

TIMEFLOW: AN IMPLICIT NEURAL REPRESENTATION APPROACH FOR CONTINUOUS TIME SERIES MODELING

Etienne Le Naour^{* 1,2}, **Louis Serrano**^{* 1}, **Léon Migus**^{* 1,3}, **Yuan Yin**¹, **Ghislain Agoua**²
Nicolas Baskiotis¹, **Patrick Gallinari**^{1,4}, **Vincent Guigue**⁵

¹ Sorbonne Université, CNRS, ISIR, 75005 Paris, France

² EDF R&D, Palaiseau, France

³ Sorbonne Université, CNRS, Laboratoire Jacques-Louis Lions, 75005 Paris, France

⁴ Criteo AI Lab, Paris, France

⁵ AgroParisTech, Palaiseau, France

{*louis.serrano, leon.migus, yuan.yin, nicolas.baskiotis, vincent.guigue*}@sorbonne-universite.fr

{*etienne.le-naour, ghislain.agoua*}@edf.fr

ABSTRACT

We introduce a novel modeling approach for time series imputation and forecasting, tailored to address the challenges often encountered in real-world data, such as irregular samples, missing data, or unaligned measurements from multiple sensors. Our method relies on a continuous-time-dependent model of the series’ evolution dynamics. It leverages adaptations of conditional, implicit neural representations for sequential data. A modulation mechanism, driven by a meta-learning algorithm, allows adaptation to unseen samples and extrapolation beyond observed time-windows for long-term predictions. The model provides a highly flexible and unified framework for imputation and forecasting tasks across a wide range of challenging scenarios. It achieves state-of-the-art performance on classical benchmarks and outperforms alternative time-continuous models.

1 INTRODUCTION

Time series analysis and modeling are ubiquitous in a wide range of fields, including industry, climate science and health. The variety, heterogeneity and increasing number of deployed sensors, raise new challenges when dealing with real-world problems for which current methods often fail. For example, data are frequently irregularly sampled, contain missing values, or are unaligned when collected from distributed sensors (Schulz and Stattegger, 1997; Clark and Bjørnstad, 2004). Recent advancements in deep learning have significantly improved state-of-the-art performance in both data imputation (Cao et al., 2018; Du et al., 2023) and forecasting tasks (Zeng et al., 2022; Nie et al., 2022). Many state-of-the-art models, such as transformers, have been primarily designed for dense and regular grids (Wu et al., 2021; Nie et al., 2022; Du et al., 2023). They struggle to handle irregular data and often suffer from significant performance degradation (Chen et al., 2001; Kim et al., 2019).

Our objective is to explore alternatives to SOTA transformers able to handle, in a unified framework, imputation and forecasting tasks for irregularly, arbitrarily sampled, and unaligned time series sources. Time-dependent continuous models (Rasmussen and Williams, 2006; Garnelo et al., 2018; Rubanova et al., 2019) offer such an alternative. However, until now, their performance has lagged significantly behind that of models designed for regular discrete grids. A few years ago, implicit neural representations (INRs) emerged as a powerful tool for representing images as continuous functions of spatial coordinates (Sitzmann et al., 2020; Tancik et al., 2020) with recent new applications such as image generation (Dupont et al., 2022) or even modeling dynamical systems (Yin et al., 2023). A more detailed related content is available in Appendix B.

* Equal contribution

In this work, we leverage the potential of conditional INR models within a meta-learning approach to introduce TimeFlow: a unified framework designed for modeling continuous time series and addressing imputation and forecasting tasks with irregular and unaligned observations. Our key contributions are the following:

- We propose a novel framework that excels in modeling time series as continuous functions of time, accepting arbitrary time step inputs, thus enabling the handling of irregular and unaligned time series for both imputation and forecasting tasks.
- We conducted an extensive comparison with state-of-the-art continuous and discrete models. It demonstrates that our approach outperforms continuous and discrete SOTA deep learning approaches for imputation. As for long-term forecasting, it outperforms existing continuous models both on regular and irregular samples. It is on par with SOTA discrete models on regularly sampled time series while allowing for a much greater flexibility for irregular samplings, allowing to cope with situations where discrete models fail.

2 THE TIMEFLOW FRAMEWORK

2.1 PROBLEM SETTING

We aim to develop a unified framework for time series imputation and forecasting that reduces dependency on a fixed sampling scheme for time series. We introduce the following notations for both tasks. During training, in the imputation setting, we have access to time series in an observation set denoted as \mathcal{T}_{in} , which is a subset of the complete time series observation set \mathcal{T} . In the forecasting setting, we observe time series within a limited past time grid, referred to as the 'look-back window' and denoted as \mathcal{T}_{in} (a subset of \mathcal{T}), as well as a future grid, the 'horizon', denoted as \mathcal{T}_{out} (also a subset of \mathcal{T}). At test time, in both cases, and given an observed subset \mathcal{T}_{in}^* included in a possibly new temporal window \mathcal{T}^* , our objective is to infer the time series values within \mathcal{T}^* .

2.2 KEY COMPONENTS

Our framework is articulated around three key components. (i) **INR-based time-continuous functions**: a time series x is represented by a time-continuous function $f: t \in \mathbb{R}_+ \rightarrow f(t) \in \mathbb{R}^d$ that can be queried at any time t . For that, we employ implicit neural representations (INRs), which are neural networks capable of learning a parameterized continuous function f_θ from discrete data by minimizing the reconstruction loss between observed data and network's outputs. (ii) **Conditional INRs with modulations**: An INR can represent only one function, whether it's an image or a time series. To effectively represent a collection of time series $(x^{(j)})_j$ using INRs, we improve their encoding by incorporating per-sample modulations, which we denote as $\psi^{(j)}$. These modulations condition the parameters θ of the INRs. We use the notation $f_{\theta, \psi^{(j)}}$ to refer to the conditioned INR with the modulations $\psi^{(j)}$. (iii) **Optimization-based encoding**: the conditioning modulation parameters $\psi^{(j)}$ are calculated as a function of codes $z^{(j)}$ that represent the individual sample series. We acquire these codes $z^{(j)}$ through a meta-learning optimization process using an auto-decoding strategy. Notably, auto-decoding has been found to be more efficient for this purpose than set encoders (Kim et al., 2019). In the following sections, we will elaborate on each component of our method.

INR-based time-continuous functions. We implement our INR with Fourier features and a feed-forward network (FFN) with ReLU activations, i.e. for a time coordinate $t \in \mathcal{T}$, the output of the INR f_θ is given by $f_\theta(t) = \text{FFN}(\gamma(t))$. The Fourier Features $\gamma(\cdot)$ are a frequency embedding of the time coordinates used to capture high-frequencies (Tancik et al., 2020; Mildenhall et al., 2021). In our case, we chose $\gamma(t) := (\sin(\pi t), \cos(\pi t), \dots, \sin(2^{N-1}\pi t), \cos(2^{N-1}\pi t))$, with N the number of fixed frequencies. For an INR with L layers, the output is computed as follows: (i) we get the frequency embedding $\phi_0 = \gamma(t)$, (ii) we update the hidden states according to $\phi_l = \text{ReLU}(\theta_l \phi_{l-1} + b_l)$ for $l = 1, \dots, L$, (iii) we project onto the output space $f_\theta(t) = \theta_{L+1} \phi_L + b_{L+1}$.

Conditional INRs with modulations. As indicated, sample conditioning of the INR is performed through modulations of its parameters. In order to adapt rapidly the model to new samples, the con-

ditioning should rely only on a small number of the INR parameters. This is achieved by modifying only the biases of the INR through the introduction of an additional bias term $\psi_l^{(j)}$ for each layer l , also known as *shift modulation*. To further limit the versatility of the conditioning, we generate the instance modulations $\psi^{(j)}$ from compact codes $z^{(j)}$ through a linear hypernetwork h with parameters w , i.e., $\psi^{(j)} = h_w(z^{(j)})$. Consequently, the approximation of a time series $x^{(j)}$, denoted globally as $f_{\theta, h_w}(z^{(j)})$, will depend on shared parameters θ and w that are common among all the INRs involved in modeling the series family and on the code $z^{(j)}$ specific to series $x^{(j)}$. The output of the l -th layer of the modulated INR is given by $\phi_l = \text{ReLU}(\theta_l \phi_{l-1} + b_l + \psi_l^{(j)})$, where $\psi_l^{(j)} = W_l z^{(j)}$, and $w := (W_l)_{l=1}^L$ are the parameters of the hypernetwork h_w . This design enables gathering information across samples into the common parameters of the INR and hypernetwork, while the codes contain only specific information about their respective time-series samples. The architecture is illustrated in Figure 1 (Appendix C).

Optimization-based encoding. We condition the INR using the data from \mathcal{T}_{in} , and learn the shared INR and hypernetwork parameters θ and w using \mathcal{T}_{in} for both imputation and forecasting, and \mathcal{T}_{out} for forecasting only. We achieve the conditioning on \mathcal{T}_{in} by optimizing the codes $z^{(j)}$ through gradient descent. The joint optimization of the codes and common parameters is challenging. In TimeFlow, it is achieved through a meta-learning approach, adapted from Dupont et al. (2022) and Zintgraf et al. (2019). The objective is to learn shared parameters so that the code $z^{(j)}$ can be adapted in just a few gradient steps for a new series $x^{(j)}$. For training, we perform parameter optimization at two levels: the inner-loop and the outer-loop. The inner-loop adapts the code $z^{(j)}$ to condition the network on the set $\mathcal{T}_{in}^{(j)}$, while the outer-loop updates the common parameters using $\mathcal{T}_{in}^{(j)}$ and also $\mathcal{T}_{out}^{(j)}$ for forecasting. We present our training optimization in Algorithm 1 (Appendix D). At each training epoch and for each batch of data \mathcal{B} composed of time series $x^{(j)}$ sampled from the training set, we first update individually the codes $z^{(j)}$ in the inner loop, before updating the common parameters in the outer loop using a loss over the whole batch. We introduce a parameter λ to weight the importance of the loss over \mathcal{T}_{out} w.r.t. the loss over \mathcal{T}_{in} for the outer-loop. In practice, when \mathcal{T}_{out} exists, i.e. for forecasting, we set $\lambda = 1$ and $\lambda = 0$ otherwise. We use an MSE loss over the observations grid $\mathcal{L}_{\mathcal{T}}(x_t, \tilde{x}_t) := \mathbb{E}_{t \sim \mathcal{T}}[(x_t - \tilde{x}_t)^2]$. We denote α and η the learning rates of the inner- and outer-loop. Using $K = 3$ steps for training and testing is sufficient for our experiments thanks to the use of second-order meta-learning as explained in Appendix G.

2.3 TIMEFLOW INFERENCE

During the inference process, we aim to infer the time series value for each timestamp in the dense grid $\mathcal{T}^{*(j)}$ based on the partial observation grid $\mathcal{T}_{in}^{*(j)} \subset \mathcal{T}^{*(j)}$. We can encounter two scenarios: (i) One where we observe the same time window as during training ($\mathcal{T}^{*(j)} = \mathcal{T}^{(j)}$) as in the imputation setting in Section 3. (ii) One, where we are dealing with a newly observed time window ($\mathcal{T}^{*(j)} \neq \mathcal{T}^{(j)}$), as in the forecasting setting in Appendix F.1. At inference, the parameters θ and w are kept fixed to their final training values. We optimize the individual parameters $z^{*(j)}$ based on the newly observed grid $\mathcal{T}_{in}^{*(j)}$ using the K inner-steps of the meta-learning algorithm as described in Algorithm 2 (Appendix D). We are then in position to query $f_{\theta, h_w}(z^{*(j)})(t)$ for any given timestamp $t \in \mathcal{T}^{*(j)}$.

