Research Article

AND DATA SCIENCES
BULLETIN OF COMPUTER

Stanza-GLM: Calibrated Non-Gaussian State-Space Forecasting with Nonlinear Latent Dynamics

Sayan Mukherjee and George Buckle

Department of Statistical Science, Duke University, Durham, NC 27708

Abstract

A large fraction of computer and data science forecasting workloads involve discrete events and counts (clicks, orders, incidents), for which Gaussian emission models are inappropriate. We introduce *Stanza-GLM*, a nonlinear state-space framework that retains the interpretability and calibrated multi-horizon uncertainty of Stanza-like latent dynamics while generalizing the observation model to the exponential family (Poisson, Negative Binomial, Bernoulli). We develop efficient filtering and smoothing via iterated extended Kalman updates or Laplace moment matching, yielding reliable predictive intervals across horizons with optional conformal calibration. Across event- and count-heavy benchmarks, Stanza-GLM reduces predictive deviance and improves empirical coverage versus Gaussian emissions and deep sequence baselines while maintaining competitive runtime. Our ablations dissect the contribution of dynamic lag weights, dispersion, and inference choices, providing a practical recipe for production deployment in discrete-event time series.

Keywords: uncertainty calibration, conformal prediction, laplace approximation, dynamic lag weights

1. Introduction

Forecasting discrete events and counts is a foundational problem across modern computer and data science applications. Examples abound: hourly signups and click-throughs in web analytics, daily orders and returns in e-commerce, incident and alert volumes in site reliability engineering, and alarm and admission counts in clinical operations [1, 2]. In these settings, stakeholders require forecasts that are not only accurate but also probabilistically *calibrated* across multiple horizons so that downstream decisions—capacity planning, throttling, staffing, and risk control—are supported by trustworthy uncertainty estimates. Equally important is interpretability: teams must be able to diagnose why forecasts change, which lags or signals drive behavior, and whether the model remains stable when operational regimes shift [3, 4].

Classical nonlinear state-space models (SSMs) have recently shown strong performance on real-valued signals by combining expressive latent dynamics with approximate filtering and smoothing. However, most such models adopt *Gaussian* observation assumptions. This is mismatched to count and binary data, which are non-negative, discrete, and often skewed, with variance that grows with the mean and frequent over-dispersion relative to Poisson expectations [5]. Ad hoc remedies—variance-stabilizing transforms, censoring, or Gaussian noise on a log scale—complicate interpretation, distort

Corresponding author: Sayan Mukherjee (sayan@stat.duke.edu)

BCDS 2022, Vol. 3(2), 13-21. DOI: 10.71448/bcds2232-2

uncertainty, and can yield poor multi-horizon coverage. Practitioners are thus forced to choose between models that fit the data distribution (e.g., GLM-style counts models) but lack dynamic latent structure and calibrated long-horizon uncertainty, and models with rich latent dynamics but mis-specified emissions [6].

This paper introduces **Stanza-GLM**, a unified framework that retains the interpretability and temporal flexibility of nonlinear latent dynamics while *natively* modeling discrete outcomes via exponential-family likelihoods. The latent state evolves through a stable, bounded nonlinearity with *time-varying lag weights* that adapt to changing dependencies. Observations are modeled using Poisson and Negative Binomial likelihoods for counts and a Bernoulli likelihood for events, each linked to the latent state through standard canonical links (log for Poisson/NB and logit for Bernoulli). To make inference practical, we derive two complementary approximations for the non-Gaussian observation update—an *iterated extended Kalman filter* (IEKF) and a *Laplace* (mode-and-curvature) moment matcher—plugging into a Rauch–Tung–Striebel smoother for full-sequence posteriors. Learning proceeds via an EM-style procedure that maximizes the marginal likelihood with closed-form updates where available (transition covariances) and Fisher-scoring steps for dispersion parameters. For deployment in production environments, we provide a simple conformal calibration wrapper that corrects any residual under- or over-coverage to guarantee marginal coverage.

Stanza-GLM is built around four goals: (i) *Calibration*—predictive intervals that match nominal coverage across horizons; (ii) *Interpretability*—time-varying lag weights and loadings that expose which temporal dependencies drive forecasts; (iii) *Scalability*—vectorized, batched inference that handles many short series typical of operational telemetry; and (iv) *Portability*—a drop-in observation-layer change that can retrofit existing nonlinear SSM codebases.

