in high-resolution gridded datasets (Houssou and Carreau, 2026).

To better exploit the spatial structure embedded in RCM simulations, Houssou and Carreau (2026) introduced the Spatial Pattern Regression (SPR) framework. SPR extracts a low-rank basis of spatial patterns from RCM data via Singular Value Decomposition (SVD, equivalent to Empirical Orthogonal Functions - EOF) (Monahan et al., 2009; Taylor et al., 2013) and reconstructs daily meteorological fields by regressing station observations onto this basis. By directly encoding the dominant spatial modes of variability derived from the RCM covariance structure, SPR provides a principled mechanism for generalizing observations to ungauged locations, naturally capturing anisotropy and non-stationarity without requiring complex spatial correlation structures. Evaluation against IDW, ordinary kriging, and KED showed that SPR performs competitively in standard conditions and provides clear advantages under very sparse network configurations (Houssou and Carreau, 2026).

A key limitation of the original SPR formulation, however, is that each time step is reconstructed independently. In practice, meteorological fields may exhibit strong temporal persistence: synoptic weather systems evolve on time scales of several days, so that temperature and precipitation fields on time step t are physically related to those of time step $t-1$ (Christensen, 2012; Pebesma, 2021; McDermott et al., 2018). This temporal coherence is a well-established property of atmospheric dynamics, reflected in the autocorrelation structures of observed meteorological time series (Chang et al., 1984; Trenberth, 1984; Zheng, 1996). Ignoring this structure means that each reconstruction is constrained only by the spatial information available on that specific time step, which can produce spurious temporal fluctuations — particularly under sparse observational conditions. For hydrological modeling, such artificial high-frequency noise can propagate into simulated catchment states, leading to unrealistic fluctuations in runoff, soil moisture, and streamflow predictions (Nijssen and Lettenmaier, 2004; Gyasi-Agyei, 2019).

A related limitation of the original SPR formulation is that it operates directly on raw meteorological fields, jointly capturing the climatological seasonal signal and the higher-frequency anomalies within a single reconstruction. A complementary line of work addresses this by decomposing meteorological fields into a climatological mean and a residual anomaly, interpolating the anomaly rather than the raw field. This approach, long used in synoptic climatology and climate downscaling (New et al., 2000; Stalenberg et al., 2018; Hamlet and Lettenmaier, 2005; Skofronick-Jackson et al., 2017), separates the large-scale seasonal signal — which can be estimated reliably from long RCM records — from the high-frequency residual variability, avoiding the need to jointly model the climatological signal and the associated fluctuations. Building on this idea, A-SPR extends SPR by reconstructing anomaly fields rather than raw observations, isolating the contribution of this climatological decomposition. However, A-SPR retains the SPR limitation of interpolating each time step independently.

The incorporation of temporal structure into spatial field reconstruction has a long history in geostatistics and spatial statistics. Spatio-temporal kriging extends the covariance model to the time dimension, enabling joint estimation across space and time (Christensen, 2012; Pravilovic et al., 2018; Takafuji et al., 2020; Kyriakidis and Journel, 1999). State-space and hierarchical Bayesian models provide a flexible framework for encoding temporal dynamics through latent processes (Wikle and Cressie, 1999; Cameletti and Biondi, 2019; McDermott et al., 2018; Hussain et al., 2010). Reduced-rank representations, in which the spatial field is projected onto a low-dimensional basis of spatial functions and the associated coefficients follow a dynamic process, offer an attractive compromise between expressiveness and computational tractability (Wikle and Cressie, 1999; McDermott et al., 2018). The Kalman filter and smoother provide optimal inference in linear Gaussian state-space models and have been applied to meteorological field reconstruction in various forms (Lorenc, 2003; Buehner et al., 2013; TSUYUKI and MIYOSHI, 2007; Wikle and Cressie, 1999). The T-SPR approach proposed in this work (Temporally-constrained Spatial Pattern Regression) adopts a similar reduced-rank representation, decomposing the reconstructed anomaly field via PCA and applying a VAR(1) temporal model to the retained component coefficients. Unlike Kalman filter/smoother-based approaches, which require specifying and propagating process and observation noise covariances through recursive filtering and smoothing, T-SPR estimates the component coefficients by minimizing a penalized least-squares objective combining data fidelity and the VAR(1) temporal penalty. Although the joint objective is not convex, each iteration admits a closed-form, OLS-like update, and convergence is achieved in a small number of iterations, offering a computationally efficient alternative to full state-space estimation.

To summarize, we introduce T-SPR, which combines the anomaly-based reconstruction of A-SPR with a first-order vector autoregressive (VAR(1)) regularization enforcing temporal persistence in the reconstructed field, consistent with the synoptic persistence of atmospheric circulation patterns. Calibration is framed as an alternating minimization, with each subproblem admitting a closed-form solution, ensuring computational efficiency. To isolate the contribution of each refinement, we evaluate SPR, A-SPR, and T-SPR on synthetic experiments across three network density levels (10%, 1%, 0.1%) and validate them on real station observations, using daily precipitation, minimum temperature, and maximum temperature over a region in southern Québec, Canada.

The remainder of this paper is organized as follows. Section 2 describes the study area and data. Section 3 presents the methodological framework, including the anomaly decomposition, the T-SPR model, the alternating minimization algorithm, and the hyperparameter selection procedure. Section 4 reports the experimental results. Section 5 concludes with a synthesis and discussion of future directions.

2 Study area and data

The study area, regional climate model data, and observational dataset used in this work follow a similar setup as in Houssou and Carreau (2026). Three standard inputs for hydrological modelling and climate impact studies are considered: daily precipitation ($\mathrm{mm\,day^{-1}}$), minimum temperature, and maximum temperature (°C).

2.1 Regional climate model data

High-resolution gridded meteorological fields were obtained from the ClimEx project (Leduc et al., 2019), based on simulations produced with the Canadian Regional Climate Model (CRCM5). The model covers North America on a regular grid of approximately 11 km horizontal resolution, resulting in a domain of 280×280 grid cells (Fig. 1). These simulations provide dynamically consistent representations of spatial climatic variability. Two distinct periods were extracted from member *kdj* of the 50-member ClimEx single-model initial-condition large ensemble (1950–2099): an auxiliary period (1980–2009) to estimate climatological means and derive spatial patterns, and an interpolation period (2000–2009) reserved for spatial interpolation and synthetic evaluation (Section 3). Although these periods overlap in time in the present study, the proposed methodology does not rely on temporal alignment between auxiliary and interpolation periods. However, in practice, the auxiliary period should represent climatic conditions similar to those of the interpolation period in order to ensure the relevance of the derived spatial patterns and climatology.

The analysis focuses on the southern Québec region illustrated in Fig. 1b. Compared to Houssou and Carreau (2026), which considered multiple regions, we restrict attention to a single domain in order to isolate the methodological contributions under a consistent climatic and geographical setting, thereby avoiding confounding effects due to spatial heterogeneity. This choice is further motivated by the availability of observational data, as the southern region provides a denser station network, enabling a meaningful validation on real observations, which would not be feasible in the more sparsely instrumented northern regions.

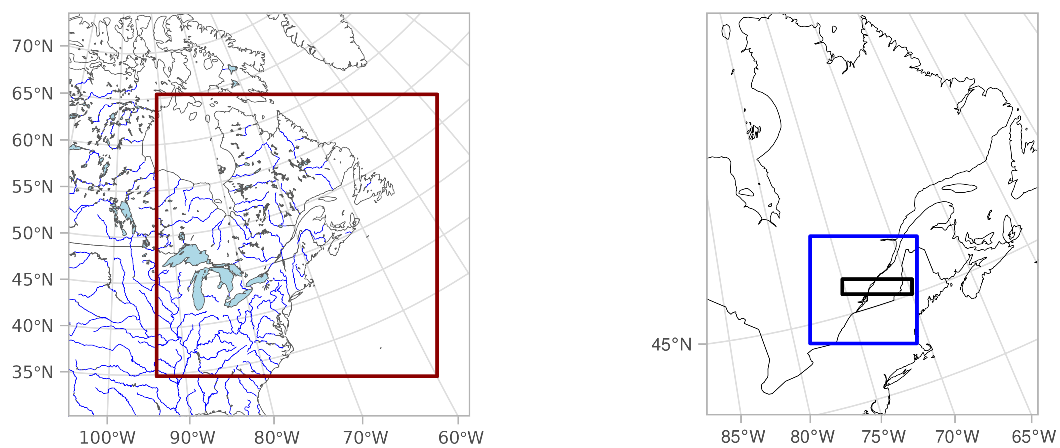

(a) North American domain of the ClimEx project (dark red rectangle).

(b) Southern Québec study region within the domain outlined in Fig. 1a (blue rectangle). The black rectangle indicates the spatial domain of the real-data experiments.

Figure 1: Spatial domains used in this study.

2.2 Observational data

Daily station observations were drawn from the national monitoring network operated by Environment and Climate Change Canada (ECCC; `https://climate.weather.gc.ca/climate_data/bulk_data_e.html`, accessed January 30, 2026). The dataset was restricted to southern Québec (latitudes 45.0–47.0°N and longitudes 75.0–70.0°W), located within the study region illustrated in Fig. 1b. To ensure a strict validation framework, only stations with complete records over a two-year period were retained, and no gap-filling procedures were applied. In order to maximize station availability, the two-year window was selected independently for each variable: October 2000 to October 2002 for precipitation, December 2000 to December 2002 for minimum temperature ($T_{\min}$), and December 2016 to December 2018 for maximum temperature ($T_{\max}$). This procedure results in 15 stations for precipitation and 44 stations for temperature forcings. Importantly, station observations are used only for performance assessment and, where applicable, hyperparameter selection, and do not influence the extraction of spatial patterns or climatological information. This guarantees a strict separation between auxiliary RCM-derived information and real-world observations.

