
Recurrent Neural Processes

Timon Willi
NNAISENSE SA
timon@nnaisense.com

Jonathan Masci
NNAISENSE SA
jonathan@nnaisense.com

Jürgen Schmidhuber
NNAISENSE SA
juergen@nnaisense.com

Christian Osendorfer
NNAISENSE SA
christian@nnaisense.com

Abstract

We extend Neural Processes (NPs) to sequential data through Recurrent NPs or RNPs, a family of conditional state space models. RNPs model the state space with Neural Processes. Given time series observed on fast real-world time scales but containing slow long-term variabilities, RNPs may derive appropriate slow latent time scales. They do so in an efficient manner by establishing conditional independence among subsequences of the time series. Our theoretically grounded framework for stochastic processes expands the applicability of NPs while retaining their benefits of flexibility, uncertainty estimation, and favorable runtime with respect to Gaussian Processes (GPs). We demonstrate that state spaces learned by RNPs benefit predictive performance on real-world time-series data and nonlinear system identification, even in the case of limited data availability.

1 Introduction

The surge of neural latent variable models for performing inference on stochastic processes is a recent development in Deep Learning. Neural Processes (NPs) [Garnelo et al., 2018b] combine the best of neural networks and GPs, estimating uncertainty and model distributions over functions, while still achieving $\mathcal{O}(n + m)$ runtime instead of $\mathcal{O}((n + m)^3)$, where n is the number of context points and m the number of target points [Garnelo et al., 2018a].

Our novel Recurrent NPs (RNPs) transfer the benefits of NPs to deep generative state-space models by modeling the state space with NPs. By introducing a hierarchy of NPs, the RNPs learn in the space-time domain and integrate information in space and time into a general model.

By modeling the state space with Neural Processes, we can capture multiple time scales. Consider a time series of 10 years of hourly temperature measurements. Daily cycles (warm days, cold nights) are superimposed by seasonal cycles (winter, summer). RNPs can derive slow latent time scales (reflecting winter and summer) despite fast real-world time scales (reflecting hourly observations). RNPs retain the benefits from NPs of (a) estimating uncertainty, (b) modeling distributions over functions, and (c) high flexibility at test time without increasing runtime. RNPs enable efficient inference of the latent space by establishing conditional independence among subsequences in a given time series. We apply RNPs to various real-world one-step-look-ahead prediction tasks, illustrating their performance and benefits.

2 Background

The name "Neural Process" has two components. The "Process" comes from the fact that we are learning a stochastic process (SP). A stochastic process $\{Z_t\}_{t \in T}$ is defined as a collection of random

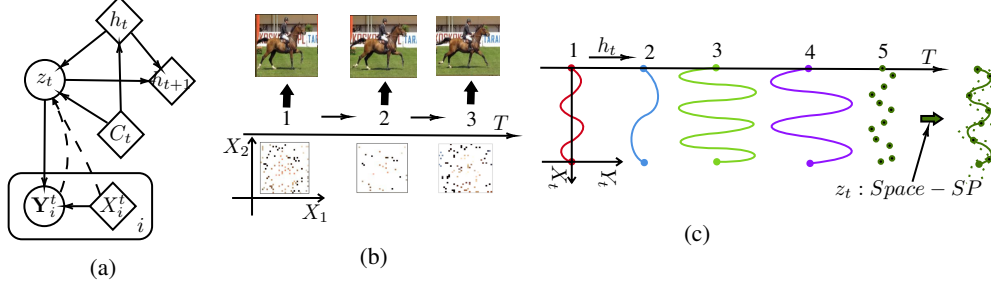


Figure 1: (a) Graphical Model of Eq. (2). Figure adapted from [Singh et al., 2019]. (b) An example application for the RNP, where we want to predict the rest of the frame at each time step of a video given some context pixels. This is just an example and does not display the real performance of the RNP. (c) An example application of the graphical model displayed in 1a. We are given a sequence of fixed-length sample functions from a stochastic process slowly changing over time. The z_t models the relations in the sample function but does not model the relations between the different time steps.

variables on a probability space (Ω, \mathcal{F}, P) [Lamperti, 2012]. The random variables Z_t are indexed by some set T . The random variables Z_t take values in a common measure space (Ξ, \mathcal{Z}) . Consider that $Z_t(\omega)$ takes two arguments: $Z(t, \omega)$, where $\omega \in \Omega$. If we fix t , we get the standard random variable $Z_t(\omega)$ from before. If we fix ω , $Z(t)$ becomes a function from T to Ξ . Basically, a stochastic process is a function-valued random variable. That is why the realization of a stochastic process is called a sample function $Z : T \rightarrow \Xi$.

The "Neural" in Neural Processes comes from the fact that we are learning the SP with a Neural Network. To learn a stochastic process means to learn a distribution over functions mapping from an input $x \in \mathbb{R}^{d_x}$ to a random variable $Y \in \mathbb{R}^{d_y}$. To learn the distribution over functions, the NPs are cast as conditional latent variable models trying to learn the conditional probability distribution

$$P(Y_T | X_T, C) = \int P(Y_T | X_T, z) P(z | C) dz \quad (1)$$

where $C = (X_C, Y_C) = (x_i, y_i)_{i \in I(C)}$ are the context and $D = (X_T, Y_T) = (x_i, y_i)_{i \in I(D)}$ are the target points, and z is some latent variable assumed to represent the parameterization of the stochastic process $P : \{y_x\}_{x \in X}$. The $I(\cdot)$ is returning the desired indices of a dataset. The points in C and D are assumed to come from the true stochastic process P .