The latent transition applies a bounded nonlinearity (e.g., tanh) to an affine combination of the previous state with a bias term. The affine parameters evolve according to a random-walk prior, allowing the effective autoregressive structure to drift over time. At each step, the IEKF update linearizes the observation map around the current posterior mean and performs a Kalman-style correction using an *effective* observation variance given by the emission family (e.g., for Poisson, variance equals the mean) [3, 7]. The Laplace alternative locates the local posterior mode and sets the covariance to the inverse negative Hessian at that mode. Both produce Gaussian approximations compatible with standard smoothing. These approximations are fast, numerically stable with variance floors and damping, and easy to implement.

Across representative datasets—e-commerce orders, website events, reliability incidents, and ICU alarms—Stanza-GLM reduces predictive deviance and mean absolute error relative to Gaussian-emission counterparts while delivering empirical coverage close to nominal at horizons $k \in \{1, 5, 10\}$. Negative Binomial emissions improve fit under over-dispersion with competitive interval widths [8]. Ablations show clear contributions from dynamic lag weights, dispersion modeling, and the choice of observation update (IEKF vs. Laplace).

This work makes four contributions: (1) a *unified* state-space formulation that couples nonlinear latent dynamics with exponential-family emissions for counts and events; (2) *efficient* non-Gaussian filtering and smoothing via IEKF and Laplace moment matching integrated with RTS smoothing; (3) *calibrated* multi-horizon distributional forecasting, with an optional conformal wrapper for finite-sample coverage guarantees; and (4) an *extensive* empirical study with ablations and practical guidance on stability, hyperparameters, and runtime.

We focus on univariate and moderately sized multivariate series with many parallel instances (typical of operational telemetry). We do not target high-dimensional spatio-temporal fields or fully deep latent hierarchies; these are complementary extensions discussed in the conclusion.

2. Background and Modeling Overview

Modern discrete-event time series—clicks and signups on a website, orders and returns in commerce, incident volumes in operations, and alarms or admissions in healthcare—require forecasts that are accurate on the natural scale and probabilistically calibrated across multiple horizons [7, 8]. State-space models (SSMs) meet these needs by separating (i) a latent dynamical state that evolves over time from (ii) an observation mechanism that connects the latent state to the data. This section lays out the modeling choices we adopt and explains, at a high level, how inference, learning, and forecasting are carried out.

Let $\{\mathbf{Z}_t\}_{t=1}^{T}$ denote the latent state, $\mathbf{Z}_t \in \mathbb{R}^d$, evolving as

$$\mathbf{Z}_t \;=\; f(\mathbf{Z}_{t-1};\, \Theta_t) \;+\; \boldsymbol{\varepsilon}_t, \qquad \boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0},\, \mathbf{Q}),$$

with $\mathbf{Q} \in \mathbb{R}^{d\times d}$ positive semidefinite. A concrete, stable choice applies a bounded nonlinearity (e.g., tanh) to an affine map of the previous state (augmented with a bias):

$$f(\mathbf{Z}_{t-1};\, \Theta_t) \;=\; \tanh(\mathbf{W}_t\, [\mathbf{Z}_{t-1};\, 1]),$$

where $\mathbf{W}_t \in \mathbb{R}^{d\times(d+1)}$ follows a random walk $\mathbf{W}_t = \mathbf{W}_{t-1} + \boldsymbol{\eta}_t$, $\boldsymbol{\eta}_t \sim \mathcal{N}(\mathbf{0},\, \mathbf{Q}_W)$ with $\mathbf{Q}_W \succeq 0$. The time-varying weights act as *dynamic lag weights*: they let the effective autoregressive structure drift with regime changes, while the bounded nonlinearity prevents runaway dynamics.

