Adaptive Sampling for Probabilistic Forecasting under Distribution Shift

Anonymous Author(s)
Affiliation
Address
email

Abstract

The world is not static: This causes real-world time series to change over time because external, and potentially disruptive, events such as macroeconomic cycles or the COVID-19 pandemic change the underlying factors that influence the time series. Once such a data distribution shift happens, it will be part of the time series history and impact future forecasting attempts. We present an adaptive sampling strategy that selects the part of the history that is relevant for the recent data distribution. We achieve this by learning a discrete distribution over relevant time steps by Bayesian optimization. We instantiate this idea with a two-step, model-agnostic method that is pre-trained with uniform sampling and then training a lightweight adaptive architecture with adaptive sampling. We show with synthetic and real-world experiments that this method adapts to distribution shift and reduces the forecasting error of the base model by 8.4%.

1 Introduction

Time series forecasting uses historical data to forecast the future evolution of a time series to assist decision making in several domains, for example retail [26, 5, 35], electric load planning [33, 16], or cloud computing management [27, 28], or labor planning [3]. Most time series forecasting algorithms make the assumption that the data distribution does not change over time. However, this is violated in practice where real-world time series data is affected by subtle drifts such as evolving customer demand or disruptive events (like the COVID-19 pandemic). These distribution shifts will inevitably be part of the time series history and therefore forecasting algorithms need to be able to account for them.

Distribution shifts can occur with different (possibly reoccurring) patterns, but likely result in higher test error because the train and test distribution differs. Thus, only a subset of the time series history might come from the same distribution as the test set. Driven by this observation, we propose an *adaptive sampling mechanism* that explicitly feeds the model with inputs from the history that are only relevant for forecasting in the current (most recent) distribution. Instead of sampling uniformly over the history, our method learns a discrete distribution over past time steps and explicitly allows continuous adaptation, so that the model can react and adapt to distribution shifts.

Our contribution is a model-agnostic adaptive sampling mechanism that learns the sampling distribution over time steps via gradient-free Bayesian optimization. We propose a two-step framework where the forecasting model is pre-trained with uniform sampling and then a lightweight architecture is trained with adaptive sampling. We evaluate the properties and performance using a combination of experiments on synthetic and real-world data.

Submitted to the Workshop on Distribution Shifts, 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

2 Related Work

While the problem of learning under distribution shift in "static" settings has seen extensive research for many years, similar ideas in the time series domain have only recently started to emerge. Kim et al. [18] employs a trainable instance-wise normalization-and-denormalization layer to address non-stationary mean and variance in time series forecasting. Arik et al. [2] proposes Self-Adaptive Forecasting to adapt the forecasting model using test-time training by "backcasting" and using the backcast errors to signal a potential distribution shift to adjust the weights of the model before inference. AdaRNN [8] splits the time series history into dissimilar segments and learns importance weights to combine the RNN hidden states over these segments. Another notion of distribution shift, adversarial attacks, has been recently considered to build more robust forecasting [38, 23].

Another class of approaches to time series forecasting use specific weighting schemes to weight the time series history to make predictions, for example exponential smoothing and the Non-Parametric Time Series (NPTS) Forecaster [12, 1], or stratified sampling [24]. In contrast to these models, our approach adaptively learns the weighting scheme over the history. Our work can be seen as a specific instantiation of weighted re-sampling (or importance sampling) for time-series forecasting. Importance sampling weights the log-likelihood in supervised learning tasks by importance weights $w(x) = p_{test}(x)/p_{train}(x)$ where p_{test} and p_{train} are the respective densities from which the train/test samples are drawn [34, 17, 37, 21]. These approaches typically differ in their approach to estimate the importance weights. Here, we propose to avoid estimating p_{train} and p_{test} and instead tune the parameters of a parametric re-sampling function using Bayesian optimization.

3 Background and Method

In this section, we state our assumptions about the distribution shift and introduce our model. Our main assumption is that we are able to observe at least some of the forthcoming shift in order to be able to react and adapt to it[1]:

Assumption 1 *Let c be a constant number of (past) time steps set a priori. We assume that there is no (marginal) distribution shift between the most recent training window of c steps and the forecasting (test) window. More precisely, let $\{\mathbf{z}_{i,t}\}$ be the given time series and let $\{\bar{\mathbf{z}}_{i,t}\}$ be the time series obtained after removing trend, seasonalities and other effects of available features, then*

$$p(\bar{\mathbf{z}}_{i,T-c}) = \cdots = p(\bar{\mathbf{z}}_{i,T}) = p(\bar{\mathbf{z}}_{i,T+1}) = \cdots = p(\bar{\mathbf{z}}_{i,T+\tau}), \quad \forall i = 1, \ldots, N. \tag{1}$$

