

Warped Dynamic Linear Models for Time Series of Counts*

Brian King[†] and Daniel R. Kowal[‡]

Abstract. Dynamic Linear Models (DLMs) are commonly employed for time series analysis due to their versatile structure, simple recursive updating, ability to handle missing data, and probabilistic forecasting. However, the options for count time series are limited: Gaussian DLMs require continuous data, while Poisson-based alternatives often lack sufficient modeling flexibility. We introduce a novel semiparametric methodology for count time series by *warping* a Gaussian DLM. The warping function has two components: a (nonparametric) transformation operator that provides distributional flexibility and a rounding operator that ensures the correct support for the discrete data-generating process. We develop conjugate inference for the warped DLM, which enables analytic and recursive updates for the state space filtering and smoothing distributions. We leverage these results to produce customized and efficient algorithms for inference and forecasting, including Monte Carlo simulation for offline analysis and an optimal particle filter for online inference. This framework unifies and extends a variety of discrete time series models and is valid for natural counts, rounded values, and multivariate observations. Simulation studies illustrate the excellent forecasting capabilities of the warped DLM. The proposed approach is applied to a multivariate time series of daily overdose counts and demonstrates both modeling and computational successes.

Keywords: discrete data, state-space model, particle filter, selection normal.

1 Introduction

Count time series data inherit all the complexities of continuous time series data: the time-ordered observations may be multivariate, seasonal, dependent on exogenous variables, and exhibit a wide variety of autocorrelation structures. At the same time, count data often present uniquely challenging distributional features, including zero-inflation, over-/underdispersion, boundedness or censoring, and heaping. Additionally, discrete data require distinct strategies for probabilistic forecasting, uncertainty quantification, and evaluation. As modern datasets commonly feature higher resolutions and lengthier time series, computational tools for both online and offline inference and forecasting are in demand. Fundamentally, the goals in count time series modeling are similar to those in the continuous setting: forecasting, trend filtering/smoothing, seasonal decomposition, and characterization of inter- and intra-series dependence, among others.

*This material is based upon work supported by the National Science Foundation: a Graduate Research Fellowship under Grant No. 1842494 (King) and SES-2214726 (Kowal).

[†]Department of Statistics, Rice University, bking@rice.edu

[‡]Department of Statistics, Rice University, daniel.kowal@rice.edu

In this paper, we develop methods, theory, and computing tools for a broad class of multivariate state space models that address each of these challenges and objectives. The core model is defined by *warping* a Gaussian Dynamic Linear Model (DLM; West and Harrison, 1997), which we refer to as a *warped DLM* (warpDLM):

$$\mathbf{y}_t = h \circ g^{-1}(\mathbf{z}_t) \quad (\text{warping}) \quad (1)$$

$$\{\mathbf{z}_t\}_{t=1}^T \sim \text{DLM} \quad (\text{see (3) and (4)}) \quad (2)$$

where $\mathbf{y}_t \in \mathbb{N}^n$ is the observed count data and $\mathbf{z}_t \in \mathbb{R}^n$ is continuous latent data. The warping has two components: a rounding operator $h : \mathcal{T} \rightarrow \mathbb{N}^n$, which ensures the correct support for the discrete data-generating process, and a monotone transformation function $g : \mathcal{T} \rightarrow \mathbb{R}^n$, which endows (possibly nonparametric) flexibility in the marginal distributions. The latent DLM enables straightforward embedding of familiar dynamic modeling structures such as local levels, seasonality, or covariates, along with a natural way of dealing with missing data. Note that this article focuses on modeling counts, but the framework is easily adaptable for general integer-valued or rounded data.

To isolate the rounding and transformation operations, the warping operation may be decomposed into $\mathbf{y}_t = h(\mathbf{y}_t^*)$ and $\mathbf{z}_t = g(\mathbf{y}_t^*)$ for a latent continuous variable \mathbf{y}_t^* . We emphasize that the warpDLM is fundamentally distinct from the “transformed DLM” strategy of (i) fitting a Gaussian DLM to (possibly transformed, e.g., logarithmically) count data and (ii) rounding the resulting (continuous) forecasts. The “transformed DLM” fails to account for the discreteness of the data in the model-fitting process and introduces a critical mismatch between the *fitted model* and the *data-generating process*. If the terminal rounding step is omitted, then the data-generating process is not discrete; yet unless the rounding step is included within the model-fitting process—as in the warpDLM—then the fitted model fails to account for crucial features in the data. In particular, the rounding operation is a nontrivial component of the model: within the warpDLM, h provides the capability to model challenging discrete distributional features, such as zero-inflation, boundedness, or censoring (see Section 2). Transformation-only models are known to be ineffective in many settings, such as low counts (O’Hara and Kotze, 2010), yet the warpDLM excels precisely in this case (see Section 5.2).

A key contribution of this article is to develop conjugate inference for the warpDLM. In particular, we show that the warpDLM likelihood is conjugate to the *selection normal distribution* (e.g., Arellano-Valle et al., 2006). Based on this result, we derive analytic and recursive updates for the warpDLM filtering and smoothing distributions. Crucially, we provide direct Monte Carlo simulators for these distributions—as well as the count-valued forecasting distribution—and construct an optimal particle filter for online inference. These models, derivations, and algorithms remain valid in the multivariate setting, and provide significant advancements over existing latent models for count time series (e.g., Jia et al., 2021). To the best of our knowledge, these results are unique for multivariate count time series models.

The warpDLM offers a unified framework for several discrete data models (dynamic as well as static) and incorporates strategies which have proven successful in other related methods. In the non-dynamic realm, Siegfried and Hothorn (2020) demonstrated

the benefits of learned transformations for discrete data linear regression, while Kowal and Wu (2021) adopted a related transformation and rounding strategy to model heaped count data. Both took a frequentist approach to estimation. Among Bayesian methods, Canale and Dunson (2011) and Canale and Dunson (2013) similarly applied rounded Gaussian and Dirichlet processes, respectively, without the transformation considerations. Kowal and Canale (2020) showcased the advantage of both the rounding *and* transformation components within a static regression setting, and Kowal and Wu (2022) extended this framework to incorporate Bayesian nonparametric estimation of the transformation. Looking to dynamic models, when g is viewed as a copula, the warpDLM resembles the count time series model proposed by Jia et al. (2021). Unlike Jia et al. (2021), we do not focus on stationary latent Gaussian processes, but rather incorporate DLMs to enable nonstationary modeling, dynamic covariates, and Bayesian inference and forecasting within a familiar setting. The warpDLM also generalizes recent work done in the binary data space, in particular the dynamic probit model of Fasano et al. (2021). In doing so, we construct novel theory utilizing a broader class of distributions and requiring distinct computations of relevant posterior quantities.

The warpDLM framework falls in the category of generalized state space models, which are traditionally separated into two classes, stemming from Cox et al. (1981): parameter-driven and observation-driven. In observation-driven models, the state process is treated as an explicit function of past data values, and methods often utilize likelihood-based inference born out of GLM theory (see Fokianos, 2015 for a review). Bayesian methods, including the warpDLM framework, are more commonly parameter-driven, which means the latent state parameter is treated as stochastic; model learning typically proceeds via recursive updating, e.g., using the Kalman filter (Kalman, 1960). The DLM is the most well-known model in this category, but relies on Gaussian assumptions unmet by count data. Dynamic Generalized Linear Models (DGLMs) were developed to adapt the state space framework for non-Gaussian data within the exponential family (West et al., 1985). For most count data, Poisson is the only available observational density that belongs to the exponential family. Binomial DGLMs are applicable for bounded data, yet perform quite poorly in the case of zero-inflation or heaping on the boundary (see Section 5.2). The negative binomial distribution with fixed dispersion parameter is also exponential family, but there is often little guidance for determining the dispersion parameter. Berry and West (2019) recently extended the DGLM family by mixing Bernoulli and Poisson DGLMs to better model zero-inflated and overdispersed count data.

The Poisson DGLM and extensions provide “data coherent” (Freeland and McCabe, 2004) inferences and forecasts, in the sense that predictions are appropriate to the type of data (count time series), but the closed-form Kalman filtering results are generally unavailable. In most DGLM specifications, the evolution equation is assumed Gaussian (see (4)), which necessitates linearized approximations or MCMC algorithms for smoothing, filtering, and forecasting. West et al. (1985) prioritized closed-form and conjugate recursions, yet require approximate and moment-based (rather than distributional) updates for state parameters. Fahrmeir (1992) designed an extended Kalman Filter to estimate posterior modes in a multivariate setting. Durbin and Koopman (2000) used importance sampling based on a linear approximation. Frühwirth-Schnatter and Wagner

(2006) constructed an approximate Gibbs sampler using data augmentation and mixture sampling. Another strategy is to replace the Gaussian evolution equation. Gamerman et al. (2013) used a Poisson model with a multiplicative state update to preserve analytic and recursive inference. However, this updating structure has limited dynamic flexibility, for example to include seasonality or covariates. Aktekin et al. (2018) proposed a multivariate extension of this multiplicative model, but the analytic updating results were not preserved and the multivariate structure only accommodated positive correlations among the series. Another exception is Bradley et al. (2018), who proposed log-Gamma processes that are conditionally conjugate to the Poisson distribution. This model still requires Gibbs sampling for all state smoothing, filtering, and forecasting distributions.

The common limitations among existing state space models for count data are (i) a lack of exact, coherent, and recursive updates for filtering, smoothing, and forecasting distributions and (ii) restricted options for count-valued distributions. The warpDLM framework directly addresses and overcomes both limitations.

The paper is organized as follows. Section 2 introduces DLMs, the proposed model, and examples for the rounding and transformation. In Section 3, we derive the smoothing, filtering, and forecasting distributions for the warpDLM. We discuss computing strategies for online and offline inference in Section 4. Finally, we present forecasting results on simulated data as well as a real-data application in Section 5 before concluding. Supplementary Material (King and Kowal, 2023) includes proofs of presented theorems, additional simulations, and further details on the application data and model specification. Code to reproduce all findings is also available on [GitHub](#). warpDLM functionality is included in the R package `countSTAR` (King and Kowal, 2023), which is available on [CRAN](#) and documented in an [online vignette](#).

2 Dynamic Linear Models and Time Series of Counts

The broad success of Bayesian time series analysis has largely been driven by Dynamic Linear Models (DLMs), also known as linear state space models. The DLM framework subsumes ARIMA models and provides decomposition of time series, dynamic regression analysis, and multivariate modeling capabilities. Additionally, the sequential updating structure provides a simple way to deal with missing observations (cf. Durbin and Koopman, 2012, Section 2.7). DLMs are widely popular not only because of their versatility, but also because the filtering, smoothing, and forecasting distributions are available in closed-form via the recursive Kalman filter.