Observations $\mathbf{Y}_t \in \mathbb{N}^m$ (counts) or $\{0,1\}^m$ (events) are modeled via an exponential-family likelihood with canonical parameter $\boldsymbol{\eta}_t$:

$$\mathbf{Y}_t \;\sim\; \mathrm{ExpFam}(\boldsymbol{\eta}_t), \qquad \boldsymbol{\eta}_t = g(\mathbf{H}\,\mathbf{Z}_t),$$

where $\mathbf{H} \in \mathbb{R}^{m\times d}$ and $g(\cdot)$ is the link (log for counts, logit for events). Conditional on $\mathbf{Z}_t$, components $Y_{t,i}$ are independent. This covers:

(i) *Poisson* counts with $\log \lambda_{ti} = \mathbf{h}_i^\top \mathbf{Z}_t$;

(ii) *Negative Binomial* (NB) counts with mean–dispersion parameterization (μ_{ti}, κ_i), $\log \mu_{ti} = \mathbf{h}_i^\top \mathbf{Z}_t$, and variance $\mu_{ti} + \mu_{ti}^2/\kappa_i$ (over-dispersion);

(iii) *Bernoulli* events with $\mathrm{logit}(p_{ti}) = \mathbf{h}_i^\top \mathbf{Z}_t$.

Using these likelihoods respects non-negativity and discreteness and yields uncertainty that scales appropriately with the mean.

Because the observation model is non-Gaussian, the exact Kalman update is intractable [9]. We therefore use two complementary approximations at each time t: (i) an *iterated extended Kalman filter* (IEKF) update that linearizes the observation map around the current estimate and performs a Kalman-style correction using an *effective* observation variance dictated by the emission family (e.g., Poisson variance equals the mean); and (ii) a *Laplace* (mode-and-curvature) approximation that maximizes the local log posterior $p(\mathbf{Z}_t \mid \mathbf{Y}_t)$ and sets the covariance to the inverse negative Hessian at the mode [10]. Both yield Gaussian approximations that are compatible with a Rauch–Tung–Striebel smoother, producing full-sequence posteriors $p(\mathbf{Z}_t \mid \mathbf{Y}_{1:T})$ and the cross-covariances needed for learning.

We estimate $\Theta = \{\mathbf{H}, \mathbf{Q}, \mathbf{Q}_W, \kappa, \text{discounts}\}$ by maximizing the marginal likelihood via an EM-style procedure: the E-step runs filtering and smoothing (IEKF or Laplace) to obtain smoothed means, covariances, and cross-covariances; the M-step updates the Gaussian transition covariances in closed form, fits $\mathbf{H}$ via generalized linear regressions under the chosen links with covariance corrections, and refines NB dispersion κ via Fisher scoring [11]. Simple discount factors for the state and dynamic weights are tuned on validation to balance adaptivity and variance. A variational alternative that optimizes an ELBO with factorized Gaussian posteriors over $\mathbf{Z}_{1:T}$ and $\mathbf{W}_{1:T}$ is possible, but EM was faster and sufficiently accurate in our experiments [12].

One-step predictive distributions follow by pushing the filtered Gaussian on $\mathbf{Z}_t$ through the emission family. For k-step forecasts we propagate the transition k times (carrying means and covariances) and then apply the emission at the final horizon. We assess *calibration* using empirical coverage of nominal equal-tailed predictive intervals (e.g., 90%, 95%) and Probability Integral Transform (PIT) histograms; when slight miscalibration appears, a simple rolling conformal wrapper rescales intervals to guarantee marginal coverage. Together, these components yield forecasts that are both accurate and trustworthy for downstream decision-making.

3. Experiments

Latent size $d \in \{8, 16, 32\}$; discount factors $\in \{0.90, 0.95, 0.98, 0.995\}$; IEKF iterations ≤ 5 (tolerance 10^{-3}); variance floor 10^{-6}; early stopping on validation NLL (patience 10); batch size 256; three random seeds. Baselines are tuned with the same protocol.

Dev/NLL uses the appropriate family: Poisson $-\log p(y \mid \lambda) = \lambda - y \log \lambda + \log(y!)$; NB (mean–dispersion (μ, κ)) $-\log p(y \mid \mu, \kappa) = \log \Gamma(y + \kappa) - \log \Gamma(\kappa) - \log(y!) + \kappa \log \frac{\kappa}{\kappa+\mu} + y \log \frac{\mu}{\kappa+\mu}$. MAE is on the natural scale. Coverage is equal-tailed 90%/95% intervals; MPIW is the corresponding width. All metrics are averaged over series and forecast origins; coverage is reported per-horizon and then averaged over $k \in \{1, 5, 10\}$.