3 EXPERIMENTS: A FOCUS ON IMPUTATION

We conducted a comprehensive evaluation of our TimeFlow framework across three different tasks, comparing its performance to state-of-the-art continuous and discrete baseline methods. In the core of the manuscript, we assess TimeFlow’s capabilities to impute sparsely observed time series under various sampling rates. In the appendixes, Appendix F.1 focuses on long-term forecasting, where we evaluate TimeFlow over standard long-term forecasting horizons. In Appendix F.2, we tackle a challenging task forecasting with incomplete look-back windows, thus combining the challenges of imputation and forecasting. This demonstrates TimeFlow’s versatility and performance.

Datasets. We tested our framework on three extensive multivariate datasets where a single phenomenon is measured at multiple locations over time, namely *Electricity*, *Traffic* and *Solar*. They are commonly used in the time series literature and are described in Appendix I.

Imputation setting. We consider the classical imputation setting where n time series are partially observed over a given time window. Using our approach, we can predict for each time series the value at any timestamp t in that time window based on partial observations.

For a time series $x^{(j)}$, we denote the set of observed points as $\mathcal{T}_{in}^{(j)}$ and the ground truth set of points as $\mathcal{T}^{(j)}$. The observed time grids may be irregularly spaced and may differ across the different time series ($\mathcal{T}_{in}^{(j_1)} \neq \mathcal{T}_{in}^{(j_2)}, \forall j_1 \neq j_2$). The model is trained for each $x^{(j)}$ following Algorithm 1. Then, we aim to infer for any unobserved $t \in \mathcal{T}^{(j)}$ the missing value $x_t^{(j)}$ conditioned on $\mathcal{T}_{in}^{(j)}$ according to Algorithm 2. For this imputation task, the TimeFlow training and inference procedures are detailed in Section 2 and illustrated in Figure 2. For comparison with the SOTA imputation discrete baselines, we assume that the ground truth time grid is the same for each sample. The subsampling rate τ is defined as the rate of observed values.

Baselines. We compare TimeFlow with various baselines, including discrete imputation methods, such as CSDI (Tashiro et al., 2021), SAITS (Du et al., 2023), BRITS (Cao et al., 2018), and TIDER (Liu et al., 2023), and continuous ones, such as Neural Process (NP, Garnelo et al., 2018), mTAN (Shukla and Marlin, 2021), and DeepTime with slight adjustments (Woo et al., 2022) (details cf. Appendix E.5). For each dataset, we divide the series into five independent time windows (consisting of 2000 timestamps for *Electricity* and *Traffic*, and 10,000 timestamps for *Solar*), perform imputation on each time window and average the performance to obtain robust results. We evaluate the quality of the models for different subsampling rates, from the easiest $\tau = 0.5$ to the most difficult $\tau = 0.05$.

Table 1: Mean MAE imputation results on the missing grid only. Each time series is divided into 5 time windows onto which imputation is performed, and the performances are averaged over the 5 windows. In the table, τ stands for the subsampling rate, i.e. the proportion of observed points considered for each time window. Bold results are best, underlined results are second best. TimeFlow improvement represents the overall percentage improvement achieved by TimeFlow compared to the specific method being considered (Appendix I).

	τ	Continuous methods				Discrete methods			
		TimeFlow	DeepTime	mTAN	Neural Process	CSDI	SAITS	BRITS	TIDER
Electricity	0.05	0.324 ± 0.013	0.379 ± 0.037	0.575 ± 0.039	0.357 ± 0.015	0.462 ± 0.021	0.384 ± 0.019	<u>0.329 ± 0.015</u>	0.427 ± 0.010
	0.10	0.250 ± 0.010	0.333 ± 0.034	0.412 ± 0.047	0.417 ± 0.057	0.398 ± 0.072	0.308 ± 0.011	<u>0.287 ± 0.015</u>	0.399 ± 0.009
	0.20	0.225 ± 0.008	<u>0.244 ± 0.013</u>	0.342 ± 0.014	0.320 ± 0.017	0.341 ± 0.068	0.261 ± 0.008	<u>0.245 ± 0.011</u>	0.391 ± 0.010
	0.30	0.212 ± 0.007	0.240 ± 0.014	0.335 ± 0.015	0.300 ± 0.022	0.277 ± 0.059	0.236 ± 0.008	<u>0.221 ± 0.008</u>	0.384 ± 0.009
	0.50	0.194 ± 0.007	0.227 ± 0.012	0.340 ± 0.022	0.297 ± 0.016	0.168 ± 0.003	0.209 ± 0.008	<u>0.193 ± 0.008</u>	0.386 ± 0.009
Solar	0.05	0.095 ± 0.015	0.190 ± 0.020	0.241 ± 0.102	<u>0.115 ± 0.015</u>	0.374 ± 0.033	0.142 ± 0.016	0.165 ± 0.014	0.291 ± 0.009
	0.10	0.083 ± 0.015	0.159 ± 0.013	0.251 ± 0.081	<u>0.114 ± 0.014</u>	0.375 ± 0.038	0.124 ± 0.018	0.132 ± 0.015	0.276 ± 0.010
	0.20	0.072 ± 0.015	0.149 ± 0.020	0.314 ± 0.035	0.109 ± 0.016	0.217 ± 0.023	<u>0.108 ± 0.014</u>	0.109 ± 0.012	0.270 ± 0.010
	0.30	0.061 ± 0.012	0.135 ± 0.014	0.338 ± 0.05	0.108 ± 0.016	0.156 ± 0.002	<u>0.100 ± 0.015</u>	<u>0.098 ± 0.012</u>	0.266 ± 0.010
	0.50	0.054 ± 0.013	0.098 ± 0.013	0.315 ± 0.080	0.107 ± 0.015	<u>0.079 ± 0.011</u>	0.094 ± 0.013	<u>0.088 ± 0.013</u>	0.262 ± 0.009
Traffic	0.05	0.283 ± 0.016	0.246 ± 0.010	0.406 ± 0.074	0.318 ± 0.014	0.337 ± 0.045	0.293 ± 0.007	<u>0.261 ± 0.010</u>	0.363 ± 0.007
	0.10	0.211 ± 0.012	<u>0.214 ± 0.007</u>	0.319 ± 0.025	0.288 ± 0.018	0.288 ± 0.017	0.237 ± 0.006	0.245 ± 0.009	0.362 ± 0.006
	0.20	0.168 ± 0.006	0.216 ± 0.006	0.270 ± 0.012	0.271 ± 0.011	0.269 ± 0.017	<u>0.197 ± 0.005</u>	0.224 ± 0.008	0.361 ± 0.006
	0.30	0.151 ± 0.007	<u>0.172 ± 0.008</u>	0.251 ± 0.006	0.259 ± 0.012	0.240 ± 0.037	0.180 ± 0.006	0.197 ± 0.007	0.355 ± 0.006
	0.50	0.139 ± 0.007	0.171 ± 0.005	0.278 ± 0.040	0.240 ± 0.021	<u>0.144 ± 0.022</u>	0.160 ± 0.008	0.161 ± 0.060	0.354 ± 0.007
TimeFlow improvement	/	24.14 %	50.53 %	31.61 %	36.12 %	20.33 %	18.90 %	53.40 %	

Results. In Table 1, we show that TimeFlow outperforms discrete and continuous models across almost all τ 's for the given datasets. The relative improvements of TimeFlow over baselines are significant, ranging from 15% to 50%. We highlight that TimeFlow outperforms all discrete baselines while inheriting the advantages of continuous modeling, such as being grid-independent and able to deal with unaligned sensors. Qualitatively, we present visual imputation results in Figure 3 (Appendix E.2) where our model demonstrates impressive imputation capabilities on the *Electricity* dataset compare to the best baseline, BRITS.

4 CONCLUSION

We have introduced a unified framework for continuous time series modeling leveraging conditional INR and meta-learning. Our experiments have demonstrated superior performance compared to other continuous methods, and better or comparable results to SOTA discrete methods. One of the standout features of our framework is its inherent continuity and the ability to modulate the INR parameters. This unique flexibility lets TimeFlow effectively tackle a wide array of challenges, including forecasting in the presence of missing values, accommodating irregular time steps, and extending the trained model’s applicability to previously unseen time series and new time windows. Our empirical results have shown TimeFlow’s effectiveness in handling homogeneous multivariate time series. As a logical next step, extending TimeFlow’s capabilities to address heterogeneous multivariate phenomena represents a promising direction for future research.

REFERENCES

- M. Bilos, K. Rasul, A. Schneider, Y. Nevmyvaka, and S. Günnemann. Modeling temporal data as continuous functions with stochastic process diffusion. In A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett, editors, *International Conference on Machine Learning, ICML*, volume 202 of *Proceedings of Machine Learning Research*, pages 2452–2470. PMLR, 2023.
- E. D. Brouwer, J. Simm, A. Arany, and Y. Moreau. Gru-ode-bayes: continuous modeling of sporadically-observed time series. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pages 7379–7390, 2019.
- W. Cao, D. Wang, J. Li, H. Zhou, Y. Li, and L. Li. Brits: bidirectional recurrent imputation for time series. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 6776–6786, 2018.
- H. Chen, S. Grant-Muller, L. Mussone, and F. Montgomery. A study of hybrid neural network approaches and the effects of missing data on traffic forecasting. *Neural Computing & Applications*, 10:277–286, 2001.
- Y. G. Cinar, H. Mirisaee, P. Goswami, É. Gaussier, and A. Aït-Bachir. Period-aware content attention rnns for time series forecasting with missing values. *Neurocomputing*, 312:177–186, 2018.
- J. S. Clark and O. N. Bjørnstad. Population time series: process variability, observation errors, missing values, lags, and hidden states. *Ecology*, 85(11):3140–3150, 2004.
- G. Corani, A. Benavoli, and M. Zaffalon. Time series forecasting with gaussian processes needs priors. In *Machine Learning and Knowledge Discovery in Databases. Applied Data Science Track: European Conference, ECML PKDD 2021, Bilbao, Spain, September 13–17, 2021, Proceedings, Part IV 21*, pages 103–117. Springer, 2021.
- W. Du, D. Côté, and Y. Liu. Saits: Self-attention-based imputation for time series. *Expert Systems with Applications*, 219:119619, 2023.
- E. Dupont, H. Kim, S. M. A. Eslami, D. J. Rezende, and D. Rosenbaum. From data to functa: Your data point is a function and you can treat it like one. In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 5694–5725. PMLR, 2022.
- R. Fathony, A. K. Sahu, D. Willmott, and J. Z. Kolter. Multiplicative filter networks. In *International Conference on Learning Representations*, 2021.
- C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pages 1126–1135. PMLR, 2017.
- E. Fons, A. Sztrajman, Y. El-Laham, A. Iosifidis, and S. Vyetrenko. Hypertime: Implicit neural representation for time series. *CoRR*, abs/2208.05836, 2022.

