

# NeuralProcesses.jl: Composing Neural Processes with Flux

Wessel P. Bruinsma<sup>1,2,\*</sup>, Jonathan Gordon<sup>1,\*</sup>,  
Richard E. Turner<sup>1</sup>

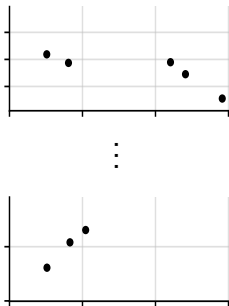
<sup>1</sup>University of Cambridge, <sup>2</sup>Invenia Labs

\*Equal contribution

JuliaCon 2020

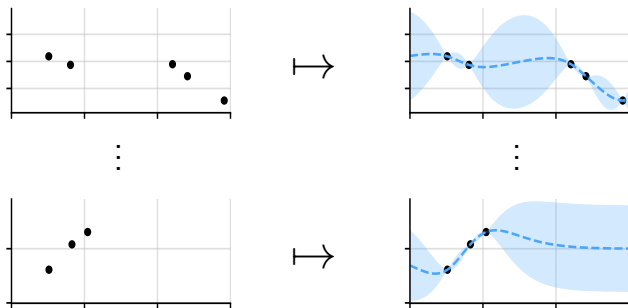
# Setting: Learning to Predict

1/7



# Setting: Learning to Predict

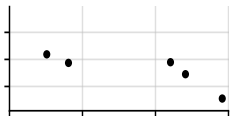
1/7



# Setting: Learning to Predict

1/7

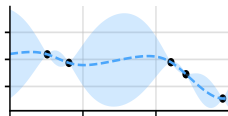
$\phi$ : data sets  $\mathcal{D}$   $\rightarrow$  predictions  $\mathcal{P}$



$\vdots$

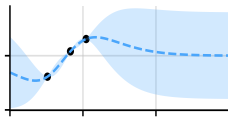


$\phi$



$\vdots$

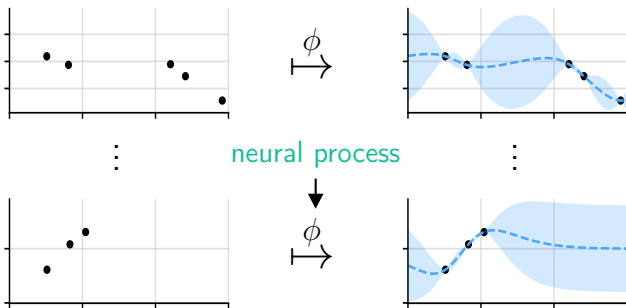
$\phi$



# Setting: Learning to Predict

1/7

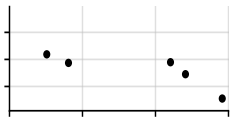
$\phi$ : data sets  $\mathcal{D}$   $\rightarrow$  predictions  $\mathcal{P}$



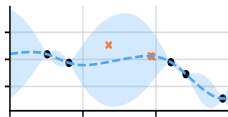
# Setting: Learning to Predict

1/7

$\phi$ : data sets  $\mathcal{D}$   $\rightarrow$  predictions  $\mathcal{P}$



$\phi$



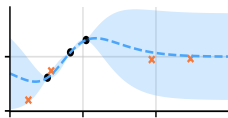
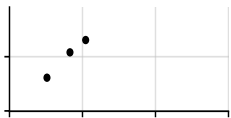
$\vdots$

neural process

$\vdots$

$\phi$

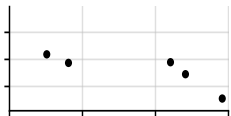
training



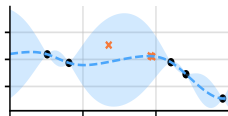
# Setting: Learning to Predict

1/7

$\phi$ : data sets  $\mathcal{D}$   $\rightarrow$  predictions  $\mathcal{P}$



$\phi$



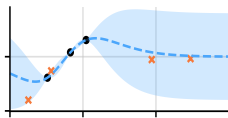
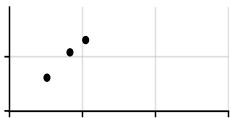
$\vdots$