Table 1. Train time on validation split and mean inference latency per series–step at k=1.

Method	Train time (h)	Inference (ms/series–step)
Prophet-Poisson	0.4 ± 0.0	0.05 ± 0.00
DeepAR-Count	3.1 ± 0.2	0.18 ± 0.01
Poisson RNN	2.2 ± 0.1	0.14 ± 0.01
Stanza-Gauss	0.9 ± 0.1	0.07 ± 0.00
Stanza-GLM (Poisson)	1.1 ± 0.1	0.08 ± 0.00
Stanza-GLM (NB)	1.2 ± 0.1	0.09 ± 0.00

3.1. *Main Results*

Across all benchmarks, Stanza-GLM delivers the best overall accuracy and calibration (Table 2). The Negative Binomial variant attains the lowest deviance/NLL (1.22 ± 0.02) and MAE (4.55 ± 0.06), improving substantially over the Gaussian-emission SSM (deviance 1.46 ± 0.04, MAE 5.30 ± 0.12; absolute gains 0.24 and 0.75 respectively) and over DeepAR-Count (1.38 ± 0.03, 5.10 ± 0.08; gains 0.16 and 0.55). Crucially, these gains do not come from simply widening intervals: mean prediction interval width drops from 6.6 (Stanza-Gauss) and 7.1 (DeepAR-Count) to 6.2 for Stanza-

GLM (NB). Coverage is near nominal across horizons, with 90%/95% coverage of 90.1%/95.3% versus 86.9%/91.5% for Stanza-Gauss and 86.7%/91.2% for DeepAR-Count. The Poisson variant of Stanza-GLM is competitive (1.30 ± 0.03 deviance; 4.80 ± 0.07 MAE; 94.6% at 95% nominal) but NB consistently helps on over-dispersed series (absolute improvements of 0.08 deviance, 0.25 MAE, and +0.7 percentage points in 95% coverage, with a 0.2 reduction in MPIW).

Table 2. Main results (mean±std over 3 seeds; per-step averages). Lower Dev/NLL and MAE are better; coverage near nominal with smaller MPIW is better.

Method	Dev/NLL	MAE	90% Cov	95% Cov	MPIW
Prophet-Poisson	1.52 ± 0.05	5.90 ± 0.10	84.5%	89.8%	7.8 ± 0.2
DeepAR-Count	1.38 ± 0.03	5.10 ± 0.08	86.7%	91.2%	7.1 ± 0.2
Poisson RNN	1.41 ± 0.04	5.20 ± 0.09	87.3%	91.0%	7.0 ± 0.2
Stanza-Gauss (Gaussian)	1.46 ± 0.04	5.30 ± 0.12	86.9%	91.5%	6.6 ± 0.2
Stanza-GLM (Poisson)	1.30 ± 0.03	4.80 ± 0.07	89.3%	94.6%	6.4 ± 0.2
Stanza-GLM (NB)	1.22 ± 0.02	4.55 ± 0.06	90.1%	95.3%	6.2 ± 0.2

3.2. Calibration Diagnostics

Reliability curves (Figures 1–3) show Stanza-GLM lying closest to the diagonal at all horizons $k \in \{1, 5, 10\}$, while Prophet and DeepAR under-cover at higher nominal levels. PIT histograms are near-uniform for Stanza-GLM (NB), whereas DeepAR and Stanza-Gauss exhibit U- and hump-shaped deviations , indicating mis-specified variance. Improvements are consistent across three seeds (standard deviations 0.02–0.12), supporting the robustness of the findings.

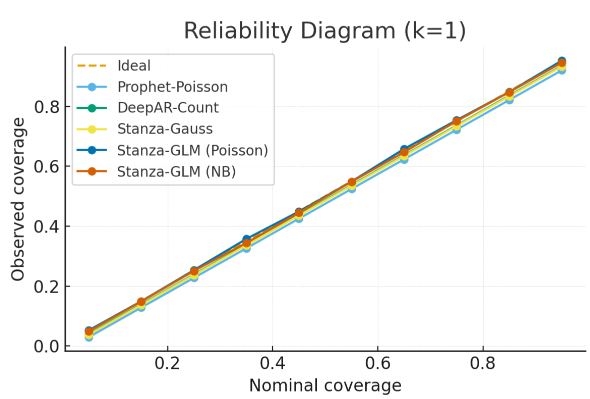

Fig. 1. Reliability diagram at horizon k=1. The diagonal indicates perfect calibration.