3 Recurrent Neural Processes

Next up we want to derive the Recurrent Neural Processes. The first step is to introduce time-dependent variants of the components of the NP. Therefore we define a time-dependent context C_t , where t is the time step. The observations in C_t , D_t are from a stochastic process P_t , instead of only P . In other words, we impose a temporal structure $P_{t-1} \rightarrow P_t$ with the underlying assumption that at each new timestep, there is a different stochastic process. Equivalently to the NPs, the processes P_t mentioned above are represented as distributions on latent variables z_t . The sequential nature of the data leads to the assumption that it is helpful to make the current z_t dependent, not only on the current context C_t , but also on previous timesteps $z_{<t}$, resulting in $P(z_t | z_{<t}, C_t)$. In order to access information from the previous step, we need to maintain a recurrent state throughout the sequence. Trivially the observation model is $P(Y_t | X_t, z_t)$. The bold letters indicate that the variables are observable. From here it follows that the generative process of the new model is

$$P(\mathbf{Y}, \mathbf{Z} | \mathbf{X}, C) = \prod_{t=1}^T P(\mathbf{Y}_t | \mathbf{X}_t, z_t) P(z_t | z_{<t}, C_t) \quad (2)$$

The graphical model of the generative process is shown in Fig. 1a. In comparison to the NP in Eq. 1 we add recurrence to the prior on z . By following the above intuitive and straightforward steps, we end up with a new NP model family able to deal with stochastic processes changing over time. Imagine the data points represent frames in a video, a sequence of images, as shown in Fig 1b. For

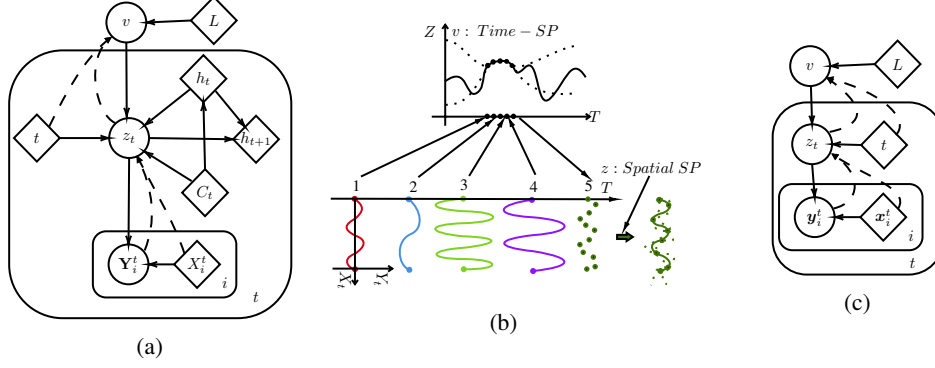


Figure 2: (a) Graphical model of the RNP with a temporal and spatial stochastic process. In this case, we have no recurrence over the temporal stochastic process. (b) For each data point, we infer its value in the temporal stochastic process. Here the observed timeline aligns perfectly with the latent timeline. This means that each $Z(t)$ parameterizes exactly one observed time step. We use the data points that we already know to infer the temporal SP. (c) In this graphical model, we have no recurrence over the spatial or the temporal SP.

each frame, we are given some context points. Our goal is to fill in the remaining pixels for each frame. This task accentuates the fact that the variables inside of the data points do not necessarily have a temporal dependency between each other. For example, in a frame of a video at a time step t , no pixel in that frame is generated before any other pixel in the same frame. The stochastic process z_t of Eq. (2) is, therefore, keeping track of the spatial relationships between the pixels of the same frame at time step t . The stochastic process z_t is not modeling the temporal relationships of pixels between different time steps. It is the recurrent state that captures all the temporal relationships between pixels at different time steps.

In theory, the recurrent state of Eq. (2) can model all necessary dynamics. Also, in practice, the model shows good performance [Singh et al., 2019]. However, a recurrent hidden state has no principled way of propagating forward uncertainty over the temporal dynamics. Furthermore, the model from Eq. (2) lacks expressivity. It is not bringing the spirit of Neural Processes to sequences. It merely uses a recurrent state to take care of the sequence dynamics and uses the stochastic process for the already known problem domain. A model true to the idea of NPs would make use of the stochastic process also for modeling the temporal dynamics. We want to add a temporal stochastic process $v : \{z_t\}_{t \in T}$ that affects all of the spatial stochastic processes z_t , as motivated in Fig. 2b. This leads to a new generative process

$$P(\mathbf{Y}, V, Z | \mathbf{X}, C, T, L) = \prod_{t=1}^T \left(\prod_{i \in I(D_t)} P(\mathbf{y}_{\mathbf{x}, i}^t | \mathbf{x}_i^t, z_t) \right) P(z_t | z_{<t}, C_t, v, t) P(v | L) \quad (3)$$

where $L = (T_L, Z_L) = (t, z_t)_{t \in I(L)}$ is the context for the temporal stochastic process v . We infer the temporal context points z_t in L from the spatial context points in C_t . The graphical model is shown in Fig. 2a. Since the temporal stochastic process v is now capturing the changing dynamics, the recurrence over the spatial SP z_t in Equation (3) might not be necessary. Leading to:

$$P(\mathbf{Y}, V, Z | \mathbf{X}, C, T, L) = \prod_{t=1}^T \left(\prod_{i \in I(D_t)} P(\mathbf{y}_{\mathbf{x}, i}^t | \mathbf{x}_i^t, z_t) \right) P(z_t | C_t, v, t) P(v | L) \quad (4)$$

The corresponding graphical model is shown in Fig. 2c.