3 Methods

This section presents Temporally-constrained Spatial Pattern Regression (T-SPR), an extension of the Spatial Pattern Regression (SPR) framework introduced in Houssou and Carreau (2026). The presentation follows the natural progression of the three methods evaluated in this work: Section 3.1 presents SPR, the classical framework operating directly on the meteorological forcing; Section 3.2 presents A-SPR, which replaces the forcing with a climatological anomaly; and Section 3.3 presents T-SPR, which additionally enforces temporal persistence in the reconstructed forcing. Each method is described in full, however A-SPR and T-SPR are presented by highlighting their key differences from the original SPR through direct reference to its equations. Finally, Section 3.4 defines the quantitative evaluation metrics, and Section 3.5 details the experimental setup and hyperparameter selection protocol.

3.1 SPR (classical framework)

3.1.1 Computation of spatial patterns

In this work, spatial patterns of meteorological forcing refer to recurrent structures that occur within a given spatial domain. We assume that the RCM simulations provide representative, physics-driven information on these spatial structures, which are thus constrained by the RCM grid resolution. Practically, these patterns are obtained as the spatial eigenvectors from the singular value decomposition (SVD) of the RCM simulation matrix over the auxiliary period. Consequently, they are dimensionless and time-invariant.

Let $\boldsymbol{Z}^{\text{aux}} \in \mathbb{R}^{T_{\text{aux}} \times D}$ be the RCM simulation matrix, where T_{aux} is the number of days in the auxiliary period — distinct from the interpolation period — and D is the number of grid cells in the spatial domain. Truncated SVD of $\boldsymbol{Z}^{\text{aux}}$ yields:

$$\boldsymbol{Z}^{\text{aux}} \approx \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^{\top}, \tag{1}$$

where $\boldsymbol{V} \in \mathbb{R}^{D \times K}$ contains the K leading spatial patterns (equivalently, Empirical Orthogonal Functions, EOF) in its columns, $\boldsymbol{\Sigma} \in \mathbb{R}^{K \times K}$ is the diagonal matrix of singular values, and $\boldsymbol{U} \in \mathbb{R}^{T_{\text{aux}} \times K}$ contains the associated temporal scores. The columns of $\boldsymbol{V}$ form an orthonormal basis capturing the dominant modes of spatial variability of the field. The number of retained patterns K is a key hyperparameter controlling the spatial complexity of the reconstruction (Section 3.5). Spatial patterns are extracted solely from the RCM auxiliary period and are independent of the station observations used for interpolation.

3.1.2 Regression

Let $\boldsymbol{Z}_t \in \mathbb{R}^d$ denote the random vector of meteorological forcing at d gauged stations on day t, with $d \ll D$, where D is the number of grid cells in the spatial domain. Given the spatial patterns $\boldsymbol{V}_d \in \mathbb{R}^{d \times K}$ restricted to the d station locations, $\boldsymbol{Z}_t$ is modeled as a linear combination of these K reduced modes:

$$\boldsymbol{Z}_t = \boldsymbol{V}_d\,\boldsymbol{\beta}_t^{\text{SPR}} + \boldsymbol{\varepsilon}_t, \qquad \boldsymbol{\varepsilon}_t \sim \mathcal{N}\big(\boldsymbol{0},\, \sigma^2 \boldsymbol{I}_d\big)\,, \tag{2}$$

The coefficients $\boldsymbol{\beta}_t^{\mathrm{SPR}}$ are estimated independently for each day t by ordinary least squares (OLS), with no coupling across time. To reconstruct forcing values across the entire gridded domain (i.e., performing spatial interpolation at ungauged locations as well), Eq. (2) is applied across the full spatial patterns $\boldsymbol{V} \in \mathbb{R}^{D\times K}$ using the OLS estimators $\hat{\boldsymbol{\beta}}_t^{\mathrm{SPR}}$:

$$\hat{\boldsymbol{Z}}_t = \boldsymbol{V}\hat{\boldsymbol{\beta}}_t^{\mathrm{SPR}}. \tag{3}$$

3.2 A-SPR

A-SPR retains the 2-step structure of SPR described in Section 3.1 but operates on the climatological anomaly rather than on the raw forcing.

3.2.1 Computation of anomalies

Anomalies are computed for both the RCM simulations and the gauged observations (real or virtual). In both cases, the baseline monthly climatology is constructed from the RCM auxiliary matrix $\boldsymbol{Z}^{\mathrm{aux}} \in \mathbb{R}^{T_{\mathrm{aux}}\times D}$. For each grid cell s, the climatological mean c_s is defined as the average of $Z_{t,s}^{\mathrm{aux}}$ over all days t within the same calendar month across the auxiliary period. This monthly climatology captures the mean seasonal cycle of the forcing. For station locations, c_s is simply extracted from the nearest grid cell; thus, the baseline climatology relies exclusively on the RCM simulations without using station observations.

Let $Z_{t,s} \in \mathbb{R}$ denote a generic forcing variable—derived either from station observations or RCM simulations—at location s and time t; the corresponding anomaly formulation (additive or multiplicative) depends on the physical nature of this variable. For temperature variables (T_{min} and T_{max}), the anomaly $Y_{t,s}$ is defined with the additive formulation:

$$Y_{t,s} = Z_{t,s} - c_s. \tag{4}$$

For precipitation, a multiplicative formulation is adopted to account for its non-negative and highly skewed distribution:

$$Y_{t,s} = \frac{Z_{t,s}}{c_s} - 1. \tag{5}$$

In both formulations, the anomaly $Y_{t,s}$ is centered at zero by construction. Given an interpolated anomaly $\widehat{Y}_{t,s}$, the final interpolated value $\widehat{Z}_{t,s}$ is obtained by inverting Eq. (4)–(5):

$$\widehat{Z}_{t,s} = \begin{cases} \widehat{Y}_{t,s} + c_s, & \text{for temperature variables } \{T_{\mathrm{min}}, T_{\mathrm{max}}\}, \\ \left(\widehat{Y}_{t,s} + 1\right) c_s, & \text{for precipitation.} \end{cases} \tag{6}$$

3.2.2 Computation of spatial patterns

Applying the anomaly decomposition of Section 3.2.1 to the RCM auxiliary data $\boldsymbol{Z}^{\mathrm{aux}}$ yields the anomaly matrix $\boldsymbol{Y}^{\mathrm{aux}}$. The spatial patterns used by A-SPR are obtained via the same truncated SVD as in SPR (Section 3.1.1, Eq. (1)), with the raw RCM matrix $\boldsymbol{Z}^{\mathrm{aux}}$ replaced by the RCM anomaly matrix $\boldsymbol{Y}^{\mathrm{aux}}$. The resulting spatial pattern $\widetilde{\boldsymbol{V}} \in \mathbb{R}^{D\times K}$ retains the same interpretation as in SPR, now capturing the dominant modes of anomaly variability over the spatial domain rather than of the raw forcing.

3.2.3 Regression

Let $\boldsymbol{Y}_t \in \mathbb{R}^d$ denote the random vector of climatological anomalies across d gauged stations on day t, obtained by applying the anomaly decomposition (Section 3.2.1) to the corresponding forcing vector

$\boldsymbol{Z}_t$. Here, $d \ll D$, where D represents the total number of grid cells in the spatial domain. The regression step follows Eq. (2), with the forcing $\boldsymbol{Z}_t$ replaced by the anomaly $\boldsymbol{Y}_t$:

$$\boldsymbol{Y}_t = \widetilde{\boldsymbol{V}}_d\, \boldsymbol{\beta}_t^{\text{A-SPR}} + \boldsymbol{\varepsilon}_t, \qquad \boldsymbol{\varepsilon}_t \sim \mathcal{N}\big(\mathbf{0},\, \sigma^2 \boldsymbol{I}_d\big)\,, \tag{7}$$

with $\widetilde{\boldsymbol{V}}_d \in \mathbb{R}^{d\times K}$ the anomaly patterns restricted to the d station locations. As in SPR, coefficients are estimated independently for each day using OLS. To reconstruct forcing values across the entire domain, Eq. (3) is applied with $\boldsymbol{V}$ replaced by $\widetilde{\boldsymbol{V}}$ and $\hat{\boldsymbol{\beta}}_t^{\text{SPR}}$ by $\hat{\boldsymbol{\beta}}_t^{\text{A}-\text{SPR}}$, followed by Eq. (6) to convert the anomalies back to physical forcing values.

3.3 T-SPR

T-SPR is the primary methodological contribution of this work. It builds directly on A-SPR, adopting its anomaly decomposition (Section 3.2.1) and spatial patterns (Section 3.2.2). Furthermore, T-SPR uses the same underlying regression model (Eq. (7)) and follows the same procedure to reconstruct the forcing field (Section 3.2.3). The two methods differ only in how the regression coefficients are estimated: rather than the independent per-day OLS of A-SPR, T-SPR estimates the coefficient sequence jointly across time under an explicit temporal regularization, described below. This regularization introduces temporal persistence in the estimated regression coefficients, which in turn ensures smoothness between consecutive reconstructed forcing fields.