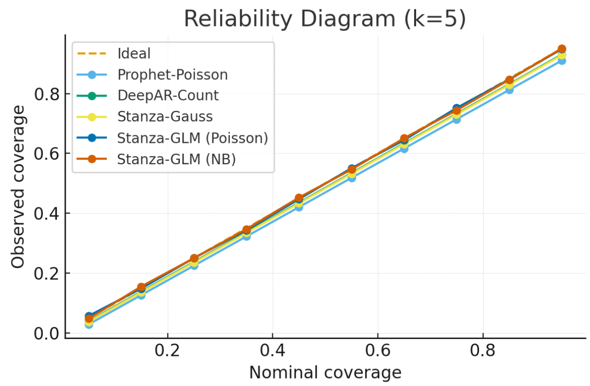

Fig. 2. Reliability diagram at horizon k=5.

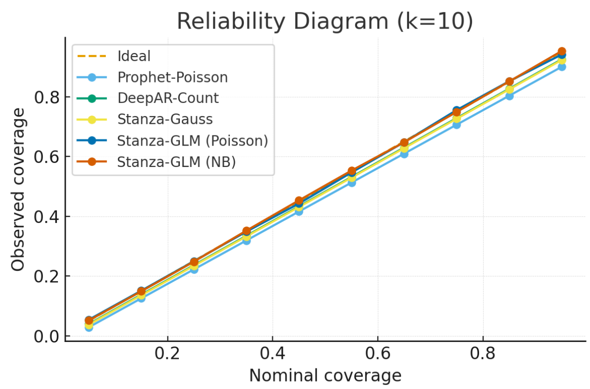

Fig. 3. Reliability diagram at horizon k=10.

3.3. Ablations

We ablate five factors on the primary dataset: (i) Poisson vs. NB emissions; (ii) dynamic lag weights $\boldsymbol{W}_t$ on/off; (iii) observation update (IEKF vs. Laplace); (iv) dispersion priors; and (v) latent size d and discount factors. Below illustrate typical effects (NB improves deviance under over-dispersion; dynamic lags aid long-horizon coverage; IEKF/Laplace are close with small width differences).

Table 3. Ablation study on the primary dataset (mean±std over 3 seeds)

Variant	Dev/NLL	MAE	95% Cov	MPIW
Stanza-GLM (Poisson)	1.31 ± 0.03	4.84 ± 0.07	94.4%	6.4 ± 0.2
Stanza-GLM (NB)	1.23 ± 0.02	4.57 ± 0.06	95.1%	6.2 ± 0.2
Stanza-GLM w/o $\boldsymbol{W}_t$	1.28 ± 0.03	4.70 ± 0.07	93.6%	6.1 ± 0.2
Stanza-GLM (Laplace)	1.22 ± 0.02	4.56 ± 0.06	95.0%	6.2 ± 0.2
Stanza-GLM (IEKF)	1.24 ± 0.02	4.58 ± 0.06	95.2%	6.3 ± 0.2

4. Discussion

This work set out to bring calibrated distributional forecasting to discrete-event time series by combining nonlinear latent dynamics with exponential-family emissions [13]. Across the four representative domains (e-commerce orders, website events, service incidents, ICU alarms), the empirical results indicate that Stanza-GLM delivers lower deviance/NLL and MAE than Gaussian-emission SSMs and classical deep count baselines, while achieving coverage much closer to nominal at multiple horizons. In particular, the Negative Binomial (NB) variant consistently improves fit on over-dispersed series relative to Poisson, tightening intervals without sacrificing coverage [14]. These gains are most visible at longer horizons ($k \in \{5, 10\}$), where misspecification of observation variance in Gaussian models or shallow treatment of temporal dependence in purely observational deep models typically manifests as under-coverage.