We consider the above to be a mild assumption, since data with $t \in [T-c, T]$ is the closest in time to the forecasting window, therefore it is reasonable to regard it as the most similar to the future. In practice, we take c to be equal to the sum of context and prediction lengths, so that we can use $\{\mathbf{z}_{i,T-c:T}\}_{i=1}^{N}$ as a validation set to learn the optimal adaptive sampling parameters. In general, we consider these two scenarios (see Figure A.1): **(a)** A new or reoccurring shift in the past and no shift between the recent training and test window and **(b)** a shift that occurred in the past but reverted back to the "standard" regime. See Section A.1 in the Appendix for a detailed definition of the cases that we consider. Note that we do not address the case where the distribution shift occurs in the test window only, which could be more of a fundamental generalization problem or data change.

Adaptive training of a lightweight architecture To address the distribution shift cases defined in Section A.1, we introduce a discrete sampling distribution $p_{\boldsymbol{\phi}}(\cdot)$ over time steps (see Section A.2). In practice, $p_{\boldsymbol{\phi}}(\cdot)$ can be any discrete distribution that is well defined over integers, such as Geometric, Negative Binomial, Poisson, mixtures of these with multiple components, parametrized by $\boldsymbol{\phi}$. This distribution is responsible for sampling time windows that are used to train our forecasting model. Ideally, the learned distribution parameters $\boldsymbol{\phi}^\star$ will avoid sampling distribution shifts in history that do not correspond to the current regime. This distribution replaces uniform sampling over the history that cannot avoid distribution shifts. However, our model *should* fall back to uniform sampling if this results in lowest validation error because this suggests that no distribution shift is present. In practice, we would expect to frequently re-estimate $\boldsymbol{\phi}^\star$ to adapt to the current distribution. To address this, we propose a two-step procedure with a lightweight adaptive architecture that is fast to train and can be frequently re-trained [2]. **i)** The model is pre-trained in the standard way using uniform

[1] We introduce the notation in Section A.1

[2] In this sense, this approach is explicitly considering a notion of time-varying parameters.

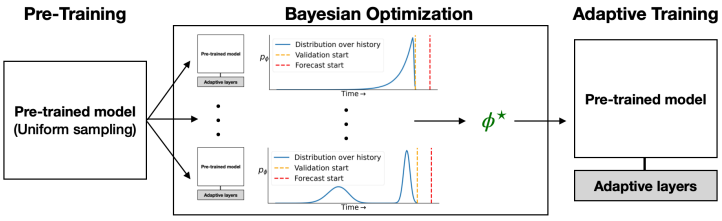

Figure 1: **Schematic description of the two-step procedure. i)** Pre-train the full model with uniform sampling; **ii.a)** use Bayesian optimization to find the optimal $\phi^\star$; **ii.b)** use $\phi^\star$ to train the lightweight architecture with adaptive sampling.

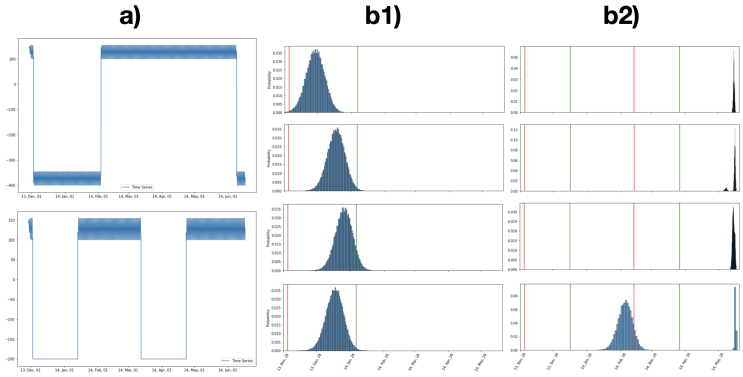

Figure 2: **Results obtained with the adaptive sampling mechanism and the lightweight adaptive architecture on synthetic data**. **a**) Synthetic data with injected shift (top) and noise (bottom). The learned sampling distribution over time steps is shown in **b1** for the top case and in **b2** for the bottom case. The red and green bars denote the start and end of a distribution shift and the histogram shows the learned distribution over time steps. **b1)** (top case in **a**): The distribution over time steps only focuses on the shifted region in the past (between the red and green bars) that is relevant for the current regime. **b2** (bottom case in **a**): The adaptive sampling distribution over time steps avoids the noise shifts (between the red and green bars) that are not relevant for the current forecast distribution.

sampling over the whole training history. Intuitively, this step aims at learning common features that are useful everywhere in time. **ii)** After freezing the weights of the pre-trained model, an adaptive lightweight architecture is attached and trained using the adaptive sampling mechanism described in Section A.2. The goal of this step is to learn time-specific features that are relevant to forecast in the current distribution. The lightweight architecture can be a short sequence of fully connected layers, or it can even be a subset of the pre-trained model.

