

# Contrastive Conditional Neural Processes

Zesheng Ye

Lina Yao

The University of New South Wales

zesheng.ye@unsw.edu.au

lina.yao@unsw.edu.au

## Abstract

*Conditional Neural Processes (CNPs) bridge neural networks with probabilistic inference to approximate functions of Stochastic Processes under meta-learning settings. Given a batch of non-i.i.d function instantiations, CNPs are jointly optimized for in-instantiation observation prediction and cross-instantiation meta-representation adaptation within a generative reconstruction pipeline. There can be a challenge in tying together such two targets when the distribution of function observations scales to high-dimensional and noisy spaces. Instead, noise contrastive estimation might be able to provide more robust representations by learning distributional matching objectives to combat such inherent limitation of generative models. In light of this, we propose to equip CNPs by 1) aligning prediction with encoded ground-truth observation, and 2) decoupling meta-representation adaptation from generative reconstruction. Specifically, two auxiliary contrastive branches are set up hierarchically, namely in-instantiation temporal contrastive learning (TCL) and cross-instantiation function contrastive learning (FCL), to facilitate local predictive alignment and global function consistency, respectively. We empirically show that TCL captures high-level abstraction of observations, whereas FCL helps identify underlying functions, which in turn provides more efficient representations. Our model outperforms other CNPs variants when evaluating function distribution reconstruction and parameter identification across 1D, 2D and high-dimensional time-series.*

## 1. Introduction

Supervised generative models learn to reconstruct joint distribution with certain prior incorporated to achieve parameter estimation. This brings in a natural fit to meta-learning [2, 12, 30] paradigm, where knowledge acquired from previous tasks can help fast adaptation in solving novel tasks drawn from the same task distribution. Conditional Neural Processes [6, 9, 16] (CNPs) lie at the intersection between these two. Mathematically, for  $T$  steps of a function instantiation  $D = \{(x_t, \mathbf{y}_t)\}_{t=1}^T$  sampled from a time-

series, where each time index  $x$  has an associated observation  $y \subseteq \mathbb{R}^d$ , CNPs learn to reconstruct the observation of a given query index  $x_q$ , with a set of time-observation pairs  $\mathcal{C} \subseteq D$  as input. Specifically, CNPs are built upon an encoder-decoder, where the encoder  $h : \mathcal{C} \rightarrow \mathbf{r}_C$  summarizes all the input as a contextual representation  $\mathbf{r}_C$ . The decoder takes  $\mathbf{r}_C$  and query index  $x_q$  and outputs estimated observation  $g : \mathbf{r}_C, x_q \rightarrow \hat{\mathbf{y}}_q \approx \mathbf{y}_q$ . In a meta-learning context, CNPs may deal with some non-i.i.d function instantiations. If, for instance, given two function instantiations  $D_1 \sim f_1 = 5x + 10$  and  $D_2 \sim f_2 = 8x + 1$ , the representation derived from  $f_1$  cannot be used directly to reconstruct observations of  $f_2$ , but rather requires a cross-instantiation adaptation step for meta-representation. In CNPs, the adaptation takes place implicitly along with data reconstruction.

The Bayesian principles enable CNPs to quantify uncertainties while performing predictive tasks under limited data volume. Moreover, CNPs are designed to model a distribution over functions whereas conventional deep learning models do not, suggesting better generalization when multiple functions are in play. These properties are desirable in applications like traffic forecasting, trajectory prediction, and activity recognition [13, 21, 24], where observations might be corrupted or sampled from different data-generation functions.

The CNP's variants attempt to improve the encoder  $h$  by introducing a range of inductive biases [9, 15, 16] to learn appropriate structures from data. Despite the performance gains achieved, the inherent generative nature of these models has raised some concerns, especially when they are considered together: 1) Correlations between observations are not modeled in CNPs. While introducing a global latent variable can yield more coherent predictions [4, 7, 16], these attempts are also fraught with intractable likelihood; 2) High-dimensional observations challenge the capacity of CNPs since generative models reconstruct low-level details but hardly form high-level abstractions [17, 23]; and 3) Supervision collapse in meta-learning [3] occurs when prediction and transfer tasks are entangled in CNPs [8].

To address this, our first motivation centers on decoupling model adaptation from generative reconstruction

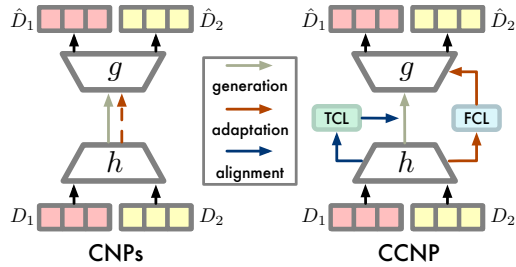


Figure 1. The motivation of the proposed CCNP against CNPs. While adaptation is **implicitly** done within generation in CNPs, CCNP **explicitly** processes it with FCL branch, and TCL branch ensures the alignment of high-level abstraction of observations.

through explicit contrastive learning, with the aim of making each objective attends to its own duties. Contrastive learning has become increasingly popular thanks to its clearer (discriminative) supervisions. Despite being implemented differently in a range of applications [1, 5, 34], the shared principle behind these contrastive methods is to create different types of tuples by utilizing certain structures in the data and training the model to identify the types. For instance in SimCLR [1], images are firstly transformed randomly with two augmentations, which are then formed as contrastive pairs labeled as either positive or negative, depending on whether two augmentations are from the same instance. Likewise, contrastive learning has renewed a surge of interest in conditional density estimation [22, 26, 27]. In this scope, discriminative pretext tasks are designed to create more efficient representations for modeling higher-dimensional data than generative counterparts [19]. One should however note that they do so at the expense of flexibility, as only downstream tasks that are directly related to the contrastive task can be accommodated.

As such, we investigate a generative-contrastive model in an effort to combine the complementary benefits of both approaches while mitigating the deficiencies of each. Existing studies linking contrastive learning to generative models spotlight the “pre-train then fine-tune” scheme [25, 32]. Contrary to this, we aim to train a hybrid model in an end-to-end way, without losing the flexibility of generative models, as well as the efficiency of contrastive models.

To this end, we propose Contrastive Conditional Neural Process (CCNP) that extends generative CNPs with two auxiliary contrastive branches, coined *in-instantiation temporal contrastive learning* (TCL) and *cross-instantiation function contrastive learning* (FCL). A particular emphasis is placed on the use of contrastive objectives to handle noisy high-dimensional observations. Concretely, TCL imposes *local* alignment between predictive representations and encoded ground-truth observations of the same timestep to model temporal correlations, enhancing the model’s scalability to higher dimensions. FCL encourages *global* consistency

across different partial views of the same instantiation to separate adaptation from reconstruction, thus improving the transferability of meta-representations.

**Contributions.** 1) We present an end-to-end generative-contrastive meta-learning model based on CNPs. 2) We demonstrate that incorporating proper contrastive objectives into generative CNPs contributes to complementary benefits, and hopefully leads to new avenues for research in CNPs. 3) We empirically verify CCNP outperforms other CNP baselines in terms of function distribution reconstruction and parameter identification across diverse tasks, including 1D few-shot regression, 2D population dynamics prediction and high-dimensional sequences reconstructions.