neural process

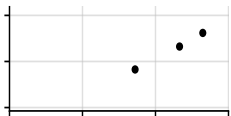
$\vdots$

$\phi$

training



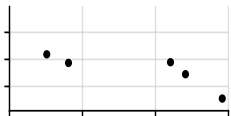
test



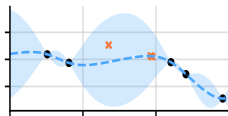
# Setting: Learning to Predict

1/7

$\phi$ : data sets  $\mathcal{D}$   $\rightarrow$  predictions  $\mathcal{P}$



$\phi$

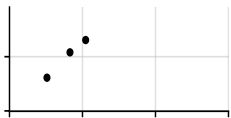


$\vdots$

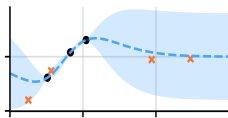
neural process

$\vdots$

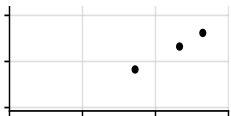
training



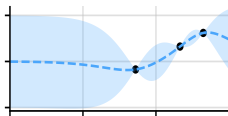
$\phi$



test



$\phi$





## Neural Processes Are Many...

2/7

- Conditional Neural Process (Garnelo, Rosenbaum, et al., 2018)

## Neural Processes Are Many...

2/7

- Conditional Neural Process (Garnelo, Rosenbaum, et al., 2018)
- Neural Process (Garnelo, Schwarz, et al., 2018)

## Neural Processes Are Many...

2/7

- Conditional Neural Process (Garnelo, Rosenbaum, et al., 2018)
- Neural Process (Garnelo, Schwarz, et al., 2018)
- Attentive Neural Process (Kim et al., 2019)

## Neural Processes Are Many...

2/7

- Conditional Neural Process (Garnelo, Rosenbaum, et al., 2018)
- Neural Process (Garnelo, Schwarz, et al., 2018)
- Attentive Neural Process (Kim et al., 2019)
- Functional Neural Process (Louizos et al., 2019)
- Sequential Neural Process (Singh et al., 2019)
- Convolutional Conditional Neural Process (Gordon et al., 2020)
- Convolutional Neural Process (Foong et al., 2020)

# Neural Processes Are Many...

2/7

- Conditional Neural Process (Garnelo, Rosenbaum, et al., 2018)
- Neural Process (Garnelo, Schwarz, et al., 2018)
- Attentive Neural Process (Kim et al., 2019)
- Functional Neural Process (Louizos et al., 2019)
- Sequential Neural Process (Singh et al., 2019)
- Convolutional Conditional Neural Process (Gordon et al., 2020)
- Convolutional Neural Process (Foong et al., 2020)

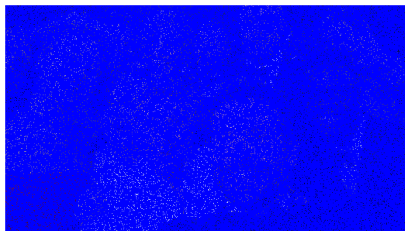


Figure from Gordon et al. (2020).

# Neural Processes Are Many...

2/7

- Conditional Neural Process (Garnelo, Rosenbaum, et al., 2018)
- Neural Process (Garnelo, Schwarz, et al., 2018)
- Attentive Neural Process (Kim et al., 2019)
- Functional Neural Process (Louizos et al., 2019)
- Sequential Neural Process (Singh et al., 2019)
- Convolutional Conditional Neural Process (Gordon et al., 2020)
- Convolutional Neural Process (Foong et al., 2020)

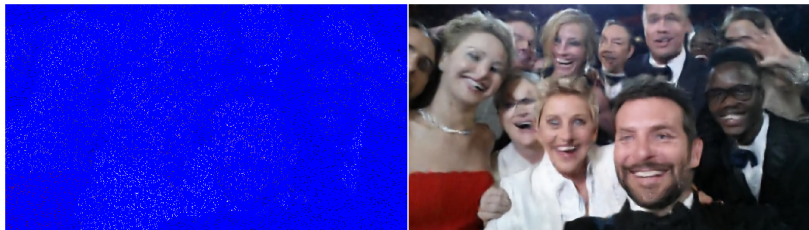


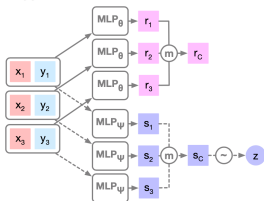
Figure from Gordon et al. (2020).