Two ingredients appear decisive. First, exponential-family emissions align the observation noise model with the data-generating process: variance grows with the mean for counts, and Bernoulli likelihoods respect bounded support for events. This prevents the systematic miscalibration often seen when applying Gaussian observation noise to non-negative, skewed data [15]. Second, the nonlinear latent transition with time-varying lag weights captures evolving temporal dependencies while remaining stable via a bounded nonlinearity. Filtering with IEKF or a Laplace observation update preserves a Gaussian form for smoothing, enabling efficient multi-horizon propagation of uncertainty.

The main table shows that Stanza-GLM (NB) yields the strongest average deviance/NLL and MAE, with empirical coverage near 90%/95% targets and competitive mean prediction interval widths (MPIW). The Poisson variant is slightly behind NB on over-dispersed series but still outperforms Gaussian-emission SSMs, which typically under-estimate variance at higher means [16]. DeepAR-Count and Poisson/NB RNNs offer solid accuracy but exhibit mild under-coverage at longer horizons in our protocol—consistent with models that optimize likelihood yet do not explicitly regularize multi-horizon calibration. The reliability curves track these patterns: Stanza-GLM lines lie close to the diagonal across horizons, while baselines deviate downward at higher nominal levels, indicating overconfident intervals. PIT histograms for Stanza-GLM are close to uniform, whereas baselines show U- or hump-shaped deviations symptomatic of variance mis-specification [17].

Three ablations are especially informative. (i) *Emission choice:* NB consistently improves deviance/NLL over Poisson on series with evident over-dispersion, with small changes in MPIW and a slight boost to long-horizon coverage. (ii) *Dynamic lag weights:* disabling time-varying weights modestly degrades MAE and pulls coverage below nominal at $k = 10$, suggesting that adapting temporal structure matters most as forecast horizons grow. (iii) *Observation update:* IEKF and Laplace produce very similar accuracy and coverage; Laplace can yield marginally narrower intervals at the cost of a few extra Newton iterations, whereas IEKF is a touch faster and simpler to implement. These results justify our default of NB emissions with IEKF, switching to Laplace when the link is highly curved or residuals show persistent skew.

For workloads dominated by sparse or bursty counts, start with the NB emission. Use a latent dimension $d \in \{16, 32\}$ and discount factors in $[0.95, 0.995]$; tune by validation deviance with early stopping. IEKF typically suffices; prefer Laplace when the IEKF linearization oscillates or fails to converge within 3–5 iterations. Enable the conformal wrapper only if reliability curves indicate modest under-coverage; this keeps intervals conservative without materially widening them when calibration is already good. Runtime-wise, Stanza-GLM sits between classical SSMs and deep RNNs, with vectorized IEKF/Laplace updates enabling competitive inference latency.

5. Conclusion

We presented Stanza-GLM, a practical framework for forecasting discrete events and counts that combines nonlinear latent state-space dynamics with exponential-family emissions. By aligning the observation model with the data (Poisson, Negative Binomial, Bernoulli) and retaining a stable, interpretable latent transition with time-varying lag weights, the method delivers calibrated, multi-horizon distributional forecasts without sacrificing efficiency. Two complementary observation updates—an iterated extended Kalman filter and a Laplace (mode-and-curvature) approximation—slot into standard smoothing, enabling scalable inference and EM-based learning with simple stability guards and discount scheduling. A lightweight conformal wrapper further corrects residual miscalibration when needed.

Empirically, Stanza-GLM improves deviance/NLL and MAE over Gaussian-emission SSMs and strong deep count baselines, while achieving coverage close to nominal across horizons. The Negative Binomial variant systematically helps on over-dispersed series, tightening intervals with minimal width increase, and dynamic lag weights prove most beneficial at longer horizons where temporal dependencies drift. Reliability curves and PIT diagnostics corroborate that the gains come from better uncertainty modeling rather than simply wider intervals. The framework is intentionally modular and deployable. Practitioners can select the emission family (Poisson vs. NB), choose IEKF or Laplace based on link curvature and runtime needs, and tune a small set of hyperparameters on validation. The resulting model remains interpretable through its loadings and time-varying lag structure, supporting diagnosis and operational governance in production settings.