Eq.(3) and (4) have a subtle aspect that turns out to be important for modelling sequences. The temporal stochastic process $v : \{z_t\}_{t \in T}$ shares the same temporal indexing T with our observed sequence of data points $\{(\mathbf{Y}_i^t, \mathbf{X}_i^t)\}_{i=1}^N\}_{t \in T}$, where N is the size of the data points. That means latent and observed time steps are aligned; one input sequence $\{(\mathbf{Y}_i^t, \mathbf{X}_i^t)\}_{i=1}^N\}_t$ corresponds to one z_t . Clearly, this does not need to be the general case. For example, let us say the input sequences $\{(\mathbf{Y}_i^t, \mathbf{X}_i^t)\}_{i=1}^N\}_{t \in T}$ correspond to sine waves and z_t corresponds to the amplitude of the sine wave. It is not unreasonable to assume that one z_t might correspond to multiple input sequences. To

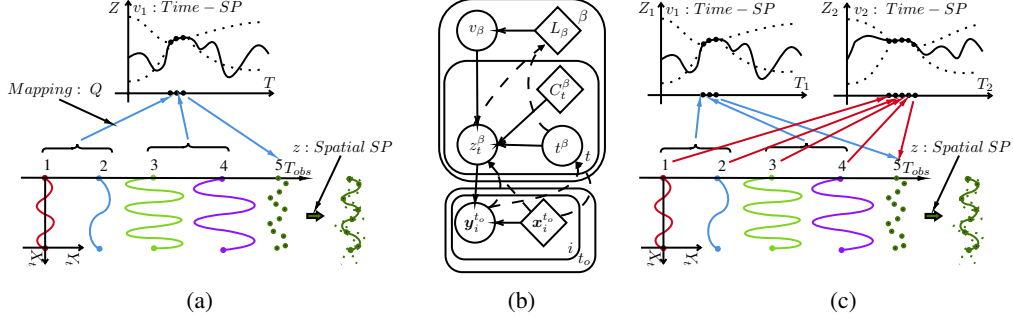


Figure 3: (a) We map two time steps of the observed time to one time step of the latent time. The mapping is indicated by the blue arrows. (b) The graphical model of the RNP accounting for multiple temporal SPs. By introducing multiple temporal SPs, we explicitly model for different latent timelines. (c) Visualization of the RNP with multiple temporal SPs.

differentiate between the observed time and the latent time, we introduce a new time line T_{obs} , which is the temporal indexing of the observed data $\{(\mathbf{Y}_i^{t_{obs}}, \mathbf{X}_i^{t_{obs}})\}_{i=1}^N\}_{t \in T_{obs}}$. The time line T stays the time line of the temporal stochastic process $v : \{z_t\}_{t \in T}$. In general, we do not know how T and T_{obs} relate to each other. As a consequence, we have to find a mapping between the different indices $T \iff T_{obs}$ to properly assign the credit of influence. We propose to do so via implicitly learning the connections in a bipartite graph Q between two index sets. One set is formed by T_{obs} and the other by T . The connections in-between the two disjoint sets have to satisfy the temporal causal constraints that when t maps to t_{obs} , $t + 1$ cannot map to $t'_{obs} \leq t_{obs}$.

$$P(Q | T, T_{obs}) = \prod_{t \in T} \prod_{t_{obs} \in T_{obs}} P(Q_{t, t_{obs}} | t_{obs}, t) \quad (5)$$

We can thus write the model as follows

$$P(\mathbf{Y}, V, Z, T | \mathbf{X}, T_{obs}, C) = \prod_{t_{obs} \in T_{obs}} \prod_{t \in T} \prod_{i \in I(D_{t_{obs}})} P(\mathbf{y}_{\mathbf{x}, i}^{t_{obs}} | \mathbf{x}_i^{t_{obs}}, z_t) P(z_t | z_{< t}, C_{t_{obs}}, v, t) P(Q | t_{obs}, t) P(v, t | L) \quad (6)$$

One can imagine the model with or without recurrence at this point. The model is displayed in Figure 3a. Interestingly we can keep the recurrence on the spatial stochastic process. The updates happen at the frequency of the corresponding latent time T and not T_{obs} . This is similar to other previously proposed methods where updates happen at different discrete time steps for different units in RNNs [Schmidhuber, 1991, 1992, El Hihi and Bengio, 1996, Koutnik et al., 2014].

Eq. (7) only accounts for one latent time scale. Again, this might be overly restrictive: In the temperature example, we would want to model multiple times. The daily temperature fluctuations tick on a different time scale than the yearly seasons. In theory, it is possible for one temporal SP to account for these factors, but again, this lacks expressivity. This leads to the generative model

$$P(\mathbf{Y}, V, Z, T | \mathbf{X}, T_{obs}, C) = \prod_{t_{obs} \in T_{obs}} \prod_{\beta \in \mathcal{V}} \prod_{i \in I(D_{t_{obs}})} \prod_{t \in T^\beta} P(\mathbf{y}_{\mathbf{x}, i}^{t_{obs}} | \mathbf{x}_i^{t_{obs}}, z_t^\beta) P(z_t^\beta | z_{< t}^\beta, C_{t_{obs}}, v^\beta, t^\beta) P(Q^\beta | t^\beta, t_{obs}) P(v^\beta, t^\beta | L^\beta) \quad (7)$$

with the graphical model shown in Figure 3b and a visualization shown in Fig. 3c. Finally, we arrived at the RNP. We were able to extend the NP to arrive at the RNP and thus show a generalization of the former. Now we will highlight some important aspects of the RNP. Then we proceed with presenting the architecture and a suitable loss function to train the model.

Whereas Eq. (2) only spans the spatial dimension with a stochastic process, the RNP additionally spans an SP over time. This has a few favorable consequences. First, a temporal SP provides more options to induce biases about the problem, which can be explained with the following example. Consider Figure 4. You are given a series of temperature measurements, one measurement per hour for a few years. The data points are one-dimensional, so technically, the spatial stochastic process consists of only one variable, and the temporal stochastic process spans over the whole sequence.