## 2. Related Works

**Conditional Neural Processes.** Conditional Neural Process (CNP) [6] was proposed to perform function approximation under a meta-learning setting. Meta-learning refers to a learning paradigm that facilitates rapid model adaptation across multiple related tasks. This is typically implemented via a bi-level learning setting, where the model solves predictive tasks within inner-level loop, while optimizing the generalization ability during outer-level learning. CNPs wrap up the inner step with outer-level optimization based on a simple encoder-decoder architecture. Recent progress includes baking a variety of inductive biases into the model so that the corresponding structure can be learned from data. AttnCNP [16] uses attention [29] to improve interpolation performance inside the range of training input. ConvCNP [9] realizes translation-equivariance and enhances extrapolation capability outside the training range. The latest evolution involves imposing symmetries on CNPs, *e.g.* group equivariance in [15] and SteerableCNP [11] to ensure spatial invariance. Even so, the majority of their efforts are devoted to encoder while neglecting supervision collapse associated with the inherent limitation of generative reconstruction. ContrNP [14] shares the motivation of preserving function consistency with us, the proposed CCNP is more scalable owing to the TCL branch.

**Contrastive Conditional Density Estimation.** The common strategy held by [8, 22, 26, 27] is to bypass likelihood-based distribution inference with noise contrastive estimation (NCE) [10]. MetaCDE [26] learns a mean embeddings to estimate conditional density of multi-modal distribution. FCLR [8] replaces the decoder of CNPs with self-supervised contrastive signals to model function representations. CReSP [20] takes a further step upon FCLR and forms a semi-supervised framework to learn representations at specific indices of time-series. While both FCLR and CReSP also seek for preserving function consistency as in ContrNP and CCNP, they stand for different tasks.

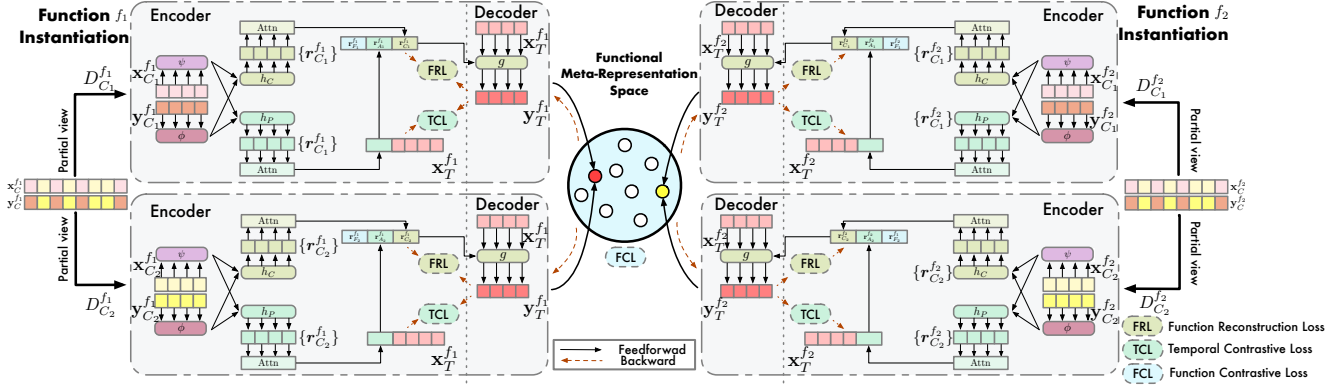


Figure 2. The proposed Contrastive Conditional Neural Processes. We showcase the scenario where there are two function instantiations (*i.e.*  $f_1$  and  $f_2$ ) with two partial views (*i.e.*  $D_{C_1}^f$  and  $D_{C_2}^f$ ) of each  $f$  for illustration only. Each instantiation  $f$  is involved with three losses, including a) Function Reconstruction Loss FRL, implemented with an encoder-decoder generative module; b) in-instantiation Temporal Contrastive Loss TCL enforcing the predictive function embedding to align with ground-truth function embedding; and c) cross-instantiation Function Contrastive Loss FCL that regularizes the transferability of meta-representation for each instantiation.

CReSP primarily accommodates downstream tasks without relying on reconstruction, whereas ContrNP and CCNP are intended to reconstruct observations ultimately.

**Summary.** Our work is orthogonal to recent signs of progress in either branch, yet combines the merits of both into one framework. To CNPs, we strip off model adaptation from the function generation process, and set up a contrastive likelihood-free objective in its place. Moreover, we learn to extract the high-level abstraction of function observations rather than wasting model’s capacity on reconstructing every low-level detail. Such the hierarchical contrastive objectives aid in scaling the model to handle high-dimensional observations and to provide robust meta-representation across multiple non-*i.i.d* instantiations. To CDEs, we investigate a hybrid generative-contrastive meta-learning model while the previous ones doing away from the reconstruction. Therefore, the proposed method may serve as a universal plug, and being able to be combined with any instance from the CNP family.

### 3. Contrastive Conditional Neural Processes

Fig. 2 overviews CCNP, wherein the model is tasked with optimizing three objectives, including a) generative function observation reconstruction loss (FRL) (Sec. 3.2), plus b) in-instantiation temporal contrastive alignment loss (TCL) (Sec. 3.3) and c) cross-instantiation function contrastive identification loss (FCL) (Sec. 3.4).

#### 3.1. Preliminaries

**Problem Statement.** We detail meta-learning for time-series prediction as modeling stochastic processes (SPs)  $\mathcal{F} : \mathcal{X} \rightarrow \mathcal{Y}$ , with  $\mathcal{X} = [0, +\infty)$  is 1-dimensional scalar time indices and  $\mathcal{Y} \subseteq \mathbb{R}^d$  represents  $d$ -dimensional obser-

vations. Let  $D_C^f = \{(x_c^f, \mathbf{y}_c^f)\}_{c \in I_C^f}$  be a set of observed sample pairs uniformly drawn from a specific underlying data-generating function  $f \sim \mathcal{F}$ , indexed by  $I_C^f$ , referred to the context set of functional instantiation  $f$ . Each  $D_C^f$  is associated with a learning task that maps context set to a contextual representation  $\mathbf{r}_C^f$  defined over  $f$ , and reconstruct the function observations on a superset of context set, denoted by the target set  $D_T^f \subset D_T^f = \{(x_t^f, \mathbf{y}_t^f)\}_{t \in I_T^f}$ .

To model the variability of SPs with meta-learning, predictions at target indices  $\mathbf{x}_T^f$  are specified as predictive posterior distribution  $p(\mathbf{y}_T^f | \mathbf{x}_T^f; D_C^f)$  that are consistent with, say, a set of  $F$  function instantiations in practice (*i.e.*  $f \in F \sim \mathcal{F}$ ), conditioning on corresponding context set. Assume, any finite sets of function evaluations of  $f$  are conditionally independent, and jointly Gaussian distributed given context set, a generative model is obtained where the output is a conditional predictive distribution

$$\begin{aligned}
 p(\mathbf{y}_T^f | \mathbf{x}_T^f; D_C^f) &= \prod_{t=1}^{|D_T^f|} p(\mathbf{y}_t^f | \Phi(D_C^f)(\mathbf{x}_t^f)) \\
 &= \prod_{t=1}^{|D_T^f|} \mathcal{N}(\mathbf{y}_t^f; \boldsymbol{\mu}_t^f, \boldsymbol{\Sigma}_t^f) \quad (1)
 \end{aligned}$$

with  $(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t) = \Phi(D_C^f)(\mathbf{x}_t) = g(h(D_C^f), \mathbf{x}_t^f) = g(\mathbf{r}_C^f, \mathbf{x}_t^f)$  and  $f(\mathbf{x}_t^f) \perp\!\!\!\perp f(\mathbf{x}_{t'}^f) | D_C^f$ ,  $t \neq t' \in I_T^f$

where  $\Phi = \{h, g\}$  is implemented as encoder-decoder that approximates the function mapping from a dataset of context set to predictive distributions over arbitrary target sets.