A DLM is defined by two equations: (i) the observation equation, which specifies how the observations are related to the latent state vector and (ii) the state evolution equation, which describes how the states are updated in a Markovian fashion. We present the Gaussian DLM for an *observable* continuous n -dimensional time series $\{\mathbf{z}_t\}_{t=1}^T$, but note that these continuous variables are latent (i.e., unobservable) within the warpDLM:

$$\mathbf{z}_t = \mathbf{F}_t \boldsymbol{\theta}_t + \mathbf{v}_t, \quad \mathbf{v}_t \sim N_n(\mathbf{0}, \mathbf{V}_t) \quad (3)$$

$$\boldsymbol{\theta}_t = \mathbf{G}_t \boldsymbol{\theta}_{t-1} + \mathbf{w}_t, \quad \mathbf{w}_t \sim N_p(\mathbf{0}, \mathbf{W}_t) \quad (4)$$

for $t = 1, \dots, T$, where $\{\mathbf{v}_t, \mathbf{w}_t\}_{t=1}^T$ are mutually independent and $\boldsymbol{\theta}_0 \sim N_p(\mathbf{a}_0, \mathbf{R}_0)$. Depending on the DLM specification, the p -dimensional state vector $\boldsymbol{\theta}_t$ could describe a local level, regression coefficients, or seasonal components, among other features. Common choices of the $n \times p$ observation matrix \mathbf{F}_t include the identity ($n = p$), a matrix of indicators that selects certain elements of $\boldsymbol{\theta}_t$ (e.g., for structural time series models), and covariate values (e.g., for dynamic regression analysis). The $p \times p$ state evolution matrix \mathbf{G}_t is frequently the identity but can be more complex, for example to capture seasonality. The observation and evolution covariance matrices are \mathbf{V}_t ($n \times n$) and \mathbf{W}_t ($p \times p$), respectively. Taken together, the quadruple $\{\mathbf{F}_t, \mathbf{G}_t, \mathbf{V}_t, \mathbf{W}_t\}_{t=1}^T$ defines the DLM. Often, these matrices will be time-invariant.

The Gaussian DLM (3)–(4) is data-incoherent for discrete time series: the forecasting distribution does not match the support of the data. DGLMs offer one resolution by replacing the Gaussian observation equation with an exponential family distribution. However, the primary option for count data is the Poisson distribution, which is often inadequate and requires additional modeling layers for common discrete data features such as zero-inflation, over/underdispersion, boundedness or censoring, and heaping. These additional layers introduce significant modeling and computational complexity. By comparisons, the warpDLM is capable of modeling *each* of these distributional features under default specifications via the warping (rounding and transformation) operation, yet maintains the useful and familiar state space formulation through the latent DLM. By leveraging Gaussian state space models, the warpDLM builds on the long history of theoretical and computational tools (West and Harrison, 1997; Petris et al., 2009; Prado et al., 2021), and operates within a familiar setting for practitioners.

The warpDLM framework links count data \mathbf{y}_t with a (latent) Gaussian DLM for \mathbf{z}_t in (3)–(4) via the warping operation (1). The rounding operation $h : \mathcal{T} \rightarrow \mathbb{N}^n$ serves as the connection mechanism between the real-valued latent space and the non-negative integers, setting $h(\mathbf{y}^*) = \mathbf{j}$ for any $\mathbf{y}^* \in \mathcal{A}_j$. These pre-image sets \mathcal{A}_j form a disjoint partition of the space \mathcal{T} . In the univariate setting, $\mathcal{A}_j = [a_j, a_{j+1})$ is simply an interval. The warpDLM likelihood can thus be written as

$$\mathbb{P}(\mathbf{y}_t = \mathbf{j} | \boldsymbol{\theta}) = \mathbb{P}\{\mathbf{z}_t \in g(\mathcal{A}_j) | \boldsymbol{\theta}\}, \quad t = 1, \dots, T \quad (5)$$

for $j \in \mathbb{N}^n$, where $\mathbf{z}_t \in g(\mathcal{A}_j)$ is defined elementwise when \mathbf{y}_t is multivariate.

Both components of the warping operation serve important purposes that lead directly to desirable model properties. The rounding function ensures that the warpDLM has the correct support for the (possibly bounded or censored) count data. For simplicity, suppose $n = 1$; generalizations occur by applying these specifications elementwise. By default, we take the rounding function to be the floor function, so $\mathcal{A}_j = [j, j + 1)$. In addition, we include the zero modification $g(\mathcal{A}_0) = (-\infty, 0)$ so that $y_t = 0$ whenever $z_t < 0$. This specification maps much of the latent space to zero, with persistence of zeros determined by the DLM (3)–(4), for example, $\mathbb{P}(y_t = 0 | y_{t-1} = 0) = \mathbb{P}(z_t < 0 | z_{t-1} < 0)$. Similarly, when there is a known upper bound y_{\max} due to natural bounds or censoring, we may simply set $\mathcal{A}_{y_{\max}} = [y_{\max}, \infty)$ so that the warpDLM has the correct support, $\mathbb{P}(y_t \leq y_{\max} | \boldsymbol{\theta}) = 1$. These useful rounding operation properties are formalized in Kowal and Canale (2020) and Kowal and Wu (2021). Importantly, the rounding operator does

not require any modification to the computing algorithms: once it is specified, inference proceeds the exact same way for all choices of h .

While the rounding operation matches the discreteness and support of the data, the transformation enables (nonparametric) distributional flexibility. We apply the transformation elementwise, $g(\mathbf{y}_t^*) = (g_1(y_{1,t}^*), \dots, g_n(y_{n,t}^*))'$, and again present the $n = 1$ case for simplicity. The only requirement of the transformation g is that it be strictly monotonic, which preserves ordering in the latent data space and ensures an inverse exists. We present both *parametric* and *nonparametric* modeling strategies. Parametric examples include classical transformations of count data, such as logarithmic, square-root, or identity transformations, and introduce no additional parameters into the model.

The nonparametric strategy uses a flexible and data-driven approach to infer the transformation based on the marginal distribution of each component of \mathbf{y} . Specifically, let $\mathcal{A}_j = [a_j, a_{j+1})$ as above and consider $n = 1$. The cumulative distribution function (CDF) of y and z are linked via $F_y(j) = F_z\{g(a_{j+1})\}$, which suggests the transformation

$$\hat{g}_0(a_{j+1}) = \bar{y} + \hat{s}_y \Phi^{-1}\{\tilde{F}_y(j)\}, \quad (6)$$

where \tilde{F}_y is an estimate or model for F_y and \bar{y} and \hat{s}_y are the sample mean and sample standard deviation, respectively, of $\{y_t\}_{t=1}^T$, to match the marginal moments of y and z (Kowal and Wu, 2021). We smoothly interpolate $(y_t, \hat{g}_0(a_{y_t+1}))$ using a monotonic spline, which ensures that the warpDLM is supported on \mathbb{N} instead of only the observed data values. Hence, a (nonparametric) model for g may be equivalently specified by a (nonparametric) model for F_y . We adopt the (rescaled) empirical CDF $\tilde{F}_y(j) = (T + 1)^{-1} \sum_{t=1}^T \mathbb{I}(y_t \leq j)$, which implies a *semiparametric* model for the warpDLM. Many other models for F_y are compatible within the warpDLM, including Bayesian nonparametric models and parametric distributions (e.g., Poisson or Negative Binomial marginals). By design, the model for the transformation g is decoupled from both the rounding operator h and the DLM (3)–(4), so the subsequent derivations and algorithms require only trivial modifications for distinct choices of g .

The nonparametric transformation serves as a reasonable default across a variety of scenarios, as suggested by our results in Section 5 and further confirmed in other (non-time series) settings (e.g. Kowal and Wu, 2021). To compare models with different transformations, there are a variety of choices. In the simulation study of Section 5.2, we select the best model based on out-of-sample forecasting performance using leave-future-out cross-validation. However, in practice, any Bayesian model comparison metric could be used, such as the widely applicable or Watanabe-Akaike information criterion (WAIC, Watanabe, 2010) or approaches more specifically tailored to time series analysis (Bürkner et al., 2020).

3 Exact Filtering and Smoothing

In this section, we derive the recursive updates for the warpDLM filtering, smoothing, and forecasting distributions. Currently, these *exact* results stemming from a *coherent* joint distribution (1)–(2) are unique among state space models for multivariate count (or rounded) data, and crucially enable MCMC-free inference and forecasting.

3.1 Selection Distributions and the warpDLM

Consider the first time step $t = 1$ of the warpDLM. Here, we omit the time subscripts for simplicity. The latent data (2) are described by the two DLM equations (3)–(4), which can be rewritten as a single equation

$$\mathbf{z} = \mathbf{F}\boldsymbol{\theta} + \mathbf{v}, \quad \mathbf{v} \sim N_n(\mathbf{0}, \mathbf{V}) \quad (7)$$

with $\boldsymbol{\theta} := \boldsymbol{\theta}_1$ and the prior $\boldsymbol{\theta} \sim N_p(\boldsymbol{\mu}_\theta = \mathbf{G}a_0, \boldsymbol{\Sigma}_\theta = \mathbf{G}\mathbf{R}_0\mathbf{G}' + \mathbf{W})$. Using the likelihood (5), the posterior distribution is

$$p(\boldsymbol{\theta}|\mathbf{y}) = p(\boldsymbol{\theta}|\mathbf{z} \in \mathcal{C}) = \frac{p(\boldsymbol{\theta})p\{\mathbf{z} \in \mathcal{C}|\boldsymbol{\theta}\}}{p\{\mathbf{z} \in \mathcal{C}\}} \quad (8)$$

which is also the first-step filtering distribution. Within the warpDLM, conditioning on \mathbf{y} is equivalent to conditioning on \mathbf{z} belonging to a set $\mathcal{C} = g(\mathcal{A}_y)$. Although the representation is general, this set is typically simple: the default floor operator for h implies that $\mathcal{A}_y = [y_1, y_1 + 1] \times \cdots \times [y_n, y_n + 1]$, so that $\mathbf{z} \in \mathcal{C}$ implies that each element of \mathbf{z} belongs to a (transformed) interval.

A distribution of the form (8) is known as a selection distribution (Arellano-Valle et al., 2006). When the two variables $\boldsymbol{\theta}$ and \mathbf{z} are jointly normal—as in the warpDLM—the resulting distribution is a *selection normal* (SLCT-N). More formally, given the joint distribution

$$\begin{pmatrix} \mathbf{z} \\ \boldsymbol{\theta} \end{pmatrix} \sim N_{n+p} \left\{ \begin{pmatrix} \boldsymbol{\mu}_z \\ \boldsymbol{\mu}_\theta \end{pmatrix}, \begin{pmatrix} \boldsymbol{\Sigma}_z & \boldsymbol{\Sigma}_{z\theta} \\ \boldsymbol{\Sigma}'_{z\theta} & \boldsymbol{\Sigma}_\theta \end{pmatrix} \right\}$$

we denote the conditional random variable $[\boldsymbol{\theta}|\mathbf{z} \in \mathcal{C}] \sim \text{SLCT-N}_{n,p}(\boldsymbol{\mu}_z, \boldsymbol{\mu}_\theta, \boldsymbol{\Sigma}_z, \boldsymbol{\Sigma}_\theta, \boldsymbol{\Sigma}_{z\theta}, \mathcal{C})$ for constraint region \mathcal{C} . This random variable has density

$$p(\boldsymbol{\theta}|\mathbf{z} \in \mathcal{C}) = \phi_p(\boldsymbol{\theta}; \boldsymbol{\mu}_\theta, \boldsymbol{\Sigma}_\theta) \frac{\bar{\Phi}_n(\mathcal{C}; \boldsymbol{\Sigma}_{z\theta}\boldsymbol{\Sigma}_\theta^{-1}(\boldsymbol{\theta} - \boldsymbol{\mu}_\theta) + \boldsymbol{\mu}_z, \boldsymbol{\Sigma}_z - \boldsymbol{\Sigma}_{z\theta}\boldsymbol{\Sigma}_\theta^{-1}\boldsymbol{\Sigma}'_{z\theta})}{\bar{\Phi}_n(\mathcal{C}; \boldsymbol{\mu}_z, \boldsymbol{\Sigma}_z)} \quad (9)$$

where $\phi_p(\cdot; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes the Gaussian density function of a Gaussian random variable with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$ and $\bar{\Phi}_n(\mathcal{C}; \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \int_{\mathcal{C}} \phi_n(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}) d\mathbf{x}$. The density (9) is somewhat unwieldy in practice, but there is also a constructive representation which allows for direct Monte Carlo simulation from the posterior density $p(\boldsymbol{\theta}|\mathbf{y})$ (see Section 4).