# Neural Processes Are Many... (2)

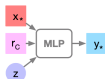
3/7

## NEURAL PROCESS

### ENCODER

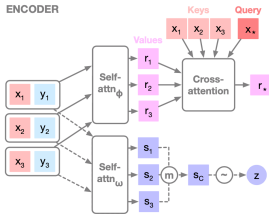


### DECODER



## ATTENTIVE NEURAL PROCESS

### ENCODER



### DECODER



Figure from Kim et al. (2019).

# Neural Processes Are Many... (2)

3/7

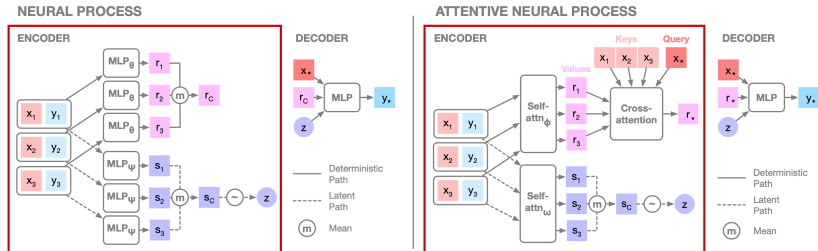


Figure from Kim et al. (2019).



# Neural Processes Are Many... (2)

3/7

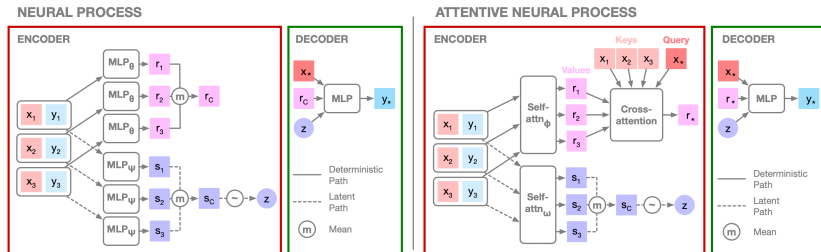


Figure from Kim et al. (2019).

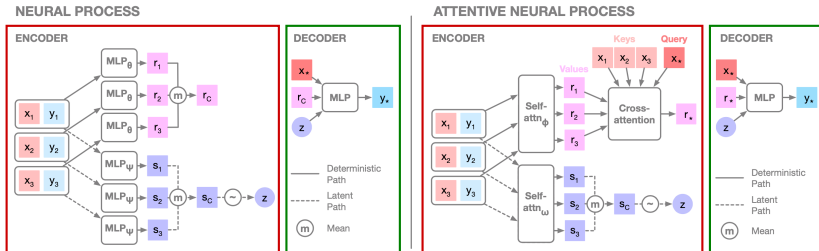


Figure from Kim et al. (2019).

- Implementations similar, but differ in details.

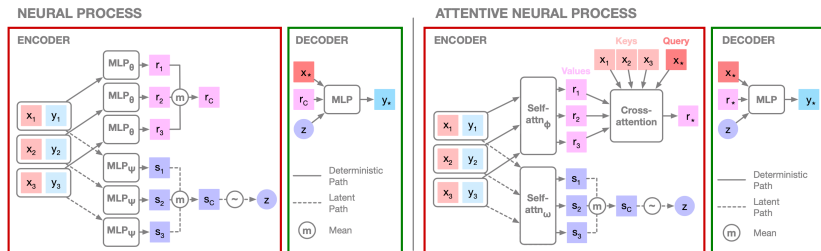
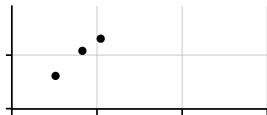


Figure from Kim et al. (2019).

- Implementations similar, but differ in details.
- Calls for a unifying framework: `NeuralProcesses.jl`.

# Encoder-Decoder Architectures

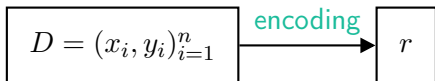
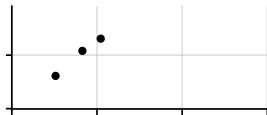
4/7



$$D = (x_i, y_i)_{i=1}^n$$

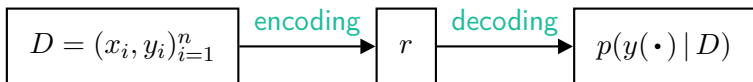
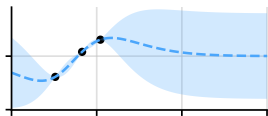
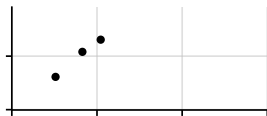
# Encoder-Decoder Architectures