**Workflow in CNPs.** Model generation in CNPs starts with passing through every context sample in  $D_C^f$  into

encoder network  $h$  to acquire a finite-dimensional representation, and performing permutation-invariant aggregation over whole context set  $\mathbf{r}_C^f = \square_{c \in \mathcal{C}}(h(x_c^f, \mathbf{y}_c^f))$ . We note that  $\square(h(\cdot, \cdot))$  is implemented with a *DeepSet* [35] structure, where the input are arranged as a set to ensure operation  $\square(\cdot)$  is invariant to various input orders and size. The decoder network  $g$  then takes the conditional prior  $\mathbf{r}_C^f$  for target indices to produce a posterior distribution  $p(\mathbf{y}_T^f | \mathbf{x}_T^f; D_C^f)$ , of which the predictive performance is measured via log-likelihood estimation on the target set.

### 3.2. Reconstructing Function Distribution

**Encoder.** The function observation reconstruction objective FRL is built upon CNPs. Unlike vanilla CNP that concatenates each pair of  $\mathbf{x}_c^f$  with  $\mathbf{y}_c^f$  to the encoder  $h_C$ , we separately pre-process  $\mathbf{x}_c^f$  and  $\mathbf{y}_c^f$  using diverse feature extraction networks  $\psi(\mathbf{x}^f)$  and  $\phi(\mathbf{y}^f)$ . Considering that  $\mathbf{y}^f$  could be in a high-dimensional setting, *i.e.*  $\mathbb{R}^{d \geq 3}$ , domain-specific feature extractors can be applied to obtain efficient embeddings of observation, *e.g.* CNN for images and GNN for graph-structured data. Following that, the encoder  $h_C(\cdot, \cdot)$  is applied to each time-observation pair of the context set  $D_C$  to obtain the corresponding *local* representation  $\mathbf{r}_c^f$

$$\mathbf{r}_c^f = h_C(\psi(\mathbf{x}_c^f), \phi(\mathbf{y}_c^f)) \quad (2)$$

**Position-Aware Aggregation.** CNP [6] applies mean aggregation over the whole context set, which, however ignores the relative position between any two context indices and results in under-fitting [9, 16]. To capture such correlation, we pass encoded representations  $\mathbf{r}_{c=1:|D_C^f|}^f$  through a multi-head self-attention [29] module to calculate the position-aware importance matrix, followed by mean aggregation, to compute the contextual representation that summarizes all the time-observation pairs from  $D_C$

$$\mathbf{r}_C^f = \frac{1}{|D_C^f|} \sum_{(\mathbf{x}_c^f, \mathbf{y}_c^f) \in D_C^f} \sum_k \frac{e^{\mathbf{W}_q \mathbf{r}_c^f \cdot \mathbf{W}_k \mathbf{r}_k^f / \sqrt{d}}}{\sum_k e^{\mathbf{W}_q \mathbf{r}_c^f \cdot \mathbf{W}_k \mathbf{r}_k^f / \sqrt{d}}} \mathbf{W}_v \mathbf{r}_k^f \quad (3)$$

where  $d$  is the dimension of  $\mathbf{r}_c^f$ ,  $\mathbf{W}_q, \mathbf{W}_k, \mathbf{W}_v$  are linear operators applied to  $\mathbf{r}_c^f$  and  $\mathbf{r}_k^f$  for  $k = 1 : |D_C^f|$ .

**Decoder.** We concatenate the query index  $\psi(\mathbf{x}_T^f)$  with contextual representation  $\mathbf{r}_C^f$ , and use decoder  $g$  to estimate observation  $\hat{\mathbf{y}}_T^f$ , where  $\mathbf{r}^f$  is a non-linear transformation of  $[\mathbf{r}_C^f, \mathbf{r}_P^f, \mathbf{r}_F^f]$ . The latter two are obtained by following the same procedure as Eq. (2), Eq. (3) with different encoder  $h_P$  and  $h_F$ , and are optimized according to Sec. 3.3 and Sec. 3.4, respectively. Following the common empirical evaluation criteria of CNPs [18], reconstructed observations of the target set  $D_T$  is assumed to be a joint Gaussian dis-

tribution.

$$\hat{\mathbf{y}}_T^f \sim \log \prod_{t=1}^{|D_T^f|} \mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\sigma}_t) \quad (4)$$

with  $\boldsymbol{\mu}_t = \text{softmax}(\mu(\mathbf{r}_t^f))$ ,

$$\boldsymbol{\sigma}_t = \text{diag}(0.9 \times \text{softplus}(\mathbf{r}_t^f) + 0.1)$$

where  $\mu(\cdot), \sigma(\cdot)$  are linear transformations of target input and output predictive mean and variance values.

### 3.3. Aligning Predictive Temporal Representation

**Likelihood-free Density Estimation.** We expect the predictive representation to be similar to the encoding of ground-truth observation [17, 20], so as to realize temporal *local* alignment. For instance, let  $\mathbf{x}_t^f$  be the query target index, we attempt to maximize the density ratio of its predictive embedding conditioned on sets of context set  $D_C^{f \in F}$  sampled from  $F$  different instantiations.

**In-Instantiation Temporal Contrastive Loss.** In practice, we formulate an in-instantiation temporal contrastive learning (TCL) objective, optimized with InfoMax [28]-based loss to estimate such ratio of likelihood. The concatenated target index and contextual representation is transformed with predictive head  $\varphi$  yielding the predictive embedding  $\hat{\varphi}_t^f$  at target index  $\mathbf{x}_t^f$ . We further apply a non-linear projection head  $\rho_p : \mathbb{R}^{\phi \cdot \varphi} \mapsto \mathbb{R}^z$  to map these embeddings to a low-dimensional space for similarity measurement [1].

$$\hat{\mathbf{z}}_t^f = \rho_p(\hat{\varphi}_t^f), \mathbf{z}_t^f = \rho_p(\varphi(\mathbf{y}_t^f)) \text{ with } \hat{\varphi}_t^f = \varphi(\mathbf{x}_t^f, \mathbf{r}_C^f) \quad (5)$$

Since  $\hat{\mathbf{z}}_t$  and  $\mathbf{z}_t$  is considered as positive pair, function embeddings of the remain indices within the batch become negative samples. The TCL loss can therefore be given by

$$\mathcal{L}_{\text{TCL}} = - \sum_{f=1}^F \sum_{t=1}^T \log \frac{e^{\text{sim}(\hat{\mathbf{z}}_t^f, \mathbf{z}_t^f) / \tau}}{e^{\text{sim}(\hat{\mathbf{z}}_t^f, \mathbf{z}_t^f) / \tau} + \sum_{f'}^F \sum_{i=1}^{T+C} \mathbb{1}_{[i \neq t]} e^{\text{sim}(\hat{\mathbf{z}}_t^f, \mathbf{z}_i^{f'}) / \tau}} \quad (6)$$

where  $\text{sim}(\mathbf{a}, \mathbf{b}) = \frac{\mathbf{a}^\top \mathbf{b}}{\|\mathbf{a}\| \|\mathbf{b}\|}$  is cosine similarity.  $\tau$  denotes the temperature coefficient.  $C, T$  and  $F$  refers to the number of sampled context set, target set and observation sequences, respectively. By minimizing Eq. (6), we ensure the predicted embedding is closer to the embedding of the true outcome than embeddings of other random outcomes.