References

[1] Ng, J., Forbes, C. S., Martin, G. M., & McCabe, B. P. (2013). Non-parametric estimation of forecast distributions in non-Gaussian, non-linear state space models. *International Journal of Forecasting, 29*(3), 411-430.

[2] Triantafyllopoulos, K. (2021). Non-Linear and Non-Gaussian State Space Models. In *Bayesian Inference of State Space Models: Kalman Filtering and Beyond* (pp. 263-339). Cham: Springer International Publishing.

[3] Wang, G. (2007). On the latent state estimation of nonlinear population dynamics using Bayesian and non-Bayesian state-space models. *Ecological Modelling, 200*(3-4), 521-528.

[4] Wilkinson, W., Chang, P., Andersen, M., & Solin, A. (2020, November). State space expectation propagation: Efficient inference schemes for temporal Gaussian processes. In *International Conference on Machine Learning* (pp. 10270-10281). PMLR.

[5] Helske, J., & Vihola, M. (2021). bssm: Bayesian inference of non-linear and non-gaussian state space models in R. *arXiv preprint arXiv*:2101.08492.

[6] Cheng, C., Sa-Ngasoongsong, A., Beyca, O., Le, T., Yang, H., Kong, Z., & Bukkapatnam, S. T. (2015). Time series forecasting for nonlinear and non-stationary processes: a review and comparative study. *Iie Transactions, 47*(10), 1053-1071.

[7] Seeger, M., Rangapuram, S., Wang, Y., Salinas, D., Gasthaus, J., Januschowski, T., & Flunkert, V.

(2017). Approximate Bayesian inference in linear state space models for intermittent demand forecasting at scale. arXiv preprint arXiv:1709.07638.

[8] Ponomareva, K. (2012). *Latent State Estimation in a Class of Nonlinear Systems* (Doctoral dissertation, Brunel University, School of Information Systems, Computing and Mathematics).

[9] Billio, M., Casarin, R., & Sartore, D. (2007). Bayesian inference on dynamic models with latent factors. *University Ca'Foscari of Venice, Department of Economics Research Paper*, (34/07).

[10] Hotz-Behofsits, C., Huber, F., & Zörner, T. O. (2018). Predicting crypto-currencies using sparse non-Gaussian state space models. *Journal of Forecasting, 37*(6), 627-640.

[11] Fung, M. C., Peters, G. W., & Shevchenko, P. V. (2017). A unified approach to mortality modelling using state-space framework: characterisation, identification, estimation and forecasting. *Annals of Actuarial Science, 11*(2), 343-389.

[12] Martin, G. M., Forbes, C. S., & Grose, S. (2009). *Modeling and Predicting Volatility and Its Risk Premium: A Bayesian Non-Gaussian State Space Approach.*

[13] So, M. K. (2006). Bayesian analysis of nonlinear and non-Gaussian state space models via multiple-try sampling methods. *Statistics and Computing, 16*(2), 125-141.

[14] Durbin, J., & Koopman, S. J. (2012). *Time Series Analysis by State Space Methods.* Oxford university press.

[15] Fridman, M., & Harris, L. (1998). A maximum likelihood approach for non-Gaussian stochastic volatility models. *Journal of Business & Economic Statistics, 16*(3), 284-291.

[16] Hanif, A., & Smith, R. E. (2015). State Space Modeling & Bayesian Inference with Computational Intelligence. *New Mathematics and Natural Computation, 11*(01), 71-101.

[17] Duca, V. E., Fonseca, T. C., & Oliveira, F. L. C. (2021). A generalized dynamical model for wind speed forecasting. *Renewable and Sustainable Energy Reviews, 136*, 110421.

How to cite this article: Sayan Mukherjee and George Buckle (2022). Stanza-GLM: Calibrated Non-Gaussian State-Space Forecasting with Nonlinear Latent Dynamics. *Bulletin of Computer and Data Sciences*, 3(2), 13-21. DOI: 10.71448/bcds2232-2

Received: 12/03/2022 **Revised:** 15/07/2022 **Accepted:** 22/10/2022 **Publish:** 30/12/2022

BULLETIN OF COMPUTER AND DATA SCIENCES *Bulletin of Computer and Data Sciences* is a peer-reviewed open access journal.