Learn ϕ via Bayesian Optimization To avoid confusion, let us denote the parameters of the pre-trained model as $\mathbf{w}_{\text{pre}}$ and those of the lightweight architecture as $\mathbf{w}_{\text{ada}}$. A crucial observation at this point is that the adaptive sampling mechanism is not directly part of the model, which implies that we cannot take gradients with respect to ϕ and train $(\mathbf{w}_{\text{ada}}, \phi)$ jointly end-to-end. We propose to treat ϕ as hyper-parameters and use Bayesian optimization [11] to find the distribution over time steps. Equation 1 plays a key role in this step because the most recent region in the training set is used as validation channel to choose $\phi^\star$ during Bayesian optimization.

4 Experiments

The instantiation of our model in this manuscript is called `Ada-DeepAR` and uses `DeepAR` [31] as a base model with a two-layer LSTM. The adaptive part of our model is the second LSTM layer. Additionally, we perform Bayesian optimization using MOBSTER, which combines asynchronous successive halving and asynchronous Hyperband with gaussian-process based Bayesian optimization [19]. We use the implementation available in `Syne Tune` [32]. For any time series i with training history T and forecast horizon τ (see Section A.2 for detailed notation), we use $\mathbf{z}_{i,0:T-\tau}$ as a validation channel for Bayesian optimization. We consider a geometric and a mixture of two negative binomial distributions as our sampling distributions $p_{\phi}(\cdot)$. Training details are given in Section A.4.

Experiments on synthetic data: To demonstrate the applicability of our method, we first evaluate it on two synthetic experiments that mimic cases **a)** and **b)** of Figure A.1. We generated an artificial

Dataset	Ada-DeepAR (ours)	Ada-DeepAR (uniform)	CRPS RevIn	DeepAR	TFT
Covid-deaths	0.074 ± 0.013	0.088 ± 0.011	0.13 ± 0.006	0.103 ± 0.053	$\mathbf{0.053} \pm 0.061$
Electricity	$\mathbf{0.048} \pm 0.002$	0.054 ± 0.005	0.063 ± 0.005	0.062 ± 0.011	0.082 ± 0.000
Traffic	0.172 ± 0.002	$\mathbf{0.171} \pm 0.004$	0.172 ± 0.003	0.176 ± 0.004	0.306 ± 0.050
Solar	$\mathbf{0.370} \pm 0.003$	0.376 ± 0.003	0.376 ± 0.004	0.388 ± 0.007	0.562 ± 0.015
Taxi	0.286 ± 0.003	$\mathbf{0.278} \pm 0.001$	0.288 ± 0.004	0.279 ± 0.004	0.509 ± 0.001
Mean nCRPS	**0.190**	0.193	0.202	0.206	0.302
Mean Rank	**1.8**	**1.8**	3.4	3.4	4.2

Table 1: The CRPS loss (lower is better) for all datasets and models. We average over ten runs and report the mean and standard deviation. We took the results for TFT from TSBench [4], which only uses two runs.

sinusoidal dataset and then injected shifts or noise in different regions. **Focus on shifted regions distant in time:** In this example we injected a shift in a distant portion of the training dataset and also in the validation and test sets (see the top of panel **a)** in Figure 2). We show that the sampling distribution over time steps can focus on specific regions that are relevant to the current domain, regardless of when they occurred in the available history (see **b1)** in Figure 2). **Exclude noisy regions:** In this example, we injected noise (a negative constant) in two different regions of the training dataset. Our adaptive sampling mechanism avoids these noisy regions from the history (**b2)** in Figure 2). Interestingly, this sometimes corresponds to sampling only in the most recent window, while in other cases the distribution exploits its bimodality to focus on two distant windows.