### 3.4. Regularizing Functional Meta-Representation

**Self-Supervised Function Identification.** We further explicitly generalize contextual representation across instantiations as in [8], rather than entangle meta-representation

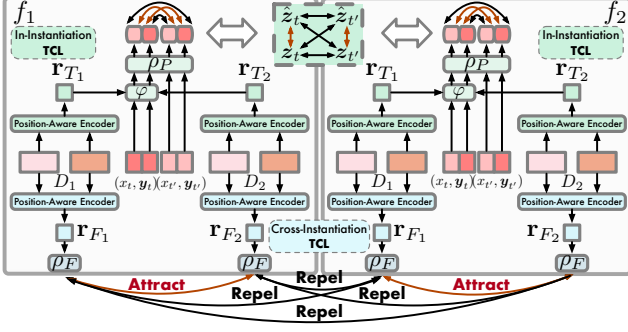


Figure 3. Conceptual illustration of TCL and FCL. We note that details of  $\phi, \psi$  are omitted for brevity. The upper half (green) depicts in-instantiation TCL with respect to Eq. (2), Eq. (3), Eq. (5) *i.e.* the objective is optimized within each instantiation  $f$ . The lower (blue) half discusses the cross-instantiation FCL regarding Eq. (2), Eq. (3), Eq. (7), Eq. (8).

acquisition with predictive meta-task together as in CNPs. Different partial observations from the same instantiation are expected to be identified to the same data-generating function. By doing so, we encourage the representations of the same partial context set to be mapped closer in the embedding space, while in contrast, different ones should be pushed away. This motivates the function contrastive learning (FCL) objective. In contrast to TCL using target indices and observations, FCL solely optimizes the training procedure with only context sets in a self-supervised spirit.

**Cross-Instantiation Function Contrastive Loss.** With accessing to  $F$  context sets drawn from corresponding functions, we halve  $D_C^f$  into two disjoint subsets by randomly selecting  $|D_C^f|/2$  samples to comprise the partial view  $D_{C_1}^f$ , then the left half makes up of another partial view  $D_{C_2}^f$ , *i.e.*

$$D_C^f = D_{C_1}^f \cup D_{C_2}^f \text{ with } D_{C_1}^f \cap D_{C_2}^f = \emptyset \quad (7)$$

We process each partial view in the similar way as TCL to obtain self-attentive representation  $\mathbf{r}_{F_i}^f, \mathbf{r}_{F_j}^f$ . Subsequently, we pass them through the projection head  $\rho_F: \mathbb{R}^r \mapsto \mathbb{R}^z$  and get the corresponding low-dimensional representations  $\mathbf{q}_i^f = \rho_F(\mathbf{r}_{F_i}^f)$  and  $\mathbf{q}_j^f = \rho_F(\mathbf{r}_{F_j}^f)$ . FCL is formulated as

$$\mathcal{L}_{\text{FCL}} = - \sum_{f=1}^F \sum_{1 \leq i < j \leq 2} \log \frac{e^{\text{sim}(\mathbf{q}_i^f, \mathbf{q}_j^f)/\tau}}{e^{\text{sim}(\mathbf{q}_i^f, \mathbf{q}_j^f)/\tau} + \sum_{f' \neq f} \sum_{1 \leq i < j \leq 2} \mathbb{1}[f \neq f'] e^{\text{sim}(\mathbf{q}_i^f, \mathbf{q}_j^{f'})/\tau}} \quad (8)$$

where  $\text{sim}(\mathbf{a}, \mathbf{b}) = \frac{\mathbf{a}^\top \mathbf{b}}{\|\mathbf{a}\| \|\mathbf{b}\|}$  refers to the cosine similarity and  $\tau$  is the temperature coefficient.

### Algorithm 1 Episodic Training for CCNP

**Require:** A dataset of  $F$  function instantiations  $\{D^f\}_{f \in F}$ .  
**Require:** context\_range =  $C$ , extra\_target\_range =  $T$ .

**Ensure:** Parameters  $\Theta = \{\phi, \psi, h_C, h_P, h_F, g, \rho_F, \rho_P, \varphi\}$   
*/\* Sampling context set and target set \*/*

- 1: **for**  $f \in F$  **do**
- 2:   Sample context size  $|D_C^f| \leftarrow C$  with  $C \sim \mathcal{U}(1, C)$
- 3:   Sample target size  $|D_T^f| \leftarrow C + T$  with  $T \sim \mathcal{U}(1, T)$
- 4:   Uniformly sample context set  $D_C^f$  with size  $|D_C^f|$
- 5:   Uniformly sample target set  $D_T^f$  with size  $|D_{C+T}^f|$
- 6: **end for**  
*/\* Optimizing FCL Objective (Section Sec. 3.4) \*/*
- 7: Obtain representations of partial views  $\mathbf{r}_{F_i}^f, \mathbf{r}_{F_j}^f \leftarrow$  Eq. (2), Eq. (3), Eq. (7), **foreach**  $f \in F$ .
- 8: Compute  $\mathcal{L}_{\text{FCL}} \leftarrow$  Eq. (8)
- 9:  $\theta \leftarrow \theta - \beta \nabla_{\theta} \mathcal{L}_{\text{FCL}}$  **foreach**  $\theta \in \{\psi, \phi, h_F, g, \rho_f\}$   
*/\* Optimizing TCL Objective (Sec. 3.3) \*/*
- 10: **for**  $f \in F$  **do**
- 11:   Obtain representation  $\mathbf{r}_P^f \leftarrow$  Eq. (2), Eq. (3)
- 12:   Obtain predictive and ground-truth embeddings of target  $\hat{\mathbf{z}}_t^f, \mathbf{z}_t^f \leftarrow$  Eq. (5) **foreach**  $t \in T$
- 13: **end for**
- 14: Compute  $\mathcal{L}_{\text{TCL}} \leftarrow$  Eq. (6)
- 15:  $\theta \leftarrow \theta - \alpha \nabla_{\theta} \mathcal{L}_{\text{TCL}}$  **foreach**  $\theta \in \{\psi, \phi, h_T, \rho_P, \varphi\}$   
*/\* Optimizing FRL Objective (Sec. 3.2) \*/*
- 16: Obtain context representation  $[\mathbf{r}_C^f, \mathbf{r}_P^f, \mathbf{r}_F^f] \leftarrow$  Eq. (2) - Eq. (3)
- 17: Compute  $\mathcal{L}_{\text{FRL}} \leftarrow$  Eq. (4) **foreach**  $f \in F$
- 18:  $\theta \leftarrow \theta - \nabla_{\theta} \mathcal{L}_{\text{FRL}}$  **foreach**  $\theta \in \{\psi, \phi, h_C, g\}$

### 3.5. Model Training

There are three targets Eq. (1), Eq. (6), Eq. (8) involved within each episode. Nonetheless, three objectives have respective goals and different parameters to learn. With the generative observation reconstruction objective  $\mathcal{L}_{\text{FRL}}$ , we maximize Eq. (1) *i.e.* the conditional expectation. In practice, we minimize the negative log-likelihood function by negating Eq. (4). At the same time, two contrastive losses  $\mathcal{L}_{\text{TCL}}$  and  $\mathcal{L}_{\text{FCL}}$  are optimized prior to estimating the conditional expectation. When adding them to our training objective with  $\alpha \geq 0$  for  $\mathcal{L}_{\text{TCL}}$  and  $\beta \geq 0$  for  $\mathcal{L}_{\text{FCL}}$  as tradeoff parameters, we summarize the overall objective

$$\arg \min_{\Theta} \mathbb{E}_{f \sim F} [\mathcal{L}_{\text{FRL}}^f_{\phi, \psi, g, h_C} - \alpha \mathcal{L}_{\text{TCL}}_{h_P, \varphi, \rho_P} - \beta \mathcal{L}_{\text{FCL}}_{h_F, \rho_F}] \quad (9)$$

where  $\Theta = \{\phi, \psi, g, h_C, h_P, h_F, \rho_F, \rho_P, \varphi\}$

The procedure of episodic training is shown in Algorithm 1. Note that forloops in the description is for illustration, we apply batch processing for parallelism in practice.