For the first-step model (7) with prior $\boldsymbol{\theta} \sim N_p(\boldsymbol{\mu}_\theta, \boldsymbol{\Sigma}_\theta)$, we report the exact posterior distribution:

Theorem 1. *Under (7), the posterior distribution is $[\boldsymbol{\theta}|\mathbf{y}] \sim \text{SLCT-N}_{n,p}(\boldsymbol{\mu}_z = \mathbf{F}\boldsymbol{\mu}_\theta, \boldsymbol{\mu}_\theta, \boldsymbol{\Sigma}_z = \mathbf{F}\boldsymbol{\Sigma}_\theta\mathbf{F}' + \mathbf{V}, \boldsymbol{\Sigma}_\theta, \boldsymbol{\Sigma}_{z\theta} = \mathbf{F}\boldsymbol{\Sigma}_\theta, \mathcal{C} = g(\mathcal{A}_y))$.*

The Gaussian prior in Theorem 1 is conjugate: Gaussian distributions can be viewed as a limiting case of a $\text{SLCT-N}_{n,p}$ where $n = 0$. However, the conjugacy suggests a more general prior for $\boldsymbol{\theta}$, namely, the SLCT-N distribution:

Lemma 1. Consider the prior $\boldsymbol{\theta} \sim \text{SLCT-}N_{n_0,p}(\boldsymbol{\mu}_{z_0}, \boldsymbol{\mu}_\theta, \boldsymbol{\Sigma}_{z_0}, \boldsymbol{\Sigma}_\theta, \boldsymbol{\Sigma}_{z_0\theta}, \mathcal{C}_0)$ with the latent-data observation equation (7). The posterior is

$$\begin{aligned} [\boldsymbol{\theta}|\mathbf{y}] \sim \text{SLCT-}N_{n_0+n,p} & \left\{ \boldsymbol{\mu}_{z_1} = \begin{pmatrix} \boldsymbol{\mu}_{z_0} \\ \mathbf{F}\boldsymbol{\mu}_\theta \end{pmatrix}, \boldsymbol{\mu}_\theta, \boldsymbol{\Sigma}_{z_1} = \begin{pmatrix} \boldsymbol{\Sigma}_{z_0} & \boldsymbol{\Sigma}_{z_0\theta}\mathbf{F}' \\ \mathbf{F}\boldsymbol{\Sigma}'_{z_0\theta} & \mathbf{F}\boldsymbol{\Sigma}_\theta\mathbf{F}' + \mathbf{V} \end{pmatrix}, \boldsymbol{\Sigma}_\theta, \right. \\ & \left. \boldsymbol{\Sigma}_{z_1\theta} = \begin{pmatrix} \boldsymbol{\Sigma}_{z_0\theta} \\ \mathbf{F}\boldsymbol{\Sigma}_\theta \end{pmatrix}, \mathcal{C}_1 = \mathcal{C}_0 \times g(\mathcal{A}_{\mathbf{y}}) \right\}. \end{aligned}$$

Thus the selection normal distribution is conjugate with the warpDLM likelihood. Moreover, since the Gaussian distribution is closed under linear transformations, this property is preserved for the selection normal distribution (Arellano-Valle et al., 2006). These results link the selection normal distribution with the warpDLM prior-to-posterior updating mechanism, and provide the key building blocks for deriving the warpDLM filtering and smoothing distributions.

3.2 Filtering and Smoothing for Warped DLMs

We now proceed to the most general setting for warpDLM, using the conjugacy and closure under linear transformation results to derive the exact forms of the filtering, smoothing, and forecasting distributions. Throughout, we assume the quadruple $\{\mathbf{F}_t, \mathbf{G}_t, \mathbf{V}_t, \mathbf{W}_t\}$ for $t = 1, \dots, T$ is known. Estimation of variance parameters is addressed in Section 4.

Filtering

Adding back the appropriate subscripts to Theorem 1, the first-step filtering distribution is $(\boldsymbol{\theta}_1|\mathbf{y}_1) \sim \text{SLCT-}N_{n,p}(\boldsymbol{\mu}_{z_1} = \mathbf{F}_1\mathbf{a}_1, \boldsymbol{\mu}_\theta = \mathbf{a}_1, \boldsymbol{\Sigma}_{z_1} = \mathbf{F}_1\mathbf{R}_1\mathbf{F}'_1 + \mathbf{V}_1, \boldsymbol{\Sigma}_\theta = \mathbf{R}_1, \boldsymbol{\Sigma}_{z_1\theta} = \mathbf{F}_1\mathbf{R}_1, \mathcal{C}_1 = g(\mathcal{A}_{\mathbf{y}_1}))$, where $\mathbf{a}_1 = \mathbf{G}_1\mathbf{a}_0$ and $\mathbf{R}_1 = \mathbf{G}_1\mathbf{R}_0\mathbf{G}'_1 + \mathbf{W}_1$. Given the first-step filtering distribution, we proceed inductively for time t :

Theorem 2. Let $(\boldsymbol{\theta}_{t-1}|\mathbf{y}_{1:t-1}) \sim \text{SLCT-}N_{n(t-1),p}(\boldsymbol{\mu}_{z_{t-1|t-1}}, \boldsymbol{\mu}_{\theta_{t-1|t-1}}, \boldsymbol{\Sigma}_{z_{t-1|t-1}}, \boldsymbol{\Sigma}_{\theta_{t-1|t-1}}, \boldsymbol{\Sigma}_{(z\theta)_{t-1|t-1}}, \mathcal{C}_{t-1|t-1})$ be the filtering distribution at time $t-1$ under the warpDLM. Then the one-step-ahead state predictive distribution at t is

$$(\boldsymbol{\theta}_t|\mathbf{y}_{1:t-1}) \sim \text{SLCT-}N_{n(t-1),p}(\boldsymbol{\mu}_{z_{t|t-1}}, \boldsymbol{\mu}_{\theta_{t|t-1}}, \boldsymbol{\Sigma}_{z_{t|t-1}}, \boldsymbol{\Sigma}_{\theta_{t|t-1}}, \boldsymbol{\Sigma}_{(z\theta)_{t|t-1}}, \mathcal{C}_{t|t-1}) \quad (10)$$

with $\boldsymbol{\mu}_{z_{t|t-1}} = \boldsymbol{\mu}_{z_{t-1|t-1}}, \boldsymbol{\mu}_{\theta_{t|t-1}} = \mathbf{G}_t\boldsymbol{\mu}_{\theta_{t-1|t-1}}, \boldsymbol{\Sigma}_{z_{t|t-1}} = \boldsymbol{\Sigma}_{z_{t-1|t-1}}, \boldsymbol{\Sigma}_{\theta_{t|t-1}} = \mathbf{G}_t\boldsymbol{\Sigma}_{\theta_{t-1|t-1}}\mathbf{G}'_t + \mathbf{W}_t, \boldsymbol{\Sigma}_{(z\theta)_{t|t-1}} = \boldsymbol{\Sigma}_{(z\theta)_{t-1|t-1}}\mathbf{G}'_t$, and $\mathcal{C}_{t|t-1} = \mathcal{C}_{t-1|t-1}$. Furthermore, the filtering distribution at time t is

$$(\boldsymbol{\theta}_t|\mathbf{y}_{1:t}) \sim \text{SLCT-}N_{n(t),p}(\boldsymbol{\mu}_{z_{t|t}}, \boldsymbol{\mu}_{\theta_{t|t}}, \boldsymbol{\Sigma}_{z_{t|t}}, \boldsymbol{\Sigma}_{\theta_{t|t}}, \boldsymbol{\Sigma}_{(z\theta)_{t|t}}, \mathcal{C}_{t|t}) \quad (11)$$

with $\boldsymbol{\mu}_{z_{t|t}} = \begin{pmatrix} \boldsymbol{\mu}_{z_{t|t-1}} \\ \mathbf{F}_t\boldsymbol{\mu}_{\theta_{t|t-1}} \end{pmatrix}, \boldsymbol{\mu}_{\theta_{t|t}} = \boldsymbol{\mu}_{\theta_{t|t-1}}, \boldsymbol{\Sigma}_{z_{t|t}} = \begin{pmatrix} \boldsymbol{\Sigma}_{z_{t|t-1}} & \boldsymbol{\Sigma}_{(z\theta)_{t|t-1}}\mathbf{F}'_t \\ \mathbf{F}_t\boldsymbol{\Sigma}'_{(z\theta)_{t|t-1}} & \mathbf{F}_t\boldsymbol{\Sigma}_{\theta_{t|t-1}}\mathbf{F}'_t + \mathbf{V}_t \end{pmatrix}, \boldsymbol{\Sigma}_{\theta_{t|t}} = \boldsymbol{\Sigma}_{\theta_{t|t-1}}, \boldsymbol{\Sigma}_{(z\theta)_{t|t}} = \begin{pmatrix} \boldsymbol{\Sigma}_{(z\theta)_{t|t-1}} \\ \mathbf{F}_t\boldsymbol{\Sigma}_{\theta_{t|t-1}} \end{pmatrix}$, and $\mathcal{C}_{t|t} = \mathcal{C}_{t|t-1} \times g(\mathcal{A}_{\mathbf{y}_t})$.

The state predictive distribution updating in (10) is analogous to that of the Kalman filter, but the filtering distribution update in (11) has a different form. In the Kalman filter, the filtering distribution parameters are updated by taking the previous best estimate for the states (from the state predictive distribution) and “correcting” them using information gained from the new observation. Under the warpDLM, the previous parameters are not corrected; instead, the dimension of the new filtering distribution increases, with the latest data point controlling the bounds over which the latent state can vary. The growing dimension ensures we are retaining all the information about the states, although it also means that sampling can become demanding for longer time series (see Section 4 for discussion). The selection region C also grows in size with each update, but from a storage perspective this is not burdensome: the region $C_{t|t}$ adds an interval for each component of \mathbf{y}_t , which simply requires appending a row to two $(t-1) \times n$ matrices. Importantly, the parameter updates depend on the new observation only through C , which implies that the system matrices do not need to be updated recursively. This point is better clarified by considering the smoothing distribution.

Smoothing

The joint smoothing distribution is $p(\boldsymbol{\theta}_{1:T}|\mathbf{y}_{1:T})$ where T is the terminal time point. Most commonly, the smoothing distribution in a DLM is obtained by first filtering “forward” recursively and then smoothing “backward” in time (Rauch et al., 1965). This strategy targets the marginal smoothing distributions $p(\boldsymbol{\theta}_t|\mathbf{y}_{1:T})$ and requires modifications to access the *joint* smoothing distribution (e.g., Durbin and Koopman, 2002).