3.3.1 Temporal regularization of the regression coefficients

Let $\boldsymbol{\beta}_t^{\text{T-SPR}} \in \mathbb{R}^K$ denote the regression coefficients for T-SPR, which adopts the same regression model as A-SPR (Eq. (7)). Rather than estimating these coefficients independently for each day via OLS, T-SPR enforces temporal persistence in the coefficients $\boldsymbol{\beta}_t^{\text{T-SPR}}$ through a first-order vector autoregressive (VAR(1)) prior:

$$\boldsymbol{\beta}_t^{\text{T-SPR}} = \boldsymbol{A}\, \boldsymbol{\beta}_{t-1}^{\text{T-SPR}} + \boldsymbol{\eta}_t, \qquad \boldsymbol{\eta}_t \sim \mathcal{N}(\mathbf{0},\, \boldsymbol{Q})\,, \tag{8}$$

where $\boldsymbol{A} \in \mathbb{R}^{K\times K}$ is the transition matrix that dictates the daily persistence (or temporal memory) of the coefficients, and $\boldsymbol{Q} \in \mathbb{R}^{K\times K}$ is the process noise covariance matrix. In the optimization formulation below (Eq. (10)), $\boldsymbol{Q}$ is not explicitly estimated; its effect is absorbed into a single scalar regularization parameter λ.

The reconstruction problem is formulated as the joint minimization over the coefficient matrix $\boldsymbol{B} = (\boldsymbol{\beta}_1, \ldots, \boldsymbol{\beta}_T)^T \in \mathbb{R}^{T\times K}$ (where the superscript T-SPR is omitted for brevity) and the transition matrix $\boldsymbol{A}$:

$$(\boldsymbol{B}^*, \boldsymbol{A}^*) = \underset{\boldsymbol{B},\, \boldsymbol{A}}{\arg\min}\;\, J(\boldsymbol{B}, \boldsymbol{A}), \tag{9}$$

The objective function is designed to combine spatial reconstruction fidelity and temporal consistency:

$$J(\boldsymbol{B}, \boldsymbol{A}) = \underbrace{\sum_{t=1}^{T} \left\| \boldsymbol{y}_t - \widetilde{\boldsymbol{V}}_d\, \boldsymbol{\beta}_t \right\|_2^2}_{\text{spatial fidelity}} + \lambda \underbrace{\sum_{t=2}^{T} \|\boldsymbol{\beta}_t - \boldsymbol{A}\, \boldsymbol{\beta}_{t-1}\|_2^2}_{\text{temporal consistency}}, \tag{10}$$

where $\boldsymbol{y}_t \in \mathbb{R}^d$ denotes the observed anomaly vector at the d gauged stations on day t, corresponding to a realization of the random vector $\boldsymbol{Y}_t$; $\widetilde{\boldsymbol{V}}_d \in \mathbb{R}^{d\times K}$ is the anomaly spatial pattern matrix restricted to the station locations, and $\lambda > 0$ controls the strength of the temporal regularization. The first term corresponds to a Gaussian likelihood of the reduced-rank observation model (Christensen, 2012), while the second is the negative log-likelihood of the VAR(1) transition (Pebesma, 2021; Christensen, 2012; Cameletti and Biondi, 2019; McDermott et al., 2018). When $\boldsymbol{A} = \mathbf{0}$, the temporal regularization reduces to a classical ridge penalty (Hastie et al., 2009) applied independently to each coefficient vector, making T-SPR a natural generalization of ridge regression that allows correlation across time through a VAR(1) prior. Note that when $\lambda = 0$, the temporal regularization vanishes and the problem reduces to the independent day-by-day anomaly regression of A-SPR (Section 3.2.3); SPR is recovered only when the anomaly decomposition is additionally removed.

The cost function Eq. (10) is quadratic in $\boldsymbol{B}$ for fixed $\boldsymbol{A}$, and quadratic in $\boldsymbol{A}$ for fixed $\boldsymbol{B}$, but is not jointly convex in $(\boldsymbol{B}, \boldsymbol{A})$ due to the bilinear interaction term $\boldsymbol{A}\boldsymbol{\beta}_{t-1}$ in the temporal penalty. Optimization is carried out via an alternating minimization strategy, iteratively updating $\boldsymbol{B}$ and $\boldsymbol{A}$ until convergence, as described below.

3.3.2 Initialization

The algorithm is initialized with $\boldsymbol{A}^{(0)} = \mathbf{0}$, corresponding to the absence of temporal constraint. Under this initialization, the temporal penalty reduces to a ridge-type regularization on $\boldsymbol{\beta}_t$, ensuring a stable and well-posed starting point for the iteration. The temporal constraint is then progressively introduced as $\boldsymbol{A}$ is learned.

3.3.3 Step I: estimation of B given B

For fixed $\boldsymbol{A}^{(k)}$, the minimization of Eq. (10) with respect to $\boldsymbol{B}$ is performed sequentially over time, from $t = 1$ to $t = T$. For $t = 1$, no past information is available; the temporal penalty in Eq. (10) reduces to $\lambda\|\boldsymbol{\beta}_1\|_2^2$, and the update is a standard ridge regression that does not depend on the iteration index k:

$$\boldsymbol{\beta}_1^{(k+1)} = \left(\widetilde{\boldsymbol{V}}_d^\top \widetilde{\boldsymbol{V}}_d + \lambda \boldsymbol{I}_K\right)^{-1} \widetilde{\boldsymbol{V}}_d^\top \boldsymbol{y}_1. \tag{11}$$

For $t \geq 2$, each coefficient vector minimizes the contribution of time step t to $J(\boldsymbol{B}, \boldsymbol{A})$:

$$\boldsymbol{\beta}_t^{(k+1)} = \arg\min_{\boldsymbol{\beta}} \left\{ \left\| \boldsymbol{y}_t - \widetilde{\boldsymbol{V}}_d \, \boldsymbol{\beta} \right\|_2^2 + \lambda \left\| \boldsymbol{\beta} - \boldsymbol{A}^{(k)} \boldsymbol{\beta}_{t-1}^{(k+1)} \right\|_2^2 \right\}, \tag{12}$$

which admits the closed-form solution:

$$\boldsymbol{\beta}_t^{(k+1)} = \left(\widetilde{\boldsymbol{V}}_d^\top \widetilde{\boldsymbol{V}}_d + \lambda \boldsymbol{I}_K\right)^{-1} \left(\widetilde{\boldsymbol{V}}_d^\top \boldsymbol{y}_t + \lambda \, \boldsymbol{A}^{(k)} \boldsymbol{\beta}_{t-1}^{(k+1)}\right). \tag{13}$$

Because $\boldsymbol{\beta}_t^{(k+1)}$ depends on $\boldsymbol{\beta}_{t-1}^{(k+1)}$ already updated at the same iteration, this update constitutes a forward pass over the interpolation period.

3.3.4 Step II: estimation of B given B

For fixed $\boldsymbol{B}^{(k+1)}$, the spatial fidelity term in Eq. (10) does not depend on $\boldsymbol{A}$, so the transition matrix is obtained by minimizing the temporal consistency term alone:

$$\boldsymbol{A}^{(k+1)} = \arg\min_{\boldsymbol{A}} \sum_{t=2}^{T} \|\boldsymbol{\beta}_t - \boldsymbol{A}\, \boldsymbol{\beta}_{t-1}\|_2^2 \,. \tag{14}$$

Defining the lagged coefficient matrices

$$\boldsymbol{B}_{1:T-1} = \begin{bmatrix} \boldsymbol{\beta}_1^\top \\ \vdots \\ \boldsymbol{\beta}_{T-1}^\top \end{bmatrix} \in \mathbb{R}^{(T-1)\times K}, \qquad \boldsymbol{B}_{2:T} = \begin{bmatrix} \boldsymbol{\beta}_2^\top \\ \vdots \\ \boldsymbol{\beta}_T^\top \end{bmatrix} \in \mathbb{R}^{(T-1)\times K}, \tag{15}$$

the multivariate OLS solution is:

$$\boldsymbol{A}^{(k+1)} = \boldsymbol{B}_{2:T}^\top \, \boldsymbol{B}_{1:T-1} \left(\boldsymbol{B}_{1:T-1}^\top \, \boldsymbol{B}_{1:T-1}\right)^+, \tag{16}$$

where $(\cdot)^+$ denotes the Moore–Penrose pseudo-inverse, which reduces to the standard matrix inverse when the Gram matrix $\boldsymbol{B}_{1:T-1}^\top \, \boldsymbol{B}_{1:T-1} \in \mathbb{R}^{K\times K}$ is full rank.