- V. Fortuin, D. Baranchuk, G. Rätsch, and S. Mandt. Gp-vae: Deep probabilistic time series imputation. In *International conference on artificial intelligence and statistics*, pages 1651–1661. PMLR, 2020.
- M. Garnelo, D. Rosenbaum, C. Maddison, T. Ramalho, D. Saxton, M. Shanahan, Y. W. Teh, D. J. Rezende, and S. M. A. Eslami. Conditional neural processes. In *Proceedings of the 35th International Conference on Machine Learning, ICML*, volume 80, pages 1690–1699. PMLR, 2018.
- L. Huang and T. Hoeffler. Compressing multidimensional weather and climate data into neural networks. In *International Conference on Learning Representations, ICLR, 2023*.
- K. Jeong and Y. Shin. Time-series anomaly detection with implicit neural representation. *CoRR*, abs/2201.11950, 2022.
- T. Kim, W. Ko, and J. Kim. Analysis and impact evaluation of missing data imputation in day-ahead pv generation forecasting. *Applied Sciences*, 9(1):204, 2019.
- S. Liu, H. Yu, C. Liao, J. Li, W. Lin, A. X. Liu, and S. Dustdar. Pyraformer: Low-complexity pyramidal attention for long-range time series modeling and forecasting. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*, 2022.
- S. Liu, X. Li, G. Cong, Y. Chen, and Y. Jiang. Multivariate time-series imputation with disentangled temporal representations. In *The Eleventh International Conference on Learning Representations, ICLR, 2023*.
- Y. Liu, R. Yu, S. Zheng, E. Zhan, and Y. Yue. Naomi: Non-autoregressive multiresolution sequence imputation. *Advances in neural information processing systems*, 32, 2019.
- Y. Luo, X. Cai, Y. Zhang, J. Xu, et al. Multivariate time series imputation with generative adversarial networks. *Advances in neural information processing systems*, 31, 2018.
- Y. Luo, Y. Zhang, X. Cai, and X. Yuan. E2gan: End-to-end generative adversarial network for multivariate time series imputation. In *Proceedings of the 28th international joint conference on artificial intelligence*, pages 3094–3100. AAAI Press, 2019.
- B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1): 99–106, 2021.
- A. Nichol, J. Achiam, and J. Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018.
- Y. Nie, N. H. Nguyen, P. Sinthong, and J. Kalagnanam. A time series is worth 64 words: Long-term forecasting with transformers. *CoRR*, abs/2211.14730, 2022.
- C. E. Rasmussen and C. K. I. Williams. *Gaussian processes for machine learning*. Adaptive computation and machine learning. MIT Press, 2006.
- Y. Rubanova, R. T. Q. Chen, and D. Duvenaud. Latent odes for irregularly-sampled time series. *CoRR*, abs/1907.03907, 2019.
- M. Schulz and K. Stattegger. Spectrum: Spectral analysis of unevenly spaced paleoclimatic time series. *Computers & Geosciences*, 23(9):929–945, 1997.
- S. N. Shukla and B. M. Marlin. Multi-time attention networks for irregularly sampled time series. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*, 2021.
- V. Sitzmann, J. N. P. Martel, A. W. Bergman, D. B. Lindell, and G. Wetzstein. Implicit neural representations with periodic activation functions. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.

- M. Tancik, P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. Barron, and R. Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems*, 33:7537–7547, 2020.
- X. Tang, H. Yao, Y. Sun, C. C. Aggarwal, P. Mitra, and S. Wang. Joint modeling of local and global temporal dynamics for multivariate time series forecasting with missing values. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI*, pages 5956–5963. AAAI Press, 2020.
- Y. Tashiro, J. Song, Y. Song, and S. Ermon. Csd: Conditional score-based diffusion models for probabilistic time series imputation. *Advances in Neural Information Processing Systems*, 34: 24804–24816, 2021.
- G. Woo, C. Liu, D. Sahoo, A. Kumar, and S. C. H. Hoi. Deeptime: Deep time-index meta-learning for non-stationary time-series forecasting. *CoRR*, abs/2207.06046, 2022.
- H. Wu, J. Xu, J. Wang, and M. Long. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. In M. Ranzato, A. Beygelzimer, Y. N. Dauphin, P. Liang, and J. W. Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 22419–22430, 2021.
- Y. Yin, M. Kirchmeyer, J.-Y. Franceschi, A. Rakotomamonjy, and P. Gallinari. Continuous pde dynamics forecasting with implicit neural representations. In *International Conference on Learning Representations, ICLR*, 2023.
- A. Zeng, M. Chen, L. Zhang, and Q. Xu. Are transformers effective for time series forecasting? *CoRR*, abs/2205.13504, 2022.
- H. Zhou, S. Zhang, J. Peng, S. Zhang, J. Li, H. Xiong, and W. Zhang. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, pages 11106–11115, 2021.
- T. Zhou, Z. Ma, X. Wang, Q. Wen, L. Sun, T. Yao, W. Yin, and R. Jin. Film: Frequency improved legendre memory model for long-term time series forecasting. *CoRR*, abs/2205.08897, 2022a.
- T. Zhou, Z. Ma, Q. Wen, X. Wang, L. Sun, and R. Jin. Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. In K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvári, G. Niu, and S. Sabato, editors, *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 27268–27286. PMLR, 2022b.
- L. Zintgraf, K. Shiarli, V. Kurin, K. Hofmann, and S. Whiteson. Fast context adaptation via meta-learning. In *International Conference on Machine Learning*, pages 7693–7702. PMLR, 2019.

A REPRODUCTIBILITY STATEMENT

Our work is entirely reproducible, and all the references to the information in order to reproduce it are in this section.

Data. The dataset description, processing and normalization are presented in Appendix I.

Model. The model and the training details are presented in Section 2 and the hyperparameter selection is available in Appendix H.1.

GPU. We used NVIDIA TITAN RTX 24Go single GPU to conduct all the experiments for our method, which is coded in PyTorch (Python 3.9.2).

B RELATED CONTENT

Discrete methods for time series imputation and forecasting. Recently, Deep Learning (DL) methods have been widely used for both time series imputation and forecasting. For imputation, BRITS (Cao et al., 2018) uses a bidirectional recurrent neural network (RNN). Alternative frameworks were later explored, e.g., GAN-based (Luo et al., 2018; 2019; Liu et al., 2019), VAE-based (Fortuin et al., 2020), diffusion-based (Tashiro et al., 2021), matrix factorization-based (TIDER, Liu et al., 2023) and transformer-based (SAITS, Du et al., 2023) approaches. These methods cannot handle irregular time series. In situations involving multiple sensors, such as those placed at different locations, incorporating new sensors necessitates retraining the entire model, thereby limiting their usability. For forecasting, most recent DL SOTA models are based on transformers. Initial approaches apply plain transformers directly to the series, each token being a series element (Zhou et al., 2021; Liu et al., 2022; Wu et al., 2021; Zhou et al., 2022b). These transformers may underperform linear models as shown in (Zeng et al., 2022). PatchTST (Nie et al., 2022) significantly improved transformers SOTA performance by considering sub-series as tokens of the series. However, all these models cannot handle properly irregularly sampled look-back windows.

Continuous methods for time series. Gaussian Processes (Rasmussen and Williams, 2006) have been a popular family of methods for modeling time series as continuous functions. They require choosing an appropriate kernel (Corani et al., 2021) and may suffer limitations in large dimensions settings. Neural Processes (NPs) (Garnelo et al., 2018; Kim et al., 2019) parameterize Gaussian processes through an encoder-decoder architecture leading to more computationally efficient implementations. NPs have been used to model simple signals for imputation and forecasting tasks, but struggle with more complex signals. Bilos et al. (2023) parameterizes a Gaussian Process through a diffusion model, but the model has difficulty adapting to a large number of timestamps. Other approaches such as Brouwer et al. (2019) and Rubanova et al. (2019) model time series continuously with latent ordinary differential equations. mTAN (Shukla and Marlin, 2021) a transformer model uses an attention mechanism to impute irregular time series. While these approaches have shown significant progress in continuous modeling for time series, their performances on regularly spaced grids are inferior compared to the aforementioned discrete models and they lack extrapolation capability when dealing with complex dynamics.

Implicit neural representations. The recent development of implicit neural representations (INRs) has led to impressive results in computer vision (Sitzmann et al., 2020; Tancik et al., 2020; Fathy et al., 2021; Mildenhall et al., 2021). INRs can represent data as a continuous function, which can be queried at any coordinate. While they have been applied in other fields such as physics (Yin et al., 2023) and meteorology (Huang and Hoefler, 2023), there has been limited research on INRs for time series analysis. Prior works (Fons et al., 2022; Jeong and Shin, 2022) focused on time series generation for data augmentation and on time series encoding for reconstruction but are limited by their fixed grid input requirement. DeepTime (Woo et al., 2022) is the closest work to our contribution. DeepTime learns a set of basis INR functions from a training set of multiple time series and combines them using a Ridge regressor. This regressor allows it to adapt to new time series. It has been designed for forecasting only. The original version cannot handle imputation properly and was adapted to do so for our comparisons. In our experiments, we will demonstrate that TimeFlow

significantly outperforms DeepTime in imputation and also in forecasting tasks when dealing with missing values in the look-back window. TimeFlow also shows a slight advantage over DeepTime in forecasting regularly sampled series.

C INR GLOBAL SCHEME

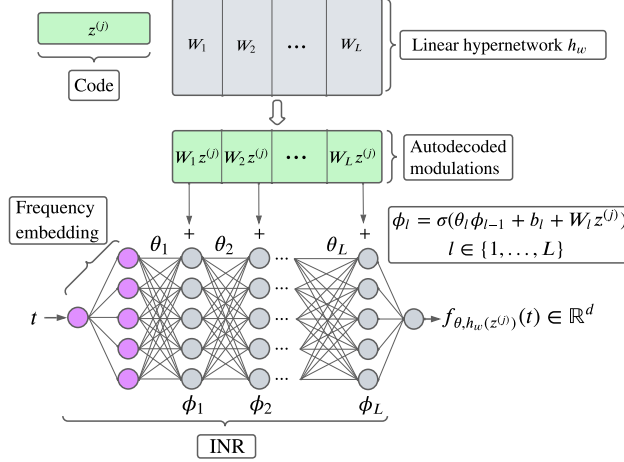


Figure 1: Overview of TimeFlow architecture. Forward pass to approximate the time series $x^{(j)}$. σ stands for the ReLU activation function.