Within the warpDLM, we instead target the joint smoothing distribution directly and analytically. Crucially, this process does not require any preliminary passes through the data: all smoothing parameters can be constructed *a priori* through matrix multiplications. Let $\mathbf{G}_{1:t} = \mathbf{G}_t \mathbf{G}_{t-1} \cdots \mathbf{G}_1$ and $\boldsymbol{\mu}_\theta = (\mathbf{G}_{1:1} \mathbf{a}_0, \dots, \mathbf{G}_{1:T} \mathbf{a}_0)$, a pT -dimensional vector. Now, let $\mathbf{R}_t = \mathbf{G}_t \mathbf{R}_{t-1} \mathbf{G}'_t + \mathbf{W}_t$ for $t = 1, \dots, T$. Then the covariance matrix of $\boldsymbol{\theta}$ is a $pT \times pT$ matrix with $p \times p$ diagonal block entries of $\boldsymbol{\Sigma}_\theta[t, t] = \mathbf{R}_t = \mathbf{G}_{1:t} \mathbf{R}_0 \mathbf{G}'_{1:t} + \sum_{q=2}^t \mathbf{G}_{q:t} \mathbf{W}_{q-1} \mathbf{G}'_{q:t} + \mathbf{W}_t$ and cross covariance entries (for $t > q$) $\boldsymbol{\Sigma}_\theta[t, q] = \boldsymbol{\Sigma}_\theta[q, t]^\top = \mathbf{G}_{(q+1):t} \boldsymbol{\Sigma}_\theta[q, q]$. We also define two block diagonal matrices, \mathfrak{F} and \mathfrak{V} , with diagonal entries of \mathbf{F}_t and \mathbf{V}_t , respectively, for $t = 1, \dots, n$, so \mathfrak{F} is a $nT \times pT$ matrix and \mathfrak{V} is a $nT \times nT$ matrix. Finally, let $\mathcal{C}_{1:T} = (g(\mathcal{A}_{\mathbf{y}_1}), \dots, g(\mathcal{A}_{\mathbf{y}_T}))$ be the matrix of selection intervals through time T . The joint smoothing distribution is given below.

Theorem 3. *Under the warpDLM, the joint smoothing distribution is*

$$(\boldsymbol{\theta}_{1:T}|\mathbf{y}_{1:T}) \sim SLCT\text{-}N_{nT,pT}\{\boldsymbol{\mu}_z = \mathfrak{F}\boldsymbol{\mu}_\theta, \boldsymbol{\mu}_\theta, \boldsymbol{\Sigma}_z = \mathfrak{V} + \mathfrak{F}\boldsymbol{\Sigma}_\theta\mathfrak{F}', \boldsymbol{\Sigma}_\theta, \boldsymbol{\Sigma}_{z\theta} = \mathfrak{F}\boldsymbol{\Sigma}_\theta, \mathcal{C} = \mathcal{C}_{1:T}\}. \quad (12)$$

These results can be applied to the time s filtering distribution $p(\boldsymbol{\theta}_{1:s}|\mathbf{y}_{1:s})$ to provide a more concise representation. Using the analogous smoothing parameters but applied up to time s instead of T , we can rewrite the parameters of the filtering distribution as $\boldsymbol{\mu}_{z_{s|s}} = \mathfrak{F}\boldsymbol{\mu}_\theta, \boldsymbol{\mu}_{\theta_{s|s}} = \mathbf{G}_s^s \mathbf{a}_0, \boldsymbol{\Sigma}_{z_{s|s}} = \mathfrak{V} + \mathfrak{F}\boldsymbol{\Sigma}_\theta\mathfrak{F}', \boldsymbol{\Sigma}_{(z\theta)_{s|s}} = \mathfrak{F}\boldsymbol{\Sigma}_\theta, \boldsymbol{\Sigma}_{\theta_{s|s}} = \mathbf{R}_s$. Although these covariance terms are quite complex in terms of notation, constructing such matrices is computationally straightforward via matrix multiplications and additions.

The marginal smoothing distribution $p(\boldsymbol{\theta}_t | \mathbf{y}_{1:T})$ at each time t is readily available from Theorem 3. In particular, the SLCT-N family is closed under marginalization, so parameters for the observations stay the same as in the joint case, and we simply pick off the correct block of the state parameters. This is formalized in the corollary below.

Corollary 1. *Under the warpDLM, the marginal smoothing distribution at time t is*

$$\begin{aligned} (\boldsymbol{\theta}_t | \mathbf{y}_{1:T}) &\sim SLCT-N_{nT,p}\{\boldsymbol{\mu}_{z_{t|T}} = \mathfrak{F}\boldsymbol{\mu}_\theta, \boldsymbol{\mu}_{\theta_{t|T}} = \boldsymbol{\mu}_\theta[t], \boldsymbol{\Sigma}_{z_{t|T}} = \mathfrak{V} + \mathfrak{F}\boldsymbol{\Sigma}_\theta\mathfrak{F}', \\ &\quad \boldsymbol{\Sigma}_{\theta_{t|T}} = \boldsymbol{\Sigma}_\theta[t, t], \boldsymbol{\Sigma}_{(z\theta)_{t|T}} = \boldsymbol{\Sigma}_{z\theta}[t, t], \mathcal{C} = \mathcal{C}_{1:T}\}, \end{aligned} \quad (13)$$

where $\boldsymbol{\Sigma}_{z\theta}[t, t]$ refers to the t -th block of p columns in $\boldsymbol{\Sigma}_{z\theta}$.

The joint smoothing distribution also enables direct computation of the warpDLM marginal likelihood:

Corollary 2. *Under the warpDLM, the marginal likelihood is*

$$p(\mathbf{y}_{1:T}) = \bar{\Phi}_{nT}(\mathcal{C}_{1:T}; \boldsymbol{\mu}_z = \mathfrak{F}\boldsymbol{\mu}_\theta, \boldsymbol{\Sigma}_z = \mathfrak{V} + \mathfrak{F}\boldsymbol{\Sigma}_\theta\mathfrak{F}'). \quad (14)$$

The marginal likelihood is useful for marginal maximum likelihood estimation of the variance parameters in (3)–(4) and other model comparison metrics.

Forecasting

A primary advantage of the warpDLM is the availability of discrete and correctly-supported forecasting distributions. Let $\boldsymbol{\mu}_{z_{t+1|t+1}}$ and $\boldsymbol{\Sigma}_{z_{t+1|t+1}}$ denote parameters of $p(\mathbf{y}_{1:(t+1)})$ analogous to (14), and let $\mathcal{C}_{1:t}$ be the vector of selection intervals up to the known time t . Then the one step forecasting distribution is

$$p(\mathbf{y}_{t+1} | \mathbf{y}_{1:t}) = \frac{p(\mathbf{y}_{1:(t+1)})}{p(\mathbf{y}_{1:t})} = \frac{\bar{\Phi}_{n(t+1)}\{\mathcal{C}_{1:t} \times g(\mathcal{A}_{y_{t+1}}); \boldsymbol{\mu}_{z_{t+1|t+1}}, \boldsymbol{\Sigma}_{z_{t+1|t+1}}\}}{\bar{\Phi}_{nt}(\mathcal{C}_{1:t}; \boldsymbol{\mu}_{z_{t|t}}, \boldsymbol{\Sigma}_{z_{t|t}})}. \quad (15)$$

The h -step ahead forecasting distribution can be defined similarly. In most applications, the forecasting distribution is defined over all non-negative integers, and thus evaluating it at every point is not possible. However, in practice most time series of counts take values in a small range (especially locally), and thus the forecasting distribution only has significant mass on a small range of values. For short to medium time series applications, evaluating these probabilities is quite fast and does not pose a computational burden.

More directly, it is straightforward to simulate from the forecasting distribution given draws from the filtering distribution. Specifically, let $\boldsymbol{\theta}_t^* \sim p(\boldsymbol{\theta}_t | \mathbf{y}_{1:t})$ denote a draw from the filtering distribution. By passing this draw through the DLM, the evolution equation (4) samples $\boldsymbol{\theta}_{t+1}^{**} \sim N_p(\mathbf{G}_{t+1}\boldsymbol{\theta}_t^*, \mathbf{W}_{t+1})$ and the observation equation (3) samples $\mathbf{z}_{t+1}^* \sim N_n(\mathbf{F}_t\boldsymbol{\theta}_{t+1}^{**}, \mathbf{V}_{t+1})$, which equivalently yields a draw from $(\mathbf{z}_{t+1} | \mathbf{y}_{1:t})$. By retaining $\tilde{\mathbf{y}}_{t+1} = h \circ g^{-1}(\mathbf{z}_{t+1}^*)$, we obtain a draw from the one-step forecasting distribution of the warpDLM; multi-step forecasting proceeds similarly. Hence, the primary computational cost is the draw from the filtering distribution, which is detailed in Section 4.

Missing Data

The warpDLM seamlessly handles missing data in any component of \mathbf{y}_t at any time t . Notably, the filtering, smoothing, and forecasting distributions depend on \mathbf{y}_t only through the implied constraint region (i.e., the collection of intervals). Hence, the absence of an observation y_{jt} implies that the interval, such as $[g(y_{jt}), g(y_{jt}+1)]$, is simply replaced by the interval $(-\infty, \infty)$: there is no information about y_{jt} and therefore no constraint for the latent z_{jt} . As a result, the modifications for the filtering, smoothing, and forecasting distributions in the presence of missingness are trivial both analytically and computationally. This result also suggests that the warpDLM is readily applicable to discrete time series data observed at mixed frequencies.

4 Computing

We develop computing strategies for offline and online inference and forecasting. For offline inference, we propose direct Monte Carlo sampling from the relevant filtering or smoothing distributions. We also introduce a Gibbs sampler for offline inference, which decreases raw computing time for large T , but sacrifices some Monte Carlo efficiency due to the autocorrelated samples. Online inference is enabled by an optimal particle filter. The *Monte Carlo* sampler and the *optimal* particle filter are uniquely available as a consequence of the results in Section 3.

4.1 Direct Monte Carlo Sampling

The analytic filtering and smoothing distributions for the warpDLM unlock the potential for direct Monte Carlo sampling. This task requires sampling from a selection normal distribution. Crucially, the selection normal distribution admits a constructive representation using a multivariate normal distribution and a multivariate truncated normal distribution (Arellano-Valle et al., 2006), which is converted to a sampler in Algorithm 1. We then apply Algorithm 1 to obtain Monte Carlo draws from the warpDLM state predictive distribution (10), filtering distribution (11), joint smoothing distribution (12), marginal smoothing distribution (13), or forecasting distribution (Section 3.2). This computing strategy avoids the need for MCMC, which often requires lengthy simulation runs and various diagnostics.

Algorithm 1: Sampling from Selection Normal Distribution.

Result: One sample $\boldsymbol{\theta}$ from SLCT-N distribution

Given $[\boldsymbol{\theta}] \sim \text{SLCT-N}_{d_1, d_2}(\boldsymbol{\mu}_z, \boldsymbol{\mu}_\theta, \boldsymbol{\Sigma}_z, \boldsymbol{\Sigma}_\theta, \boldsymbol{\Sigma}_{z\theta}, \mathcal{C})$

1. **Sample truncated multivariate normal:** Sample \mathbf{V}_0 from $N_{d_1}(\mathbf{0}, \boldsymbol{\Sigma}_z)$ truncated to region $\mathcal{C} - \boldsymbol{\mu}_z$
2. **Sample multivariate normal:** Sample \mathbf{V}_1 from $N_{d_2}(\mathbf{0}, \boldsymbol{\Sigma}_\theta - \boldsymbol{\Sigma}'_{z\theta} \boldsymbol{\Sigma}_z^{-1} \boldsymbol{\Sigma}_{z\theta})$
3. **Combine results:** Compute $\boldsymbol{\theta} = \boldsymbol{\mu}_\theta + \mathbf{V}_1 + \boldsymbol{\Sigma}'_{z\theta} \boldsymbol{\Sigma}_z^{-1} \mathbf{V}_0$

These distributions treat the variance components \mathbf{V}_t and \mathbf{W}_t as fixed and known. Unknown variance components may be accommodated using point estimates (e.g., marginal maximum likelihood estimators using (14)) or within a blocked Gibbs sampler that iterates draws of the state parameters (Algorithm 1) and the variance components.