To ensure stationarity of the VAR(1) process, the spectral radius of $\boldsymbol{A}^{(k+1)}$, denoted $\rho(\boldsymbol{A}^{(k+1)})$, is constrained to be strictly less than one. If $\rho(\boldsymbol{A}^{(k+1)}) \geq 1$, the matrix is rescaled as:

$$\boldsymbol{A}^{(k+1)} \leftarrow \frac{\boldsymbol{A}^{(k+1)}}{\rho(\boldsymbol{A}^{(k+1)}) + \varepsilon}, \qquad \varepsilon = 10^{-4}. \tag{17}$$

3.3.5 Convergence and field reconstruction

Steps I (Section 3.3.3) and II (Section 3.3.4) are repeated until the relative decrease in the objective function Eq. (10) falls below a predefined threshold:

$$\frac{\left|J^{(k+1)} - J^{(k)}\right|}{J^{(k)}} < \delta, \tag{18}$$

with $\delta = 10^{-4}$, or until a maximum number of iterations is reached. In practice, convergence is observed within a small number of iterations. At convergence, the anomaly field is reconstructed over the full domain using Eq. (3) with $\boldsymbol{V}$ replaced by $\widetilde{\boldsymbol{V}}$ and $\hat{\boldsymbol{\beta}}_t^{\mathrm{SPR}}$ by $\hat{\boldsymbol{\beta}}_t^{\mathrm{T-SPR}}$. The final interpolated field $\widehat{\boldsymbol{Z}}_t$ is then recovered by applying Eq. (6).

3.4 Evaluation metrics

To quantitatively evaluate spatial field reconstruction performance, model predictions are compared against ground-truth reference values at unobserved locations (i.e., grid cells containing no virtual or real stations). Five complementary metrics jointly evaluate spatial fidelity, temporal consistency, and day-to-day dynamics:

RMSE – pointwise interpolation accuracy. Daily RMSE is computed across all evaluation locations and subsequently averaged over the test period.

Structural Similarity Index Measure (SSIM) – for each evaluation location, the reconstructed and ground-truth time series over the full test period are compared through their agreement in mean level, variability, and temporal correlation, combined following the luminance–contrast–structure decomposition; the reported SSIM is the average of these location-wise scores over all evaluation locations.

Power Spectral Density (PSD) – for each evaluation location, the power spectral density of the reconstructed and ground-truth time series is estimated using Welch's method (periodogram averaging over overlapping segments). The median spectrum and the corresponding interquartile range (IQR) across evaluation locations are then reported at each frequency, summarizing dominant variability scales, high-frequency behavior, and spatial variability in temporal dynamics.

Auto-Correlation Function (ACF) – for each evaluation location, the lag-dependent autocorrelation of the reconstructed and ground-truth series is computed; the mean ACF across evaluation locations is reported at each lag, together with its IQR, and close correspondence with the ground truth indicates that the reconstruction successfully captures the underlying temporal structure.

Inter-day variability – the absolute day-to-day change $|\Delta_t|$ is first averaged across all evaluation locations for each day, and the resulting daily series is compared, over the test period, between the reconstructed and ground-truth fields; a distribution close to that of the ground truth indicates faithful day-to-day dynamics, while under-dispersion signals excessive smoothing and over-dispersion signals spurious fluctuations.

Together, these five metrics provide a comprehensive characterization of spatial fidelity, structural realism, temporal consistency, and day-to-day dynamics of the reconstructed fields.

3.5 Experimental framework and hyperparameter selection

To compare SPR, A-SPR, and T-SPR under controlled and real-world conditions, we establish two distinct experimental settings alongside a standardized hyperparameter selection protocol.

3.5.1 Synthetic experimental framework

The synthetic experiments rely on an idealized setting where a virtual station network is defined as a sub-sample of the RCM grid cells. By emulating observational networks of varying sizes, the virtual stations yield pseudo-observations that are used in the same way as real station data. The objective is to evaluate how well each variant can reconstruct a known spatially complete field from sparse observations, rather than to assess realism with respect to actual meteorological observations.

Three controlled experiments are designed by setting the density of the virtual station network to 10%, 1%, and 0.1% of the total grid cells (3025). The 0.1% configuration—retaining only three virtual stations over the full domain—corresponds to a stress-test scenario representative of severely under-instrumented regions such as northern Québec and other remote environments. Interpolation is assessed at all remaining grid cells (i.e., excluding the virtual stations) by comparing the reconstructed fields directly to the reference RCM fields, which serve as ground truth by design.

For each density level, a single virtual-station layout is defined. This design choice provides a computationally feasible framework for evaluating multiple variables, density scenarios, and hyperparameter grids without requiring repeated spatial resampling. Because the underlying RCM field provides a spatially complete reference, evaluating models across the thousands of remaining ungauged grid cells over the full temporal domain yields a massive sample size ($d_{\text{eval}} \times T$). The breadth of this evaluation ensures that observed performance differences between methods are statistically robust under each station density.

3.5.2 Real-observation experimental framework

To evaluate model performance under realistic observational conditions, an experiment is conducted using the daily station observation dataset (Section 2), which spans a two-year period. As detailed in Section 2, the available physical network comprises 44 stations for temperature ($T_{\max}$ and $T_{\min}$) and 15 stations for precipitation (P_r).

Unlike the synthetic setting, where station placement can be controlled across the RCM grid, the real station network is sparse and fixed by physical geography. To mirror the most challenging stress-test density from the synthetic framework (0.1% density), each experiment uses only three stations for model training ($d = 3$).

Because this evaluation set is relatively small, performance metrics are sensitive to the spatial arrangement of the chosen training stations. To mitigate spatial sampling noise and quantify topology sensitivity, a repeated random sampling protocol is applied. At each repetition, three training stations are randomly drawn without replacement from the available network, and the remaining stations are used for evaluation. This selection process is repeated 100 times using independent random draws for each variable, and performance metrics are aggregated across all repetitions.

3.5.3 Temporal splitting strategy

To structure independent validation and test sets without sacrificing spatial coverage, the available data are split across two distinct time periods. For a given time period, the regression steps for all method variants are calibrated exclusively at the gauged stations using data from that specific period. Meanwhile, the spatial patterns—whether derived directly from forcings or anomalies—are pre-computed over an independent auxiliary period (Section 3). Model performance is then assessed over the remaining ungauged locations during the same period. For the synthetic experiments, these evaluation sets range from several hundred grid cells at 10% density up to 3022 grid cells in the sparsest 0.1% experiment. For the real-observation experiment (Section 2), where 3 stations are selected for training, the evaluation set comprises the remaining 41 ungauged stations for temperature ($T_{\max}$ and $T_{\min}$) and 12 for precipitation (P_r).

To maintain a strict firewall between hyperparameter selection and final performance assessment, this protocol is executed across two distinct temporal phases obtained by partitioning the dataset along the time dimension: a validation period (2000–2004 for synthetic experiments; Year 1 for real observations) and a test period (2005–2009 for synthetic experiments; Year 2 for real observations), as detailed in Section 2. The ungauged evaluation locations are used first during the validation phase to select optimal hyperparameters, and subsequently during the held-out test phase to evaluate final reconstruction skill. This unified design allows all available spatial locations to be fully leveraged while guaranteeing independent validation and testing.

3.5.4 Hyperparameter selection and constraints

All three method variants share the structural hyperparameter K, which dictates the number of retained spatial patterns or EOF (Section 3.1.1). Candidate values for K are specified as proportions of the maximum available spatial patterns $K_{\max}$ and rounded to the nearest integer. The temporally regularized variant (T-SPR) introduces an additional hyperparameter: the temporal regularization parameter $\lambda \geq 0$ (Eq. (10)). The unregularized variants (SPR and A-SPR) correspond to the special case where $\lambda = 0$, requiring no optimization along this dimension.

Hyperparameter configurations are selected by minimizing the spatial reconstruction RMSE (Section 3.4) over the ungauged locations during the validation phase (Section 3.5.3), aligning model selection directly with the primary goal of field reconstruction.

For sparse network configurations, the admissible range of K is bounded by the number of available stations d. Specifically, at the 0.1% synthetic density level and in the real-observation experiment—both of which feature $d = 3$ training stations—the regression system requires $K < d$ to remain well posed. Setting $K = 3$ results in an exactly determined system that perfectly fits the station observations but lacks spatial generalization, whereas $K = 1$ retains only the primary spatial mode and fails to capture local gradients. To maintain numerical stability while maximizing spatial expressiveness, K is fixed to 2 ($K = 2$) across all $d = 3$ settings. For T-SPR, the regularization parameter λ is subsequently selected via the temporal validation procedure.

For the real-observation experiment, the optimal regularization strengths selected on the validation year are $\lambda_{P_r} = 0.04$, $\lambda_{T_{\min}} = 0.001$, and $\lambda_{T_{\max}} = 0.003$ for T-SPR. A summary of the selected hyperparameter configurations across the synthetic experimental settings is provided in Table 1.

Table 1: Optimal proportion (and absolute count in parentheses) of retained spatial patterns and VAR(1) regularization parameter λ selected on the validation sub-period (2000–2004) for the synthetic experiments. The proportion of spatial patterns is defined relative to the total number of available patterns $K_{\max}$ for a given station density extracted from the RCM anomaly matrix. At 0.1% density, K is fixed at proportion 0.67 ($K = 2$) for all methods.