D TRAINING AND INFERENCE ALGORITHMS

Algorithm 1: TimeFlow Training

while no convergence do

 Sample batch \mathcal{B} of data $(x^{(j)})_{j \in \mathcal{B}}$;

 Set codes to zero $z^{(j)} \leftarrow 0, \forall j \in \mathcal{B}$;

 // inner loop for encoding:

for $j \in \mathcal{B}$ and $step \in \{1, \dots, K\}$ **do**

$z^{(j)} \leftarrow z^{(j)} - \alpha \nabla_{z^{(j)}} \mathcal{L}_{\mathcal{T}_{in}^{(j)}}(f_{\theta, h_w(z^{(j)})}, x^{(j)})$;

 // outer loop step:

$[\theta, w] \leftarrow [\theta, w] - \eta \nabla_{[\theta, w]} \frac{1}{|\mathcal{B}|} \sum_{j \in \mathcal{B}} [\mathcal{L}_{\mathcal{T}_{in}^{(j)}}(f_{\theta, h_w(z^{(j)})}, x^{(j)}) + \lambda \mathcal{L}_{\mathcal{T}_{out}^{(j)}}(f_{\theta, h_w(z^{(j)})}, x^{(j)})]$;

Algorithm 2: TimeFlow Inference with trained θ, w

For the j -th series $(x^{(j)})$, set code to zero $z^{*(j)} \leftarrow 0$;

for $step \in \{1, \dots, K\}$ **do**

$z^{*(j)} \leftarrow z^{*(j)} - \alpha \nabla_{z^{*(j)}} \mathcal{L}_{\mathcal{T}_{in}^{*(j)}}(f_{\theta, h_w(z^{*(j)})}, x_t)$

Query $f_{\theta, h_w(z^{*(j)})}(t)$ for any $t \in \mathcal{T}^{*(j)}$

E IMPUTATION EXPERIMENTS

E.1 MAIN SETTING DETAILS

For the imputation task in Section 3, the TimeFlow training and inference procedures are detailed and illustrated in Figure 2.




Figure 2: Training and inference procedures of TimeFlow for imputation. (i) During training, for each time series $x^{(j)}$, our observations (red dots ●) are restricted to the sparsely sampled grid, denoted as $\mathcal{T}_{in}^{(j)}$. (ii) During inference, our objective is to infer the values over the dense grids $\mathcal{T}^{(j)}$, on the unobserved data points (such as the blue dots ● on the figure).

E.2 IMPUTATION VISUAL RESULTS

Qualitatively, we see on example series in Figure 3 that our model shows significant imputation capabilities, with on a subsampling rate at $\tau = 0.1$ on the *Electricity* dataset. It captures well different frequencies and amplitudes in a challenging case (sample 35), although it underestimates the amplitude of some peaks. In a more challenging scenario (sample 25), where the series exhibit additional trend changes and frequency variations within the data, TimeFlow correctly imputes most timestamps, outperforming BRITS, which is the best-performing method for the *Electricity* dataset.

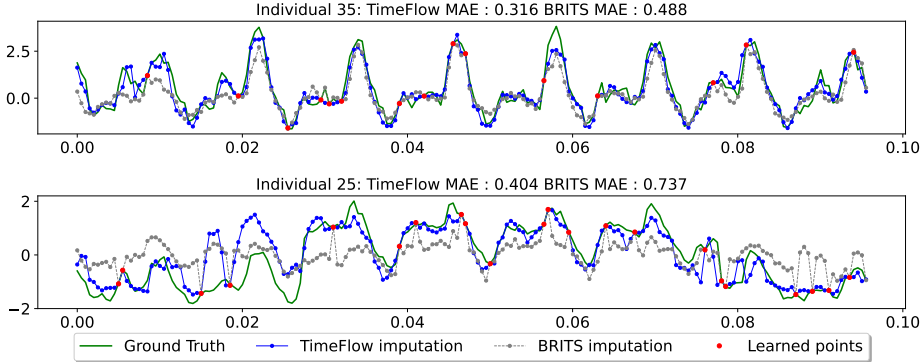


Figure 3: *Electricity* dataset. TimeFlow imputation (blue line) and BRITS imputation (gray line) with 10% of known point (red points) on the eight first days of samples 35 (top) and 25 (bottom).

E.3 MODELS COMPLEXITY

We can see in Table 2 that our method has fewer parameters than SOTA imputation methods, 10 times less than BRITS and 20 times less than SAITS. It is mainly due to their modelisation of interaction between samples. SAITS, which is based on transformers has the highest number of parameters when mTAN has the lowest number of parameters.

E.4 IMPUTATION FOR PREVIOUSLY UNSEEN TIME SERIES

In more practical scenarios, such as cases involving the installation of new sensors, we often encounter new time series originating from the same underlying phenomenon. In such instances, it becomes crucial to make inferences for these previously unseen time series. Thanks to efficient

Table 2: Number of parameters for each DL methods on the imputation task on the *Electricity* dataset.

	TimeFlow	DeepTime	NeuralProcess	mTAN	SAITS	BRITS	TIDER
Number of parameters	602k	1315k	248k	113k	11 137k	6 220k	1 034k

adaptation in latent space, our model can easily be applied to these new time series, contrasting with SOTA methods like SAITS and BRITS, which require full model retraining on the whole set of time series.

Setting In this section we analyze in details the imputations results for previously unseen time series described in Section 3. Specifically, TimeFlow is trained on a given set of time series within a defined time window and then used for inference on new time series. We train TimeFlow on 50 % of the samples and consider the remaining 50 % as the new time series.

We compare in Table 3 observed grid fit scores and missing grid inference scores for time series known at training and time series unknown at training.

Table 3: TimeFlow MAE imputation errors results for imputation previously unseen time series.

	τ	Known time series		New time series	
		Fit	Inference	Fit	Inference
Electricity	0.05	0.060 ± 0.010	0.402 ± 0.021	0.142 ± 0.083	0.413 ± 0.026
	0.10	0.046 ± 0.006	0.302 ± 0.010	0.144 ± 0.098	0.309 ± 0.016
	0.20	0.067 ± 0.015	0.285 ± 0.014	0.154 ± 0.089	0.291 ± 0.022
	0.30	0.093 ± 0.022	0.266 ± 0.010	0.163 ± 0.073	0.271 ± 0.017
	0.50	0.108 ± 0.012	0.236 ± 0.010	0.167 ± 0.061	0.245 ± 0.017
Solar	0.05	0.014 ± 0.002	0.104 ± 0.015	0.050 ± 0.037	0.109 ± 0.016
	0.10	0.017 ± 0.002	0.092 ± 0.015	0.052 ± 0.036	0.099 ± 0.017
	0.20	0.028 ± 0.008	0.078 ± 0.014	0.058 ± 0.031	0.089 ± 0.017
	0.30	0.038 ± 0.009	0.072 ± 0.013	0.063 ± 0.028	0.084 ± 0.018
	0.50	0.045 ± 0.011	0.066 ± 0.013	0.067 ± 0.025	0.080 ± 0.019
Traffic	0.05	0.044 ± 0.003	0.291 ± 0.013	0.094 ± 0.051	0.291 ± 0.012
	0.10	0.033 ± 0.001	0.209 ± 0.010	0.093 ± 0.060	0.216 ± 0.012
	0.20	0.037 ± 0.006	0.175 ± 0.008	0.095 ± 0.058	0.186 ± 0.013
	0.30	0.048 ± 0.005	0.164 ± 0.006	0.098 ± 0.051	0.175 ± 0.013
	0.50	0.068 ± 0.004	0.159 ± 0.007	0.110 ± 0.042	0.169 ± 0.012

Results The results presented in Table 3 indicate that the inference MAE for missing grids shows consistency between known and new samples, regardless of the data or sampling rate. However, it is worth noting that there is a slight drop in performance compared to the results in table Table 1. This decrease is because in Table 3, the shared architecture is trained on only half the samples, affecting its overall performance.

E.5 DETAILS ON DEEPTIME ADAPTATION FOR IMPUTATION

As DeepTime was proposed to address the forecasting task with a deeptime-index model, the authors did not tackle the task of imputation and left it out for future work. Given the success of this method and the motivation of our work, we wanted to explore its capabilities to impute time series with several subsampling rates. Following our current framework, we first tried to train the model in a self-supervised way, i.e. trying to reconstruct observations $x^{(j)} \in \mathcal{T}^{(j)}$ after the INR has been conditioned with the Ridge Regressor on the same set of observations, but discovered failure cases for $\tau \leq 0.20$. To be faithful to the original supervised training of DeepTime, we therefore randomly mask out 50% of the observations that we use as context for the Ridge Regressor and try to infer the other 50% (the targets) to train the INR.

We provide a qualitative comparison of the model’s performance with these two different training procedures in Figure 4. We can notice that the model that results from the self-supervised training perfectly fits the observations but completely misses the important patterns of the series. On the other hand, when DeepTime is trained to infer target values based on observations, it is able to

capture the general trends. We think that in the small subsampling regime ($\tau \leq 0.20$), the Ridge Regressor easily fits very well all the observations which hinders the training of the INR’s basis.




Figure 4: *Electricity dataset*. Self supervised DeepTime imputation (blue line) and supervised DeepTime imputation (black line) with 5% of known point (red points) on the eight first days of samples 11 (top) and 29 (bottom).

E.6 IMPUTATION AGAINST NON DEEP LEARNING METHODS

Setting In addition to the deep learning methods presented in Table 1, we evaluate TimeFlow against two classic machine learning baselines, K-Nearest Neighbours (KNN) and linear interpolation, which are valuable for getting an idea of the complexity of the problem.

Table 4: Mean MAE imputation results on the missing grid only over five different time window. τ stands for the subsampling rate. Bold results are best, underline results are second best.

	τ	TimeFlow	Linear interpolation	KNN (k=3)
Electricity	0.05	0.324 ± 0.013	0.828 ± 0.045	<u>0.531 ± 0.033</u>
	0.10	0.250 ± 0.010	0.716 ± 0.039	<u>0.416 ± 0.020</u>
	0.20	0.225 ± 0.008	0.518 ± 0.029	<u>0.363 ± 0.019</u>
	0.30	0.212 ± 0.007	0.396 ± 0.022	<u>0.342 ± 0.017</u>
	0.50	0.194 ± 0.007	<u>0.275 ± 0.015</u>	0.323 ± 0.016
Solar	0.05	0.095 ± 0.015	0.339 ± 0.031	<u>0.151 ± 0.017</u>
	0.10	0.083 ± 0.015	0.170 ± 0.014	<u>0.128 ± 0.017</u>
	0.20	0.072 ± 0.015	<u>0.088 ± 0.010</u>	0.110 ± 0.016
	0.30	0.061 ± 0.012	<u>0.063 ± 0.009</u>	0.103 ± 0.017
	0.50	<u>0.054 ± 0.013</u>	0.044 ± 0.008	0.096 ± 0.016
Traffic	0.05	0.283 ± 0.016	0.813 ± 0.027	<u>0.387 ± 0.014</u>
	0.10	0.211 ± 0.012	0.701 ± 0.026	<u>0.293 ± 0.012</u>
	0.20	0.168 ± 0.006	0.508 ± 0.022	<u>0.249 ± 0.010</u>
	0.30	0.151 ± 0.007	0.387 ± 0.018	<u>0.228 ± 0.009</u>
	0.50	0.139 ± 0.007	0.263 ± 0.013	<u>0.204 ± 0.009</u>
TimeFlow improvement	/	/	49.06 %	35.95 %

Results KNN imputation uses information from other individuals and gives satisfactory results at all sampling rates. On the other hand, the purely univariate approach of linear interpolation struggles at low sampling rates but performs well at high sampling rates. TimeFlow significantly outperforms both baselines by a large margin.

F FORECASTING EXPERIMENTS

F.1 CLASSICAL FORECASTING

In this section, we are interested in the conventional long-term forecasting scenario. It consists in predicting the phenomenon in a specific future period, the horizon, based on the history of a limited past period, the look-back window. The forecaster is trained on a set of n observed time series for a given time window (train period) and tested on new time windows.




