

# Quant GANs: Deep Generation of Financial Time Series

Hyelin Choi

Department of Mathematics  
Sungkyunkwan University

Dec 11, 2024

# Table of Contents

- I. Introduction
- II. Generative Adversarial Networks (GANs)
- III. Temporal Convolutional Networks (TCNs)
- IV. Log-return Neural Process
- V. Numerical Results

# Table of Contents

I. Introduction


II. Generative Adversarial Networks (GANs)

III. Temporal Convolutional Networks (TCNs)

IV. Log-return Neural Process

V. Numerical Results

# What is S&P 500?



TESLA



# S&P 500 Stock Price Path



S&P 500 index data, Google Finance

# Return

Let  $S_t$  be the stock price at time  $t$ . There are two types of returns.


- Relative return

$$R_t = (S_t - S_{t-1})/S_{t-1}$$

- Log-return

$$R_t = \log \frac{S_t}{S_{t-1}}$$

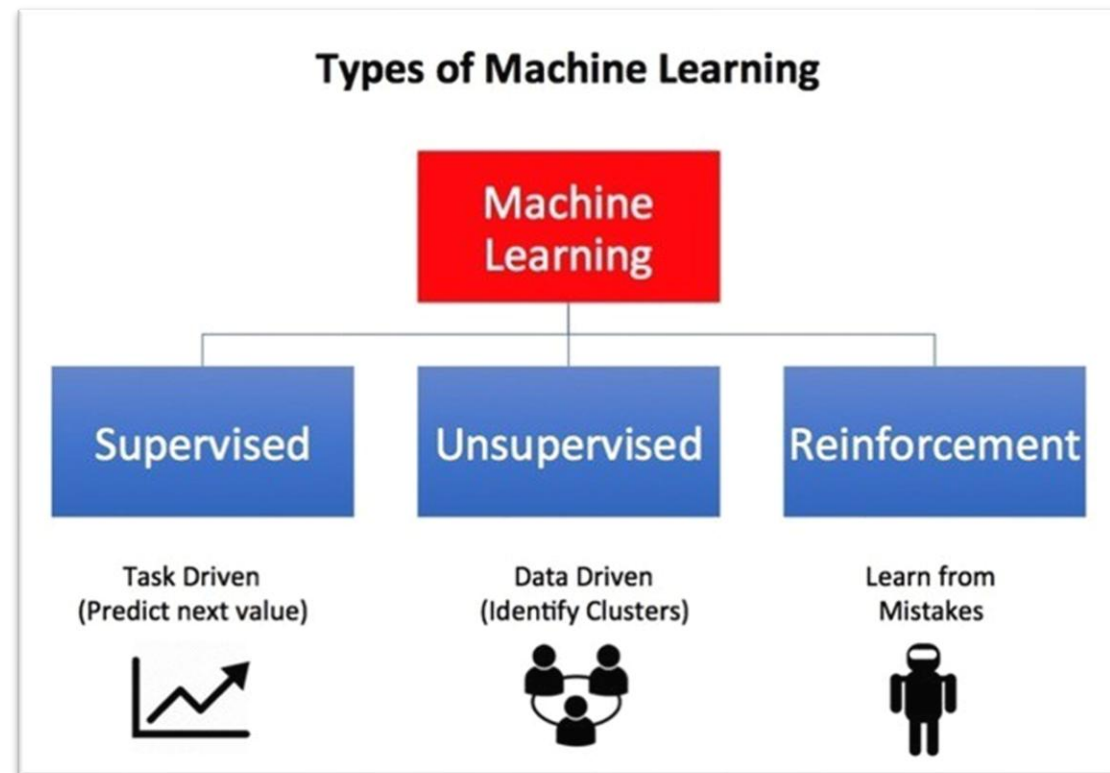
# S&P 500 Returns



S&P 500 Historical Annual Returns

# Applications of QuantGANs

1. QuantGANs enable us to generate realistic log-return paths for the S&P 500 across multiple scenarios.





# Applications of QuantGANs

2. We can predict **future prices** or **trends** of assets based on past behavior.




# Long-range dependency

What key information should QuantGANs capture from log-return data?