The computational bottleneck of the Monte Carlo sampler is the draw (a) from the multivariate truncated normal. Efficient sampling from multivariate truncated normal distributions is an active area of research; we apply the minimax tilting strategy from Botev (2017), which is implemented in the R package `TruncatedNormal`. This method is highly efficient and accurate for dimensions up to about $d_1 \approx 500$. For the warpDLM, the dimension d_1 corresponds to nt for the filtering distribution at time t or nT for the smoothing distributions. Hence, when the combined data dimension (nT) is very large, other computational approaches are needed; we provide these in the subsequent sections.

4.2 Gibbs Sampling

For lengthy time series, we develop a Gibbs sampler for offline inference that circumvents the computational bottlenecks of direct Monte Carlo sampling. The Gibbs sampler iterates between a *latent data augmentation* draw and a *state parameter* draw (plus the variance components ψ , if unknown). This is represented in Algorithm 2. Although not shown, one can also easily sample from the posterior predictive and forecast distributions for \mathbf{z}_t in each pass. Draws from the corresponding distributions for \mathbf{y}_t are given by computing $h(g^{-1}(\mathbf{z}_t))$.

Algorithm 2: Gibbs Sampler for WarpDLM.

1. **Sample the latent data:** draw $[\mathbf{z}_t | \{\mathbf{y}_t, \boldsymbol{\theta}_t\}_{t=1}^T, \psi]$ from $N_n(\mathbf{F}_t \boldsymbol{\theta}_t, \mathbf{V}_t)$ truncated to $\mathbf{z}_t \in g(\mathcal{A}_{y_t})$ for $t = 1, \dots, T$
2. **Sample the states:** draw $[\boldsymbol{\theta}_{1:T} | \mathbf{z}_{1:T}, \psi]$
3. **Sample the parameters:** draw $[\psi | \{\mathbf{z}_t, \boldsymbol{\theta}_t\}_{t=1}^T]$ (from the appropriate full conditionals depending on priors)

Crucially, the first step of latent data augmentation can be decomposed into T *independent* draws from an n -dimensional truncated normal distribution; hence the Gibbs sampling approach doesn't suffer from the same dimensionality problems as the direct sampling method. Furthermore, the latent state parameter full conditional in step 2 is exactly the smoothing distribution for a Gaussian DLM, and we can thus rely on already-developed Kalman filter-based simulation techniques such as forward filtering backward sampling (FFBS; Frühwirth-Schnatter, 1994; Carter and Kohn, 1994) or the simulation smoother (Durbin and Koopman, 2002). Importantly, these sampling steps are *linear* in the length of the time series T . Additionally, both methods have R implementations which can be readily leveraged: the package `d1m` (Petris, 2010) performs FFBS and the `KFAS` package (Helske, 2017) applies the simulation smoother.

In a fully Bayesian framework, we place priors on all variances, and sample from the appropriate full conditionals in step 3. For univariate models, we put a $\text{Uniform}(0, A)$ prior with A large on all standard deviations, leading to relatively simple Gibbs updates. In the multivariate context, we place independent inverse-Wishart priors on variance matrices as in Petris et al. (2009).

The sampler for the warpDLM framework is similar to that developed in Kowal and Canale (2020) for the static count regression scenario. The Bayesian dynamic probit sampler (Albert and Chib, 1993; Fasano et al., 2021) can be regarded as a special case. In particular, if our data y_t were binary, we could let our transformation g be the identity while defining our rounding operator such that $\mathcal{A}_{y_t=0} = (-\infty, 0)$ and $\mathcal{A}_{y_t=1} = (0, \infty)$, thus recovering the probit Gibbs sampling algorithm.

Overall, the Gibbs sampler provides scalable offline inference for the warpDLM in the case of large T , although at the expense of Monte Carlo efficiency; naturally, the direct Monte Carlo sampler (Algorithm 1) is preferred when the data dimension permits.

4.3 Optimal Particle Filtering

We design an *optimal* particle filter for online inference with the warpDLM. The key sampling steps do *not* depend on T , which provides scalability for lengthy time series. Online particle filtering algorithms for state space models usually apply Sequential Importance Sampling with or without resampling (Doucet et al., 2000). For each time point, sample trajectories or particles are drawn from an importance function, and each particle has an associated importance weight. For the resampling step, draws from the filtering distribution are obtained via a weighted resampling from the previously drawn particles. The purpose of this step is to slow the effect of particle degeneration.

A well-known drawback of particle filtering methods is that the variance of the importance weights grows over time, and thus the trajectories will degenerate, eventually leaving only one particle with nonzero weight. To avoid quick degeneration, it is important to choose a good importance function; the *optimal* function is that which minimizes the conditional variance of weights. Applying Doucet et al. (2000) to the warpDLM, the optimal importance function is $p(\boldsymbol{\theta}_t | \boldsymbol{\theta}_{t-1}, \mathbf{y}_t)$ with weights given by $p(\mathbf{y}_t | \boldsymbol{\theta}_{t-1})$. For many state space models, this function is not known analytically; however, our results from Section 3 produce an exact form for the warpDLM, formalized in the following corollary.

Corollary 3. *Under the warpDLM, we have*

$$\begin{aligned} (\boldsymbol{\theta}_t | \boldsymbol{\theta}_{t-1}, \mathbf{y}_t) &\sim \text{SLCT-}N_{n,p}(\boldsymbol{\mu}_z = \mathbf{F}_t \mathbf{G}_t \boldsymbol{\theta}_{t-1}, \boldsymbol{\mu}_\theta = \mathbf{G}_t \boldsymbol{\theta}_{t-1}, \boldsymbol{\Sigma}_z = \mathbf{V}_t + \mathbf{F}_t \mathbf{W}_t \mathbf{F}'_t, \\ &\quad \boldsymbol{\Sigma}_\theta = \mathbf{W}_t, \boldsymbol{\Sigma}_{z\theta} = \mathbf{F}_t \mathbf{W}_t, C = g(\mathcal{A}_{y_t})) \end{aligned} \quad (16)$$

and

$$p(\mathbf{y}_t | \boldsymbol{\theta}_{t-1}) = \bar{\Phi}_n(g(\mathcal{A}_{y_t}); \boldsymbol{\mu}_z = \mathbf{F}_t \mathbf{G}_t \boldsymbol{\theta}_{t-1}, \boldsymbol{\Sigma}_z = \mathbf{V}_t + \mathbf{F}_t \mathbf{W}_t \mathbf{F}'_t) \quad (17)$$

We apply Corollary 3 to design an *optimal* particle filter for the warpDLM in Algorithm 3. Crucially, the draw from the selection normal in (16) is not dependent on the length of the time series, which provides scalability in T .

Algorithm 3: Optimal Particle Filter for warpDLM.

Result: Draws $(\boldsymbol{\theta}_t^{(1)}, \dots, \boldsymbol{\theta}_t^{(S)})$ from filtering distribution $p(\boldsymbol{\theta}_t | \mathbf{y}_{1:t})$ for $t = 1, \dots, T$

Given S particle trajectories at time $t = 0$ denoted by $\boldsymbol{\theta}_0^{(s)}$;

for $t \leftarrow 1$ **to** T **do**

for $s \leftarrow 1$ **to** S **do**

Sample $\tilde{\boldsymbol{\theta}}_t^{(s)}$ from (16) conditional on $\boldsymbol{\theta}_{t-1} = \boldsymbol{\theta}_{t-1}^{(s)}$;

Calculate weight up to normalizing constant using (17):

$w_t^{(s)} \propto p(\mathbf{y}_t | \boldsymbol{\theta}_{t-1}^{(s)})$

end

Normalize importance weights: $\tilde{w}_t^{(s)} = \frac{w_t^{(s)}}{\sum_{s=1}^S w_t^{(s)}}$;

Obtain final draws $\boldsymbol{\theta}_t^{(1)}, \dots, \boldsymbol{\theta}_t^{(S)}$ by resampling $\tilde{\boldsymbol{\theta}}_t^{(1)}, \dots, \tilde{\boldsymbol{\theta}}_t^{(S)}$ with weights $\tilde{w}_t^{(s)}$

end

The (unnormalized) weights can be used to estimate the marginal likelihood $p(\mathbf{y}_{1:T})$:

$$\hat{p}(\mathbf{y}_{1:T}) := \hat{p}(\mathbf{y}_1) \prod_{t=2}^T \hat{p}(\mathbf{y}_t | \mathbf{y}_{1:t-1}) \quad (18)$$

where $\hat{p}(\mathbf{y}_1) = \frac{1}{S} \sum_{s=1}^S w_1^{(s)}$ and $\hat{p}(\mathbf{y}_t | \mathbf{y}_{1:t-1}) = \frac{1}{S} \sum_{s=1}^S w_t^{(s)}$. As with equation (14), this can be used for model comparisons or estimation of unknown parameters (Kantas et al., 2015). Such an approximation is also key when designing more complex sequential Monte Carlo schemes or sampling based on particle methods (Andrieu et al., 2010).

As both the cross-section dimension n and the length of the time series T increase, we expect that the quality of the particles will decline, despite the optimality of our transition density. This effect is typically more pronounced for the smoothing distribution than the filtering distribution. In that case, we may seek to improve the particle filter (Algorithm 3) by applying relevant techniques from sequential Monte Carlo sampling, such as a mutation step, or resort to particle smoothing (Doucet et al., 2009) when T is especially large. However, we show empirically that the proposed particle filter is capable of accurate inference for $n = 5$ and $T = 4000$ (see the Supplementary Material).

5 Warped DLMs in Action

In this section, we showcase several facets of the warpDLM framework. We first perform a forecasting analysis on simulated data, showing that the warpDLM offers better distributional forecasting than competing Bayesian methods. We then apply the warpDLM to model a bivariate time series of drug overdose counts in Cincinnati, and highlight both the offline and online computing capabilities.

5.1 Evaluating Probabilistic Forecasts

A prominent advantage of the warpDLM is that it provides a forecasting *distribution* that matches the (discrete) support of the data. Distributional forecasts are especially important for discrete data, whereas point forecasts are less informative. Evaluation of probabilistic forecasts should first and foremost verify that forecast distributions are *calibrated* in the sense that observed values are consistent with samples from the distribution (Czado et al., 2009). To assess calibration, we apply the randomized Probability Integral Transform (rPIT). Let H_t denote the true (discrete) cumulative forecasting distribution, from which a count y_t is observed. Each model produces an estimate \hat{H}_t of the forecasting distribution. The rPIT is drawn from $\tilde{p}_t \stackrel{iid}{\sim} \text{Uniform}[\hat{H}_t(y_t - 1), \hat{H}_t(y_t)]$. If the model is well-calibrated, then $\hat{H}_t \approx H_t$ and consequently $\tilde{p}_t \stackrel{iid}{\sim} \text{Uniform}(0, 1)$. Thus, forecasts are evaluated by (i) drawing \tilde{p}_t for each t and (ii) comparing the draws to a standard uniform distribution.

Among calibrated forecasts, we evaluate the *sharpness* of the forecasting distribution, such as the tightness of the prediction intervals. Czado et al. (2009) propose to evaluate the sharpness via scoring rules, in particular the logarithmic scoring rule $-\log(p_x)$ where p_x is the probability mass of the predictive distribution at the observed count x . We will utilize these diagnostics to evaluate various warpDLMs and DGLM competitors.