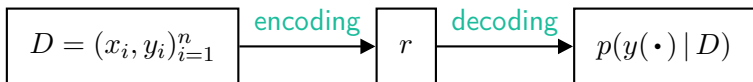
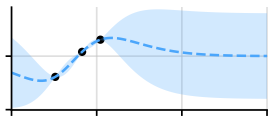
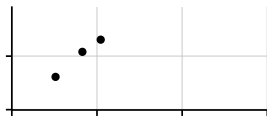
4/7



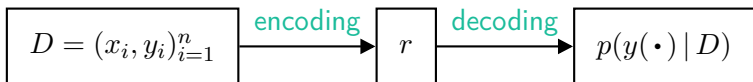
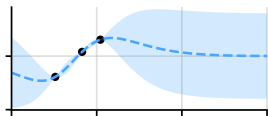
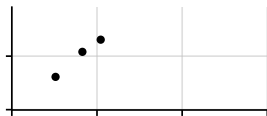
# Encoder-Decoder Architectures

4/7



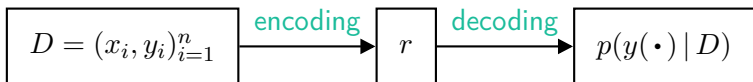
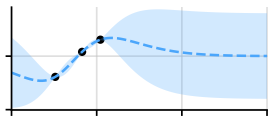
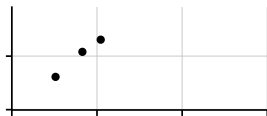


- Need a common representation.

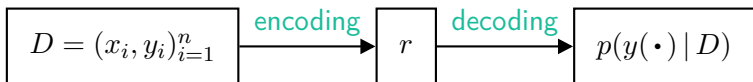
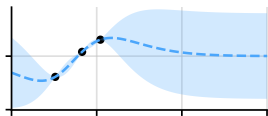
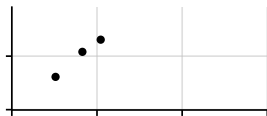


- Need a common representation.
- Data is a function:  $D: \{x_1, \dots, x_n\} \rightarrow \{y_1, \dots, y_n\}$ .

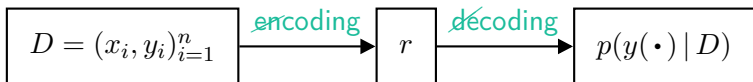
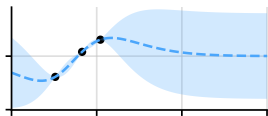
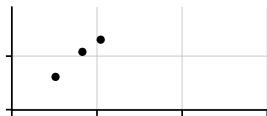




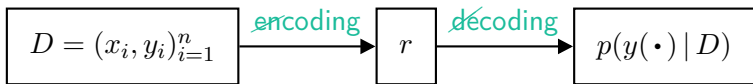
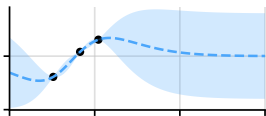
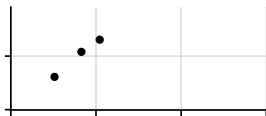
- Need a common representation.
- Data is a function:  $D: \{x_1, \dots, x_n\} \rightarrow \{y_1, \dots, y_n\}$ .
- Prediction is a function:  $p(y(\cdot) | D): \mathbb{R} \rightarrow \mathbb{R} \times \mathbb{R}_{>0}$  ( $\mu$  and  $\sigma$ ).



- Need a common representation: **functions!**
- Data is a function:  $D: \{x_1, \dots, x_n\} \rightarrow \{y_1, \dots, y_n\}$ .
- Prediction is a function:  $p(y(\cdot) | D): \mathbb{R} \rightarrow \mathbb{R} \times \mathbb{R}_{>0}$  ( $\mu$  and  $\sigma$ ).