## 4. Empirical Studies

We empirically validate the proposed CCNP on a broad range of time-series data encompassing 1D, 2D function regression and multi-object trajectory prediction. Our main interests fall onto the following three questions:

- i. Do introduce auxiliary contrastive losses help to increase the predictive performance of CNP?
- ii. If so, does explicitly model FCL beneficial for model adaptation?
- iii. Is CCNP capable of handling against high-dimensional observations than CNP?

We attempt to address these research questions by reporting the quantitative experimental results and discussing associated explanations. In Sec. 4.2 we compare the predictive regression performance on three scopes of datasets to answer question i). Then, we discuss if the meta-transferability can be improved by FCL in Sec. 4.3. Moreover, in Sec. 4.4 we look into if the introduced TCL makes CNP scalable to higher-dimensional temporal sequences.

### 4.1. Experiments Setup

**Datasets.** We briefly describe the datasets covered in this work. See Appendix A.1 for detailed dataset description.

For *1D Functions*, we practice few-shot regression over four different function families, including sinusoids, exponentials, damped oscillators and straight lines. Each function family is depicted with the linear combination of a amplitude  $\alpha \sim \mathcal{U}(\alpha_{\text{floor}}, \alpha_{\text{ceil}})$  and a phase  $\beta \sim \mathcal{U}(\beta_{\text{floor}}, \beta_{\text{floor}})$  in a closed-form. For each function family, we construct a meta-dataset by randomly sampling different  $\alpha$  and  $\beta$  for each instantiation, with each dataset is composed of 490 training, 10 validation and 10 test sequences. In evaluation, the models are tested 5-shot, 10-shot and 20-shot setting (*i.e.* the size of context set  $|D_C|$ ), respectively. All models are trained for 25 episodes.

For *2D Population dynamics*, we study a predator-prey system, where the dynamics can be modeled by the Lotka-Volterra (LV) Equations [31]. The populations of predator  $y_1$  and prey  $y_2$  are mutually influenced after each time increment, with the interactions between two species being dominated by the greeks coefficients  $\alpha, \beta, \delta, \gamma$ . Expressly,  $\alpha$  and  $\gamma$  represent predator and prey’s birth rates, while the  $\beta$  and  $\delta$  are associated with the death rates. We conducted the experiments with two sets of system configurations. For one setting where initial populations  $y_1 \in [50, 100)$  and  $y_2 \in [100, 150)$  are provided, we fix a set of greeks coefficients  $(\alpha, \beta, \delta, \gamma) = (\frac{2}{3}, \frac{4}{3}, 1, 1)$ . In another setting, the initial populations are fixed to 160 and 80, while the greeks are randomly drawn from the predefined range to form each function instantiation. For both configurations, we generate 180 training, 10 validation and 10 test sequences with

100 times evolvment in each instantiation. We run 200 episodes for model training, where the sizes of context set and extra target set within each episode are randomly drawn from  $[1, 80]$  and  $[0, 20]$  therein.

For *Higher-dimensional sequence prediction*, we study a synthetic bouncing ball system and rotating MNIST data<sup>1</sup>, derived from generative temporal modeling scenarios. In the bouncing ball system, each trajectory (samples sequence) contains the movements of three interacting balls within a rectangular box. Each timestep is framed as a  $32 \times 32$  image. The models are tasked with inferring the locations of the balls as well as interaction rules between them, and reconstructing the pixels in unobserved timesteps, without prior assumption on details of the scene (e.g., ball count and velocities). We randomly grab 10000 and 500 trajectories for training and testing, where each trajectory consists of 20 simulated time steps of motion. The number of context samples and extra target samples are fixed to 5 and 10. In RotMNIST data, we predict 16 rotation angles of handwritten 3 digits in the shape of  $28 \times 28$  grayscale pixels sequences. We follow the same settings as in [33] where the 400 sequences with a 9:1 ratio of train-test split to run the experiments. Similarly, in each sequence, the model learns from 5 context samples randomly drawn from 16 frames is tasked with predicting the remaining frames.

**Baselines.** Resembling most of related studies, we consider vanilla CNP (CNP) [6], Attentive CNP (AttnCNP) [16], and Convolutional CNP (ConvCNP) [9] as baselines.

**Evaluation Metrics.** To answer question i), we evaluate the models’ performances in terms of predictive mean log-likelihood and reconstruction error measured by mean squared error (MSE) as usual for CNPs. In Sec. 4.3 we additionally validate the transferability of meta-representation by predicting the function coefficients of a specific data-generating function. MSE is adopted as the measurement.

### 4.2. Exact Reconstruction Results

**1D Few-Shot Function Regression.** Tab. 1 summarizes our finding in running few-shot function regression over four groups of data-generating functions. In this task, the predictive model takes a context set  $D_C^f$  of an instantiation  $f$ . The contextual information is extracted to a global representation  $\mathbf{r}^f$ , which is used then to produce the Gaussian mean of  $\mathbf{y}_t^f$  and the predictive variance in the target index  $x_t$ , with  $x_t \in D_T^f$ . Sec. 4.2 illustrates that CCNP can make reasonable predictions with higher confidence (lower variance) even only 5 context points are provided. CCNP remarkably outperforms all the competitive baselines in relating to both predictive log-likelihood on context sets and reconstruction error on target sets.