## Long-range dependency

Long-range dependency means correlations persist across distant time lags.

# Volatility, Drift, and Innovation



S&P 500 index data, Google Finance

# Geometric Brownian Motion(GBM)

Geometric Brownian Motion (GBM) is a mathematical model used to describe the random behavior of asset prices over time.

$$dS_t = \mu S_t dt + \sigma S_t dW_t \quad \text{where } W_t: \text{Brownian motion, } \mu: \text{drift and } \sigma: \text{volatility}$$



$$R_t = \log \frac{S_t}{S_{t-1}} = \mu^- + \sigma(W_t - W_{t-1}) \quad \text{where } \mu^- = \mu - \frac{\sigma^2}{2}$$

Drift  
 $\mu_{t,\theta}$

Volatility  
 $\sigma_{t,\theta}$

Innovation  
 $\varepsilon_{t,\theta} \sim N(0,1)$

# Table of Contents

I. Introduction


**II. Generative Adversarial Networks (GANs)**

III. Temporal Convolutional Networks (TCNs)

IV. Log-return Neural Process

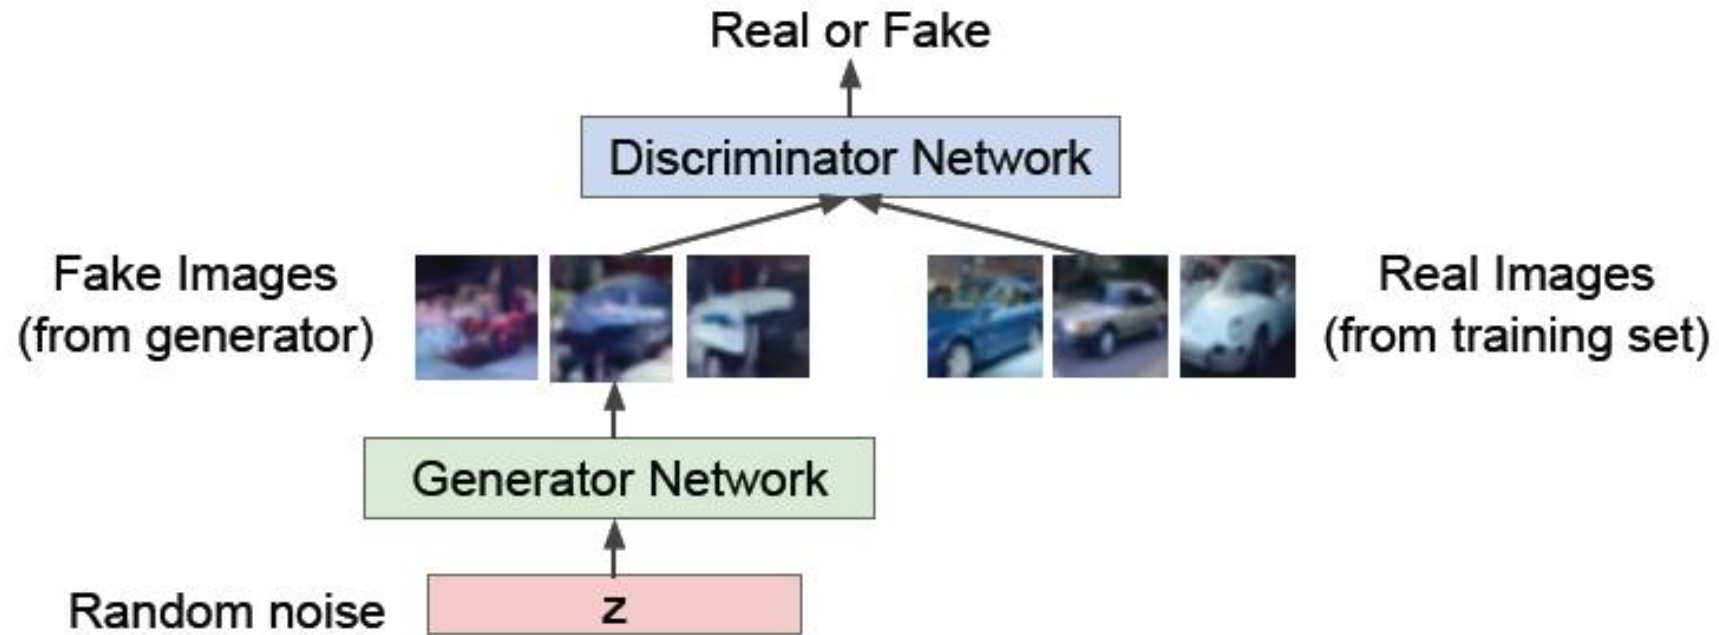
V. Numerical Results

# Construction of QuantGANs

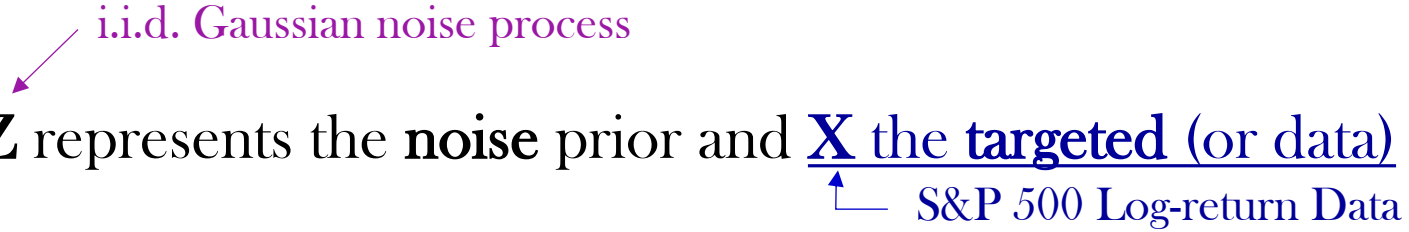


# Generative Adversarial Networks

대립적인



# Generative Adversarial Networks

- The random variable  $\mathbf{Z}$  represents the **noise** prior and  $\mathbf{X}$  the targeted (or data) random variable.  