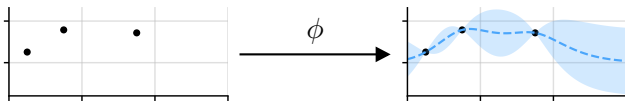
- Need a common representation: **functions!**
- Data is a function:  $D: \{x_1, \dots, x_n\} \rightarrow \{y_1, \dots, y_n\}$ .
- Prediction is a function:  $p(y(\cdot) | D): \mathbb{R} \rightarrow \mathbb{R} \times \mathbb{R}_{>0}$  ( $\mu$  and  $\sigma$ ).
- \***coding** is then a transformation of functions



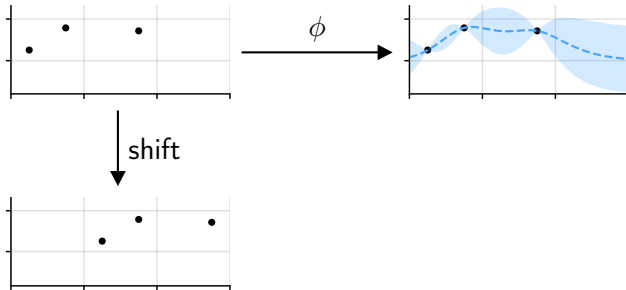
- Need a common representation: functions!
  - Data is a function:  $D: \{x_1, \dots, x_n\} \rightarrow \{y_1, \dots, y_n\}$ .
  - Prediction is a function:  $p(y(\cdot) | D): \mathbb{R} \rightarrow \mathbb{R} \times \mathbb{R}_{>0}$  ( $\mu$  and  $\sigma$ ).
  - \*coding is then a transformation of functions
- $\Rightarrow$  fundamental abstraction of `NeuralProcesses.jl`.

- A formulation that incorporates **translation equivariance**:

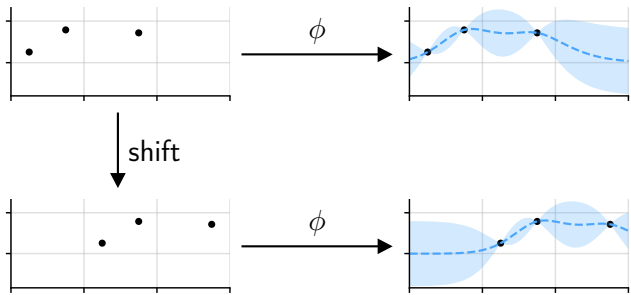
- A formulation that incorporates **translation equivariance**:



- A formulation that incorporates **translation equivariance**:

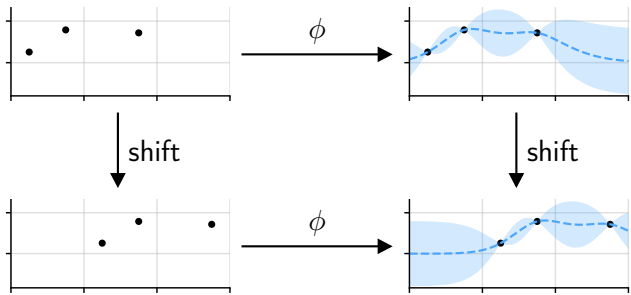


- A formulation that incorporates **translation equivariance**:

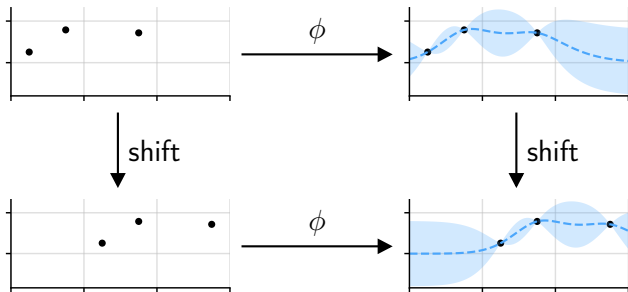




- A formulation that incorporates **translation equivariance**:



- A formulation that incorporates **translation equivariance**:

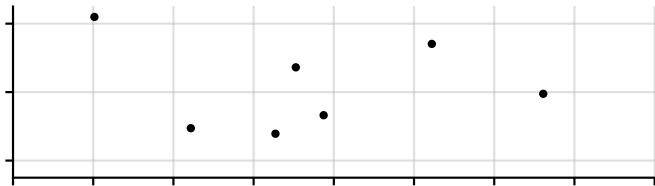


- Leverages the parameter efficiency of CNNs: Flux.jl!