5.2 Forecasting on Simulated Data

Existing models for count time series are often unable to capture multiple challenging discrete distributional features, such as zero-inflation, boundedness, and heaping. By comparison, the warpDLM is particularly well-suited for these features. At the same time, the warpDLM is sufficiently flexible to model simpler count data settings for which existing models have been designed, such as Poisson data. To showcase this breadth of modeling, we designed simulation experiments for both scenarios and compared the forecasting performance of various warpDLMs against Bayesian competitors.

In the first experiment, we simulated zero-inflated and bounded count time series data. One real-world example of such data is analyzed in Ensom et al. (2014), where the quantity of interest is the number of hours per day which exceed a certain pollutant threshold. Many days show no pollution, leading to zero-inflation, but occasionally every hour in a day exceeds the threshold, resulting in heaped values at the upper bound of 24. Specifically, we sampled from a zero-inflated Poisson with time varying mean λ_t , where $\lambda_1 \sim \text{Uniform}(5, 15)$ and $\lambda_{t+1} \sim N(\lambda_t, 0.2)$ for $t = 2, \dots, T$. The zero-inflation component was drawn uniformly from $[0.1, 0.3]$. Any sampled value above 24 was rounded down to the bound.

In the second scenario, we simulated data from the INGARCH class using the R package `tscount` (Liboschik et al., 2017). An INGARCH(p, q) model assumes $y_t \sim \text{Poisson}(\lambda_t)$ with $\lambda_t = \beta_0 + \sum_{k=1}^p \beta_k Y_{t-k} + \sum_{\ell=1}^q \alpha_{\ell} \lambda_{t-\ell}$. In our simulation, we chose relatively simple dynamics of $p = 1$ and $q = 1$ and set $\beta_0 = 0.3$, $\beta_1 = 0.6$, and $\alpha_1 = 0.2$, with coefficients selected such that simulated series were low-count time series.

For both simulation scenarios, 30 time series of length $T = 200$ were simulated. Among all competing state space models, we applied the same evolution equation: a local level DLM with the univariate state θ_t and the time-invariant system $\{1, 1, V, W\}$ with initial state $\theta_0 \sim N(0, 3)$. The warpDLMs included *parametric* transformations (identity and square-root) and the *nonparametric* transformation (6). In each case, the variance components V and W were estimated via maximization of the marginal likelihood (14) and then Monte Carlo samples from the joint smoothing distribution (12) were drawn using Algorithm 1. These draws were used to simulate from the forecasting distribution, as outlined in Section 3.2. Competing methods included Poisson, negative binomial, and (for the bounded data scenario) binomial DGLMs, as implemented in the R package **KFAS** (Helske, 2017). We estimated variance components with maximum likelihood and used the **KFAS** default value of the dispersion parameter for the negative binomial DGLM.

For the DGLM models, it took several attempts to identify an initial value that resulted in convergence for the maximum likelihood estimation. In general, the likelihoods of state space models often have complex landscapes which can make parameter estimation difficult and sensitive to starting inputs (Helske, 2017). For the warpDLM methods, a suitable starting point was found by running a Gibbs sampler for 2000 iterations (treating the first 500 as burn-in) and then using the posterior means as the input to the optimization algorithm.

In order to evaluate the forecasting performances, we used time series cross-validation for one-step-ahead predictions. Starting at $t = 100$, each of the five methodologies (three warpDLM methods and two DGLM methods) were fitted and one-step-ahead forecasts were drawn. For 50 equally dispersed time points from $t = 100$ to $t = 200$, the process of training and drawing one-step ahead forecasts was repeated. Thus, we have 50 observations to compare to 50 one-step-ahead forecasts for each of the 30 time series.

Since we evaluate a multitude of series, we follow Kolassa (2016) in testing the uniformity of the rPIT values using the data-driven smooth test of Ledwina (1994). Hence, the output is a p-value, where small values indicate poor calibration. We then consider the sharpness by evaluating the logarithmic scoring rule. Our output for these models is 5000 draws from the predictive distribution, so we approximate the probability mass by calculating the proportion of draws which predicted that value. If a value was not sampled in any posterior draw, but was observed, then the logarithmic scoring rule would be infinite. To avoid this, we instead set the mass in that case to 0.0001. For each series, we find the mean log score across the 50 time points and compute the percent difference over a baseline method, chosen to be the Poisson DGLM. The logarithmic scoring rule is negatively oriented, so a negative percent difference equates to improved forecasting.

Figure 1 presents the calibration and sharpness across simulations for the zero-inflated and bounded Poisson example. The binomial DGLM exhibited extremely poor forecasting performance and is omitted for clarity of presentation; results including the binomial DGLM are in the Supplement. A clear takeaway is that the warpDLM with nonparametric transformation is well-calibrated. The identity warpDLM is also reasonably well-calibrated for many simulations, highlighting the importance of the rounding

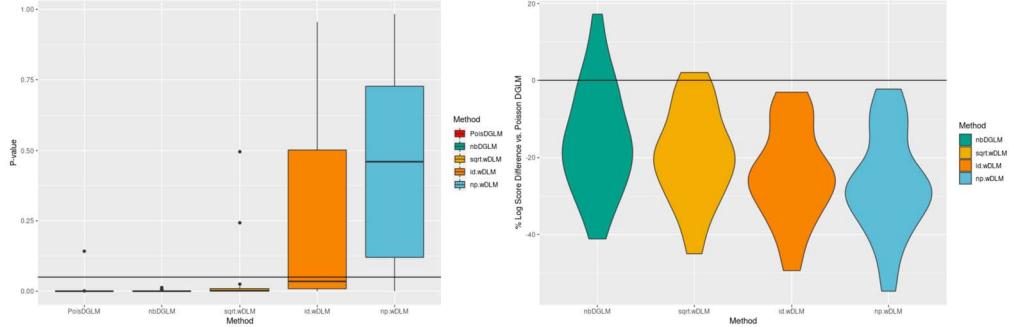


Figure 1: Zero-inflated Poisson simulations; Left: p-values measuring calibration (larger is better) across simulations, with a line at $p = 0.05$; Right: percent difference in log score compared to baseline (negative values imply improved forecasting relative to Poisson DGLM).

operator. The competing models produce extremely low p-values for almost every simulation, which decisively shows that these models are not well-calibrated; Czado et al. (2009) argue that such poor calibration should automatically disqualify these models. The results for sharpness are consistent: the warpDLM with nonparametric transformation produces the best distributional forecasts (on average 30% better than the Poisson DGLM), while the warpDLM with identity transformation is the second best. Modifications for the Poisson and negative binomial DGLMs to enforce the upper bound of 24 on forecasts did not substantively change the results. Clearly, the warpDLM offers significant forecasting gains over competing methods, and advertises both good calibration and maximal sharpness.

The second simulation scenario is presented in Figure 2. This simpler scenario should be more favorable for the Poisson DGLM. Nonetheless, the warpDLM has the flexibility to match or improve upon the DGLM forecasts. All five models show proper calibration for nearly all simulations. Each of warpDLMs exhibits sharpness similar to the Poisson DGLM, and the warpDLM with square-root transformation actually appears to improve on the forecasting for the majority of the series. The negative binomial DGLM performs uniformly worse than all other methods.

Additional simulation results for higher-dimensional data ($n = 5$, $T = 4000$) are included in the Supplementary Material.

5.3 Real Data Analysis

We apply the warpDLM to a multivariate time series of drug overdose (OD) counts in Cincinnati. The OD reports are contained within a dataset of all fire and EMS incidents in the [city's open access data repository](#), and include daily counts of EMS calls for heroin and non-heroin ODs from January 1st, 2018 to January 31st, 2021; see the Supplement for data wrangling and cleaning details. Notably, Cincinnati has

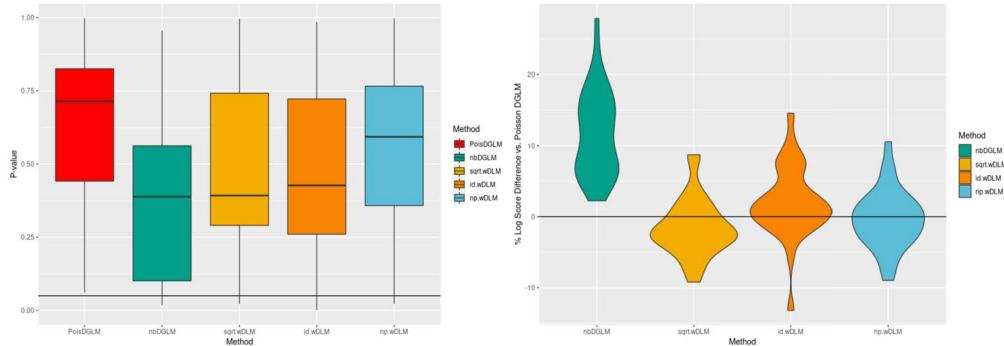


Figure 2: INGARCH simulations; Left: p-values measuring calibration (larger is better) across simulations, with a line at $p = 0.05$; Right: percent difference in log score compared to baseline (negative values imply improved forecasting relative to Poisson DGLM).

struggled with a severe problem of drug ODs in the city, and in particular heroin (Li et al., 2019). From a statistical modeling perspective, these data are discrete, time-ordered, multivariate ($n = 2$), and moderately lengthy ($T = 1127$); see Figure 3.

We specify a warpDLM with the nonparametric transformation from (6), the default rounding operator, and a linear growth model for each series $i \in \{1, 2\}$:

$$z_{i,t} = \mu_{i,t} + v_{i,t} \quad (19)$$

$$\mu_{i,t} = \mu_{i,t-1} + \beta_{i,t-1} + w_{i,t-1}^\mu \quad (20)$$

$$\beta_{i,t} = \beta_{i,t-1} + w_{i,t-1}^\beta \quad (21)$$

where $\mu_{i,t}$ is a local level, $\beta_{i,t}$ is a slope, and (20)–(21) comprise the DLM evolution equation (4). The linear growth model allows for more flexibility than the local level specification previously introduced, but also can effectively reduce to the simpler local level if the data supports it (i.e. if the slope variance parameters are estimated to be very small). Dependence between the $n = 2$ series is induced by the observation error distribution $\mathbf{v}_t \sim N_2(\mathbf{0}, \mathbf{V})$ and the state error distributions $\mathbf{w}_t^\mu \sim N_2(\mathbf{0}, \mathbf{W}_\mu)$ and $\mathbf{w}_t^\beta \sim N_2(\mathbf{0}, \mathbf{W}_\beta)$ for $\mathbf{v}_t = (v_{1,t}, v_{2,t})'$, $\mathbf{w}_t^\mu = (w_{1,t}^\mu, w_{2,t}^\mu)'$, and $\mathbf{w}_t^\beta = (w_{1,t}^\beta, w_{2,t}^\beta)'$, with $\mathbf{V}, \mathbf{W}_\mu, \mathbf{W}_\beta$ assigned inverse-Wishart priors. Model (19)–(21) is expressed in the usual DLM form of (3)–(4) in the Supplement. As an alternative approach, we also constructed a seasonal warpDLM, but its performance was indistinguishable from the linear growth model (see the Supplementary Material).