i.i.d. Gaussian noise process

S&P 500 Log-return Data
- The goal of GANs is to train a network  $g: \mathbb{R}^{N_z} \times \Theta^{(g)} \rightarrow \mathbb{R}^{N_x}$  such that the induced random variable  $g_\theta(Z) := g_\theta \circ Z$  for some parameter  $\theta \in \Theta^{(g)}$  and the targeted random variable  $X$  have the same distribution, i.e.  $g_\theta(\mathbf{Z}) \stackrel{d}{=} \mathbf{X}$ .



# Generative Adversarial Networks

## Loss function of GANs

$$\begin{aligned}\mathcal{L}(\theta, \eta) &:= \mathbb{E} [\log(d_\eta(X))] + \mathbb{E} [\log(1 - d_\eta(g_\theta(Z)))] \\ &= \mathbb{E} [\log(d_\eta(X))] + \mathbb{E} [\log(1 - d_\eta(\tilde{X}_\theta))] .\end{aligned}$$

$\eta$ : parameter of discriminator  
 $\theta$ : parameter of generator  
 $d_\eta(\cdot)$ : function of discriminator  
 $X$ : targeted r.v.  
 $\tilde{X}_\theta$ : generated r.v.

## Step 1

Let the 1 represent real data and 0 represent fake data.

The discriminator's parameter  $\eta \in \Theta^{(d)}$  are chosen to maximize the function  $\mathcal{L}(\theta, \cdot), \theta \in \Theta^{(g)}$ .