Density	Variable	Methods			
		SPR	A-SPR	T-SPR	
		K	K	K	λ
10%	Pr	0.25 (K=76)	0.25 (K=76)	0.30 (K=91)	0.001
	$T_{\min}$	0.25 (K=76)	0.25 (K=76)	0.25 (K=76)	0.001
	$T_{\max}$	0.25 (K=76)	0.25 (K=76)	0.25 (K=76)	0.001
1%	Pr	0.25 (K=8)	0.30 (K=9)	0.30 (K=9)	0.0001
	$T_{\min}$	0.30 (K=9)	0.30 (K=9)	0.30 (K=9)	0.0003
	$T_{\max}$	0.30 (K=9)	0.30 (K=9)	0.30 (K=9)	0.0002
0.1%	Pr	0.67 (K=2)	0.67 (K=2)	0.67 (K=2)	0.001
	$T_{\min}$	0.67 (K=2)	0.67 (K=2)	0.67 (K=2)	0.0001
	$T_{\max}$	0.67 (K=2)	0.67 (K=2)	0.67 (K=2)	0.0002

4 Results

We evaluate three methods: SPR, which operates directly on raw observations following Houssou and Carreau (2026) (Section 3.1); A-SPR, which applies the same framework to anomaly fields (Section 3.2); and T-SPR, which extends A-SPR with a VAR(1) temporal prior on the pattern coefficients (Section 3.3) to promote temporal persistence in the reconstructed fields. Reconstructed fields are back-transformed, where applicable, to the original scale before evaluation (Section 3.2.1, Eq. (6)), ensuring comparability across methods. Results are presented for precipitation (P_r), minimum temperature ($T_{\min}$), and maximum temperature ($T_{\max}$).

The results are organized in two parts. Section 4.1 presents the synthetic experiments, in which virtual station networks of controlled density (10%, 1%, and 0.1%) are sampled from the RCM grid, and reconstruction accuracy is evaluated against the known RCM ground truth. For each criterion considered, the three variants are discussed jointly, so that the respective contributions of the anomaly decomposition (SPR $\rightarrow$ A-SPR) and of the temporal regularization (A-SPR $\rightarrow$ T-SPR) can be assessed together. Section 4.2 then evaluates the three variants on real station observations, under a station configuration analogous to the sparsest synthetic experiment.

4.1 Synthetic experiments

4.1.1 Overall accuracy

At 10% and 1% station densities, all three variants perform comparably well across variables (Table 2), demonstrating that observational density alone is sufficient to capture spatial structure under moderate to dense coverage. Because performance differences at these densities are marginal, subsequent detailed analysis focuses on the 0.1% stress-test regime (3 virtual stations across the domain), where methodological differences become highly pronounced.

Table 2: Synthetic experiments: Average daily RMSE and SSIM for SPR, A-SPR, and T-SPR across density levels and variables, computed over the ungauged locations during the test period (2005–2009). Best value per row is shown in bold. Higher SSIM and lower RMSE indicate better reconstruction quality.

Density	Variable	RMSE			SSIM		
		SPR	A-SPR	T-SPR	SPR	A-SPR	T-SPR
10%	P_r	**1.7806**	1.7880	1.7862	**0.9436**	0.9409	0.9418
	$T_{\min}$	0.8438	0.8401	**0.8371**	0.9960	**0.9960**	0.9960
	$T_{\max}$	0.5260	0.5217	**0.5193**	0.9985	0.9985	**0.9986**
1%	P_r	**3.0476**	3.1203	3.0987	**0.8127**	0.8056	0.8061
	$T_{\min}$	1.5372	1.5215	**1.4993**	**0.9867**	0.9850	0.9848
	$T_{\max}$	1.1672	1.0982	**1.0904**	0.9931	0.9941	**0.9942**
0.1%	P_r	6.5263	5.8417	**4.8230**	**0.4780**	0.4698	0.4281
	$T_{\min}$	6.5917	3.0880	**2.9710**	0.7661	0.9163	**0.9291**
	$T_{\max}$	3.8260	2.6412	**2.4375**	0.9201	0.9585	**0.9669**

For $T_{\min}$ and $T_{\max}$, reconstruction accuracy follows a clear progressive hierarchy (SPR $\rightarrow$ A-SPR $\rightarrow$ T-SPR). Transitioning to A-SPR yields a massive leap in skill, cutting $T_{\min}$ RMSE by over 50% and $T_{\max}$ RMSE by 31% compared to baseline SPR, while substantially boosting spatial structural fidelity (SSIM). T-SPR provides a consistent secondary refinement, delivering the top overall performance across both temperature variables by preventing unnatural day-to-day fluctuations. Overall, these results show that while baseline SPR suffers severe degradation under extreme sparsity, anomaly-based centering stabilizes the spatial regression, and temporal persistence ensures physically realistic continuity across complex terrain.

For P_r at 0.1% density, the transition across variants reveals a distinct trade-off between magnitude accuracy and spatial structure. A-SPR improves magnitude estimation over baseline SPR, while T-SPR delivers a further substantial error reduction, lowering overall precipitation RMSE by 26% relative to baseline SPR. However, this gain in magnitude precision comes at the cost of spatial structural fidelity: baseline SPR retains a slightly higher SSIM than A-SPR and T-SPR. Because precipitation fields under severe sparsity lack smooth background gradients, enforcing temporal persistence dampens erratic daily magnitude spikes at ungauged locations, but also penalizes abrupt, highly localized convective events that occur independently from day to day.

4.1.2 Spatial distribution of grid-cell RMSE

Figure 2 shows the per-grid-cell daily RMSE maps at 0.1% density for all three variants. The maps reveal three consistent patterns. First, SPR produces large localized errors in areas remote from the three virtual station locations, with the spatial distribution of errors showing both a high overall level and substantial heterogeneity across the domain. Second, A-SPR substantially reduces these errors – particularly for temperature variables; for precipitation, the region of high RMSE away from the virtual stations remains, but its spatial extent is substantially reduced. Third, T-SPR further reduces errors across the domain for all variables, with the largest gains visible in the contraction of high-error regions; this effect is especially marked for precipitation, where T-SPR effectively limits the extreme local errors that persist under SPR and A-SPR.

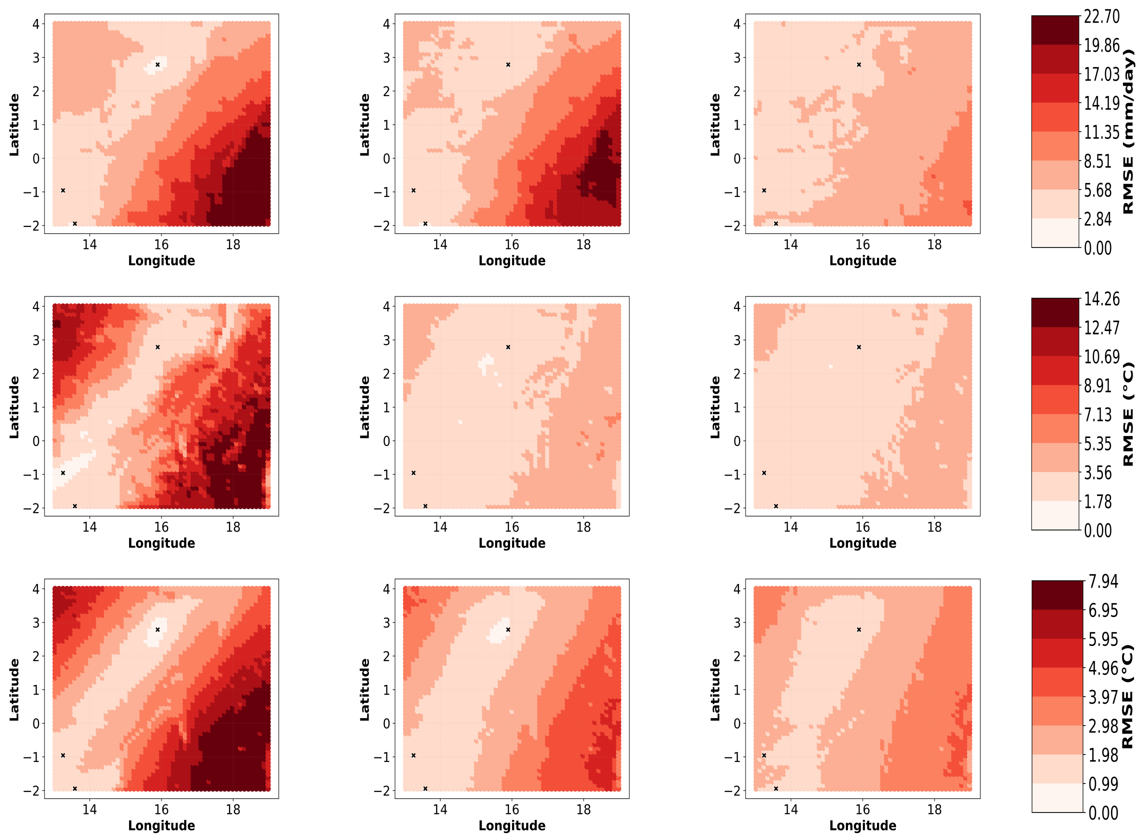

(a) SPR (b) A-SPR (c) T-SPR

Figure 2: Synthetic experiments: spatial distribution of per-grid-cell daily RMSE at 0.1% station density for SPR (left column), A-SPR (center column), and T-SPR (right column). Rows correspond to P_r (top), $T_{\min}$ (middle), and $T_{\max}$ (bottom). Each panel shows the RMSE averaged over ungauged locations during the test period (2005–2009), evaluated on the original scale. The three synthetic station locations used for interpolation are indicated by black crosses ($\times$).

4.1.3 Temporal consistency: ACF and PSD

Figure 3 and Figure 4 show the ACF and PSD of reconstructed forcing fields at 0.1% station density for all three variants, compared against the ground truth (RCM reference field). At 10% and 1% station density (not shown), all variants closely reproduce the ground-truth ACF and PSD structure, indicating that temporal dynamics are already well constrained when observational coverage is sufficient.