Experiments on real-world datasets: Finally, we show preliminary results on five real-world datasets (see Section A.3). We compare the following models: Our two-step training prodcedure with adaptive sampling (`Ada-DeepAR`) and with uniform sampling (`Ada-DeepAR-uniform`), reversible instance normalization (`RevIn`) [18], and the attention-based Temporal Fusion Transformer `TFT` [22]. For `Ada-DeepAR`, `Ada-DeepAR-uniform` and `RevIn` we use `DeepAR`[31] as the base model. We evaluate these probabilistic methods using the continuous-ranked probability score (CRPS) [14, 25]. Table 1 summarizes the results. Averaged over all datasets, `Ada-DeepAR` improves the base model `DeepAR` by 8.4% (relative improvement). For `Covid-deaths` and `Electricity`, we observe significant improvements (39.2% and 29.2%). `Ada-DeepAR` selected the mixture of two negative binomial distributions for all datasets except `Covid-deaths` (for which the geometric distribution was selected). `RevIn` also improves over `DeepAR`, but with a much smaller margin. `TFT` performs best on `Covid-deaths`, but overall worse than the other methods on these datasets. `Ada-DeepAR-uniform` can be seen as an ablation of `Ada-DeepAR`, because it uses the two-step architecture but only uniform sampling. Interestingly, `Ada-DeepAR-uniform` also improves the error by 6.7% over its base model. In particular, `Ada-DeepAR-uniform` is the best model on the `Traffic` and `Taxi` datasets. Our hypothesis is that adaption to distribution shift is not necessary, and perhaps even counterproductive, for these datasets. This is also suggested by the performance of `RevIn`, which performs similar to `Ada-DeepAR` on these datasets. We would like to point out that a uniform distribution could be seen as one of the possible distributions $p_{\boldsymbol{\phi}}(\cdot)$ that our adaptive architecture should consider and if no other distribution reduces the validation error, the uniform distribution should be selected. However, further work is required to confirm whether our method behaves like this robustly. Further investigating two-step training approaches in forecasting could be another interesting research direction.

5 Conclusion

We introduced a model-agnostic adaptive sampling mechanism over the training history, whose parameters are learned using Bayesian optimization. Our experiments show that our method improves the forecasting error over its base model. While we find Bayesian optimization to be a powerful technique to find good adaptive sampling parameters, this approach is computationally expensive (because many parallel trials need to be evaluated) and could degrade as the number of parameters increases. Re-framing our approach as an input-selection layer that uses the re-parametrization trick for sampling could allow for cheaper end-to-end learning of the adaptive model parameters and $p_{\boldsymbol{\phi}}(\cdot)$.

References

[1] Alexander Alexandrov, Konstantinos Benidis, Michael Bohlke-Schneider, Valentin Flunkert, Jan Gasthaus, Tim Januschowski, Danielle C Maddix, Syama Sundar Rangapuram, David Salinas, Jasper Schulz, et al. Gluonts: Probabilistic and neural time series modeling in python. *J. Mach. Learn. Res.*, 21(116):1–6, 2020.

[2] Sercan O Arik, Nathanael C Yoder, and Tomas Pfister. Self-adaptive forecasting for improved deep learning on non-stationary time-series. *arXiv preprint arXiv:2202.02403*, 2022.

[3] Michael Bohlke-Schneider, Shubham Kapoor, and Tim Januschowski. Resilient neural forecasting systems. In *Proceedings of the Fourth International Workshop on Data Management for End-to-End Machine Learning*, DEEM'20, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450380232. doi: 10.1145/3399579.3399869. URL `https://doi.org/10.1145/3399579.3399869`.

[4] Oliver Borchert, David Salinas, Valentin Flunkert, Tim Januschowski, and Stephan Günnemann. Multi-objective model selection for time series forecasting, 2022. URL `https://arxiv.org/abs/2202.08485`.

[5] J. D. Croston. Forecasting and stock control for intermittent demands. *Journal of the Operational Research Society*, 23(3):289–303, Sep 1972. ISSN 1476-9360. doi: 10.1057/jors.1972.50. URL `https://doi.org/10.1057/jors.1972.50`.

[6] Pierpaolo De Blasi, Stefano Favaro, Antonio Lijoi, Ramsés H Mena, Igor Prünster, and Matteo Ruggiero. Are gibbs-type priors the most natural generalization of the dirichlet process? *IEEE transactions on pattern analysis and machine intelligence*, 37(2):212–229, 2013.