## Step 2

The generator's parameters  $\theta \in \Theta^{(g)}$  are trained to minimize the probability of generated samples being identified as such and not from the data distribution.

# Generative Adversarial Networks

We get the min-max problem

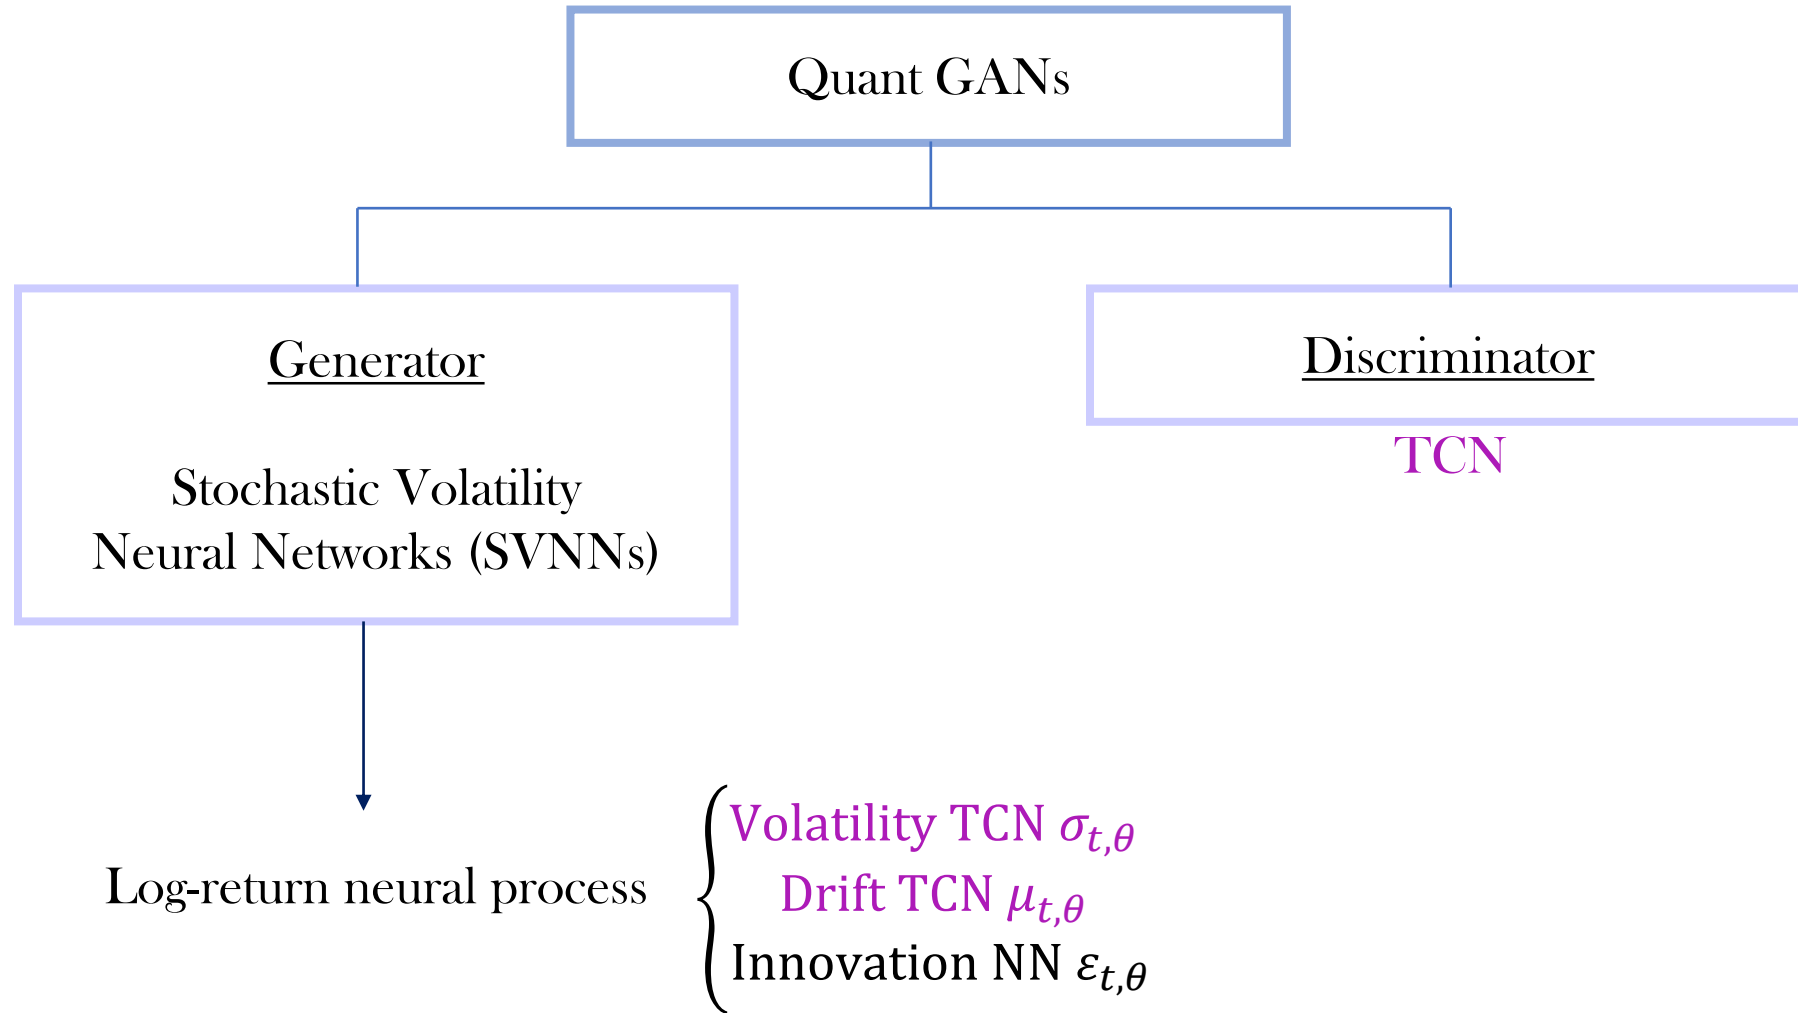
$$\min_{\theta \in \Theta^{(g)}} \max_{\eta \in \Theta^{(d)}} \mathcal{L}(\theta, \eta)$$