# The ConvCNP with NeuralProcesses.jl

6/7



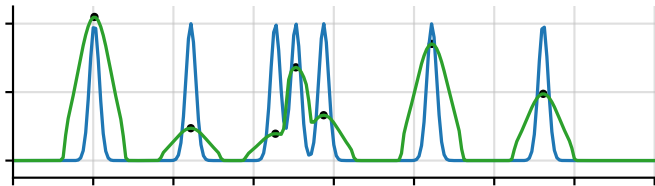


```
encoder = FunctionalCoder(  
    UniformDiscretisation1D(...),
```

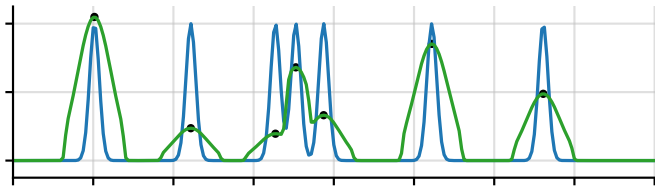
```
)
```



```
encoder = FunctionalCoder(  
    UniformDiscretisation1D(...),  
    Chain(  
        # See Gordon et al. (2020).  
        set_conv(...),  
        Deterministic()  
    )  
)
```

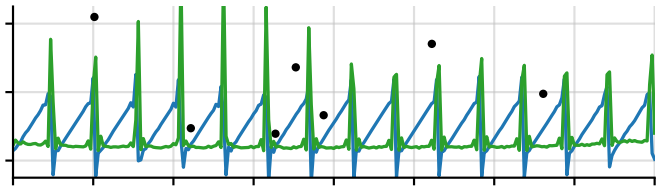


```
encoder = FunctionalCoder(  
    UniformDiscretisation1D(...),  
    Chain(  
        # See Gordon et al. (2020).  
        set_conv(...),  
        Deterministic()  
    )  
)
```



```
encoder = FunctionalCoder(  
    UniformDiscretisation1D(...),  
    Chain(  
        # See Gordon et al. (2020).  
        set_conv(...),  
        Deterministic()  
    )  
)
```

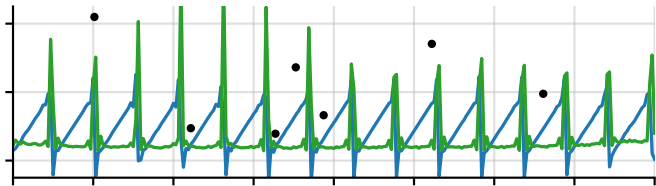
```
decoder = Chain(  
    build_conv(...),  
  
)
```



```
encoder = FunctionalCoder(  
  UniformDiscretisation1D(...),  
  Chain(  
    # See Gordon et al. (2020).  
    set_conv(...),  
    Deterministic()  
  )  
)
```

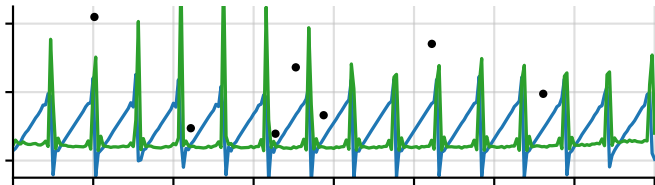
```
decoder = Chain(  
  build_conv(...),  
  
)
```





```
encoder = FunctionalCoder(  
  UniformDiscretisation1D(...),  
  Chain(  
    # See Gordon et al. (2020).  
    set_conv(...),  
    Deterministic()  
  )  
)
```

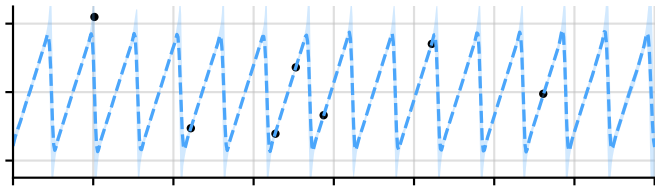
```
decoder = Chain(  
  build_conv(...),  
  set_conv(...),  
  HeterogeneousGaussian()  
)
```



```
encoder = FunctionalCoder(  
    UniformDiscretisation1D(...),  
    Chain(  
        # See Gordon et al. (2020).  
        set_conv(...),  
        Deterministic()  
    )  
)
```

```
decoder = Chain(  
    build_conv(...),  
    set_conv(...),  
    HeterogeneousGaussian()  
)
```

```
convcnp = Model(encoder, decoder)
```