[7] Ensheng Dong, Hongru Du, and Lauren Gardner. An interactive web-based dashboard to track COVID-19 in real time. *The Lancet Infectious Diseases*, 20(5):533–534, May 2020. ISSN 1473-3099. doi: 10.1016/S1473-3099(20)30120-1. URL `https://doi.org/10.1016/S1473-3099(20)30120-1`. Publisher: Elsevier.

[8] Yuntao Du, Jindong Wang, Wenjie Feng, Sinno Pan, Tao Qin, Renjun Xu, and Chongjun Wang. Adarnn: Adaptive learning and forecasting of time series. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pages 402–411, 2021.

[9] Dheeru Dua and Casey Graff. UCI machine learning repository, 2017. URL `http://archive.ics.uci.edu/ml`.

[10] Thomas S Ferguson. A bayesian analysis of some nonparametric problems. *The annals of statistics*, pages 209–230, 1973.

[11] Peter I Frazier. A tutorial on bayesian optimization. *arXiv preprint arXiv:1807.02811*, 2018.

[12] Everette S Gardner Jr. Exponential smoothing: The state of the art. *Journal of forecasting*, 4(1): 1–28, 1985.

[13] Samuel J Gershman and David M Blei. A tutorial on bayesian nonparametric models. *Journal of Mathematical Psychology*, 56(1):1–12, 2012.

[14] Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007. doi: 10.1198/016214506000001437. URL `https://doi.org/10.1198/016214506000001437`.

[15] Rakshitha Godahewa, Christoph Bergmeir, Geoffrey I. Webb, Rob J. Hyndman, and Pablo Montero-Manso. Monash time series forecasting archive. In *Neural Information Processing Systems Track on Datasets and Benchmarks*, 2021.

[16] Tao Hong, Jingrui Xie, and Jonathan Black. Global energy forecasting competition 2017: Hierarchical probabilistic load forecasting. *International Journal of Forecasting*, 35(4):1389 – 1399, 2019.

[17] Jiayuan Huang, Arthur Gretton, Karsten Borgwardt, Bernhard Schölkopf, and Alex Smola. Correcting sample selection bias by unlabeled data. In B. Schölkopf, J. Platt, and T. Hoffman, editors, *Advances in Neural Information Processing Systems*, volume 19. MIT Press, 2006. URL https://proceedings.neurips.cc/paper/2006/file/a2186aa7c086b46ad4e8bf81e2a3a19b-Paper.pdf.

[18] Taesung Kim, Jinhee Kim, Yunwon Tae, Cheonbok Park, Jang-Ho Choi, and Jaegul Choo. Reversible instance normalization for accurate time-series forecasting against distribution shift. In *International Conference on Learning Representations*, 2021.

[19] Aaron Klein, Louis C. Tiao, Thibaut Lienart, Cedric Archambeau, and Matthias Seeger. Model-based asynchronous hyperparameter and neural architecture search, 2020. URL https://arxiv.org/abs/2003.10865.

[20] Guokun Lai, Wei-Cheng Chang, Yiming Yang, and Hanxiao Liu. Modeling long- and short-term temporal patterns with deep neural networks. *CoRR*, abs/1703.07015, 2017. URL http://arxiv.org/abs/1703.07015.

[21] Fengpei Li, Henry Lam, and Siddharth Prusty. Robust importance weighting for covariate shift. In Silvia Chiappa and Roberto Calandra, editors, *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 352–362. PMLR, 26–28 Aug 2020. URL https://proceedings.mlr.press/v108/li20b.html.

[22] Bryan Lim, Sercan Ö. Arık, Nicolas Loeff, and Tomas Pfister. Temporal fusion transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting*, 37(4): 1748–1764, 2021. ISSN 0169-2070. doi: https://doi.org/10.1016/j.ijforecast.2021.03.012. URL https://www.sciencedirect.com/science/article/pii/S0169207021000637.

[23] Linbo Liu, Youngsuk Park, Trong Nghia Hoang, Hilaf Hasson, and Jun Huan. Towards robust multivariate time-series forecasting: Adversarial attacks and defense mechanisms. *arXiv preprint arXiv:2207.09572*, 2022.

[24] Yucheng Lu, Youngsuk Park, Lifan Chen, Yuyang Wang, Christopher De Sa, and Dean Foster. Variance reduced training with stratified sampling for forecasting models. In *International Conference on Machine Learning*, pages 7145–7155. PMLR, 2021.

[25] James E. Matheson and Robert L. Winkler. Scoring rules for continuous probability distributions. *Management Science*, 22(10):1087–1096, 1976. ISSN 00251909, 15265501. URL http://www.jstor.org/stable/2629907.