Figure 5: Training and inference procedure of TimeFlow for forecasting. (i) During training (top-figure), for each time series $x^{(j)}$, we observe some look-back window/horizon drawing pairs in the trained period. TimeFlow is trained with Algorithm 1 to predict all observed timestamps in this drawing pairs while being conditioned by the observed look-back window. (ii) Once TimeFlow is optimized, the objective during inference (bottom-figure) is to infer the horizon over new time windows (blue dots \bullet) while being conditioned by the newly observed look-back window (red dots \bullet).

Setting For a given time series $x^{(j)}$, $\mathcal{T}_{in}^{(j)}$ denotes the look-back window and $\mathcal{T}_{out}^{(j)}$ the horizon of H points. During training, at each epoch, we train $f_{\theta, h_w}(z^{(j)})$ following Algorithm 1 with randomly drawn pairs of look-back window and horizon $(\mathcal{T}_{in}^{(j)} \cup \mathcal{T}_{out}^{(j)})_{j \in \mathcal{B}}$ within the observed train period. Then, for a new time window $\mathcal{T}^{*(j)}$, given a look-back window $\mathcal{T}_{in}^{*(j)}$ we forecast future values any $t \in \mathcal{T}^{*(j)}$, the horizon interval, following Algorithm 2. We illustrate the training and inference of TimeFlow for the forecasting task in Figure 5.

Baselines. To evaluate the quality of our model in long-term forecasting, we compare it to the discrete baselines PatchTST (Nie et al., 2022), DLinear (Zeng et al., 2022), AutoFormer (Wu et al., 2021), and Informer (Zhou et al., 2021). We also include continuous baselines DeepTime and Neural Process (NP). In Table 5, we present the forecasting results for standard horizons in long-term forecasting: $H \in \{96, 192, 336, 720\}$. The look-back window length is fixed to 512.

Results. The results in Table 5 show that our approach ranks in the top two across all datasets and horizons and is the overall best continuous method. TimeFlow’s performance is comparable to the current SOTA model PatchTST, with only 2% relative difference. Moreover, TimeFlow shows consistent results across the three datasets, whereas the other best discrete and continuous baselines, i.e. PatchTST and DeepTime, performance drops for some datasets. We also note that, despite the great performance of the SOTA PatchTST, other transformer-based baselines (discrete methods in Table 5) perform poorly. Overall, although this evaluation setting favors discrete methods because the time series are observed at evenly distributed time steps, TimeFlow consistently performs as well as PatchTST and outperforms all the other methods, whether discrete or continuous. It is the first time that a continuous model has achieved the same level of performance as discrete methods within their specific setting. We observe that, despite comparable results between TimeFlow and the

Table 5: Mean MAE forecast results averaged over different time windows. Each time, the model is trained on one time window and tested on the others (there are 2 windows for *SolarH* and 5 for *Electricity* and *Traffic*). H stands for the horizon. Bold results are best, and underlined results are second best. TimeFlow improvement represents the overall percentage improvement achieved by TimeFlow compared to the specific method being considered.

	H	Continuous methods			Discrete methods			
		TimeFlow	DeepTime	Neural Process	Patch-TST	DLinear	AutoFormer	Informer
Electricity	96	0.228 ± 0.028	0.244 ± 0.026	0.392 ± 0.045	0.221 ± 0.023	0.241 ± 0.030	0.546 ± 0.277	0.603 ± 0.255
	192	<u>0.238 ± 0.020</u>	0.252 ± 0.019	0.401 ± 0.046	0.229 ± 0.020	0.252 ± 0.025	0.500 ± 0.190	0.690 ± 0.291
	336	<u>0.270 ± 0.031</u>	0.284 ± 0.034	0.434 ± 0.076	0.251 ± 0.027	0.288 ± 0.038	0.523 ± 0.188	0.736 ± 0.271
	720	<u>0.316 ± 0.055</u>	0.359 ± 0.051	0.607 ± 0.150	0.297 ± 0.039	0.365 ± 0.059	0.631 ± 0.237	0.746 ± 0.265
SolarH	96	0.190 ± 0.013	0.190 ± 0.020	0.221 ± 0.048	0.262 ± 0.070	0.208 ± 0.014	0.245 ± 0.045	0.248 ± 0.022
	192	0.202 ± 0.020	<u>0.204 ± 0.028</u>	0.244 ± 0.048	0.253 ± 0.051	0.217 ± 0.022	0.333 ± 0.107	0.270 ± 0.031
	336	0.209 ± 0.017	0.199 ± 0.026	0.240 ± 0.006	0.259 ± 0.071	0.217 ± 0.026	0.334 ± 0.079	0.328 ± 0.048
	720	0.218 ± 0.041	<u>0.229 ± 0.024</u>	0.403 ± 0.147	0.267 ± 0.064	0.249 ± 0.034	0.351 ± 0.055	0.337 ± 0.037
Traffic	96	0.217 ± 0.032	0.228 ± 0.032	0.283 ± 0.027	0.203 ± 0.037	0.228 ± 0.033	0.319 ± 0.059	0.372 ± 0.078
	192	<u>0.212 ± 0.028</u>	0.220 ± 0.022	0.292 ± 0.024	0.197 ± 0.030	0.221 ± 0.023	0.368 ± 0.057	0.511 ± 0.247
	336	<u>0.238 ± 0.034</u>	0.245 ± 0.038	0.305 ± 0.039	0.222 ± 0.039	0.250 ± 0.040	0.434 ± 0.061	0.561 ± 0.263
	720	<u>0.279 ± 0.050</u>	0.290 ± 0.052	0.339 ± 0.038	0.269 ± 0.057	0.300 ± 0.057	0.462 ± 0.062	0.638 ± 0.067
TimeFlow improvement	/	3.74 %	29.06 %	3.23 %	6.92 %	42.09 %	48.57 %	

DLinear, DeepTime baselines for the 96, 192, and 336 time horizons, the difference becomes more significant when the longest horizon is considered.

Discussion. Our method consistently performs as well as PatchTST, in the setting with evenly distributed time steps. This achievement is significant as our continuous model can attain the same level of performance as discrete methods within their specific setting.

Forecasting on previously unseen time series. TimeFlow considers that the series observed at different locations are independent, similar to PatchTST, NP, and DeepTime. This allows it to generalize to previously unseen time series from the same phenomenon. Note that this is not the case for most discrete methods. We show in Appendix F.5, Table 9 that TimeFlow is able to generalize to previously unseen time series with no significant performance drop.

F.1.1 FORECASTING ON NEW PERIODS FOR UNKNOWN SAMPLES WITH MISSING VALUES

Setting After undergoing training for a specific duration on a given set of samples, our model demonstrates continuity, enabling us to extend its capabilities to model a continuum for new samples across different time windows. To accomplish this, we maintain the same training configuration as described in Appendix F.1, but apply inference to new samples within unseen time windows.

Baseline In the above scenario, models such as PatchTST (and other DL models) must proceed in two steps: (i) completing the look-back window on a dense regular grid using imputation; (ii) apply the model on the completed window to predict the future. We compared TimeFlow with the following two-step processing baseline: linear interpolation handling the missing values within the partially observed look-back window, and PatchTST handling the forecasting task. Due to the unavailability of data from other samples for the same period, when dealing with unseen time windows for new samples, we are restrained to interpolation as imputation method. We conducted experiments on the *Traffic* and *Electricity* datasets, focusing on the 96 and 192 horizons. In Table 6, we present the results at different sampling rates $\tau \in \{0.5, 0.2, 0.1\}$ within the look-back window.

Results Table 6 shows that our method significantly outperforms the baseline (linear interpolation + PatchTST). This poor performance of the baseline can be attributed to the imputation errors induced by the linear interpolation. Although our method was not explicitly trained for imputation, it demonstrates a remarkable capability to recover the ground truth values for the partially observed historical window, leading to precise forecasts. We show the extended results for *Traffic* and *Electricity* datasets in the supplementary materials. In addition, our method exhibits effectiveness even at significantly low sampling rates. When it comes to the results presented in Table 5, our method shows a relatively modest degradation when the subsampling rate drops, in comparison to the results

obtained in the optimal setting. We illustrate the qualitative results in Figure 6 using a look-back window with a 20 % sampling rate for a horizon of 96 time steps.

F.2 CHALLENGING TASK: FORECAST WHILE IMPUTING INCOMPLETE LOOK-BACK WINDOWS

In real-world scenarios, it is common to encounter missing or irregularly sampled series when making predictions on new time windows (Cinar et al., 2018; Tang et al., 2020). Continuous methods can handle these cases, as they are designed to accommodate irregular sampling within the look-back window. In this section, we formulate a task to simulate these real-world scenarios. It’s worth noting that this task is often encountered in practice but is rarely considered in the DL literature.

Setting and baselines. This scenario is similar to the forecast setting in Appendix F.1 and illustrated in Figure 5. The difference is that during inference, the look-back window is subsampled at a rate τ smaller than the one used for the training phase. This simulates a situation with missing observations in the look back window. Consequently, two distinct tasks emerge during the inference phase: imputing missing points within the sparsely observed look-back window, and forecasting over the horizon with this degraded context. In Table 6, we compare to the two other continuous baselines, DeepTime and NP on *Electricity* and *Traffic* for different τ s and horizons.

Table 6: MAE results for forecasting with missing values in the look-back window. τ stands for the percentage of observed values in the look-back window. Best results are in bold. TimeFlow improvement represents the overall percentage improvement (for each task) achieved by TimeFlow compared to the specific method being considered.

	H	τ	TimeFlow		DeepTime		Neural Process	
			Imputation error	Forecast error	Imputation error	Forecast error	Imputation error	Forecast error
Electricity	96	0.5	0.151 ± 0.003	0.239 ± 0.013	0.209 ± 0.004	0.270 ± 0.019	0.460 ± 0.048	0.486 ± 0.078
		0.2	0.208 ± 0.006	0.260 ± 0.015	0.249 ± 0.006	0.296 ± 0.023	0.644 ± 0.079	0.650 ± 0.095
		0.1	0.272 ± 0.006	0.295 ± 0.016	0.284 ± 0.007	0.324 ± 0.026	0.740 ± 0.083	0.737 ± 0.106
	192	0.5	0.149 ± 0.004	0.235 ± 0.011	0.204 ± 0.004	0.265 ± 0.018	0.461 ± 0.045	0.498 ± 0.070
		0.2	0.209 ± 0.006	0.257 ± 0.013	0.244 ± 0.007	0.290 ± 0.023	0.601 ± 0.075	0.626 ± 0.101
		0.1	0.274 ± 0.010	0.289 ± 0.016	0.282 ± 0.007	0.315 ± 0.025	0.461 ± 0.045	0.724 ± 0.090
Traffic	96	0.5	0.180 ± 0.016	0.219 ± 0.026	0.272 ± 0.028	0.243 ± 0.030	0.436 ± 0.025	0.444 ± 0.047
		0.2	0.239 ± 0.019	0.243 ± 0.027	0.335 ± 0.026	0.293 ± 0.027	0.596 ± 0.049	0.597 ± 0.075
		0.1	0.312 ± 0.020	0.290 ± 0.027	0.385 ± 0.025	0.344 ± 0.027	0.734 ± 0.102	0.731 ± 0.132
	192	0.5	0.176 ± 0.014	0.217 ± 0.017	0.241 ± 0.027	0.234 ± 0.021	0.477 ± 0.042	0.476 ± 0.043
		0.2	0.233 ± 0.017	0.236 ± 0.021	0.286 ± 0.027	0.276 ± 0.020	0.685 ± 0.109	0.678 ± 0.108
		0.1	0.304 ± 0.019	0.277 ± 0.021	0.331 ± 0.025	0.324 ± 0.021	0.888 ± 0.178	0.877 ± 0.174
TimeFlow improvement			/	/	18.97 %	11.87 %	61.88 %	58.41 %

Results. In Table 6, the results show that TimeFlow consistently outperforms other methods in imputation and forecasting for every scenarios. When comparing with the complete look-back windows observations scenario from Table 5, one observes that at a 0.5 sampling rate, TimeFlow presents only a slight reduction in performance, whereas other baseline methods experience more significant drops. For instance, when we compare forecast results between a complete window and a $\tau = 0.5$ subsampled window for *Electricity* with a forecasting horizon of $H = 96$, TimeFlow’s error increases by a mere 4.6% (from 0.228 to 0.239). In contrast, DeepTime’s error grows by over 10% (from 0.244 to 0.270), and NP experiences a rise of around 25% (from 0.392 to 0.486). For lower sampling rates, TimeFlow still delivers correct predictions. Qualitatively, we see on the series example in Figure 6 that despite observing only 10% of the look-back window, the model can correctly infer both the complete look-back window and the horizon. Both quantitative and qualitative results show the robustness and efficiency of TimeFlow on this particularly challenging setting.