We highlight the capabilities of the warpDLM for both offline and online analysis of multivariate count time series data. Specifically, we perform an offline analysis for the years of 2018 and 2019 ($T = 730$) and switch to an online particle filter for the 2020 and January 2021 data ($T = 397$). Given the lengthy time series, offline inference was conducted using the Gibbs sampler (Algorithm 2), which returned 10,000 draws from the state filtering distribution of $\boldsymbol{\theta}_t$ for $t = 730$ and the posterior means of the variance

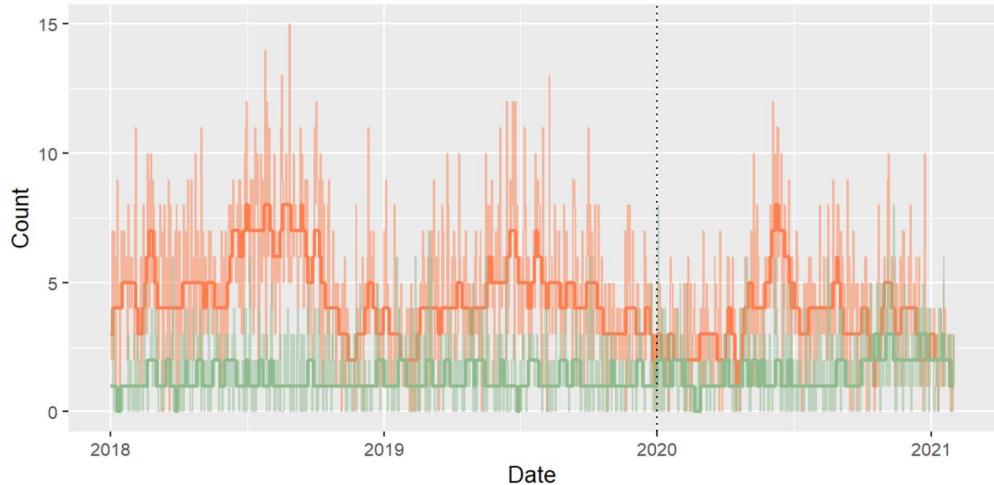


Figure 3: Bivariate time series of daily counts of EMS calls for heroin (larger values) and non-heroin ODs. The dotted vertical line represents the transition between the online and offline datasets. The solid lines are the median smoothed predictions from the warpDLM.

components. These quantities were input into the particle filter of Algorithm 3 and run sequentially through each of the $T = 397$ time points in the online dataset, resulting in 10,000 sequential paths of the smoothed states. These smoothed states were iterated one step ahead to obtain draws from the one-step forecasting distribution as in Section 3.2. We also considered longer forecast horizons (see the Supplementary Material). Unsurprisingly, the forecast distribution accuracy decreased for longer horizons.

To summarize the trends in Figure 3, we compute the pointwise medians of the smoothing predictive distributions $p(\tilde{\mathbf{y}}_t|\mathbf{y}_{1:T})$ under the warpDLM. For each sampled path of states $\boldsymbol{\theta}_{1:T}^* \sim p(\boldsymbol{\theta}_{1:T}|\mathbf{y}_{1:T})$, we draw a corresponding path $\tilde{\mathbf{z}}_{1:T}$ via the DLM observation equation (3), and set $\tilde{\mathbf{y}}_{1:T} = h \circ g^{-1}(\tilde{\mathbf{z}}_{1:T})$, similar to the forecasting procedure in Section 3.2. Notably, this quantity provides a count-valued point estimate that “smooths” each time series conditional on the complete data $\mathbf{y}_{1:T}$. The heroin ODs exhibit both higher values and greater variability over time, although the two time series appear to converge to similar levels in early 2021.

To evaluate the fitness of the warpDLM, we report calibration of the one-step forecasts in Figure 4. The sorted rPIT values of both series are plotted against standard uniform quantiles; the 45 degree line indicates calibration. Since the rPIT values are random quantities, some variation is natural. We illustrate this inherent variability by sampling 100 draws of size 396 from a standard uniform distribution, and plot their sorted values against the uniform quantiles.

The warpDLM one-step forecasts are well-calibrated for non-heroin ODs. For the heroin ODs, the calibration shows minor deviations from uniformity. The histogram

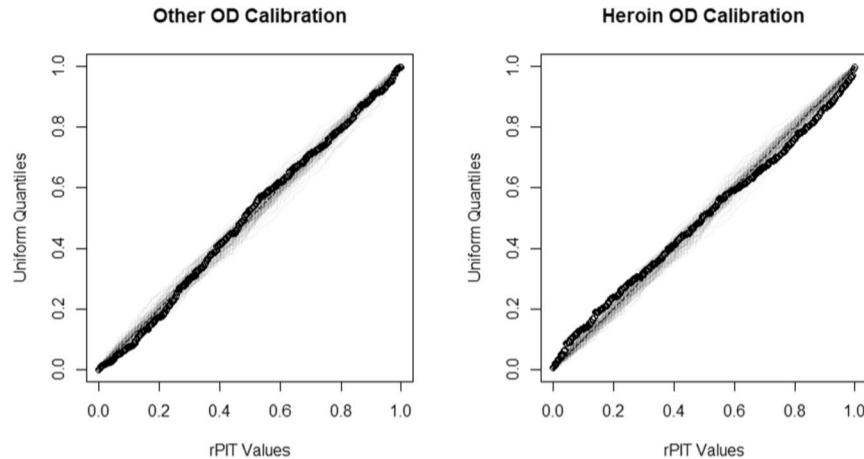


Figure 4: rPIT plots from one-step-ahead particle filter forecasts for both OD count series. The dashed 45 degree line indicates perfect calibration. Gray lines are computed from standard uniform draws, to illustrate inherent variability in the randomized metric.

of the rPIT values (not shown here) shows a slightly inverse-U shape, which implies forecasts are overdispersed relative to the actual values. A closer investigation of the data shows that the online dataset showed considerably less variation than the offline dataset. Since the nonparametric transformation (6) was inferred using only the *offline* dataset, this result suggests the need to occasionally update the transformation during online inference, for example by updating the estimate of the marginal CDF F_y .

Lastly, we evaluate the computational performance of the optimal particle filter, both for Monte Carlo efficiency and raw computing time. Figure 5 presents the effective sample size, $ESS = 1/\sum_{s=1}^S (\tilde{w}_t^{(s)})^2$ using the normalized weights \tilde{w}_t , along with the number of seconds needed to update the model at each time point (on a laptop with an Intel Core i5-6200U CPU with 8 GB RAM). The *optimal* particle filter—uniquely obtained via the warpDLM distributional results from Section 3—produces consistently large ESS values. Crucially, the computation time does *not* increase with the time index and hovers around 13–14 seconds for most updates. Given this computation time, the warpDLM may be applied to more demanding streaming data, such as minute-by-minute data. In all, the optimal particle filter for the warpDLM enabled (i) modeling of multivariate count time series data, (ii) well-calibrated forecasts, and (iii) sufficient scalability for moderate- to high-frequency online analysis.

6 Conclusion

We introduced a state space modeling framework for a multivariate time series of counts by *warping* a dynamic linear model (warpDLM). The warpDLM advertises (i) a data-coherent model (and forecasting distribution) that matches the discrete support of the

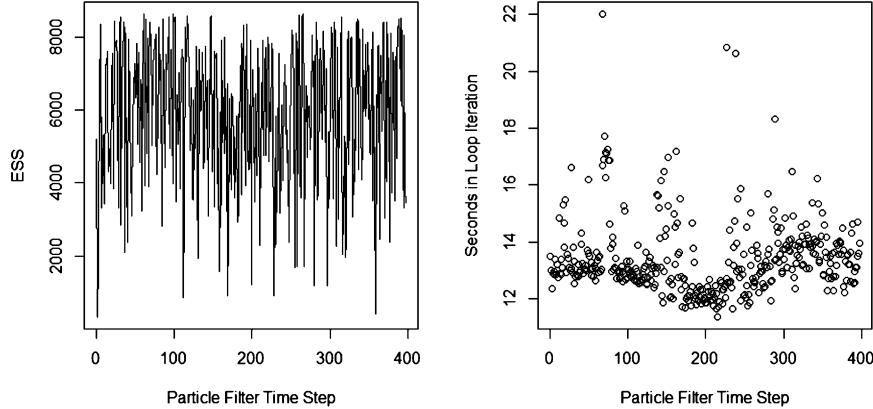


Figure 5: Effective sample sizes (left; maximum is 10,000) and raw computing time in seconds (right) for each time point in the online dataset.

data; (ii) semiparametric modeling capabilities via an inferred transformation; (iii) the familiarity and dynamic flexibility of DLMs, including the ability to handle missing data; (iv) analytic filtering and smoothing recursions; and (v) customized algorithms for forecasting and inference, including direct Monte Carlo sampling and a Gibbs sampler for offline inference and an optimal particle filter for online inference. These results and methods also apply for integer-valued and rounded data.

Using simulated data, we showed that the warpDLM offers better distributional forecasting than Bayesian competitors, particularly when the count data exhibit multiple complexities such as zero-inflation and boundedness. Finally, we showcased the online inference and forecasting capabilities of the warpDLM for a multivariate time series of drug overdose counts.

The generality of the warpDLM enables several useful extensions. Successful tools from (Gaussian) DLMs, such as scale-free variance modeling or discount factors, may be incorporated into the warpDLM framework. These specifications often result in t -distributed state updates and predictions, which may be linked to our results via the selection- t distribution. Similarly, other increases in modeling complexity can be accompanied by appropriate algorithmic advancements, such as sequential Monte Carlo methods that sample both the unknown parameters and the states concurrently (Chopin et al., 2012). Despite the inherent challenges in these more complex modeling and computing environments, the analytic and recursive updates derived for the warpDLM offer a promising pathway for efficient and perhaps optimal implementations.

The warpDLM can also be thought of as a (nonlinear) hierarchical state-space model (HSSM; Gamerman and Migon, 1993), with the latent data z_t introducing a first-level hierarchy. In particular, connections may be made to the very general HSSM structure presented in Katzfuss et al. (2020), with the warping operation presented here acting as the transformation layer in their setup. Exploring this relationship further

might open up new possibilities for applying the warpDLM in high-dimensional problems.

Supplementary Material

Supplement to “Warped Dynamic Linear Models for Time Series of Counts” (DOI: [10.1214/23-BA1394SUPP](https://doi.org/10.1214/23-BA1394SUPP); .pdf).