Figure 4: We can build a hierarchy of Neural Processes. In the temperature example, we could use one stochastic process to model the intra-day dynamics, another one to model the intra-week dynamics. The next stochastic process of the hierarchy could capture the intra-month dynamics, and so on.

The temporal SP is, therefore, tasked to model all of the dynamics, and the spatial SP is idle. However, given our prior knowledge about the temperature data, we can offload computation from the temporal SP to the spatial SP by splitting the time series into same-sized chunks, e.g., we split them into days. One chunk is considered one data point. Now the spatial SP is responsible for modeling the intra-day dynamics of temperature and the temporal SP of the inter-day dynamics. We cannot only chunk the observed sequence, but we can also chunk the temporal stochastic process and impose another stochastic process onto those chunks. We could split the observed sequence into daily chunks, have a first level of temporal SP that models the intra-month dynamics, and a second level of temporal SP that models the inter-month dynamics and so on. This leads to an architecture we term the Hierarchical Neural Process (HNP), which is the extension of Figure 2c to an arbitrary number of SPs. If the chunks are self-similar, one could even think of a Recursive Neural Process. If we were to model recursive vegetables, we could even introduce the Romanesco Neural Process. Although the HNP was just derived from a time series example, it is not restricted to problems with a temporal structure.

The above application is also a good example to show another limitation of Eq. (2). What if we had no intuition on how to split the temperature data? An example shown in Figure 5b.

The spatial SP of Equation (2) would be tasked with modeling only one value, whereas the recurrent hidden state would have to capture all the dynamics over time. In the RNP, all of the dynamics could be modeled by the temporal SP or recurrence. To have two components available to model the dynamics is advantageous. The temporal SP might be better suited to capture uncertainty than the recurrent state. The recurrent state can help to distinguish the true signal from noise.

The second consequence of a temporal SP is the ability to use information from future timesteps, if available, in a meaningful way. For example, imagine you are in a setup such as displayed in Figure 5a. You are given information about your targeted time step from the past and from the future.

First, how would you include information from future timesteps in Eq. (2)? One could, for example, first process the available data points 1, 7, 18 and 26, appended with the time step, and then predict datapoint 13. The hidden state is then additionally tasked with keeping track of time jumps. This goes directly against the model design where we assume to receive data points in the sequence in which they were generated $P_{t-1} \rightarrow P_t$. Additionally, it is overloading the spatial stochastic process with uncertainty in time and space. However, in the RNP, the explicit indexing with time T_{obs} enables a principled way of querying the temporal stochastic process for specific positions in time. This allows using future time steps as inducing points. Therefore, when we predict timestep 13 in Figure 5a, information from the inducing points 18 and 26 of the temporal SP flows into the current prediction. Now we proceed with the architectural details of the RNP.

3.1 Architecture

We are given a sequence of one-dimensional data points, as displayed in Fig. 5b and we infer temporal stochastic processes by selecting a set of N subsequences $\{(x_j, y_j), \dots, (x_{j+J}, y_{j+J})\}_{i=1}^N$. The J is the length of the subsequences. The index j is the starting point of the subsequence and can be any legal value given the length of the time series. The selection of the subsequences is problem dependent. It is a convenient way of introducing prior knowledge about the problem into the model. The subsequences do not necessarily need to have the same length J . The formulation of the mapping allows variable length subsequences. For example, in Fig. 6a, we select subsequences as contexts (on

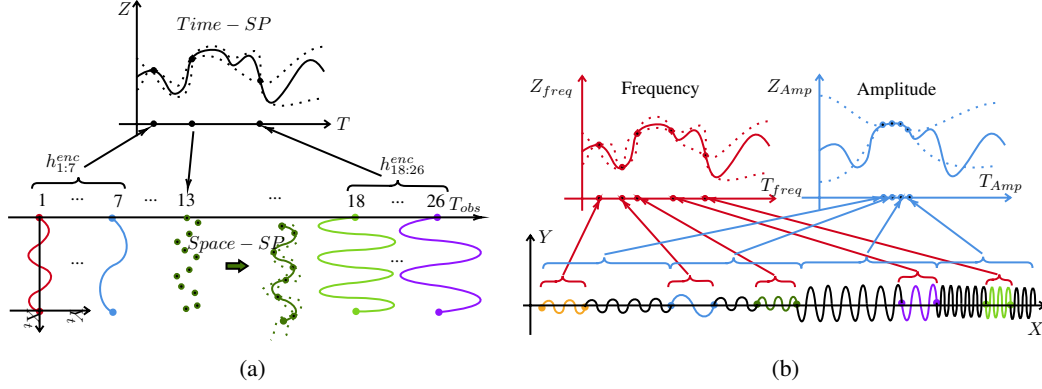


Figure 5: (a) If we are provided with observed time steps that are further apart, intuitively this should give better coverage of the temporal stochastic process. The idea of temporal SP also allows by design to account for future time steps. (b) If we do not know when the dynamics change or we have no prior knowledge about the problem setting, we have to deal with the one-dimensional time series. In this case, the spatial SP is underutilized. For example, in a sine wave we know it is parameterized through, e.g. Frequency and Amplitude. The RNP models the dynamics of the Frequency and Amplitude with a stochastic process inferred from previously seen time steps.

the left) and predict our target sequence (on the right).