[26] Srayanta Mukherjee, Devashish Shankar, Atin Ghosh, Nilam Tathawadekar, Pramod Kompalli, Sunita Sarawagi, and Krishnendu Chaudhury. ARMDN: Associative and recurrent mixture density networks for retail demand forecasting. *arXiv preprint arXiv:1803.03800*, 2018.

[27] Youngsuk Park, Kanak Mahadik, Ryan A Rossi, Gang Wu, and Handong Zhao. Linear quadratic regulator for resource-efficient cloud services. In *Proceedings of the ACM Symposium on Cloud Computing*, pages 488–489, 2019.

[28] Youngsuk Park, Ryan Rossi, Zheng Wen, Gang Wu, and Handong Zhao. Structured policy iteration for linear quadratic regulator. In *International Conference on Machine Learning*, pages 7521–7531. PMLR, 2020.

[29] Mihael Perman, Jim Pitman, and Marc Yor. Size-biased sampling of poisson point processes and excursions. *Probability Theory and Related Fields*, 92(1):21–39, 1992.

[30] Kashif Rasul, Abdul-Saboor Sheikh, Ingmar Schuster, Urs Bergmann, and Roland Vollgraf. Multivariate probabilistic time series forecasting via conditioned normalizing flows. *arXiv preprint arXiv:2002.06103*, 2020.

[31] David Salinas, Valentin Flunkert, Jan Gasthaus, and Tim Januschowski. Deepar: Probabilistic forecasting with autoregressive recurrent networks. *International Journal of Forecasting*, 36(3): 1181–1191, 2020.

[32] David Salinas, Matthias Seeger, Aaron Klein, Valerio Perrone, Martin Wistuba, and Cedric Archambeau. Syne tune: A library for large scale hyperparameter tuning and reproducible research. In *First Conference on Automated Machine Learning (Main Track)*, 2022.

[33] Harshit Saxena, Omar Aponte, and Katie T. McConky. A hybrid machine learning model for forecasting a billing period's peak electric load days. *International Journal of Forecasting*, 35 (4):1288 – 1303, 2019. ISSN 0169-2070. doi: https://doi.org/10.1016/j.ijforecast.2019.03.025. URL `http://www.sciencedirect.com/science/article/pii/S016920701930144X`.

[34] Masashi Sugiyama, Shinichi Nakajima, Hisashi Kashima, Paul Buenau, and Motoaki Kawanabe. Direct importance estimation with model selection and its application to covariate shift adaptation. In J. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc., 2007. URL `https://proceedings.neurips.cc/paper/2007/file/be83ab3ecd0db773eb2dc1b0a17836a1-Paper.pdf`.

[35] Elham Taghizadeh. Utilizing artificial neural networks to predict demand for weather-sensitive products at retail stores. *arXiv preprint arXiv:1711.08325*, 2017.

[36] NYC Taxi and Limousine Commission. TLC trip record data. `https://www1.nyc.gov/site/tlc/about/tlc-trip-record-data.page`, 2015.

[37] Olivia Wiles, Sven Gowal, Florian Stimberg, Sylvestre-Alvise Rebuffi, Ira Ktena, Krishnamurthy Dj Dvijotham, and Ali Taylan Cemgil. A fine-grained analysis on distribution shift. In *International Conference on Learning Representations*, 2022. URL `https://openreview.net/forum?id=Dl4LetuLdyK`.

[38] TaeHo Yoon, Youngsuk Park, Ernest K Ryu, and Yuyang Wang. Robust probabilistic time series forecasting. In *International Conference on Artificial Intelligence and Statistics*, pages 1336–1358. PMLR, 2022.