<sup>1</sup><https://github.com/cagatayildiz/ODE2VAE>

Model / Data	Sinusoids			Exponential			Linear			Oscillators		
	Predictive Log-likelihood ( $\times 10^{-2}$ ) $\uparrow\uparrow$											
	5-shot	10-shot	20-shot	5-shot	10-shot	20-shot	5-shot	10-shot	20-shot	5-shot	10-shot	20-shot
CNP	0.872 $\pm$ 0.292	1.024 $\pm$ 0.117	1.084 $\pm$ 0.061	1.298 $\pm$ 0.028	1.313 $\pm$ 0.036	1.335 $\pm$ 0.019	0.393 $\pm$ 0.140	0.428 $\pm$ 0.172	0.531 $\pm$ 0.070	1.188 $\pm$ 0.083	1.233 $\pm$ 0.031	1.243 $\pm$ 0.051
AttnCNP	<u>0.902 <math>\pm</math> 0.228</u>	1.004 $\pm$ 0.143	1.079 $\pm$ 0.120	1.343 $\pm$ 0.024	1.353 $\pm$ 0.010	1.361 $\pm$ 0.011	<u>0.545 <math>\pm</math> 0.141</u>	<u>0.620 <math>\pm</math> 0.250</u>	<u>0.706 <math>\pm</math> 0.198</u>	1.242 $\pm$ 0.064	1.285 $\pm$ 0.039	1.312 $\pm$ 0.026
ConvCNP	0.370 $\pm$ 0.326	0.694 $\pm$ 0.281	<b>1.331 <math>\pm</math> 0.226</b>	<b>1.807 <math>\pm</math> 0.310</b>	<b>2.069 <math>\pm</math> 0.234</b>	<b>2.488 <math>\pm</math> 0.294</b>	-0.646 $\pm$ 0.331	-0.348 $\pm$ 0.534	0.160 $\pm$ 0.342	<b>1.345 <math>\pm</math> 0.101</b>	<b>1.625 <math>\pm</math> 0.193</b>	<b>2.113 <math>\pm</math> 0.313</b>
CCNP (Ours)	<b>1.073 <math>\pm</math> 0.086</b>	<b>1.144 <math>\pm</math> 0.137</b>	<u>1.240 <math>\pm</math> 0.067</u>	1.356 $\pm$ 0.013	1.364 $\pm$ 0.020	<u>1.372 <math>\pm</math> 0.007</u>	<b>1.021 <math>\pm</math> 0.120</b>	<b>1.130 <math>\pm</math> 0.075</b>	<b>1.178 <math>\pm</math> 0.123</b>	1.344 $\pm$ 0.014	1.364 $\pm$ 0.009	<u>1.372 <math>\pm</math> 0.002</u>
	Reconstruction Error ( $\times 10^2$ ) $\downarrow\downarrow$											
CNP	1.135 $\pm$ 0.803	<u>0.714 <math>\pm</math> 0.255</u>	0.571 $\pm$ 0.089	0.164 $\pm$ 0.058	0.129 $\pm$ 0.062	0.090 $\pm$ 0.041	2.657 $\pm$ 0.760	2.726 $\pm$ 0.934	2.316 $\pm$ 0.447	0.393 $\pm$ 0.182	0.291 $\pm$ 0.077	0.279 $\pm$ 0.108
AttnCNP	<u>1.105 <math>\pm</math> 0.832</u>	0.784 $\pm$ 0.400	0.549 $\pm$ 0.182	<u>0.075 <math>\pm</math> 0.056</u>	<u>0.055 <math>\pm</math> 0.025</u>	<u>0.039 <math>\pm</math> 0.021</u>	<u>1.688 <math>\pm</math> 0.57</u>	<u>1.915 <math>\pm</math> 0.836</u>	<u>1.357 <math>\pm</math> 0.697</u>	<u>0.25 <math>\pm</math> 0.12</u>	<u>0.169 <math>\pm</math> 0.083</u>	0.127 $\pm$ 0.047
ConvCNP	3.788 $\pm$ 2.222	1.490 $\pm$ 1.213	<u>0.353 <math>\pm</math> 0.183</u>	1.289 $\pm$ 0.731	0.541 $\pm$ 0.347	0.132 $\pm$ 0.072	48.367 $\pm$ 59.953	21.610 $\pm$ 24.210	6.777 $\pm$ 7.635	0.897 $\pm$ 0.426	0.362 $\pm$ 0.165	<u>0.091 <math>\pm</math> 0.045</u>
CCNP (Ours)	<b>0.479 <math>\pm</math> 0.231</b>	<b>0.420 <math>\pm</math> 0.312</b>	<b>0.230 <math>\pm</math> 0.112</b>	<b>0.041 <math>\pm</math> 0.029</b>	<b>0.033 <math>\pm</math> 0.041</b>	<b>0.014 <math>\pm</math> 0.009</b>	<b>0.519 <math>\pm</math> 0.260</b>	<b>0.333 <math>\pm</math> 0.097</b>	<b>0.277 <math>\pm</math> 0.190</b>	<b>0.068 <math>\pm</math> 0.028</b>	<b>0.036 <math>\pm</math> 0.017</b>	<b>0.019 <math>\pm</math> 0.005</b>

Table 1. Comparison between CCNP and the baseline CNP models. The upper half shows the predictive log-likelihood ( $\uparrow\uparrow$  the higher the better) of each model on the target set for each function family. The lower half shows reconstruction error ( $\downarrow\downarrow$  the lower the better) on the target indices. Results are reported with  $\pm 1$  standard deviation over 6 different random seeds. Results in **Bold** indicate the best performances, and underlined refer to the second-best.

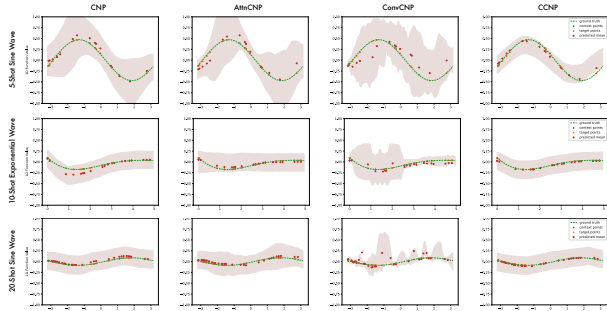
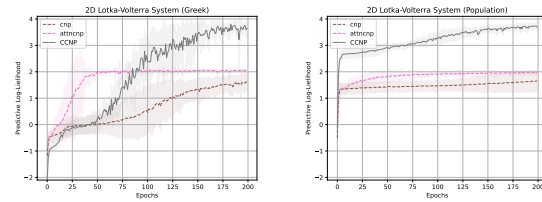


Figure 4. The qualitative comparison of CCNP and CNPs baselines on three different function families with 5-shot, 10-shot and 20-shot regression. The shown results in each column are model predictions on the shared target samples, where the plottings from left to right are of CNP, AttnCNP, ConvCNP and CCNP therein, respectively. Dashed line is the ground-truth function value, with red diamonds representing the predictive mean values  $\mu$  and the shaded regions indicating 95% confidence interval ( $\mu \pm 2\sigma$ ).

**2D Population Dynamics Prediction.** We also compare the function reconstruction capabilities on the Lotka-Volterra system of CNP, AttnCNP and CCNP. On a par with 1D tasks, we report the predictive log-likelihood of target indices on the validation set. Under both configurations, three models can converge within 100 epochs, whilst CCNP can fit LV equations with fewer epochs when the initial populations are provided. For qualitative comparisons, the maximum likelihood of CNP reaches  $1.626 \pm 0.24$  and  $1.652 \pm 0.13$  under two modes. AttnCNP performs faster adaptation in the former case, which, however is not compared favorably with the ones achieved by CCNP after 100 epochs, indicated by  $2.059 \pm 0.11(\text{AttnCNP}) < 3.794 \pm 0.58(\text{CCNP})$ .

**Higher-dimensional Sequence Prediction.** We provide both qualitative and quantitative comparisons among CNP, AttnCNP and CCNP for predicting the evolvement of higher-dimensional image sequences (see Fig. 6 and Tab. 2). With the same encoding structure and training condition,



(a) Log-likelihood with random initial populations. (b) Log-likelihood with fixed initial populations.

Figure 5. Predictive performances of 2D predator-prey dynamics. Left reveals the simulation with randomly initialized population and fixed greeks coefficients, and vice versa in the right. Grey solid line represents the predictive mean of CCNP over 6 seeds. Pink and brown dashed lines correspond to CNP and AttnCNP, respectively. We denote 95% confidence interval with shaded areas.