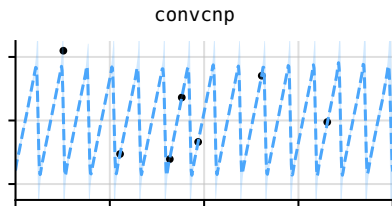


```
encoder = FunctionalCoder(  
    UniformDiscretisation1D(...),  
    Chain(  
        # See Gordon et al. (2020).  
        set_conv(...),  
        Deterministic()  
    )  
)
```

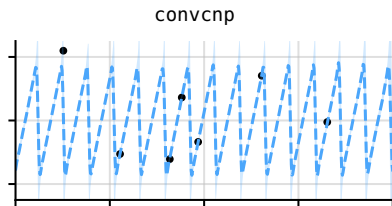
```
decoder = Chain(  
    build_conv(...),  
    set_conv(...),  
    HeterogeneousGaussian()  
)
```

```
convcnp = Model(encoder, decoder)
```

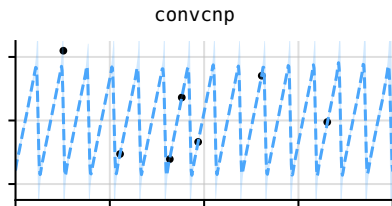
```
# Gordon et al. (2020)
convcnp = Model(
  FunctionalCoder(
    UniformDiscretisation1D(...),
    Chain(
      set_conv(...),
      Deterministic()
    )
  ),
  Chain(
    build_conv(...),
    set_conv(...),
    HeterogeneousGaussian()
  )
)
```



```
# Gordon et al. (2020)
convcnp = Model(
  FunctionalCoder(
    UniformDiscretisation1D(...),
    Chain(
      set_conv(...),
      Deterministic()
    )
  ),
  Chain(
    build_conv(...),
    set_conv(...),
    HeterogeneousGaussian()
  )
)
```



```
# Foong et al. (2020)
convnp = Model(
  FunctionalCoder(
    UniformDiscretisation1D(...),
    Chain(
      set_conv(...),
      build_conv(...),
      HeterogeneousGaussian()
    )
  ),
  Chain(
    build_conv(...),
    set_conv(...),
    HeterogeneousGaussian()
  )
)
```

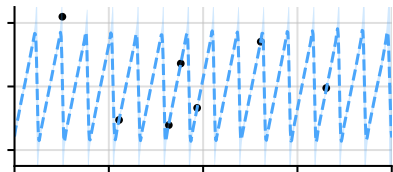


```

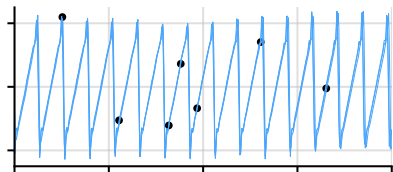
# Foong et al. (2020)
convnp = Model(
  FunctionalCoder(
    UniformDiscretisation1D(...),
    Chain(
      set_conv(...),
      build_conv(...),
      HeterogeneousGaussian()
    )
  ),
  Chain(
    build_conv(...),
    set_conv(...),
    HeterogeneousGaussian()
  )
)

```

convcnp



convnp



```
# Foong et al. (2020)
convnp = Model(
  FunctionalCoder(
    UniformDiscretisation1D(...),
    Chain(
      set_conv(...),
      build_conv(...),
      HeterogeneousGaussian()
    )
  ),
  Chain(
    build_conv(...),
    set_conv(...),
    HeterogeneousGaussian()
  )
)
```

- Composable building blocks



```
# Garnelo et al. (2018)
```

```
np = Model(  
  Parallel(  
    Chain(  
      InputsEncoder(),  
      Deterministic()  
    ),  
    Chain(  
      MLPDecoder(...),  
      Deterministic()  
    ),  
    Chain(  
      MLPDecoder(...),  
      HeterogeneousGaussian()  
    )  
  ),  
  Chain(  
    Materialise(),  
    batched_mlp(...),  
    HeterogeneousGaussian()  
  )  
)
```

- Composable building blocks

```
# Garnelo et al. (2018)
```

```
np = Model(  
  Parallel(  
    Chain(  
      InputsEncoder(),  
      Deterministic()  
    ),  
    Chain(  
      MLPDecoder(...),  
      Deterministic()  
    ),  
    Chain(  
      MLPDecoder(...),  
      HeterogeneousGaussian()  
    )  
  ),  
  Chain(  
    Materialise(),  
    batched_mlp(...),  
    HeterogeneousGaussian()  
  )  
)
```