The subsequences are processed by an LSTM, which will result in a sequence of hidden states $\{\{h_j^i\}_{j=1}^J\}^i = f^{\text{LSTM_enc}}(\{(x_j, y_j), \dots, (x_{j+J}, y_{j+J})\}^i)$. The hidden states are used to produce a representation $z_i = f^{\text{MLP}}(h_j^i)$ of the subsequence. There are many options on how to produce the representation. One could, for example, simply take the last hidden state h_j^i as a representation of the sequence. The representations z_i are aggregated, $\text{Agg}(z_i) = z$. A temporal SP v is sampled, $v \sim \mathcal{N}(\mu_v, \sigma_v)$, where $[\mu_v, \sigma_v] = f^v(z)$. Since we do not have multidimensional inputs, there is no spatial SP, which means there are no spatial context points C_* . So we can proceed with predicting our target value y^* at x^* with our decoder: $y_* \mid v, x_*, h_{t-1} \sim \mathcal{N}(\mu_y, \sigma_y)$, where $[\mu_y, \sigma_y] = f^{\text{LSTM_dec}}(h_{t-1}, v, x_*)$. The decoder could be either an MLP or an LSTM, because the temporal stochastic process v should be able to model all of the dynamics. We can even one-shot predict a whole sequence of length m if we are certain about the governing dynamics of the next m time steps. The architecture is visually described in the Appendix in Fig. 9.

This is the problem setting we are interested in for our experiments. That is why we provide a loss function for this model.

3.2 ELBO

We now proceed with deriving the ELBO for this specific RNP model. Note that we use indexing for both z and t , although t is already the indexing of z . This is done to stay consistent with previous literature, and it makes it clearer to distinguish between the context points and the target points. However, note that the index of t represents the value that t assumes.

Let $q(v \mid z_{1:L}, t_{1:L})$ be the variational posterior with which we can learn the non-linear decoder $g(\cdot)$. The Evidence Lower Bound then looks as follows:

$$\log p(z_{1:L} \mid t_{1:L}) \geq \mathbb{E}_{q(v \mid z_{1:L}, t_{1:L})} \left[\sum_{i=1}^L \log p(z_i \mid v, t_i) + \log \frac{p(v)}{q(v \mid z_{1:L}, t_{1:L})} \right] \quad (8)$$

Equivalently to the NPs we split the dataset into a context set $z_{1:k}, t_{1:k}$ and a target set $z_{k+1:L}, t_{k+1:L}$

$$\log p(z_{k+1:L} \mid t_{1:L}, z_{1:k}) \geq \mathbb{E}_{q(v \mid z_{1:L}, t_{1:L})} \left[\sum_{i=k+1}^L \log p(z_i \mid v, t_i) + \log \frac{p(v \mid z_{1:k}, t_{1:k})}{q(v \mid z_{1:L}, t_{1:L})} \right] \quad (9)$$

As $p(v \mid z_{1:k}, t_{1:k})$ is intractable we will use $q(v \mid z_{1:k}, t_{1:k})$

$$\log p(z_{k+1:L} \mid t_{1:L}, z_{1:k}) \geq \mathbb{E}_{q(v \mid z_{1:L}, t_{1:L})} \left[\sum_{i=k+1}^L \log p(z_i \mid v, t_i) + \log \frac{q(v \mid z_{1:k}, t_{1:k})}{q(v \mid z_{1:L}, t_{1:L})} \right] \quad (10)$$

Since we cannot directly observe the latent random functions we have to express this in terms of the observed data. We can express it with the observed data because we have a mapping Q from the observed time steps and the latent time steps. Thanks to the mapping Q we can approximate $(z_{1:k}, t_{1:k})$ with $(\mathbf{y}_{\{\mathbf{x} \in Q_{t_{1:k}}\}}, Q_{t_{1:k}})$, where $Q_{t_{1:k}} = \{\mathbf{x} \mid z \text{ at } t \text{ parameterizes } \mathbf{y} \text{ at } \mathbf{x}\}$. Basically the latent sequence provides a factorization for the joint distribution of the observed sequence. In other words, we have independent subsequences in the observed sequence data conditioned on the hidden sequence, as visually motivated in Fig. 5b.

$$\begin{aligned} & \log p(\mathbf{y}_{\{\mathbf{x} \in Q_{t_{k+1:L}}\}} \mid Q_{t_{1:L}}, \mathbf{y}_{\{\mathbf{x} \in Q_{t_{1:k}}\}}) \\ & \geq \mathbb{E}_{q(v \mid \mathbf{y}_{\{\mathbf{x} \in Q_{t_{1:L}}\}})} \left[\sum_{j=k+1}^L \left[\sum_{\mathbf{x}=\min\{\mathbf{x} \in Q_{t_j}\}}^{\max\{\mathbf{x} \in Q_{t_j}\}} \log p(\mathbf{y}_{\mathbf{x}} \mid \mathbf{x}, v) \right] + \log \frac{q(v \mid \mathbf{y}_{\{\mathbf{x} \in Q_{t_{1:k}}\}}, Q_{t_{1:k}})}{q(v \mid \mathbf{y}_{\{\mathbf{x} \in Q_{t_{1:L}}\}}, Q_{t_{1:L}})} \right] \end{aligned} \quad (11)$$

The jump from Eq. (10) to Eq. (11) is again justified via the mapping. Basically, we approximate $(z_{1:k}, t_{1:k})$ with $(\mathbf{y}_{\{\mathbf{x} \in Q_{t_{1:k}}\}}, Q_{t_{1:k}})$, where $Q_{t_{1:k}} = \{\mathbf{x} \mid z \text{ at } t \text{ parameterizes } \mathbf{y} \text{ at } \mathbf{x}\}$.

The difficult aspect of Eq. (11) is that we also have to learn the mapping to approximate the right ELBO. As stated in Eq. (11), we need to iterate over independent subsequences of our observed time series to calculate the correct loss. In theory, we would have to select the subsequences corresponding to the mapping to the hidden variables. In practice, however, we do not know what the mapping looks like. An option is to choose all possible subsequences. Another option is to select subsequences by hand.