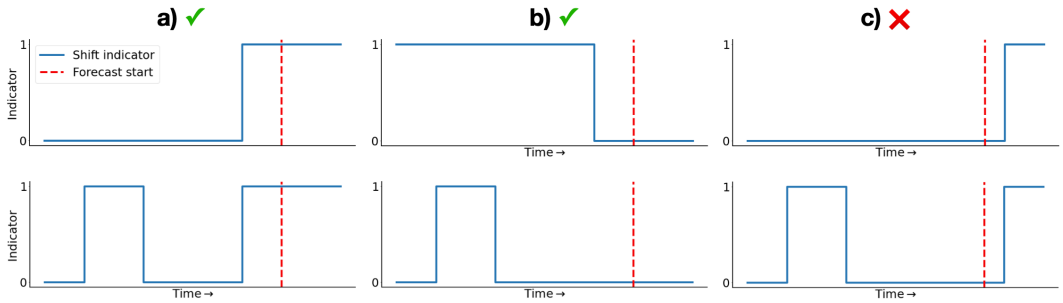

Figure A.1: **Equation 1 implies that we can handle all cases where there is no shift between the most recent region before the forecast start date and the forecast region itself.** Each frame shows an indicator function on the y-axis which is equal to 1 if a shift with respect to an assumed "standard" distribution is happening. Time is on the x-axis. The red dashed line separates the past (i.e., the training set) from the future (i.e. test set). **a)** New or reoccurring shift: the distribution over time steps should focus on the shifted regions. **b)** This case is complementary to the previous one: here the distribution over time steps should exclude the shifted regions. **c)** The shift occurs in the forecasting window but not in the most recent region before the forecast start date: our model needs to wait until part of the shift is in the training set in order to be able to adapt to it.

A Appendix

A.1 Scenarios of distribution shift occurrence

Equation 1 allows us to provide a precise breakdown of the specific cases of distribution shift that our model can handle. This is important because it determines the exact applicability of the work described in this paper. As Figure A.1 shows, we identify three domains:

New or reoccurring shift: When a distribution shift is new or reoccurring in the past and there is no shift between $t \in [T-c, T]$ and $t \in [T+1, T+\tau]$. In this case the adaptive sampling mechanism (see Section A.2) is expected to focus only on the shifted regions that are relevant for the current regime.

Reverted shift: When a shift occurred in the past but the distribution eventually reverted back to the "standard" regime. This case is complementary to the previous one: here the adaptive sampling mechanism provides a way to exclude the shifted regions. One can see the shifts in this case as regions with noisy or corrupted data: excluding them during training would guarantee robustness against noise.

Future shift: When the shift occurs in the forecasting window but not in the most recent region before the forecast start date. Our adaptive sampling mechanism cannot work if we have no information regarding the distribution shift happening in the test set (no free lunch).[3]

A.2 Background

Notation We denote the value of time series i at time t by $\mathbf{z}_{i,t}$, where $\mathbf{z} \in \mathbb{R}^k$, and corresponding time-varying covariates as $\mathbf{x}_{i,t} \in \mathbb{R}^d$.The past of each time series (i.e., training set) is denoted as $\mathbf{z}_{i,1:T} := (\mathbf{z}_{i,1}, \ldots, \mathbf{z}_{i,T})$, and the future (i.e., forecasting window or test set) as $\mathbf{z}_{i,T+1:T+\tau} := (\mathbf{z}_{i,T+1}, \ldots, \mathbf{z}_{i,T+\tau})$. The number of time series in the dataset is denoted by N.

Consider any forecasting algorithm $f_{\mathbf{w}}(\cdot)$, parameterized by weights $\mathbf{w}$, that models the conditional distribution[4] of a general time series instance $\mathbf{z}_{T+1}$ at a future time step given its past target $\mathbf{z}_{\leq T}$ and input covariate $\mathbf{x}_{\leq T}$) as:

$$p(\mathbf{z}_{T+1}|\mathbf{z}_{\leq T}, \mathbf{x}_{\leq T}) = f_{\mathbf{w}}(\mathbf{z}_{\leq T}, \mathbf{x}_{\leq T}). \tag{2}$$

[3]Note that in our setting we consider the forecasting window to be a mere continuation of the training set, which implies that the context window used at test time was part of training inputs. In this sense, our procedure cannot handle this case if and only if the shift happens solely in the forecasting range.

[4]Although we lay out our exposition using a probabilistic forecasting framework, there is technically no impediment in using our procedure with an algorithm that produces point forecasts.

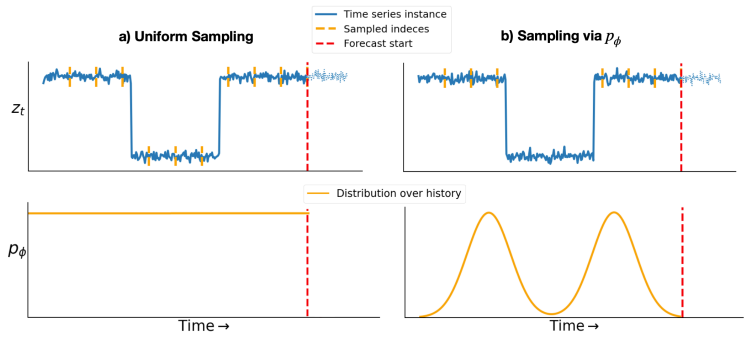

Figure A.2: **Adaptive sampling uses p_ϕ to sample training windows to come from distributions that are similar to the forecasting window.** Both panels show a time series instance (blue) at the top, with context and prediction windows (green) sampled during training, and the distribution over time steps at the bottom. **a)** Uniform sampling. **b)** Adaptive sampling.