CCNP returns lower MSE in reconstruction error across multiple prediction steps on both datasets. This can also be verified in quantitative results, where the second row consists of predictive reconstruction of CNP, while CCNP is reported within the third row. While CNP produces blur results at some query indices, CCNP can reconstruct more accurate values more confidently. For both CNP and CCNP, there is no significant performance drop when prediction sizes are extended by 3-4 times longer, reflecting that CNPs have a great potential in handling low data problems.

We further provide experiments in terms of additional 1D datasets and running efficiency (See Appendix B.).

### 4.3. Transferability of Meta-representation

Besides the evaluation of reconstructing observations, we are also interested in how well the model can recognize the particular generating function, which is associated with the transferability of meta-representation [8]. We study the following experiment *i.e.* function coefficient inference in closed-formed 1D sinusoid functions to discuss question ii).

The goal of this task is to predict the amplitude  $\alpha$  and phase  $\beta$  of the sinusoid function  $f$ , given a randomly sampled context set  $D_C^f$ . For each model, we fix the parameters

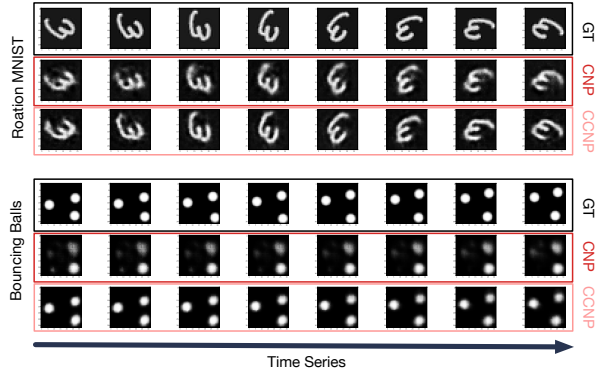
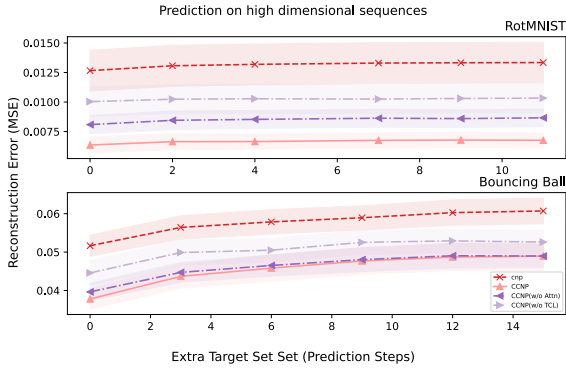


Figure 6. The qualitative and quantitative comparisons of CNP and CCNP on two image sequences datasets.

Model / Data	RotMNIST		BouncingBall	
	10	16	10	20
CNP	0.0132 ± 0.002	0.0133 ± 0.002	0.0579 ± 0.004	0.0607 ± 0.004
CCNP (-Attn)	<b>0.0085 ± 0.003</b>	<b>0.0087 ± 0.001</b>	0.0465 ± 0.003	0.0490 ± 0.003
CCNP (-TCL)	0.0103 ± 0.001	0.0103 ± 0.001	0.0505 ± 0.003	0.0526 ± 0.003
CCNP (Full)	<b>0.0066 ± 0.001</b>	<b>0.0068 ± 0.001</b>	<b>0.0458 ± 0.004</b>	<b>0.0489 ± 0.004</b>

Table 2. Prediction performance on high-dimensional images sequences with **ablation studies**. We report MSE ( $\downarrow$ ) obtained on RotMNIST and BouncingBall data, whereby we also conduct ablation studies for removing a) position-aware aggregation (indicated by CCNP (-Attn)); b) TCL branch (CCNP (-TCL)). Results in **Bold** indicate the best performances and underlined refer to the second-best, with  $\pm 1$  standard deviation over 6 runs for each.

of contextual representation  $\mathbf{r}^f$  pre-trained on function regression tasks, of which we append a classifier on top to perform coefficient inference. Fig. 7a describes performances of CNP, AttnCNP and CCNP on this task, where CCNP produces the lowest MSE in evaluating the prediction of  $\alpha$  and  $\beta$ , with  $-0.5$  (to CNP) and  $-0.4$  (to AttnCNP) of gap over 6 seeds. We conclude that CCNP demonstrates higher transferability of meta-representation, benefited by FCL regularizing its capability of aligning function instantiations. For quantitative illustration, we showcase that AttnCNP struggles at adapting to shifted  $\alpha$  within the same function family while CCNP maintains decent capability (Fig. 7b).

#### 4.4. Scalability to Higher-Dimensional Data

We investigate question iii) *i.e.* whether TCL is beneficial to the scalability of CNPs. The results in Tab. 2 suggest the prediction performances given by different design choices of CCNP. Removing either encoding self-attention or TCL leads to poor performance in predicting high-dimensional data. Stripping TCL from CCNP brings in decreased results in predicting both datasets, which can be attributed to additional efficiency provided by contrastive loss against high-dimensional data. The position-aware encoding also boosts the performance of TCL as the combination of the two is better than if either were applied alone.

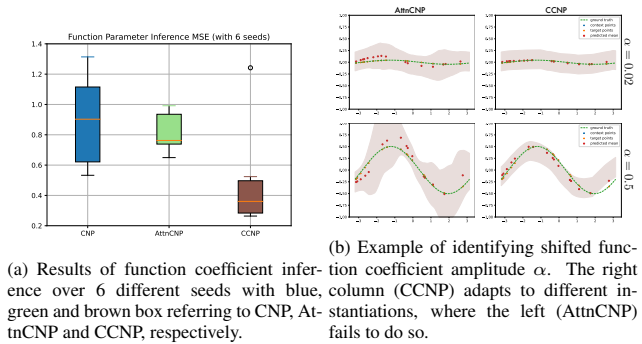


Figure 7. Comparisons of CNP, AttnCNP and CCNP on the transferability of meta-representations.

## 5. Conclusion

We present Contrastive Conditional Neural Processes, a hybrid generative-contrastive model that explores the complementary advantages of generative and contrastive approaches in meta-learning, towards scaling CNPs to handle high-dimensional and noisy time-series. Generative CNPs are augmented by two contrastive, hierarchically organized branches. While the in-instantiation temporal contrastive branch revolves around extracting high-level abstraction of observations, the cross-instantiation function contrastive branch cope with supervision-collapse of common generative meta-learning models. Improved reconstruction performances on diverse time-series datasets demonstrate that, adding contrastive regularization to CNPs not only directly facilitates discriminative down-stream tasks, but also contributes significantly to generative reconstruction tasks.

**Limitation.** In this work, we focus on improving the efficacy of CNPs with two auxiliary contrastive regularizers. The costs associated with constructing tuples and comparing pairwise similarity are, however, causing a loss in computation efficiency. Such tradeoffs may be addressed by exploring more efficient contrastive sampling methods.