- Composable building blocks
- Use Chain and Parallel to glue things together

```
# Garnelo et al. (2018)
```

```
np = Model(  
  Parallel(  
    Chain(  
      InputsEncoder(),  
      Deterministic()  
    ),  
    Chain(  
      MLPDecoder(...),  
      Deterministic()  
    ),  
    Chain(  
      MLPDecoder(...),  
      HeterogeneousGaussian()  
    )  
  ),  
  Chain(  
    Materialise(),  
    batched_mlp(...),  
    HeterogeneousGaussian()  
  )  
)
```

- Composable building blocks
- Use Chain and Parallel to glue things together
- Inference and learning are automatic

# NeuralProcesses.jl

7/7

```
# Garnelo et al. (2018)
```

```
np = Model(  
  Parallel(  
    Chain(  
      InputsEncoder(),  
      Deterministic()  
    ),  
    Chain(  
      MLPEncoder(...),  
      Deterministic()  
    ),  
    Chain(  
      MLPEncoder(...),  
      HeterogeneousGaussian()  
    )  
  ),  
  Chain(  
    Materialise(),  
    batched_mlp(...),  
    HeterogeneousGaussian()  
  )  
)
```

- Composable building blocks
- Use Chain and Parallel to glue things together
- Inference and learning are automatic

```
# Kim et al. (2019)
anp = Model(
  Parallel(
    Chain(
      InputsEncoder(),
      Deterministic()
    ),
    Chain(
      attention(...),
      Deterministic()
    ),
    Chain(
      MLPDecoder(...),
      HeterogeneousGaussian()
    )
  ),
  Chain(
    Materialise(),
    batched_mlp(...),
    HeterogeneousGaussian()
  )
)
```

- Composable building blocks
- Use Chain and Parallel to glue things together
- Inference and learning are automatic

```
# Kim et al. (2019)
anp = Model(
  Parallel(
    Chain(
      InputsEncoder(),
      Deterministic()
    ),
    Chain(
      attention(...),
      Deterministic()
    ),
    Chain(
      MLPDecoder(...),
      HeterogeneousGaussian()
    )
  ),
  Chain(
    Materialise(),
    batched_mlp(...),
    HeterogeneousGaussian()
  )
)
```

- Composable building blocks
- Use Chain and Parallel to glue things together
- Inference and learning are automatic

Thank you for listening :)

# Appendix

## References

- Foong, A. Y. K., Bruinsma, W. P., Gordon, J., Dubois, Y., Requeima, J., & Turner, R. E. (2020). Meta-learning stationary stochastic process prediction with convolutional neural processes. *arXiv preprint arXiv:2007.01332*. eprint: <https://arxiv.org/abs/2007.01332>
- Garnelo, M., Rosenbaum, D., Maddison, C. J., Ramalho, T., Saxton, D., Shanahan, M., ... Eslami, S. M. A. (2018). Conditional neural processes. *arXiv preprint arXiv:1807.01613*. eprint: <https://arxiv.org/abs/1807.01613>
- Garnelo, M., Schwarz, J., Rosenbaum, D., Viola, F., Rezende, D. J., Eslami, S. M. A., & Teh, Y. W. (2018). Neural processes. *arXiv preprint arXiv:1807.01622*. eprint: <https://arxiv.org/abs/1807.01622>



## References (2)

- Gordon, J., Bruinsma, W. P., Foong, A. Y. K., Requeima, J., Dubois, Y., & Turner, R. E. (2020). Convolutional conditional neural processes. *International Conference on Learning Representations (ICLR), 8th*. Retrieved from <https://openreview.net/forum?id=Skey4eBYPS>
- Kim, H., Mnih, A., Schwarz, J., Garnelo, M., Eslami, A., Rosenbaum, D., ... Teh, Y. W. (2019). Attentive neural processes. *arXiv preprint arXiv:1901.05761*. eprint: <https://arxiv.org/abs/1901.05761>
- Louizos, C., Shi, X., Schutte, K., & Welling, M. (2019). The functional neural process. *arXiv preprint arXiv:1906.08324*. eprint: <https://arxiv.org/abs/1906.08324>
- Singh, G., Yoon, J., Son, Y., & Ahn, S. (2019). Sequential neural processes. *arXiv preprint arXiv:1906.10264*. eprint: <https://arxiv.org/abs/1906.10264>