References

Aktekin, T., Polson, N., and Soyer, R. (2018). “Sequential Bayesian analysis of multivariate count data.” *Bayesian Analysis*, 13(2). [MR3780428](https://doi.org/10.1214/17-BA1054). doi: <https://doi.org/10.1214/17-BA1054>. 4

Albert, J. H. and Chib, S. (1993). “Bayes inference via Gibbs sampling of autoregressive time series subject to Markov mean and variance shifts.” *Journal of Business and Economic Statistics*, 11(1): 1–15. 13

Andrieu, C., Doucet, A., and Holenstein, R. (2010). “Particle Markov chain Monte Carlo methods.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(3): 269–342. [MR2758115](https://doi.org/10.1111/j.1467-9868.2009.00736.x). doi: <https://doi.org/10.1111/j.1467-9868.2009.00736.x>. 14

Arellano-Valle, R. B., Branco, M. D., and Genton, M. G. (2006). “A unified view on skewed distributions arising from selections.” *Canadian Journal of Statistics*, 34(4): 581–601. [MR2347047](https://doi.org/10.1002/cjs.5550340403). doi: <https://doi.org/10.1002/cjs.5550340403>. 2, 7, 8, 11

Berry, L. R. and West, M. (2019). “Bayesian forecasting of many count-valued time series.” *Journal of Business & Economic Statistics*, 38(4): 872–887. [MR4154894](https://doi.org/10.1080/07350015.2019.1604372). doi: <https://doi.org/10.1080/07350015.2019.1604372>. 3

Botev, Z. I. (2017). “The normal law under linear restrictions: Simulation and estimation via minimax tilting.” *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 79(1): 125–148. [MR3597967](https://doi.org/10.1111/rssb.12162). doi: <https://doi.org/10.1111/rssb.12162>. 12

Bradley, J. R., Holan, S. H., and Wikle, C. K. (2018). “Computationally efficient multivariate spatio-temporal models for high-dimensional count-valued data (with discussion).” *Bayesian Analysis*, 13(1). [MR3773410](https://doi.org/10.1214/17-BA1069). doi: <https://doi.org/10.1214/17-BA1069>. 4

Bürkner, P.-C., Gabry, J., and Vehtari, A. (2020). “Approximate leave-future-out cross-validation for Bayesian time series models.” *Journal of Statistical Computation and Simulation*, 90(14): 2499–2523. [MR4145352](https://doi.org/10.1080/00949655.2020.1783262). doi: <https://doi.org/10.1080/00949655.2020.1783262>. 6

Canale, A. and Dunson, D. B. (2011). “Bayesian kernel mixtures for counts.” *Journal of the American Statistical Association*, 106(496): 1528–1539. [MR2896854](https://doi.org/10.1198/jasa.2011.tm10552). doi: <https://doi.org/10.1198/jasa.2011.tm10552>. 3

Canale, A. and Dunson, D. B. (2013). “Nonparametric Bayes modelling of count processes.” *Biometrika*, 100(4): 801–816. [MR3142333](#). doi: <https://doi.org/10.1093/biomet/ast037>. 3

Carter, C. K. and Kohn, R. (1994). “On Gibbs sampling for state space models.” *Biometrika*, 81(3): 541–553. [MR1311096](#). doi: <https://doi.org/10.1093/biomet/81.3.541>. 12

Chopin, N., Jacob, P. E., and Papaspiliopoulos, O. (2012). “SMC2: an efficient algorithm for sequential analysis of state space models.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 75(3): 397–426. [MR3065473](#). doi: <https://doi.org/10.1111/j.1467-9868.2012.01046.x>. 21

Cox, D. R., Gudmundsson, G., Lindgren, G., Bondesson, L., Harsaae, E., Laake, P., Juselius, K., and Lauritzen, S. L. (1981). “Statistical analysis of time series: some recent developments [with discussion and reply].” *Scandinavian Journal of Statistics*, 8(2): 93–115. URL <http://www.jstor.org/stable/4615819> [MR0623586](#). 3

Czado, C., Gneiting, T., and Held, L. (2009). “Predictive model assessment for count data.” *Biometrics*, 65(4): 1254–1261. [MR2756513](#). doi: <https://doi.org/10.1111/j.1541-0420.2009.01191.x>. 15, 17

Doucet, A., Godsill, S., and Andrieu, C. (2000). “On sequential Monte Carlo sampling methods for Bayesian filtering.” *Statistics and Computing*, 10(3): 197–208. 13

Doucet, A., Johansen, A. M., et al. (2009). “A tutorial on particle filtering and smoothing: Fifteen years later.” *Handbook of Nonlinear Filtering*, 12(656-704): 3. [MR2884612](#). 14

Durbin, J. and Koopman, S. J. (2000). “Time series analysis of non-Gaussian observations based on state space models from both classical and Bayesian perspectives.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 62(1): 3–56. URL <https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/1467-9868.00218> [MR1745604](#). doi: <https://doi.org/10.1111/1467-9868.00218>. 3

Durbin, J. and Koopman, S. J. (2002). “A simple and efficient simulation smoother for state space time series analysis.” *Biometrika*, 89(3): 603–615. [MR1929166](#). doi: <https://doi.org/10.1093/biomet/89.3.603>. 9, 12

Durbin, J. and Koopman, S. J. (2012). *Time Series Analysis by State Space Methods*. Oxford University Press. [MR3014996](#). doi: <https://doi.org/10.1093/acprof:oso/9780199641178.001.0001>. 4

Ensor, K. B., Ray, B. K., and Charlton, S. J. (2014). “Point source influence on observed extreme pollution levels in a monitoring network.” *Atmospheric Environment*, 92: 191–198. doi: <https://doi.org/10.1016/j.atmosenv.2014.04.017>. 15

Fahrmeir, L. (1992). “Posterior mode estimation by extended Kalman filtering for multivariate dynamic generalized linear models.” *Journal of the American Statistical Association*, 87(418): 501–509. URL <https://www.tandfonline.com/doi/abs/10.1080/01621459.1992.10475232>. 3

Fasano, A., Rebaudo, G., Durante, D., and Petrone, S. (2021). “A closed-form filter for binary time series.” *Statistics and Computing*, 31(4). [MR4274691](#). doi: <https://doi.org/10.1007/s11222-021-10022-w>. 3, 13

Fokianos, K. (2015). “Statistical analysis of count time series models: A GLM perspective.” in: *Handbook of Discrete-Valued Time Series, Handbooks of Modern Statistical Methods*, 3–28. [MR3699398](#). 3

Freeland, R. and McCabe, B. (2004). “Forecasting discrete valued low count time series.” *International Journal of Forecasting*, 20(3): 427–434. doi: [https://doi.org/10.1016/S0169-2070\(03\)00014-1](https://doi.org/10.1016/S0169-2070(03)00014-1). 3

Frühwirth-Schnatter, S. and Wagner, H. (2006). “Auxiliary mixture sampling for parameter-driven models of time series of counts with applications to state space modelling.” *Biometrika*, 93(4): 827–841. [MR2285074](#). doi: <https://doi.org/10.1093/biomet/93.4.827>. 3

Frühwirth-Schnatter, S. (1994). “Data augmentation and dynamic linear models.” *Journal of Time Series Analysis*, 15(2): 183–202. [MR1263889](#). doi: <https://doi.org/10.1111/j.1467-9892.1994.tb00184.x>. 12

Gamerman, D., dos Santos, T. R., and Franco, G. C. (2013). “A non-Gaussian family of state-space model with exact marginal likelihood.” *Journal of Time Series Analysis*, 34(6): 625–645. [MR3127211](#). doi: <https://doi.org/10.1111/jtsa.12039>. 4

Gamerman, D. and Migon, H. S. (1993). “Dynamic hierarchical models.” *Journal of the Royal Statistical Society: Series B (Methodological)*, 55(3): 629–642. [MR1223932](#). 21

Helske, J. (2017). “KFAS: exponential family state space models in R.” *Journal of Statistical Software*, 78(10). doi: <https://doi.org/10.18637/jss.v078.i10>. 12, 16

Jia, Y., Kechagias, S., Livsey, J., Lund, R., and Pipiras, V. (2021). “Latent Gaussian count time series.” *Journal of the American Statistical Association*, 1–11. [MR4571144](#). doi: <https://doi.org/10.1080/01621459.2021.1944874>. 2, 3

Kalman, R. E. (1960). “A new approach to linear filtering and prediction problems.” *Transactions of the ASME-Journal of Basic Engineering*, 82(Series D): 35–45. [MR3931993](#). 3

Kantas, N., Doucet, A., Singh, S. S., Maciejowski, J., and Chopin, N. (2015). “On particle methods for parameter estimation in state-space models.” *Statistical Science*, 30(3). [MR3383884](#). doi: <https://doi.org/10.1214/14-STS511>. 14

Katzfuss, M., Stroud, J. R., and Wikle, C. K. (2020). “Ensemble Kalman methods for high-dimensional hierarchical dynamic space-time models.” *Journal of the American Statistical Association*, 115(530): 866–885. [MR4107685](#). doi: <https://doi.org/10.1080/01621459.2019.1592753>. 21

King, B. and Kowal, D. R. (2023). *countSTAR: Flexible Modeling of Count Data*. R package version 1.0.1.9000. URL <https://bking124.github.io/countSTAR/>. 4

King, B. and Kowal, D. R. (2023). “Supplementary Material for “Warped Dynamic

Linear Models for Time Series of Counts”” *Bayesian Analysis*. doi: <https://doi.org/10.1214/23-BA1394SUPP>. 4

Kolassa, S. (2016). “Evaluating predictive count data distributions in retail sales forecasting.” *International Journal of Forecasting*, 32(3): 788–803. doi: <https://doi.org/10.1016/j.ijforecast.2015.12.004>. 16

Kowal, D. R. and Canale, A. (2020). “Simultaneous Transformation and Rounding (STAR) models for integer-valued data.” *Electronic Journal of Statistics*. MR4083734. doi: <https://doi.org/10.1214/20-EJS1707>. 3, 5, 13

Kowal, D. R. and Wu, B. (2021). “Semiparametric count data regression for self-reported mental health.” *Biometrics*. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/biom.13617>. 3, 5, 6

Kowal, D. R. and Wu, B. (2022). “Semiparametric discrete data regression with Monte Carlo inference and prediction.” URL <https://arxiv.org/abs/2110.12316>. 3

Ledwina, T. (1994). “Data-driven version of Neyman’s smooth test of fit.” *Journal of the American Statistical Association*, 89(427): 1000–1005. MR1294744. doi: <https://doi.org/10.1080/01621459.1994.10476834>. 16

Li, Z. R., Xie, E., Crawford, F. W., Warren, J. L., McConnell, K., Copple, J. T., Johnson, T., and Gonsalves, G. S. (2019). “Suspected heroin-related overdoses incidents in Cincinnati, Ohio: A spatiotemporal analysis.” *PLoS Medicine*, 16(11): e1002956. 18

Liboschik, T., Fokianos, K., and Fried, R. (2017). “Tscount: An R package for analysis of count time series following generalized linear models.” *Journal of Statistical Software*, 82(5): 1–51. MR3699398. 15

O’Hara, R. and Kotze, J. (2010). “Do not log-transform count data.” *Nature Precedings*, 1–1. doi: <https://doi.org/10.1038/npre.2010.4136.1>. 2

Petris, G. (2010). “An R package for dynamic linear models.” *Journal of Statistical Software*, 36(12). doi: <https://doi.org/10.18637/jss.v036.i12>. 12

Petris, G., Petrone, S., and Campagnoli, P. (2009). *Dynamic Linear Models with R*. Springer. MR2730074. doi: <https://doi.org/10.1007/b135794>. 5, 13

Prado, R., Ferreira, M. A. R., and West, M. (2021). *Time Series: Modeling, Computation, and Inference*. Chapman and Hall/CRC. doi: <https://doi.org/10.1201/9781351259422>. 5

Rauch, H. E., Tung, F., and Striebel, C. T. (1965). “Maximum likelihood estimates of linear dynamic systems.” MR0181489. doi: <https://doi.org/10.2514/3.3166>. 9

Siegfried, S. and Hothorn, T. (2020). “Count transformation models.” *Methods in Ecology and Evolution*, 11(7): 818–827. doi: <https://doi.org/10.1111/2041-210X.13383>. 2

Watanabe, S. (2010). “Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory.” *Journal of Machine Learning Research*, 11(Dec): 3571–3594. MR2756194. 6

West, M. and Harrison, J. (1997). *Bayesian Forecasting and Dynamic Models* (2nd Ed.). Berlin, Heidelberg: Springer-Verlag. [MR1482232](#). 2, 5

West, M., Harrison, P. J., and Migon, H. S. (1985). “Dynamic generalized linear models and Bayesian forecasting.” *Journal of the American Statistical Association*, 80(389): 73–83. [MR0786598](#). doi: <https://doi.org/10.1080/01621459.1985.10477131>. 3