4 Experiments

Electricity: The model is evaluated on the UCI Individual Household Electric Power Consumption Dataset [Dua and Graff, 2017]. This is a time series of 2’075’259 measurements collected at 1-minute intervals from a single household. Each measurement consists of 7 features related to power consumption. In order to compare to other works, we choose the feature “active power” as our target and the other features as inputs. We compare to earlier results [Lim et al., 2019].

Drives: The Coupled Electric Drives is a dataset for nonlinear system identification [Wigren, 2010]. It consists of 500 time steps and was produced by two motors driving a pulley using a flexible belt, where the input is the sum of the voltages on the motors, and the target is the speed of the belt. We compare our model to the results reported earlier [Mattos et al., 2015].

Metrics: We evaluate model accuracy using the Mean Squared Error (**MSE**) between the predicted mean and the target value. The MSEs are normalized based on LSTM performance, which allows for comparing to previous work. The prediction interval coverage probability (**PICP**) is a way of measuring the quality of the predicted uncertainty. We measure the performance on a 90% prediction interval. It is defined as $\text{PICP} = \frac{1}{T} \sum_{t=1}^T c_t$, where $c_t = 1$ if $\psi(0.05, t) < y_t < \psi(0.95, t)$ and $c_t = 0$ otherwise. $\psi(0.05, t)$ is the 5th percentile derived from predictions sampled from $N(f(\mathbf{x}_t), \sigma)$ [Lim et al., 2019].

5 Results

The results in Table 1 and Table 2 suggest that the performance of the RNP model is comparable to that of other SSMs. The outcome for the Electricity dataset indicates that the conditioning of the LSTM decoder was deceptive, rather than informative. This could be due to multiple factors (we did not choose an exhaustive amount of subsequences or sample them intelligently). For Drives, RNPs outperform LSTM and suggest informative uncertainty measures. For a qualitative analysis, see Figures 6a-6b. The model learns to capture the time series. Figures 6a-6c show that the RNP can predict the start of the target sequence and adapt its uncertainty.

Table 1: Normalized MSE for One-Step Predictions

	Electricity	Drives
LSTM	1.000	1.000
VRNN	1.902	-
DKF	1.252	-
DSSM	1.131	-
RNF-LG	0.918	-
RNF-NP	0.856	-
MLP-NARX	-	1.017
GP-NARX	-	0.953
REVARB	-	0.462
RNP	1.111	0.238

Table 2: PICP for One-Step Prediction

	Electricity	Drives
VRNN	0.986	-
DKF	1.000	-
DSSM	0.964	-
RNF-LG	0.960	-
RNF-NP	0.927	-
RNP	0.947	0.874

6 Conclusion

One significant benefit of Neural Processes is its improved computational cost over Gaussian Processes. This benefit becomes especially prominent when stacking the stochastic processes like in the Hierarchical Neural Process, which we hinted at in this thesis. The meta-learning background of the Neural Process also helps in time series analysis. Some regular events happen very rarely. The switch from spring to summer happens only once a year, or the crash of the stock market happens every decade or so. Since the Neural Process was designed to deal with a lack of data, these regular but rare events are, at least in theory, more easily captured. The good performance on the Drives dataset also indicates that the RNP framework can profit from the meta-learning background.

To conclude, we introduce a family of models called Recurrent Neural Processes (RNPs), a generalization of Neural Processes (NPs) to sequences by introducing a notion of latent time through modeling the state space with NPs, with a wide range of applications. RNPs can derive appropriate slow latent time scales from long sequences of quickly changing observations hiding specific long term patterns.

The framework derived for RNPs enables efficient inference of temporal context by establishing conditional independence among subsequences in a given time series. It also provides an appropriate loss function for training RNPs. RNPs are also not restricted to modeling only two stochastic processes. The framework can be arbitrarily expanded to a hierarchy of stochastic processes, as shown in Fig. 4.

Acknowledgments

We would like to thank Florian Triffterer, Pranav Shyam, Giorgio Giannone, Jan Eric Lenssen, David Ackermann and Heng Xin Fun for insightful discussions, and everyone at NNAISENSE for being part of such a conducive research environment.