At 0.1% density, larger differences emerge between methods. SPR generally exhibits the widest IQR and the largest departures from the ground-truth ACF, indicating high variability under extreme sparsity. A-SPR reduces this variability, while T-SPR further increases temporal consistency by producing narrower IQR and smoother autocorrelation structures. T-SPR slightly overestimates persistence, consistent with a mild temporal smoothing effect.

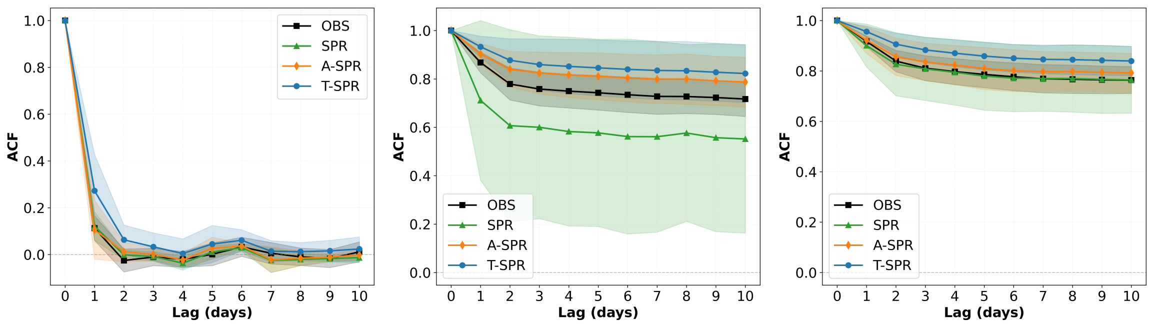

(a) Precipitation (b) Minimum temperature (c) Maximum temperature

Figure 3: Synthetic experiment: autocorrelation function (ACF) of reconstructed fields at 0.1% station density for SPR, A-SPR, T-SPR, and the ground truth. Each curve represents the median ACF across test grid cells; the shaded bands show the IQR across test cells for each method and the ground truth. A closer alignment of a method's ACF with the ground truth indicates better temporal consistency. Results at 10% and 1% density are not shown as all methods closely reproduce the ground-truth ACF at those levels.

The PSD analysis (Figure 4) leads to similar conclusions. Under extreme sparsity, SPR tends to exhibit greater spatial dispersion and excess high-frequency variability, while A-SPR and T-SPR produce spectra that are more consistent across grid cells. However, T-SPR systematically underestimates the PSD across most frequencies, most notably for precipitation, with its spectrum lying below both the other methods and the ground-truth reference. This behavior indicates a loss of variance across a wide range of scales and suggests that T-SPR introduces excessive smoothing, particularly at higher frequencies.

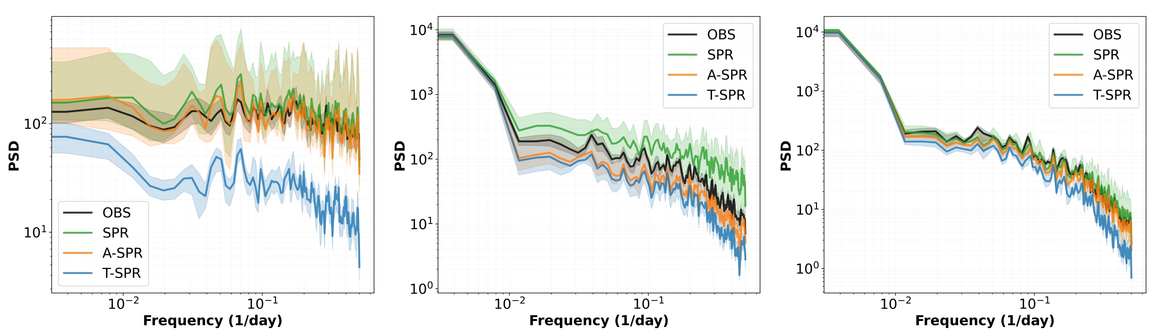

(a) Precipitation (b) Minimum temperature (c) Maximum temperature

Figure 4: Synthetic experiment: power spectral density (PSD) of reconstructed fields at 0.1% station density for SPR, A-SPR, T-SPR, and the ground truth. Each curve represents the median PSD across test grid cells; the shaded bands show the IQR across test cells. A method whose PSD curve lies within or close to the ground-truth IQR reproduces the observed temporal dynamics more faithfully. Layout as in Figure 3.

Taken together, the ACF and PSD diagnostics indicate that the main contribution of temporal regularization provided by T-SPR is not a systematic improvement in the median temporal structure, but rather a reduction in spatial variability and spurious spectral fluctuations. The resulting reconstructions are more stable across grid cells and more temporally coherent, particularly under extremely sparse observational conditions.

4.1.4 Day-to-day variability

Figure 5 shows the distribution of the absolute lag-1 differences $|\Delta_t| = |\hat{Z}_t - \hat{Z}_{t-1}|$ across ungauged grid cells over the test period at 0.1% station density, compared against the ground truth.

The effects of temporal regularization are most pronounced for precipitation. SPR and A-SPR produce day-to-day changes whose distribution is substantially wider and shifted toward larger values than the ground truth, indicating excessive volatility in the reconstructed forcing fields. T-SPR, by contrast, over-corrects in the opposite direction: its distribution is markedly narrower and centered below the ground truth, producing smoother day-to-day transitions than the observed dynamics. This behavior reflects the smoothing effect of the VAR(1) prior, which penalizes abrupt temporal changes and therefore reduces high-frequency variability. As a consequence, the temporal variability of the reconstructed precipitation fields is underestimated, consistent with the decrease in SSIM noted in Section 4.1.1.

For temperature variables, the picture is more balanced. SPR overestimates inter-day variability, particularly for $T_{\min}$, where its distribution is markedly broader than the ground truth. A-SPR substantially corrects this bias, producing distributions close to the ground truth for both $T_{\min}$ and $T_{\max}$. T-SPR further contracts the distributions, remaining slightly below the ground-truth level for both temperature variables, consistent with the smoothing bias noted for precipitation but with a substantially smaller magnitude.

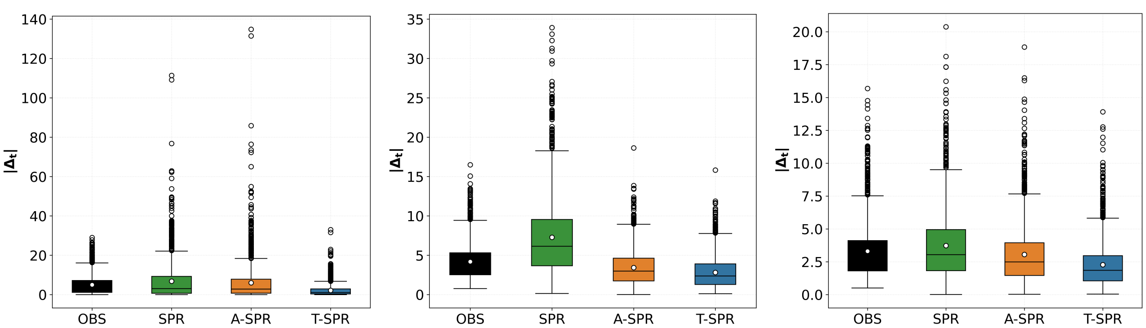

(a) Precipitation **(b) Minimum temperature** **(c) Maximum temperature**

Figure 5: Synthetic experiment: distribution of absolute inter-day differences $|\Delta_t| = |\hat{Z}_t - \hat{Z}_{t-1}|$ at 0.1% station density for the ground truth (OBS), SPR, A-SPR, and T-SPR. Variables: P_r (left), $T_{\min}$ (center), $T_{\max}$ (right). Each box shows the IQR of this daily, spatially-averaged series across the full test period; outliers are shown as circles and the average as a white marker. A distribution closer to that of the ground truth indicates that the method reproduces observed day-to-day temporal dynamics more faithfully.

4.2 Experiment on real observations

Evaluation in the real-observation experiment encompasses magnitude accuracy (RMSE), spatial agreement (SSIM), inter-day variability ($|\Delta_t|$), and temporal persistence (ACF_1) across the 100 station configurations. Figure 6 displays the resulting distributions over the 100 repetitions as boxplots for each metric, meteorological forcing, and method.

Across all variables, SPR RMSE exhibits the largest variability, with wide distributions and frequent high-error outliers, indicating strong sensitivity to the choice of training stations. A-SPR substantially reduces both median error and dispersion, while T-SPR achieves median RMSE comparable to that of A-SPR with the most compact distribution. SSIM values are generally similar across variants for precipitation and $T_{\min}$, but for $T_{\max}$, T-SPR consistently attains the highest values with reduced dispersion, indicating more stable preservation of spatial structure.