	Freq.	Horizon	Number of Series	Avg. Length	Number of observations
`Covid-deaths`	D	30	227	182	41,314
`Electricity`	H	24	321	21,044	6,755,124
`Traffic`	H	48	862	17,496	15,081,552
`Solar`	H	24	137	7,009	960,233
`Taxi`	30 MIN	24	1,214	1,488	1,806,432

Table A.1: Dataset statistics for the used datasets, taken from Borchert et al. [4].

This conditional distribution can be modelled with a specific output distribution parameterized by $\boldsymbol{\psi}(\mathbf{z}_{\leq T}, \mathbf{x}_{\leq T})$ – see, e.g., [31] – or with flexible density estimators – see, e.g, [30] for a recent example using normalizing flows.

Sampling using a distribution over time steps Standard training in time series settings is usually done by sampling fixed-size windows *uniformly at random* over the available time series history (i.e., the past). Since distribution shifts happen with different (possibly reoccurring) patterns over time, uniform sampling over the training set would likely feed the algorithm with irrelevant or even misleading data to forecast in the current distribution. Therefore, we propose the following simple and intuitive idea: introduce a distribution $p_{\boldsymbol{\phi}}(\cdot)$ over time steps, to be learned in conjunction with $f_{\mathbf{w}}(\cdot)$, which allows the sampling mechanism to be adaptive with respect to the current domain. In practice, $p_{\boldsymbol{\phi}}(j)$ can be any discrete distribution that is well defined over $[0, T]$, such as Geometric, Negative Binomial, Poisson, mixtures of these with multiple components, and so on. Stochastic processes such as the Dirichlet process [10, 13], the Pitman-Yor process [29] and their generalizations [6] represent a flexible alternative, since they naturally induce a partition of the space where the number of components is automatically determined.

A.3 Datasets

We use datasets from the Monash Time Series Forecasting Repository [15] and GluonTS [1] for our experiments. We use the same train/test splits as described in TSBench [4]. See Table A.1 for dataset statistics. The frequencies of our used datasets are 30 minutes (30 MIN), hourly (H), and daily (D).

In particular, we use the following datasets:

`Covid-deaths`: Daily COVID-19 deaths of different countries between January 22 and August 21 2020 [7].

`Electricity`: Hourly household consumption data between January 2012 and June 2014 [9].

`Traffic`: Hourly occupancy rates of freeways in the San Francisco Bay area between 2015 and 2017 [9].

`Solar`: Hourly power output of 137 photovoltaic power stations in the US [20].

`Taxi`: Number of taxi rides in different locations in New York sampled at 30 minute windows [36].

A.4 Training Details

The parameters of the distribution over history are optimized via Bayesian Optimization [11] using the recent scalable implementation provided in `Syne-Tune` [32]. We leveraged an AWS EC2 instance with ≈ 42 cores and set the number of random initializations to 44 to allow for sufficient exploration of the parameter space. That is, the optimizer starts with 44 trials, each with a different parameter configuration, and sequentially stops some of them or instantiate new ones given the information from the Gaussian Process posterior. The optimizer stopping criterion was chosen to be the maximum number of trials taken to completion, which was set to 200. We used the CRPS used to decide which trials were stopped or continued was the mean quantile loss over the validation set.

The lightweight architecture is supposed to be simple enough to allow frequent re-training through adaptive sampling, so that the forecasting model can quickly adapt to new distributions. Clearly, different architectures will work better with different datasets. In all our experiments with `DeepAR` [31], we found that using the default two-layer LSTM with 40 hidden units for the pre-trained model, and freezing the first layer only for adaptive training yielded the best results. To clarify, this means that $\mathbf{w}_{ada} \subset \mathbf{w}_{pre}$. Both pre-training and adaptive training were done over 100 epochs with a small dropout probability equal to 0.1.

For `TFT`, we selected the hyperparameter setting from the TSBench [4] that resulted in lowest CRPS loss on the validation set for each dataset. The hyperparameters tuned in TSBench are the context length, the number of hidden states, and the number of heads.

For the `DeepAR` and `TFT` models, we use the implementation available in GluonTS [1].