which refer to as the GAN objective.

# Table of Contents

- I. Introduction
- II. Generative Adversarial Networks (GANs)
- III. Temporal Convolutional Networks (TCNs)**
- IV. Log-return Neural Process
- V. Numerical Results

# Construction of QuantGANs



# Advantages of TCNs

- 1) TCNs are able to capture long-range dependencies in sequences.
- 2) TCNs with skip connections have an advantage of avoiding exponentially vanishing gradients.

# Construction of TCNs


TCNs are neural network models primarily designed to efficiently handle sequential data, such as time series.

- Constructions

Dilated causal convolutions = Causal convolutions + Dilated convolutions  
인과 확장

# Causal Convolution


Causal convolutions are convolutions, where output only depends on past sequence elements.  
인과



# Dilated Convolution

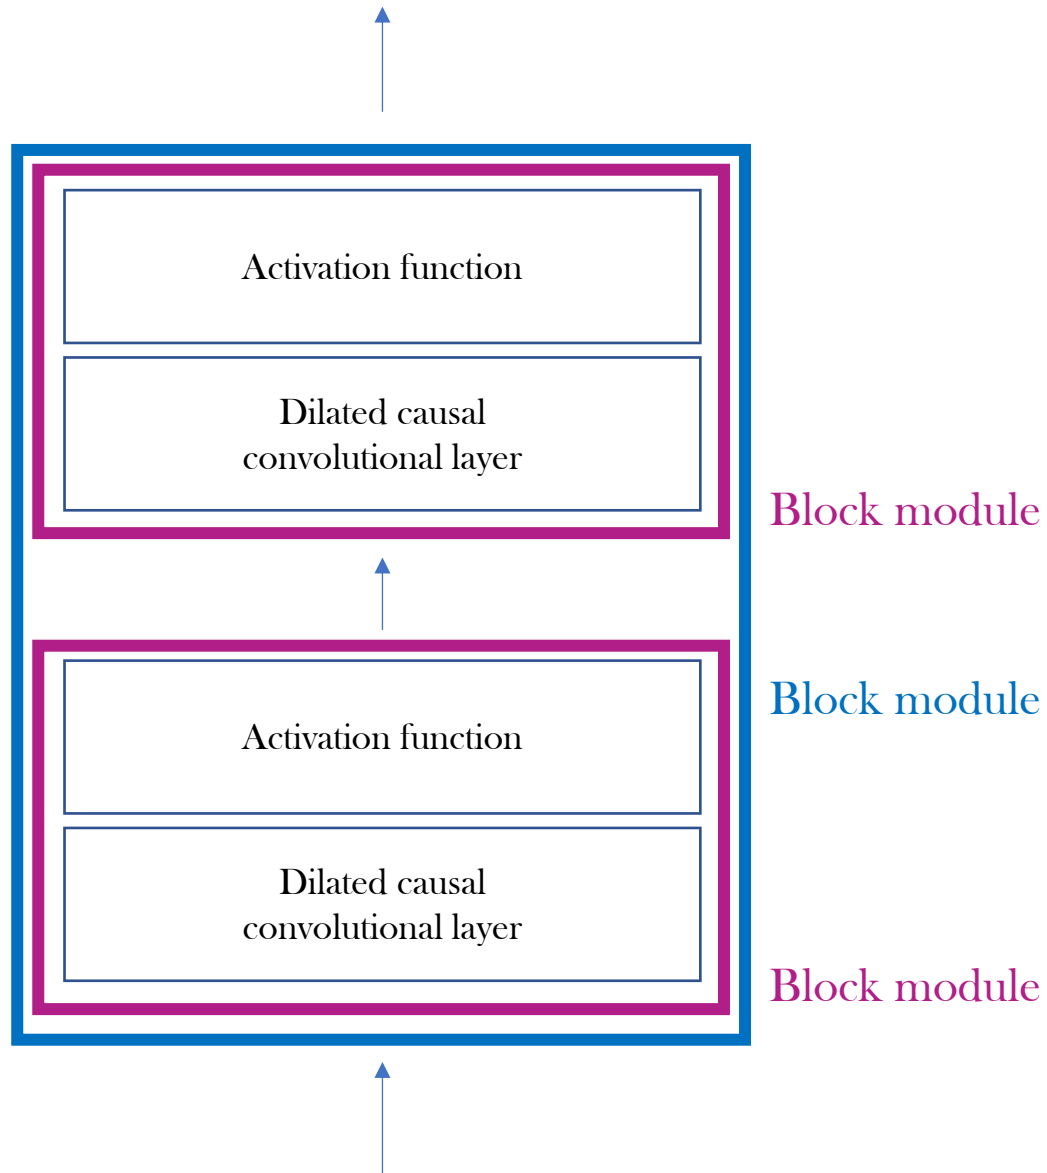
Dilated convolutions are convolutions ‘with holes’.