## References

- [1] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. [2](#), [4](#)
- [2] Yutian Chen, Matthew W Hoffman, Sergio Gómez Colmenarejo, Misha Denil, Timothy P Lillicrap, Matt Botvinick, and Nando Freitas. Learning to learn without gradient descent by gradient descent. In *International Conference on Machine Learning*, pages 748–756. PMLR, 2017. [1](#)
- [3] Carl Doersch, Ankush Gupta, and Andrew Zisserman. Crosstransformers: spatially-aware few-shot transfer. *Advances in Neural Information Processing Systems*, 33:21981–21993, 2020. [1](#)
- [4] Andrew Foong, Wessel Bruinsma, Jonathan Gordon, Yann Dubois, James Requeima, and Richard Turner. Meta-learning stationary stochastic process prediction with convolutional neural processes. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 8284–8295. Curran Associates, Inc., 2020. [1](#)
- [5] Tianyu Gao, Xingcheng Yao, and Danqi Chen. Simcse: Simple contrastive learning of sentence embeddings. *arXiv preprint arXiv:2104.08821*, 2021. [2](#)
- [6] Marta Garnelo, Dan Rosenbaum, Christopher Maddison, Tiago Ramalho, David Saxton, Murray Shanahan, Yee Whye Teh, Danilo Rezende, and SM Ali Eslami. Conditional neural processes. In *International Conference on Machine Learning*, pages 1704–1713. PMLR, 2018. [1](#), [2](#), [4](#), [6](#)
- [7] Marta Garnelo, Jonathan Schwarz, Dan Rosenbaum, Fabio Viola, Danilo J. Rezende, S. M. Ali Eslami, and Yee Whye Teh. Neural processes, 2018. [1](#)
- [8] Muhammad Waleed Gondal, Shruti Joshi, Nasim Rahman, Stefan Bauer, Manuel Wuthrich, and Bernhard Schölkopf. Function contrastive learning of transferable meta-representations. In *International Conference on Machine Learning*, pages 3755–3765. PMLR, 2021. [1](#), [2](#), [4](#), [7](#)
- [9] Jonathan Gordon, Wessel P. Bruinsma, Andrew Y. K. Foong, James Requeima, Yann Dubois, and Richard E. Turner. Convolutional conditional neural processes. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. [1](#), [2](#), [4](#), [6](#)
- [10] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 297–304. JMLR Workshop and Conference Proceedings, 2010. [2](#)
- [11] Peter Holdrieth, Michael J Hutchinson, and Yee Whye Teh. Equivariant learning of stochastic fields: Gaussian processes and steerable conditional neural processes. In *International Conference on Machine Learning*, pages 4297–4307. PMLR, 2021. [2](#)
- [12] Timothy Hospedales, Antreas Antoniou, Paul Micaelli, and Amos Storkey. Meta-learning in neural networks: A survey. *arXiv preprint arXiv:2004.05439*, 2020. [1](#)
- [13] Boris Ivanovic and Marco Pavone. The trajectron: Probabilistic multi-agent trajectory modeling with dynamic spatiotemporal graphs. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2375–2384, 2019. [1](#)
- [14] Konstantinos Kallidromitis, Denis Gudovskiy, Kozuka Kazuki, Ohama Iku, and Luca Rigazio. Contrastive neural processes for self-supervised learning. In *Asian Conference on Machine Learning*, pages 594–609. PMLR, 2021. [2](#)
- [15] Makoto Kawano, Wataru Kumagai, Akiyoshi Sannai, Yusuke Iwasawa, and Yutaka Matsuo. Group equivariant conditional neural processes. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. [1](#), [2](#)
- [16] Hyunjik Kim, Andriy Mnih, Jonathan Schwarz, Marta Garnelo, Ali Eslami, Dan Rosenbaum, Oriol Vinyals, and Yee Whye Teh. Attentive neural processes. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019. [1](#), [2](#), [4](#), [6](#)
- [17] Thomas N. Kipf, Elise van der Pol, and Max Welling. Contrastive learning of structured world models. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. [1](#), [4](#)
- [18] Tuan Anh Le, Hyunjik Kim, Marta Garnelo, Dan Rosenbaum, Jonathan Schwarz, and Yee Whye Teh. Empirical evaluation of neural process objectives. In *NeurIPS workshop on Bayesian Deep Learning*, 2018. [4](#)
- [19] Xiao Liu, Fanjin Zhang, Zhenyu Hou, Li Mian, Zhaoyu Wang, Jing Zhang, and Jie Tang. Self-supervised learning: Generative or contrastive. *IEEE Transactions on Knowledge and Data Engineering*, 2021. [2](#)
- [20] Emile Mathieu, Adam Foster, and Yee Teh. On contrastive representations of stochastic processes. *Advances in Neural Information Processing Systems*, 34, 2021. [2](#), [4](#)
- [21] Ross Messing, Chris Pal, and Henry Kautz. Activity recognition using the velocity histories of tracked keypoints. In *2009 IEEE 12th international conference on computer vision*, pages 104–111. IEEE, 2009. [1](#)
- [22] Lorenzo Pacchiardi and Ritabrata Dutta. Score matched neural exponential families for likelihood-free inference. *Journal of Machine Learning Research*, 23(38):1–71, 2022. [2](#)
- [23] Evan Racah and Sarath Chandar. Slot contrastive networks: A contrastive approach for representing objects. *CoRR*, abs/2007.09294, 2020. [1](#)
- [24] Tim Salzmann, Boris Ivanovic, Punarjay Chakravarty, and Marco Pavone. Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVIII 16*, pages 683–700. Springer, 2020. [1](#)
- [25] Jihoon Tack, Sangwoo Mo, Jongheon Jeong, and Jinwoo Shin. Csi: Novelty detection via contrastive learning on distributionally shifted instances. *arXiv preprint arXiv:2007.08176*, 2020. [2](#)

- [26] Jean-Francois Ton, Lucian CHAN, Yee Whye Teh, and Dino Sejdinovic. Noise contrastive meta-learning for conditional density estimation using kernel mean embeddings. In Arindam Banerjee and Kenji Fukumizu, editors, *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pages 1099–1107. PMLR, 13–15 Apr 2021. 2
- [27] Yao-Hung Hubert Tsai, Tianqin Li, Martin Q. Ma, Han Zhao, Kun Zhang, Louis-Philippe Morency, and Ruslan Salakhutdinov. Conditional contrastive learning with kernel. In *International Conference on Learning Representations*, 2022. 2
- [28] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding, 2019. 4
- [29] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008, 2017. 2, 4
- [30] Ricardo Vilalta and Youssef Drissi. A perspective view and survey of meta-learning. *Artificial intelligence review*, 18(2):77–95, 2002. 1
- [31] Darren J Wilkinson. *Stochastic modelling for systems biology*. Chapman and Hall/CRC, 2018. 6
- [32] Jim Winkens, Rudy Bunel, Abhijit Guha Roy, Robert Stanforth, Vivek Natarajan, Joseph R Ledsam, Patricia MacWilliams, Pushmeet Kohli, Alan Karthikesalingam, Simon Kohl, et al. Contrastive training for improved out-of-distribution detection. *arXiv preprint arXiv:2007.05566*, 2020. 2
- [33] Cagatay Yildiz, Markus Heinonen, and Harri Lahdesmaki. Ode2vae: Deep generative second order odes with bayesian neural networks. *Advances in Neural Information Processing Systems*, 32:13412–13421, 2019. 6
- [34] Yuning You, Tianlong Chen, Yongduo Sui, Ting Chen, Zhangyang Wang, and Yang Shen. Graph contrastive learning with augmentations. *Advances in Neural Information Processing Systems*, 33:5812–5823, 2020. 2
- [35] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabás Póczos, Ruslan Salakhutdinov, and Alexander J. Smola. Deep sets. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 3391–3401, 2017. 4