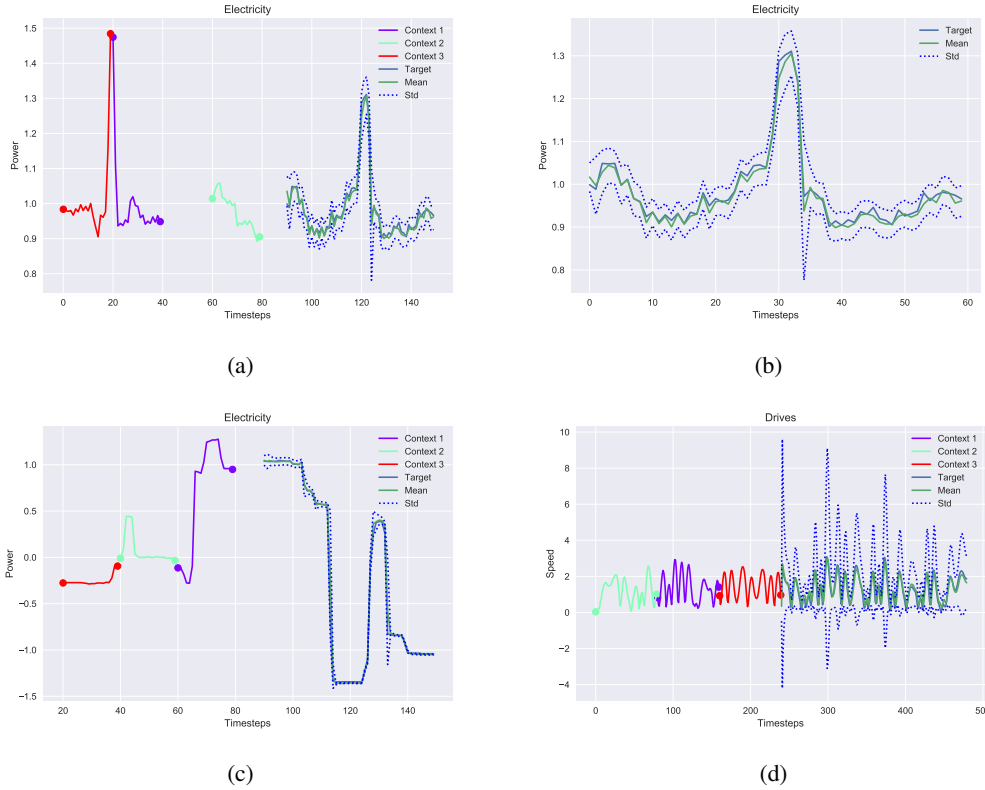


Figure 6: (a) A sample of the predictive performance on the Electricity dataset. The green line depicts the mean; the blue dotted line depicts one standard deviation. The blue line is the target. The multi-colored segments are randomly sampled subsequences used as context. (b) Close-up of the predicted sequence from the previous figure. (c) A second sample of the predictive performance on the Electricity dataset. The model is able to predict the start of the target sequence, given the context sequences. (d) The model is able to approximate the second half of the Drives dataset. However, it is relatively uncertain about the start point.

References

- Dheeru Dua and Casey Graff. UCI machine learning repository - individual household electric power consumption data set, 2017., 2017. URL <http://archive.ics.uci.edu/ml>.
- Salah El Hihi and Yoshua Bengio. Hierarchical recurrent neural networks for long-term dependencies. In *Advances in neural information processing systems*, pages 493–499, 1996.
- Marta Garnelo, Dan Rosenbaum, Chris J. Maddison, Tiago Ramalho, David Saxton, Murray Shannahan, Yee Whye Teh, Danilo J. Rezende, and S. M. Ali Eslami. Conditional Neural Processes. jul 2018a. URL <http://arxiv.org/abs/1807.01613>.
- Marta Garnelo, Jonathan Schwarz, Dan Rosenbaum, Fabio Viola, Danilo J. Rezende, S. M. Ali Eslami, and Yee Whye Teh. Neural Processes. jul 2018b. URL <http://arxiv.org/abs/1807.01622>.
- Greff, Aaron Klein, Martin Chovanec, Frank Hutter, and Jürgen Schmidhuber. The Sacred Infrastructure for Computational Research. In Katy Huff, David Lippa, Dillon Niederhut, and M Pacer, editors, *Proceedings of the 16th Python in Science Conference*, pages 49 – 56, 2017. doi: 10.25080/shinma-7f4c6e7-008.
- Hyunjik Kim, Andriy Mnih, Jonathan Schwarz, Marta Garnelo, Ali Eslami, Dan Rosenbaum, Oriol Vinyals, and Yee Whye Teh. Attentive Neural Processes. jan 2019. URL <http://arxiv.org/abs/1901.05761>.

- Jan Koutník, Klaus Greff, Faustino Gomez, and Juergen Schmidhuber. A clockwork rnn. *arXiv preprint arXiv:1402.3511*, 2014.
- John Lamperti. *Stochastic processes: a survey of the mathematical theory*, volume 23. Springer Science & Business Media, 2012.
- Tuan Anh Le, Hyunjik Kim, Marta Garnelo, Dan Rosenbaum, Jonathan Schwarz, and Yee Whye Teh. Empirical Evaluation of Neural Process Objectives. Technical report, 2018. URL <http://bayesiandeeplearning.org/2018/papers/92.pdf>.
- Bryan Lim, Stefan Zohren, and Stephen Roberts. Recurrent Neural Filters: Learning Independent Bayesian Filtering Steps for Time Series Prediction. jan 2019. URL <https://arxiv.org/abs/1901.08096>.
- César Lincoln C. Mattos, Zhenwen Dai, Andreas Damianou, Jeremy Forth, Guilherme A. Barreto, and Neil D. Lawrence. Recurrent Gaussian Processes. nov 2015. URL <https://arxiv.org/abs/1511.06644>.
- J. Schmidhuber. Learning complex, extended sequences using the principle of history compression. *Neural Computation*, 4(2):234–242, 1992.
- Jürgen H. Schmidhuber. Neural sequence chunkers. Technical Report FKI-148-91, Technische Universität München, Munich, Germany, 1991.
- Gautam Singh, Jaesik Yoon, Youngsung Son, and Sungjin Ahn. Sequential Neural Processes. June 2019. URL <https://arxiv.org/abs/1906.10264>.
- Torbjörn Wigren. Input-output data sets for development and benchmarking in nonlinear identification. 2010.

Appendix

Overview

The Supplementary Material presents more details about the RNP’s implementation, the parameter search, and the experiments.

Experimental Details

For all three experiments, we performed a coarse hyperparameter search in a grid with the following possible value ranges. Minibatch size was always 256 except for the smaller Drives set, where it was 8. We trained the RNP using teacher forcing. The experiments were conducted on p3.16xlarge instances on the Amazon Web Services.

Table 3: Parameter space searched during hyperparameter search.