F.3 PLOTS COMPARISON: TIMEFLOW VS PATCHTST

Table 5 demonstrates the similar forecasting performance of TimeFlow and PatchTST across all horizons. To visually represent their predictions, the figures below showcase the forecasted outcomes of these methods for two samples (24 and 38) and two horizons (96 and 192) on the *Electricity*, *SolarH*, and *Traffic* datasets.




Figure 6: *Traffic dataset, sample 95*. In this figure, TimeFlow simultaneously imputes and forecasts at horizon 96 with a 10% partially observed look-back window of length 512.




Figure 7: Qualitative comparisons of TimeFlow vs PatchTST on the *Electricity* dataset for new time windows

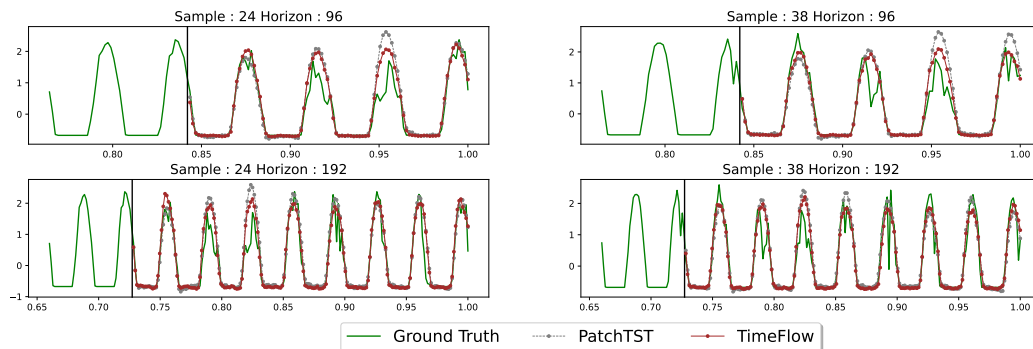


Figure 8: Qualitative comparisons of TimeFlow vs PatchTST on the *SolarH* dataset for new time windows

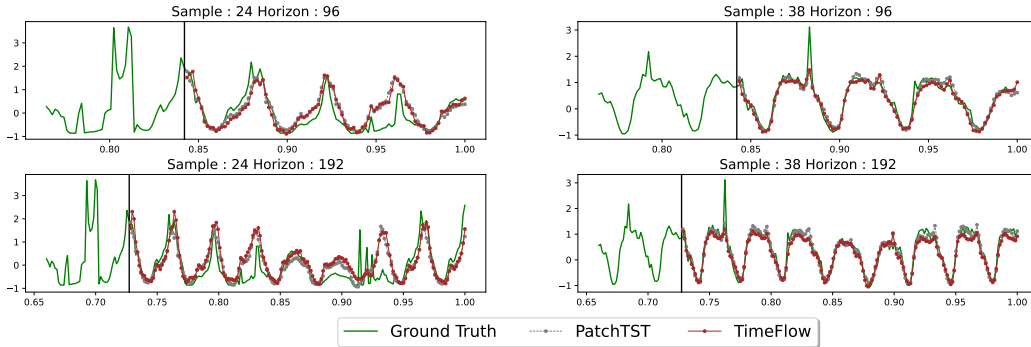


Figure 9: Qualitative comparisons of TimeFlow vs PatchTST on the *Traffic* dataset for new time windows

Results The visual analysis of the figures above reveals that the predictions of TimeFlow and PatchTST are remarkably similar. For instance, when examining sample 24 and horizon 192 of the *Traffic* dataset, both forecasters exhibit similar error patterns. The only noticeable distinction emerges in the *SolarH* dataset, where PatchTST tends to overestimate certain peaks.

F.4 MODELS COMPLEXITY

In this section, we present the parameter counts and the inference time for the main forecasting baselines. Except for TimeFlow and DeepTime, the number of parameters varies with the number of samples, the look-back window, and the horizon. Thus, we report the number of parameters for two specific configurations, including a fixed dataset, a fixed look-back window, and a fixed horizon. In Table 7, we see that for PatchTST and DLinear, the larger the horizon, the more the number of parameters increases. In Table 8, it is shown that all methods’ computational time increases with the horizon, which is expected. Moreover, TimeFlow is slower than the baselines that use forward computations only. Still, on the *Electricity* dataset, for example, the method can infer for 321 samples a horizon of 720 values with a look-back window of 512 timestamps in less than 0.2s, which does not look prohibitive for many real-world usages. This is mainly due to the small number of gradient steps at inference.

Table 7: The number of parameters for main baselines on the forecasting task on the *Electricity* dataset for horizons 96 and 720. The look-back window size is 512.

	TimeFlow	DeepTime	Neural Process	Patch-TST	DLinear	Informer	Autoformer
96	602k	1 315k	480k	1 194k	98k	984k	1 005k
720	602k	1 315k	480k	6 306k	739k	984k	1 005k

Table 8: Inference time (in seconds) for the forecasting task on the *Electricity* dataset with horizons 96 and 720 and a look-back window of length 512. The statistics are computed over 10 runs using an NVIDIA TITAN RTX GPU.

	TimeFlow	Patch-TST	DLinear	DeepTime	AutoFormer	Informer
96	0.147 ± 0.007	0.016 ± 0.002	0.007 ± 0.003	0.006 ± 0.002	0.027 ± 0.001	0.0191 ± 0.002
720	0.176 ± 0.009	0.020 ± 0.001	0.009 ± 0.001	0.010 ± 0.002	0.034 ± 0.001	0.0251 ± 0.002

F.5 FORECASTING FOR PREVIOUSLY UNSEEN TIME SERIES

Setting and baseline. As mentioned in Appendix F.1, most forecasters explicitly model the dependencies between samples, which limits their ability to generalize to new time series without

retraining the entire model. However, TimeFlow, PatchTST, and DeepTime have the advantage of being reusable for new samples. In Table 9, we present the results of TimeFlow and PatchTST for new periods, considering both known samples and new samples. We train TimeFlow and PatchTST on 50 % of the samples and consider the remaining 50 % as the new time series.

Table 9: MAE results over horizon for the forecasting task in the context of generalization to new time series.

	H	Known time series		New time series	
		TimeFlow MAE error	PatchTST MAE error	TimeFlow MAE error	PatchTST MAE error
Electricity	96	0.228 ± 0.023	0.211 ± 0.007	0.241 ± 0.023	0.224 ± 0.020
	192	0.244 ± 0.022	0.225 ± 0.014	0.254 ± 0.024	0.238 ± 0.024
	336	0.269 ± 0.036	0.267 ± 0.019	0.277 ± 0.033	0.285 ± 0.005
	720	0.331 ± 0.058	0.310 ± 0.026	0.333 ± 0.059	0.331 ± 0.045
Traffic	96	0.226 ± 0.035	0.208 ± 0.036	0.222 ± 0.031	0.203 ± 0.037
	192	0.217 ± 0.028	0.202 ± 0.029	0.215 ± 0.026	0.199 ± 0.030
	336	0.242 ± 0.036	0.228 ± 0.041	0.240 ± 0.031	0.224 ± 0.036
	720	0.283 ± 0.053	0.275 ± 0.059	0.283 ± 0.049	0.272 ± 0.055
SolarH	96	0.237 ± 0.077	0.256 ± 0.055	0.236 ± 0.081	0.256 ± 0.062
	192	0.238 ± 0.051	0.251 ± 0.239	0.239 ± 0.058	0.250 ± 0.050
	336	0.220 ± 0.027	0.255 ± 0.663	0.220 ± 0.034	0.255 ± 0.066
	720	0.240 ± 0.039	0.267 ± 0.062	0.240 ± 0.042	0.267 ± 0.063
TimeFlow improvement		/	-1.05 %	/	-0.39 %

Results Table 9 demonstrates the good adaptability of both methods to new samples, as the difference in MAE between known and new samples is marginal.

G DISCUSSION ON IMPLEMENTATION CHOICES

As indicated before, adapting the components and enhancing their synergy for the tasks of imputation and forecasting is not trivial and requires careful choices. We conducted several ablation studies to provide a comprehensive examination of key implementation choices of our framework.

Our findings can be summarized as follows. **Choice of INR:** An FFN with Fourier Features outperformed other popular INRs for the tasks considered in this study. Unlike SIREN, which does not explicitly incorporate frequencies but uses sine activation functions, the Fourier features network can more effectively capture a broader range of frequencies, especially at low sampling rates. This is crucial for accurately capturing high frequencies in sparsely observed time series. Our experiments, detailed in Section H.2.1 and Table 10, demonstrate this superiority across various datasets. **Choice of encoding / meta-learning:** TimeFlow with a set encoder for learning the compact conditioning codes z in place of the auto-decoding strategy used here, proved much less effective on complex datasets. This is further elaborated in Section H.2.4 and Table 17. Additionally, replacing the 2nd-order optimization for a 1st-order one, such as REPTILE, led to unstable training, as shown in Table 16. **Choice of modulations:** Complexifying the modulation by introducing scaling parameters in addition to shift parameters did not provide performance gains. Our experiments on the *Electricity* dataset, detailed in Section H.2.5 and Table 18, indicate that shift-only modulation is more efficient.

For TimeFlow, across all experiments, we used a code dimension of 128, an FFN with a depth of 5 and a width of 256, and 64 Fourier features. We used 3 inner steps and a learning rate of 0.01 for the inner-loop, and a learning rate of 5×10^{-4} for the outer-loop. We performed a comprehensive analysis to understand notably the **influence of the z dimension:** a latent code dimension of 128 was suitable for our tasks; this is supported by results in Section H.2.2 and Table 11 - and the **influence of the number of inner steps:** using 3 inner steps for training and inference struck a favorable balance between reconstruction capabilities and computational efficiency, as detailed in Section H.2.3. and Tables 12, 13, 14, and 15.