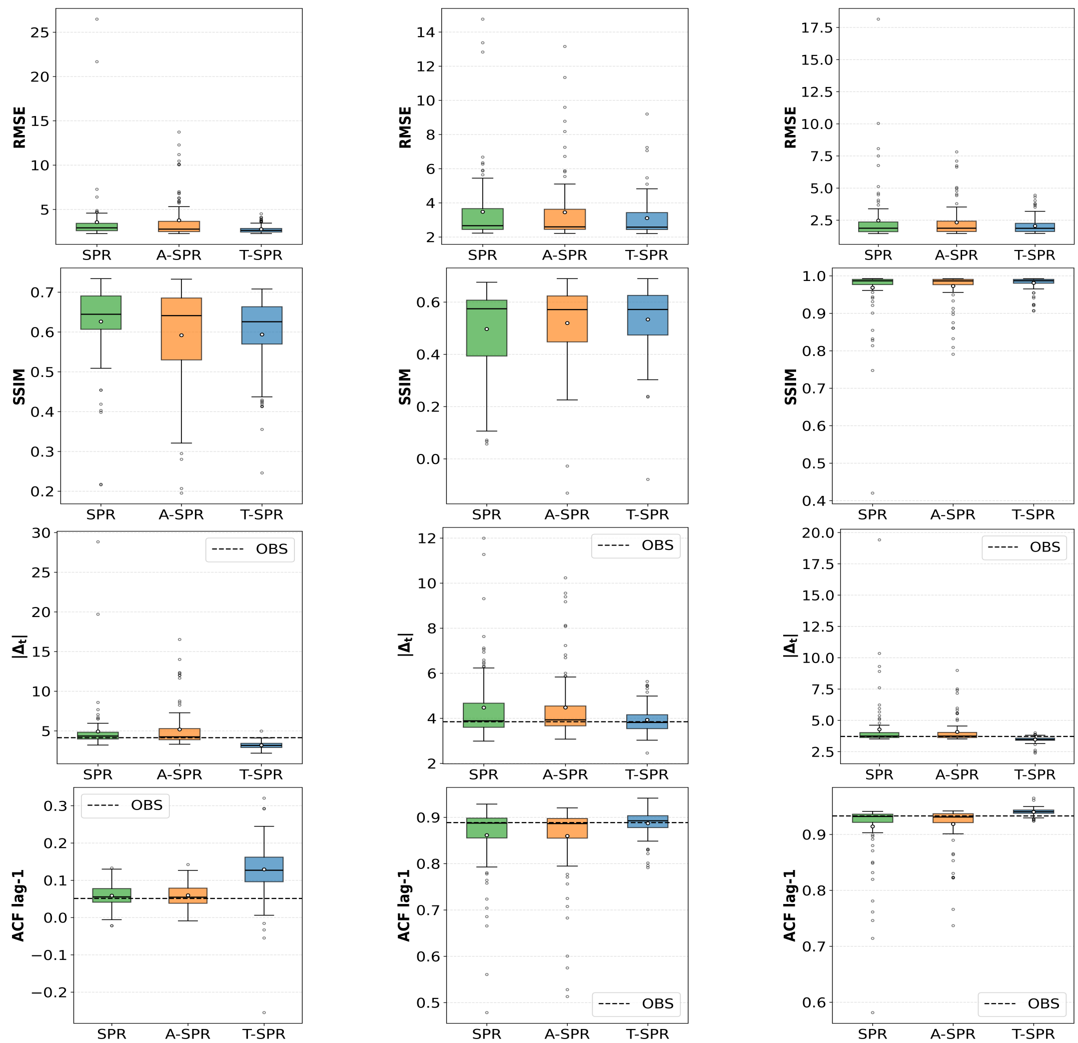

(a) Precipitation **(b) Minimum temperature** **(c) Maximum temperature**

Figure 6: Real-observation experiment: distribution, over 100 random selections of three training stations, of RMSE (first row), SSIM (second row), day-to-day variability $|\Delta_t|$ (third row), and lag-1 autocorrelation (fourth row), for SPR, A-SPR, and T-SPR, for precipitation (left), minimum temperature (center), and maximum temperature (right). Each box shows the interquartile range across the 100 repetitions; outliers are shown as circles. More compact distributions and metric values closer to the ground-truth reference indicate more robust and accurate reconstructions.

Differences between A-SPR and T-SPR are most evident in temporal behavior. For inter-day variability ($|\Delta_t|$), T-SPR produces distributions that are more tightly concentrated around the ground-truth (OBS) reference for all variables, reflecting improved consistency of day-to-day dynamics. For lag-1 autocorrelation, all variants perform similarly for precipitation, but for both temperature variables the distributions obtained with A-SPR and T-SPR are more compact and closely centered around the ground-truth (OBS) values, with fewer outliers than SPR, indicating a more stable representation of temporal variablity.

Overall, while gains in median accuracy remain modest, T-SPR provides the most robust performance by substantially limiting worst-case errors. These results demonstrate that temporal regularization reduces sensitivity to highly sparse and noisy observations, yielding more stable spatial reconstructions and more temporally consistent dynamics.

5 Conclusion and discussion

This study introduced two extensions of the Spatial Pattern Regression (SPR) framework of Houssou and Carreau (2026): the anomaly-based variant A-SPR and the temporally constrained variant T-SPR. Both variants were evaluated on synthetic experiments across three observational densities, as well as on real-data experiments, using daily precipitation, minimum temperature, and maximum temperature. The results show that the two methodological modifications address complementary limitations of the original SPR framework. The anomaly decomposition improves the spatial reconstruction problem by separating the large-scale climatological signal from the day-to-day variability. This allows the spatial patterns to focus on the residual variability rather than simultaneously representing seasonal and weather-scale signals. The benefits are particularly evident for temperature variables, whose variability is dominated by a strong and spatially coherent climatological cycle. Under sparse observational conditions, A-SPR substantially improves reconstruction accuracy and robustness relative to the original SPR formulation. The second extension concerns temporal persistence. By introducing a VAR(1) constraint on the pattern coefficients, T-SPR explicitly accounts for the persistence of meteorological fields from one day to the next. Across both synthetic and real-data experiments, temporal regularization produces more stable reconstructions, reduces sensitivity to sparse observations, and limits the occurrence of extreme local errors. While improvements in average accuracy are generally moderate, gains in robustness become increasingly important as observational density decreases. These results suggest that temporal regularization provides a valuable additional constraint when the spatial information available on a given day is insufficient to reliably estimate the field.

By addressing spatial climatology and temporal persistence simultaneously, T-SPR yields reconstructions that are both more physically coherent and less sensitive to observation sparsity than the original SPR formulation. Crucially, these gains are achieved with minimal computational overhead, as the framework leverages fast spatial SVD decompositions and an alternating minimization scheme with closed-form update steps during regression. Integrating climatological structure and temporal dependence within SPR thus provides an accurate, robust, and computationally tractable solution for sparse field reconstruction.

The main limitation identified in this study concerns precipitation. Compared with temperature, precipitation exhibits weaker temporal persistence and stronger intermittency at the daily scale. Consequently, the VAR(1) assumption smooths out higher frequencies, which we hypothesize may correspond to an attenuation of high-frequency variability. A further limitation of this study is that the real-data evaluation is restricted to a single region in southern Québec. Since the climatological mean and the principal components underlying the anomaly decomposition are derived from RCM simulations, the spatial resolution at which T-SPR can reconstruct forcing fields is inherently tied to that of the underlying RCM. Generalizing the approach to other regions therefore requires access to RCM simulations at sufficient spatial resolution, and the method's performance may vary depending on the quality and resolution of RCM products available for a given domain. Further evaluation across regions with different climatic regimes and RCM data availability would help establish the broader applicability of T-SPR. Lastly, while this study evaluates forcing reconstruction accuracy directly, propagating T-SPR outputs through a hydrological model to quantify downstream impacts on streamflow simulation is left for future work.

Several directions for future research emerge from this work. First, building on the limitations observed for precipitation, adaptive temporal regularization strategies could be developed to account for seasonal or weather-regime-dependent variations in temporal persistence. Second, extending the framework to jointly reconstruct multiple meteorological variables may provide additional benefits by exploiting cross-variable dependencies. Finally, replacing the linear EOF representation with more flexible nonlinear spatial representations may further improve the reconstruction of complex spatial structures.

Code and data availability

All analyses were performed using the R programming language version 4.6.1 and the Python programming language version 3.10.12. The analysis scripts developed for this study are not publicly archived but can be provided by the corresponding author on request. The primary data source consists of daily climate simulations from the ClimEx project (https://climex-data.srv.lrz.de/Public/). We used one ensemble member (kdj) from the CanESM2-driven simulations. Real station observations were obtained from the national monitoring network operated by Environment and Climate Change Canada (ECCC). Data were downloaded automatically through the official bulk data API (https://climate.weather.gc.ca/climate_data/bulk_data_e.html).

References

Behnke, R., Vavrus, S., Allstadt, A., Albright, T., Thogmartin, W. E., and Radeloff, V. C. (2016). Evaluation of downscaled, gridded climate data for the conterminous united states. Ecological Applications, 26(5):1338–1351.

Bokke, A. (2017). Comparative evaluation of spatial interpolation methods for estimation of missing meteorological variables over ethiopia. Journal of Water Resource and Protection, 09:945–959.

Buehner, M., Morneau, J., and Charette, C. (2013). Four-dimensional ensemble-variational data assimilation for global deterministic weather prediction. Nonlinear Processes in Geophysics, 20(5):669–682.

Cameletti, M. and Biondi, F. (2019). Hierarchical modeling of space-time dendroclimatic fields: Comparing a frequentist and a bayesian approach. Arctic, Antarctic, and Alpine Research, 51(1):115–127.