	Possible Values
LSTM Hidden State Size	3, 16, 32, (64)
Context Sequence Length	5, 50, 150 / 20, 60, 80
Latent Vector Size	4, 32, 128
Bidirectional LSTM-Encoder	Yes, No
LSTM Layers	1, 2

To manage the experiments and ensure reproducibility we used Sacred [Greff et al., 2017].

Electricity


We reproduced the LSTM baseline model of previous work [Lim et al., 2019] as accurately as possible, to compare our model to the reported normalized performances. For dataset preprocessing and architectural details of the other models we refer to the appendix of the aforementioned paper. We tried multiple context sequence lengths, {50, 150, 300}, to train and test the network, but length did not have a significant impact on performance. Experiments on this dataset profited from an increased hidden state and a decreased latent vector size. We trained the model for 130 epochs.

Drives


For the Drives dataset, we relied on results reported by previous work [Mattos et al., 2015]. Parameter search favored a small hidden state size and a larger latent vector size. Due to the dataset size, we trained the model on shorter context and target sequences of length 5 and 15 respectively. For testing, we used longer sequences to which the model was able to adapt. We trained it for 100’000 epochs.

Table 4: We used these hyperparameters for our final performance comparison.

	LSTM		RNP	
	Electricity	Drives	Electricity	Drives
Hidden State Size	50	up to 2048	64	32
Context Sequence Length	-	-	20	80
Test Sequence Length	50	250	60	250
Latent Vector Size	-	-	4	32
Bidirectional LSTM-Encoder	-	-	Yes	No
LSTM Layers	1	up to 3	2	2
Learning Rate	0.1	-	0.01	0.0001



(a) Learning curve of RNP on training and validation data of the Electricity dataset. Keep in mind that we do not optimize for MSE, but for log likelihood.




(b) Learning curve of LSTM on training and validation data of the Electricity dataset, optimized for MSE.




(c) Learning curve of RNP on training and validation data of the Drives dataset.

Figure 7: Training and Validation Metrics for RNP and LSTM on the Electricity ((a), (b)) and Drives dataset ((c))



(a) We can see that the network is able to capture the time series in its bounds of one standard deviation.



(b) The model captures the system dynamics; the apparent lag in prediction is possibly due to training by teacher forcing.

Figure 8: Zooming in on RNP's predictive performance on Electricity ((a)) and Drives ((b))

Architecture

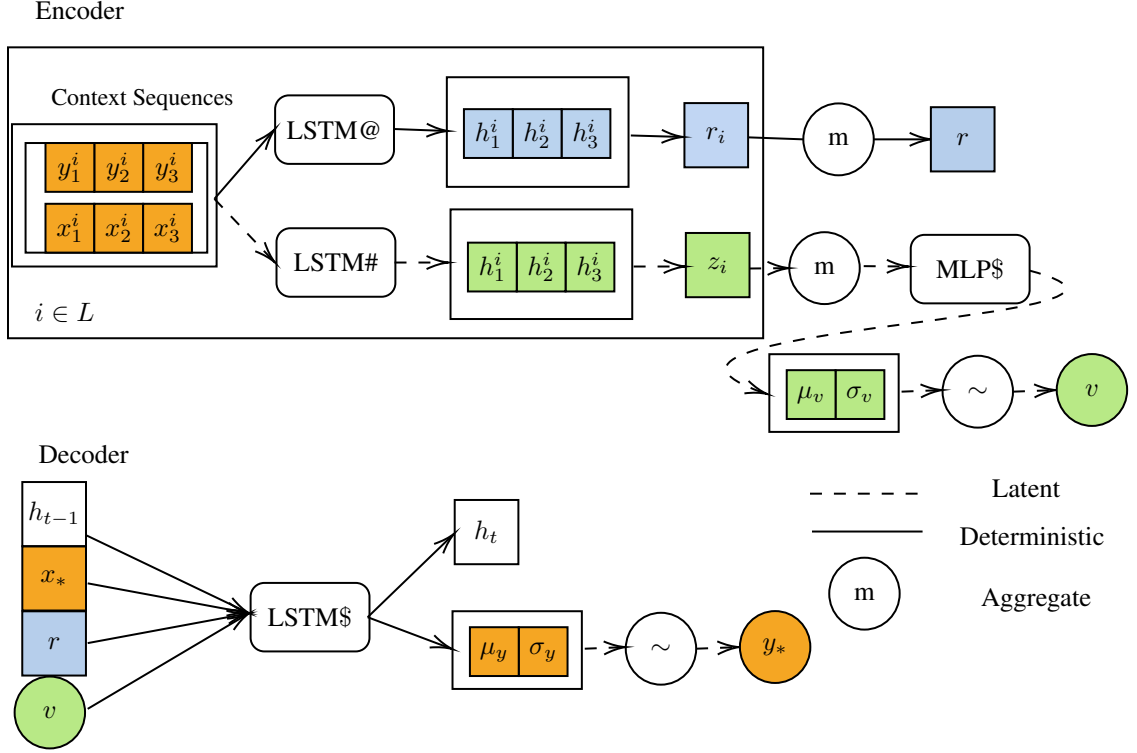


Figure 9: The architecture displayed here was used in the experiments, as explained in the Architecture section. We infer the temporal stochastic process from hand-selected subsequences. Adapted from Kim et al. [2019]

The architecture depicted in Figure 9 exhibits a stochastic and a deterministic path. Having both paths was reported to be beneficial [Le et al., 2018]. Each path has its own encoder. Subsequences fed into the left will be encoded by LSTMs providing sequences of hidden states.

Bidirectional paths yield two sequences of hidden states that one could encode in various ways. Our final model uses only the last hidden state and the corresponding input.

The codes are aggregated and fed into the LSTM decoder at each time step. Inferring the first hidden state of the decoder from the representation was found to be helpful.