Comparison with other INRs We found that an FFN with Fourier Features was the best performing INR backbone for TimeFlow for time-series reconstruction, imputation and forecasting. We hypothesize that the FFN with ReLU activations correctly learns the different frequencies of

time series with multi-temporal patterns by switching on or off the Fourier embedding frequencies. We conducted a direct comparison with SIREN (Sitzmann et al., 2020) on the imputation task in Table 10.

TimeFlow with an encoder. We also tested our modulated INR in combination with the set-encoder from the original neural process architecture. Though this kind of encoder could be leveraged to tackle the challenges we aim to address, we found it underfitted the data on complex datasets. We provide quantitative results for the forecasting task in Table 17.

Comparison of meta-learning algorithms. Several options are possible to tackle the joint optimization of the codes and the common parameters of TimeFlow. We noticed that first-order methods such as REPTILE Nichol et al. (2018) were experiencing instabilities during training and as a result were not able to properly fit the data, especially when the number of observation in the temporal support was high. Our comparison for the imputation task is presented in Table 16.

Other types of modulations. We found no benefits in using both shift and scale modulations (Zhou et al., 2022a) compared to shift-only modulations for TimeFlow. Even though it doubles the number of modulations and therefore increases the computational cost of training and inference, it does not improve performance. Results on the forecasting task are shown in Table 18.

Number of inner steps and size of latent space. For all datasets and tasks, we use 3 inner-steps for training the model and 3 inner-steps at inference. We found that it provided a good trade-off between reconstruction capabilities and computational cost. We do not use more inner-steps at test time as it does not improve the performance, as shown in Table 12. We set the dimension of the latent code to 128 for all datasets and tasks. A larger code dimension facilitated training but generalized well, and a smaller one smoothed the time-series too much.

H ARCHITECTURE DETAILS AND ABLATION STUDIES

H.1 ARCHITECTURE DETAILS

For all imputation and forecasting experiments we choose the following hyperparameters :

- z dimension: 128
- Number of layers: 5
- Hidden layers dimension: 256
- $\gamma(t) \in \mathbb{R}^{2 \times 64}$
- z code learning rate (α in Algorithm 1): 10^{-2}
- Hypernetwork and INR learning rate: 5×10^{-4}
- Number of steps in inner loop: $K = 3$
- Number of epochs: 4×10^4
- Batch size: 64

It is worth noting that the hyperparameters mentioned above remain consistent across all experiments conducted in the paper. We chose to maintain a fixed set of hyperparameters for our model, while other imputation and forecasting approaches commonly fine-tune hyperparameters based on a validation dataset. The obtained results exhibit high robustness across various settings, suggesting that the selected hyperparameters are already effective in achieving reliable outcomes.

H.2 ABLATION STUDIES

H.2.1 FOURIER FEATURES VS SIREN ON IMPUTATION TASK

Baseline The SIREN network differs from the Fourier features network because it does not explicitly incorporate frequencies as input. Instead, it is a multi-layer perceptron network that utilizes sine activation functions. An adjustable parameter, denoted ω_0 , is multiplied with the input matrices of the preceding layers to capture a broader range of frequencies. For this comparison, we adopt the

same hyperparameters described in Appendix H.1, selecting $\omega_0 = 30$ to align with Sitzmann et al. (2020). Furthermore, we set the learning rate of both the hypernetwork and the INR to 5×10^{-5} to enhance training stability. In Table 10, we compare the imputation results obtained by the Fourier features network and the SIREN network, specifically focusing on the first time window from the *Electricity*, *Traffic* and *Solar* datasets.

Table 10: MAE imputation errors on the first time window of each dataset. Best results are bold.

	τ	TimeFlow	TimeFlow w SIREN
Electricity	0.05	0.323	0.466
	0.10	0.252	0.350
	0.20	0.224	0.242
	0.30	0.211	0.222
	0.50	0.194	0.209
Solar	0.05	0.105	0.114
	0.10	0.083	0.094
	0.20	0.065	0.079
	0.30	0.061	0.072
	0.50	0.056	0.066
Traffic	0.05	0.292	0.333
	0.10	0.220	0.252
	0.20	0.168	0.191
	0.30	0.152	0.163
	0.50	0.141	0.154

Results According to the results presented in Table 10, the Fourier features network outperforms the SIREN network in the imputation task on these datasets. Notably, the performance gap between the two network architectures are more pronounced at low sampling rates. This disparity can be attributed to the SIREN network’s difficulty in accurately capturing high frequencies when the time series is sparsely observed. We hypothesize that the MLP with ReLU activations correctly learns the different frequencies of time series with multi-temporal patterns by switching on or off the Fourier embedding frequencies.

H.2.2 INFLUENCE OF THE LATENT CODE DIMENSION

The dimension of the latent code z is a crucial parameter in our architecture. If it is too small, it underfits the timeseries. Consequently, this adversely affects the performance of both the imputation and forecasting tasks. On the other hand, if the dimension of z is too large, it can lead to overfitting, hindering the model’s ability to generalize to new data points.

Baselines To investigate the impact of z dimensionality on the performance of TimeFlow, we conducted experiments on the three considered datasets, specifically focusing on the forecasting task. We varied the sizes of z within $\{32, 64, 128, 256\}$. The other hyperparameters are set as presented in Appendix H.1. The obtained results for each z dimension are summarized in Table 11.

Table 11: MAE error for different z dimension.

	H	32	64	128	256
Electricity	96	0.232 ± 0.016	0.222 ± 0.017	0.222 ± 0.018	0.215 ± 0.019
	192	0.245 ± 0.020	0.239 ± 0.018	0.230 ± 0.026	0.233 ± 0.017
	336	0.254 ± 0.029	0.244 ± 0.028	0.262 ± 0.031	0.243 ± 0.032
	720	0.295 ± 0.027	0.284 ± 0.028	0.303 ± 0.041	0.283 ± 0.029
SolarH	96	0.182 ± 0.009	0.181 ± 0.012	0.179 ± 0.003	0.225 ± 0.047
	192	0.195 ± 0.014	0.195 ± 0.016	0.193 ± 0.015	0.197 ± 0.029
	336	0.181 ± 0.011	0.182 ± 0.011	0.189 ± 0.013	0.183 ± 0.012
	720	0.201 ± 0.027	0.199 ± 0.025	0.209 ± 0.029	0.200 ± 0.030
Traffic	96	0.223 ± 0.024	0.215 ± 0.028	0.215 ± 0.037	0.210 ± 0.033
	192	0.214 ± 0.018	0.217 ± 0.025	0.206 ± 0.023	0.203 ± 0.024
	336	0.238 ± 0.029	0.231 ± 0.029	0.226 ± 0.030	0.229 ± 0.029
	720	0.272 ± 0.040	0.269 ± 0.035	0.259 ± 0.038	0.262 ± 0.040

Results The results presented in Table 11 suggest that a z dimension of 128 is a reasonable compromise but only optimal for some settings. Moreover, even though the choice of z dimension seems important, it doesn’t critically impact the MAE error for the forecasting task.

H.2.3 INFLUENCE OF THE NUMBER OF GRADIENT STEPS

As can be seen in Table 12, using three gradient steps at inference yield an inference time of less than 0.2 seconds. The latter can still be reduced by doing only one step at the cost of an increase in the forecasting error. As observed in Table 12, increasing the number of gradient steps above 3 steps during inference does not improve forecasting performance.

Table 12: Inference time (in seconds) and MAE error on the forecasting task on the *Electricity* dataset for a horizon of length 720, a look-back window of length 512, and a varying number of adaptation gradient steps. The statistics are computed over 10 runs using an NVIDIA TITAN RTX GPU.

Gradient descent steps	1	3	10	50	500	5000
Inference time (s)	0.109 ± 0.003	0.176 ± 0.009	0.427 ± 0.031	3.547 ± 0.135	17.722 ± 0.536	189.487 ± 8.060
MAE	0.351 ± 0.038	0.303 ± 0.041	0.300 ± 0.040	0.299 ± 0.039	0.302 ± 0.038	0.308 ± 0.037

Table 13: MAE error on the forecasting task using 1 inner-step during training and a varying number of adaptation gradient steps at inference. Best results are in bold and / symbol means that the MAE score is very high (≥ 1).

	H	1	3	10	50
Electricity	96	0.244 ± 0.017	0.246 ± 0.017	0.261 ± 0.016	/
	192	0.253 ± 0.024	0.253 ± 0.022	0.261 ± 0.020	0.265 ± 0.019
	336	0.267 ± 0.032	0.268 ± 0.030	0.277 ± 0.028	0.281 ± 0.027
	720	0.302 ± 0.030	0.306 ± 0.029	0.310 ± 0.028	0.301 ± 0.029
SolarH	96	0.192 ± 0.023	0.623 ± 0.397	/	/
	192	0.175 ± 0.006	0.252 ± 0.068	/	/
	336	0.192 ± 0.016	0.471 ± 0.029	/	/
	720	0.216 ± 0.034	0.465 ± 0.063	/	0.550 ± 0.187
Traffic	96	0.215 ± 0.029	0.329 ± 0.039	/	/
	192	0.208 ± 0.019	0.310 ± 0.033	0.312 ± 0.032	/
	336	0.237 ± 0.028	0.307 ± 0.038	/	/
	720	0.263 ± 0.038	0.320 ± 0.040	/	/

Table 14: MAE error on the forecasting task using 3 inner-steps during training and a varying number of adaptation gradient steps at inference. Best results are in bold.

	H	1	3	10	50
Electricity	96	0.259 ± 0.020	0.222 ± 0.018	0.222 ± 0.017	0.228 ± 0.019
	192	0.269 ± 0.020	0.230 ± 0.026	0.232 ± 0.020	0.233 ± 0.026
	336	0.273 ± 0.033	0.262 ± 0.031	0.264 ± 0.032	0.268 ± 0.032
	720	0.351 ± 0.038	0.303 ± 0.041	0.300 ± 0.040	0.299 ± 0.039
SolarH	96	0.487 ± 0.196	0.179 ± 0.003	0.181 ± 0.003	0.186 ± 0.003
	192	0.411 ± 0.088	0.193 ± 0.015	0.195 ± 0.014	0.199 ± 0.013
	336	0.435 ± 0.153	0.189 ± 0.013	0.203 ± 0.006	0.223 ± 0.012
	720	0.394 ± 0.173	0.209 ± 0.029	0.203 ± 0.006	0.209 ± 0.027
Traffic	96	0.320 ± 0.038	0.215 ± 0.037	0.219 ± 0.043	0.226 ± 0.046
	192	0.299 ± 0.023	0.206 ± 0.023	0.209 ± 0.026	0.214 ± 0.027
	336	0.345 ± 0.038	0.226 ± 0.030	0.228 ± 0.031	0.233 ± 0.032
	720	0.321 ± 0.034	0.259 ± 0.038	0.260 ± 0.038	0.266 ± 0.039

Table 15: MAE error on the forecasting task using 10 inner-steps during training and a varying number of adaptation gradient steps at inference. Best results are in bold.