Chang, T. J., Kavvas, M. L., and Delleur, J. W. (1984). Daily precipitation modeling by discrete autoregressive moving average processes. Water Resources Research, 20(5):565–580.

Christensen, O. F. (2012). Statistics for spatio-temporal data by cressie, n. and wikle, c. k. Biometrics, 68(4):1328–1329.

Faghih, M. and Brissette, F. (2023). Temporal and spatial amplification of extreme rainfall and extreme floods in a warmer climate. Journal of Hydrometeorology, 24(8):1331–1347.

Gasset, N., Fortin, V., Dimitrijevic, M., Carrera, M., Bilodeau, B., Muncaster, R., Gaborit, E., Roy, G., Pentcheva, N., Bulat, M., Wang, X., Pavlovic, R., Lespinas, F., Khedhaouiria, D., and Mai, J. (2021). A 10 km north american precipitation and land-surface reanalysis based on the gem atmospheric model. Hydrology and Earth System Sciences, 25(9):4917–4945.

Gyasi-Agyei, Y. (2019). Propagation of uncertainties in interpolated rainfields to runoff errors. Hydrological Sciences Journal, 64(5):587–606.

Hamlet, A. F. and Lettenmaier, D. P. (2005). Production of temporally consistent gridded precipitation and temperature fields for the continental united states. Journal of Hydrometeorology, 6(3):330–336.

Hartkamp, A., de Beurs, K., Stein, A., and White, J. (1999). Interpolation techniques for climate variables. Geographic Information Systems Series 99-01. International Maize and Wheat Improvement Center (CIMMYT), Mexico 1999.

Hastie, T., Tibshirani, R., and Friedman, J. (2009). The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Second Edition (Springer Series in Statistics).

Hengl, T., Heuvelink, G., and Stein, A. (2003). Comparison of kriging with external drift and regression-kriging. Technical Note.

Houssou, V. and Carreau, J. (2026). Spatial pattern regression for meteorological fields interpolation. EGUsphere, 2026:1–23.

Hussain, I., Spöck, G., Pilz, J., and Yu, H.-L. (2010). Spatio-temporal interpolation of precipitation during monsoon periods in pakistan. Advances in Water Resources, 33(8):880–886.

Hwang, Y., Clark, M., Rajagopalan, B., and Leavesley, G. (2012). Spatial interpolation schemes of daily precipitation for hydrologic modeling. Stochastic Environmental Research and Risk Assessment, 26:295–320.

Kidd, C., Huffman, G., Maggioni, V., Chambon, P., and Oki, R. (2021). The global satellite precipitation constellation: Current status and future requirements. Bulletin of the American Meteorological Society, 102(10):E1844–E1861.

Kyriakidis, P. C. and Journel, A. G. (1999). Geostatistical space–time models: A review. Mathematical Geology, 31(6):651–684.

Leduc, M., Mailhot, A., Frigon, A., et al. (2019). The ClimEx project: A 50-member ensemble of climate change projections at 12-km resolution over europe and northeastern north america with the canadian regional climate model (CRCM5). J. of App. Meteo. & Clim., 58(4):663–693.

Li, J. and Heap, A. D. (2008). A review of spatial interpolation methods for environmental scientists. Geoscience Australia.

Li, J. and Heap, A. D. (2014). Spatial interpolation methods applied in the environmental sciences: A review. Environmental Modelling & Software, 53:173–189.

Livneh, B., Bohn, T. J., Pierce, D. W., Munoz-Arriola, F., Nijssen, B., Vose, R., Cayan, D. R., and Brekke, L. (2015). A spatially comprehensive, hydrometeorological data set for mexico, the u.s., and southern canada 1950–2013. Scientific Data, 2(1):150042.

Looper, J. P. and Vieux, B. E. (2012). An assessment of distributed flash flood forecasting accuracy using radar and rain gauge input for a physics-based distributed hydrologic model. Journal of Hydrology, 412-413:114–132.

Lorenc, A. C. (2003). The potential of the ensemble kalman filter for nwp—a comparison with 4d-var. Quarterly Journal of the Royal Meteorological Society, 129(595):3183–3203.

Lucas-Picher, P., Arsenault, R., Poulin, A., Ricard, S., Lachance-Cloutier, S., and Turcotte, R. (2020). Application of a high-resolution distributed hydrological model on a u.s.-canada transboundary basin: Simulation of the multiyear mean annualhydrograph and 2011 flood of therichelieu river basin. Journal of Advances in Modeling Earth Systems, 12(4):e2019MS001709.

Ly, S., Charles, C., and Degré, A. (2013). Different methods for spatial interpolation of rainfall data for operational hydrology and hydrological modeling at watershed scale: A review. Biotechnology, Agronomy, Society and Environment, 17(2):392–406.

Maurer, E. P., Wood, A. W., Adam, J. C., Lettenmaier, D. P., and Nijssen, B. (2002). A long-term hydrologically based dataset of land surface fluxes and states for the conterminous united states. Journal of Climate, 15(22):3237–3251.

McDermott, P. L., Wikle, C. K., and Millspaugh, J. (2018). A hierarchical spatiotemporal analog forecasting model for count data. Ecology and Evolution, 8(1):790–800. PMC5756884.

Monahan, A. H., Fyfe, J. C., Ambaum, M. H. P., Stephenson, D. B., and North, G. R. (2009). Empirical orthogonal functions: The medium is the message. Journal of Climate, 22(24):6501–6514.

New, M., Hulme, M., and Jones, P. (2000). Representing twentieth-century space–time climate variability. part ii: Development of 1901–96 monthly grids of terrestrial surface climate. Journal of Climate, 13(13):2217–2238.

Nijssen, B. and Lettenmaier, D. P. (2004). Effect of precipitation sampling error on simulated hydrological fluxes and states: Anticipating the global precipitation measurement satellites. Journal of Geophysical Research: Atmospheres, 109(D2).

Pebesma, E. (2021). Spatio-temporal statistics with R by Christopher K. Wikle, Andrew Zammit-Mangion and Noel Cressie. Statistical Papers, 62(2):1093–1094.

Pravilovic, S., Appice, A., and Malerba, D. (2018). Leveraging correlation across space and time to interpolate geophysical data via cokriging. International Journal of Geographical Information Science, 32(1):191–212.

Segond, M.-L., Wheater, H. S., and Onof, C. (2007). The significance of spatial rainfall representation for flood runoff estimation: A numerical evaluation based on the lee catchment, uk. Journal of Hydrology, 347(1):116–131.

Skofronick-Jackson, G., Petersen, W. A., Berg, W., Kidd, C., Stocker, E. F., Kirschbaum, D. B., Kakar, R., Braun, S. A., Huffman, G. J., Iguchi, T., Kirstetter, P. E., Kummerow, C., Meneghini, R., Oki, R., Olson, W. S., Takayabu, Y. N., Furukawa, K., and Wilheit, T. (2017). The global precipitation measurement (gpm) mission for science and society. Bulletin of the American Meteorological Society, 98(8):1679–1695.

Spangler, K., Weinberger, K., and Wellenius, G. (2019). Suitability of gridded climate datasets for use in environmental epidemiology. Journal of Exposure Science & Environmental Epidemiology, 29.

Stalenberg, E., Hutchinson, M. F., and Foley, W. J. (2018). Using historical normals to improve modern monthly climate normal surfaces for madagascar. International Journal of Climatology, 38(15):5746–5765.

Takafuji, E. H. d. M., da Rocha, M. M., and Manzione, R. L. (2020). Spatiotemporal forecast with local temporal drift applied to weather patterns in patagonia. SN Applied Sciences, 2(6):1001.

Tarek, M., Brissette, F. P., and Arsenault, R. (2020). Large-scale analysis of global gridded precipitation and temperature datasets for climate change impact studies. Journal of Hydrometeorology, 21(11):2623–2640.

Taylor, M. H., Losch, M., Wenzel, M., and Schröter, J. (2013). On the sensitivity of field reconstruction and prediction using empirical orthogonal functions derived from gappy data. Journal of Climate, 26(22):9194–9205.

Trenberth, K. E. (1984). Some effects of finite sample size and persistence on meteorological statistics. part i: Autocorrelations. Monthly Weather Review, 112(12):2359–2368.

TSUYUKI, T. and MIYOSHI, T. (2007). Recent progress of data assimilation methods in meteorology. Journal of the Meteorological Society of Japan. Ser. II, 85B:331–361.

Varentsov, M., Esau, I., and Wolf, T. (2020). High-resolution temperature mapping by geostatistical kriging with external drift from large-eddy simulations. Monthly Weather Review, 148(3):1029–1048.

Warren, F. et al. (2022). Canada in a Changing Climate: Regional Perspectives Report. Government of Canada, Ottawa, ON. Accessed: 2025-04-07 at https://changingclimate.ca/regional-perspectives/.

Werner, A., Schnorbus, M., Shrestha, R., Cannon, A., Zwiers, F., Dayon, G., and Anslow, F. (2019). A long-term, temporally consistent, gridded daily meteorological dataset for northwestern north america. Scientific Data, 6(1):1–16.

Wikle, C. and Cressie, N. (1999). A dimension-reduced approach to space-time kalman filtering. Biometrika, 86(4):815–829.

Zheng, X. (1996). Unbiased estimation of autocorrelations of daily meteorological variables. Journal of Climate, 9(9):2197–2203.