	H	1	3	10	50
Electricity	96	0.381 ± 0.030	0.249 ± 0.024	0.236 ± 0.024	0.238 ± 0.024
	192	0.448 ± 0.045	0.273 ± 0.019	0.244 ± 0.014	0.244 ± 0.013
	336	0.514 ± 0.053	0.283 ± 0.033	0.241 ± 0.025	0.242 ± 0.024
	720	0.647 ± 0.068	0.400 ± 0.051	0.286 ± 0.023	0.287 ± 0.021
SolarH	96	0.605 ± 0.029	0.380 ± 0.018	0.188 ± 0.012	0.199 ± 0.015
	192	0.382 ± 0.072	0.250 ± 0.012	0.202 ± 0.034	0.204 ± 0.035
	336	0.745 ± 0.105	0.431 ± 0.221	0.201 ± 0.033	0.208 ± 0.032
	720	0.745 ± 0.082	0.477 ± 0.039	0.205 ± 0.030	0.205 ± 0.029
Traffic	96	0.450 ± 0.023	0.273 ± 0.026	0.225 ± 0.028	0.230 ± 0.034
	192	0.506 ± 0.028	0.318 ± 0.021	0.233 ± 0.022	0.236 ± 0.026
	336	0.500 ± 0.042	0.320 ± 0.021	0.247 ± 0.028	0.249 ± 0.031
	720	0.511 ± 0.035	0.323 ± 0.022	0.266 ± 0.027	0.272 ± 0.024

We conduct more extensive experiments in Table 13, Table 14, Table 15 to quantify the MAE score variation according to different number of gradient steps during training and inference. The tables show that using the same number of steps in training and inference leads to better results. This is expected since using different gradient steps for training and inference makes the inference model slightly different from the training model. In addition, using 3 gradient steps instead of 1 clearly improves the performances, but using 10 instead of 3 does not. Indeed, it usually leads to overall better results for longer horizon, but the gain is not clear for smaller horizons. Hence using 3 gradient steps is a suitable choice.

H.2.4 TIMEFLOW VARIANTS WITH OTHER META-LEARNING TECHNIQUES

Baselines Before converging to the current architecture and optimization of TimeFlow, we explored different options to condition the INR with the observations. The first one was inspired by the neural process architecture, which uses a set encoder to transform a set of observations $(t_i, x_{t_i})_{i \in \mathcal{I}}$ into a latent code z by applying a pooling layer after a feed forward network. We observed that this encoder in combination with the modulated fourier features network was able to achieve relatively good results on the forecasting task but suffered of underfitting on more complex datasets such as *Electricity*.

This led us to consider auto-decoding methods instead, i.e. encoder-less architectures for conditioning the weights of the coordinate-based network. We trained TimeFlow with the REPTILE algorithm (Nichol et al., 2018), which is a first-order meta-learning technique that adapts the code in a few steps of gradient descent. In contrast with a second-order method, we observed that REPTILE was less costly to train but struggled to escape sub optimal minima, which led to unstable training and underfitting.

From an implementation point of view, the only difference between second order and first order, is that in the latter the code is detached from the computation graph before taking the outer-loop parameter update. When the code is not detached, it remains a function of the common parameters $z = z(\theta, w)$, which means that the computation graph for the outer-loop also includes the inner-loop updates to the codes. Therefore the outer-loop gradient update involves a gradient through a gradient and requires an additional backward pass through the INR to compute the Hessian. Please refer to Finn et al. (2017) for more technical details.

Table 16: Comparison of second-order and first-order (REPTILE) meta learning for TimeFlow on the imputation task. Mean MAE results on the missing grid over five different time windows. τ stands for the subsampling rate. Bold results are best.

	τ	TimeFlow	TimeFlow w REPTILE
Electricity	0.05	0.324 ± 0.013	0.363 ± 0.062
	0.10	0.250 ± 0.010	0.343 ± 0.036
	0.20	0.225 ± 0.008	0.312 ± 0.043
	0.30	0.212 ± 0.007	0.308 ± 0.035
	0.50	0.194 ± 0.007	0.305 ± 0.046
Solar	0.05	0.095 ± 0.015	0.125 ± 0.025
	0.10	0.083 ± 0.015	0.123 ± 0.032
	0.20	0.072 ± 0.015	0.108 ± 0.021
	0.30	0.061 ± 0.012	0.105 ± 0.027
	0.50	0.054 ± 0.013	0.102 ± 0.021
Traffic	0.05	0.283 ± 0.016	0.304 ± 0.026
	0.10	0.211 ± 0.012	0.264 ± 0.009
	0.20	0.168 ± 0.006	0.242 ± 0.019
	0.30	0.151 ± 0.007	0.218 ± 0.020
	0.50	0.139 ± 0.007	0.216 ± 0.017

Results In Table 16, we show the performance of first-order TimeFlow on the imputation task. In low sampling regimes the difference with TimeFlow is less perceptible, but its performance plateaus when the number of points increases. This is not surprising. Indeed, as though the task is actually simpler when τ increases, the optimization is made more difficult with the increased number of observations. We provide the performance of TimeFlow with a set encoder on the Forecasting task in Table 17. We observed that this version failed to generalize well for complex datasets.

Table 17: Comparison of optimization-based and set-encoder-based meta learning for TimeFlow on the forecasting task. Mean MAE forecast results over different time windows. H stands for the horizon. Bold results are best.

	H	TimeFlow	TimeFlow w set encoder
Electricity	96	0.228 ± 0.026	0.362 ± 0.032
	192	0.238 ± 0.020	0.360 ± 0.028
	336	0.270 ± 0.031	0.382 ± 0.038
	720	0.316 ± 0.055	0.431 ± 0.059
SolarH	96	0.190 ± 0.013	0.251 ± 0.071
	192	0.202 ± 0.020	0.239 ± 0.058
	336	0.209 ± 0.017	0.235 ± 0.040
	720	0.218 ± 0.048	0.231 ± 0.032
Traffic	96	0.217 ± 0.036	0.276 ± 0.031
	192	0.212 ± 0.028	0.281 ± 0.034
	336	0.238 ± 0.034	0.297 ± 0.042
	720	0.279 ± 0.050	0.333 ± 0.048

H.2.5 INFLUENCE OF THE MODULATION

In TimeFlow, we apply shift modulations to the parameters of the INR, i.e. for each layer l we only modify the biases of the network with an extra bias term $\phi_l^{(j)}$. We generate these bias terms with a linear hypernetwork that maps the code $z^{(j)}$ to the modulations. The output of the l -th layer of the modulated INR is thus given by $\phi_{l+1} = \text{ReLU}(\theta_l \phi_{l-1} + b_l + \psi_l^{(j)})$, where $\psi_l^{(j)} = W_l z^{(j)}$ and $(W_l)_{l=1}^L$ are parameters of the hypernetwork. However, another common modulation is the combination of the scale and shift modulation, which leads to the output of the l -th layer of the modulated INR being given by $\phi_{l+1} = \text{ReLU}((S_l z^{(j)}) \circ (\theta_l \phi_{l-1} + b_l) + \psi_l^{(j)})$, where $\psi_l^{(j)} = W_l z^{(j)}$, and $(W_l)_{l=1}^L$ and $(S_l)_{l=1}^L$ are parameters of the hypernetwork and \circ is the Hadamard product.

In Table 18, we conduct additional experiments on the *Electricity* dataset in the forecasting setting with different time horizons. In these experiments, we compare two scenarios: one where the INR is modulated only by a shift factor and the other where the INR is modulated by both a shift and a scale factor. We kept the architecture and hyperparameters consistent with those described in Appendix H.1. The experiments shown in Table 18 indicate that the INR is longer to train with shift and scale modulations due to the increased number of parameters involved. Furthermore, we observe that the shift and scale modulated INR performed similarly or even worse than the INR with only shift modulation. These two drawbacks, namely an increased computational time and similar or worse performances, motivate modulating the INR only by a shift factor.

Table 18: Ablation on modulations for the forecasting task on *Electricity* dataset for different horizons. Models are trained on a given time window and tested on four new time windows. Models are trained on a single NVIDIA TITAN RTX GPU.

	96		192		336		720	
	MAE	Training time	MAE	Training time	MAE	Training time	MAE	Training time
Shift	0.233 ± 0.014	2h30	0.245 ± 0.016	2h31	0.264 ± 0.020	2h33	0.303 ± 0.041	2h46
Shift and scale	0.257 ± 0.019	3h29	0.263 ± 0.014	3h32	0.268 ± 0.025	3h45	0.308 ± 0.037	4h14

H.2.6 DISCUSSION ON OTHER HYPERPARAMETERS

While the dimension of z is indeed a crucial hyperparameter, it is important to note that other hyperparameters also play a significant role in the performance of the INR. For example, the number of layers in the FFN directly affects the ability of the model to fit the time series. In our experiments, we have observed that using five or more layers yields good performance, and including additional layers can lead to slight improvements in the generalization settings.

Similarly, the number of frequencies used in the frequency embedding is another important hyperparameter. Using too few frequencies can limit the network’s ability to capture patterns, while using too many frequencies can hinder its ability to generalize accurately.

The choice of learning rate is critical for achieving stable convergence during training. Therefore, in practice, we use a low learning rate combined with a cosine annealing scheduler to ensure stable and effective training.

I DATASETS, SCORES AND NORMALIZATION

For the complete datasets, *Electricity* dataset is available [here](#), *Traffic* dataset [here](#) and *Solar* data set [here](#).

Datasets information The *Electricity* dataset comprises hourly electricity load curves of 321 customers in Portugal, spanning the years 2012 to 2014. The *Traffic* dataset is composed of hourly road occupancy rates from 862 locations in San Francisco during 2015 and 2016. Lastly, the *Solar* dataset contains measurements of solar power production from 137 photovoltaic plants in Alabama, recorded at 10-minute intervals in 2006. Additionally, we have created an hourly version, *SolarH*, for the sake of consistency in the forecasting section. These datasets exhibit diversity in various characteristics: • They exhibit diverse temporal frequencies, including daily and weekly seasonality observed in the *Traffic* and *Electricity* datasets, while the *Solar* dataset possesses only daily frequency. • There is individual variability across data samples and more pronounced trends in the *Electricity* dataset compared to the *Traffic* and *Solar* datasets. Table 19 provides a concise overview of the main information about the datasets used for forecasting and imputation tasks.

Table 19: Summary of datasets information

Dataset name	Number of samples	Number of time steps	Sampling frequency	Location	Years
Electricity	321	26 304	hourly	Portugal	2012 – 2014
Traffic	862	17 544	hourly	San Francisco bay	2015 – 2016
Solar	137	52 560	10 minutes	Alabama	2006
SolarH	137	8 760	hourly	Alabama	2006

How TimeFlow relative improvement score is computed In many paper tables, the TimeFlow improvement score appears in the last row of the table. Its purpose is to quantify the average marginal gain of TimeFlow over the method under consideration. It is computed as follows:

$$\text{TimeFlow improvement} = \frac{1}{L} \sum_{l=1}^L \frac{s_l(\text{baseline}) - s_l(\text{TimeFlow})}{s_l(\text{baseline})}$$

- s stands for the Mean Absolute Error score of the considered method against the ground truth at line l (for instance in Table 1, $s_1(\text{TimeFlow})= 0.324$, $s_2(\text{TimeFlow})= 0.250$ etc).
- L stands for the number of line in the table (for instance 15 in Table 1).

z-normalization To preprocess each dataset, we apply the widely used z-normalization technique per-sample j on the entire series: $x_{norm}^{(j)} = \frac{x^{(j)} - \text{mean}(x^{(j)})}{\text{std}(x^{(j)})}$.