확장

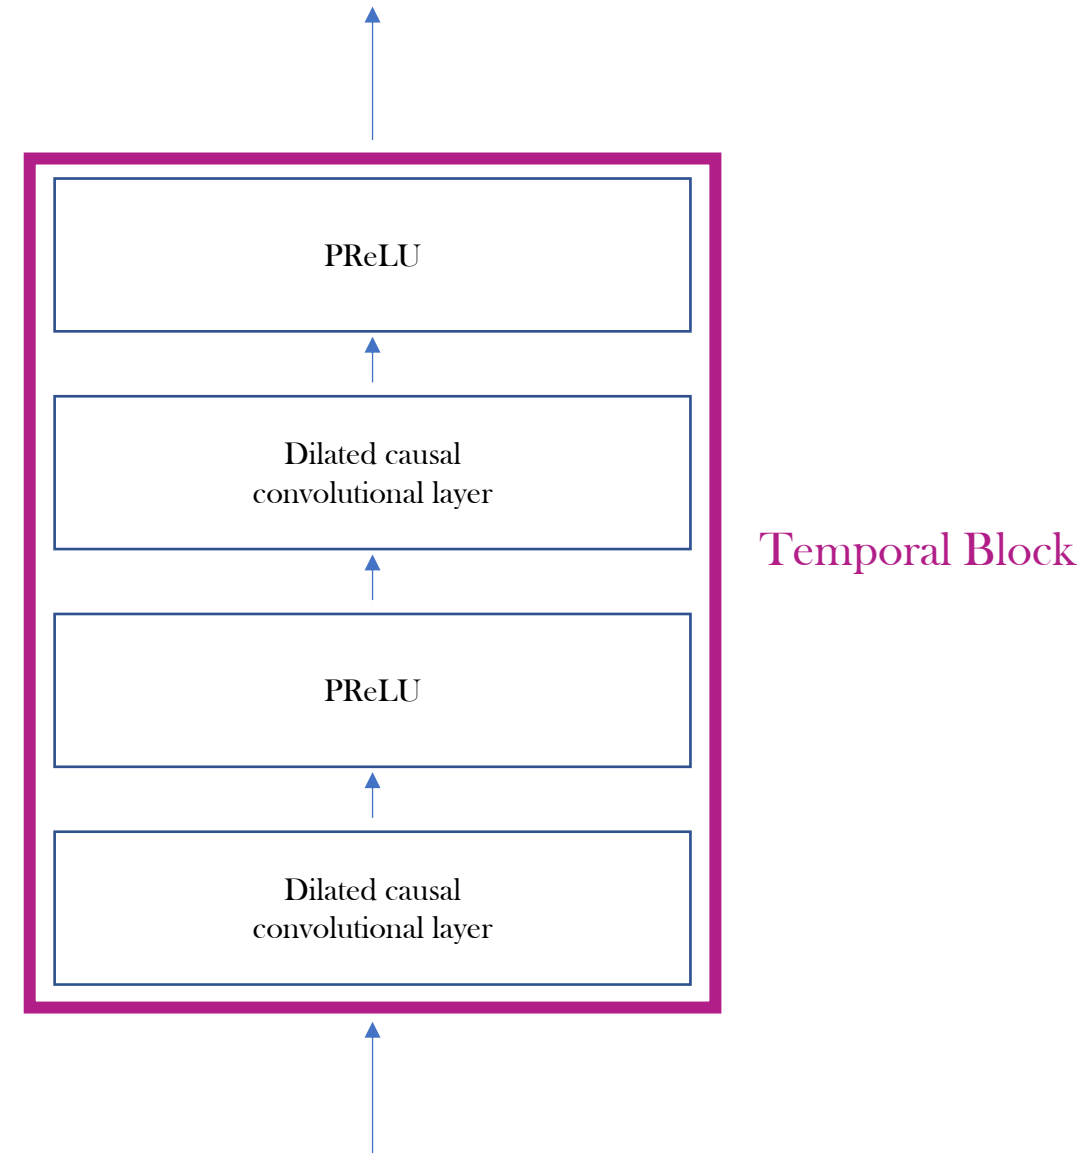




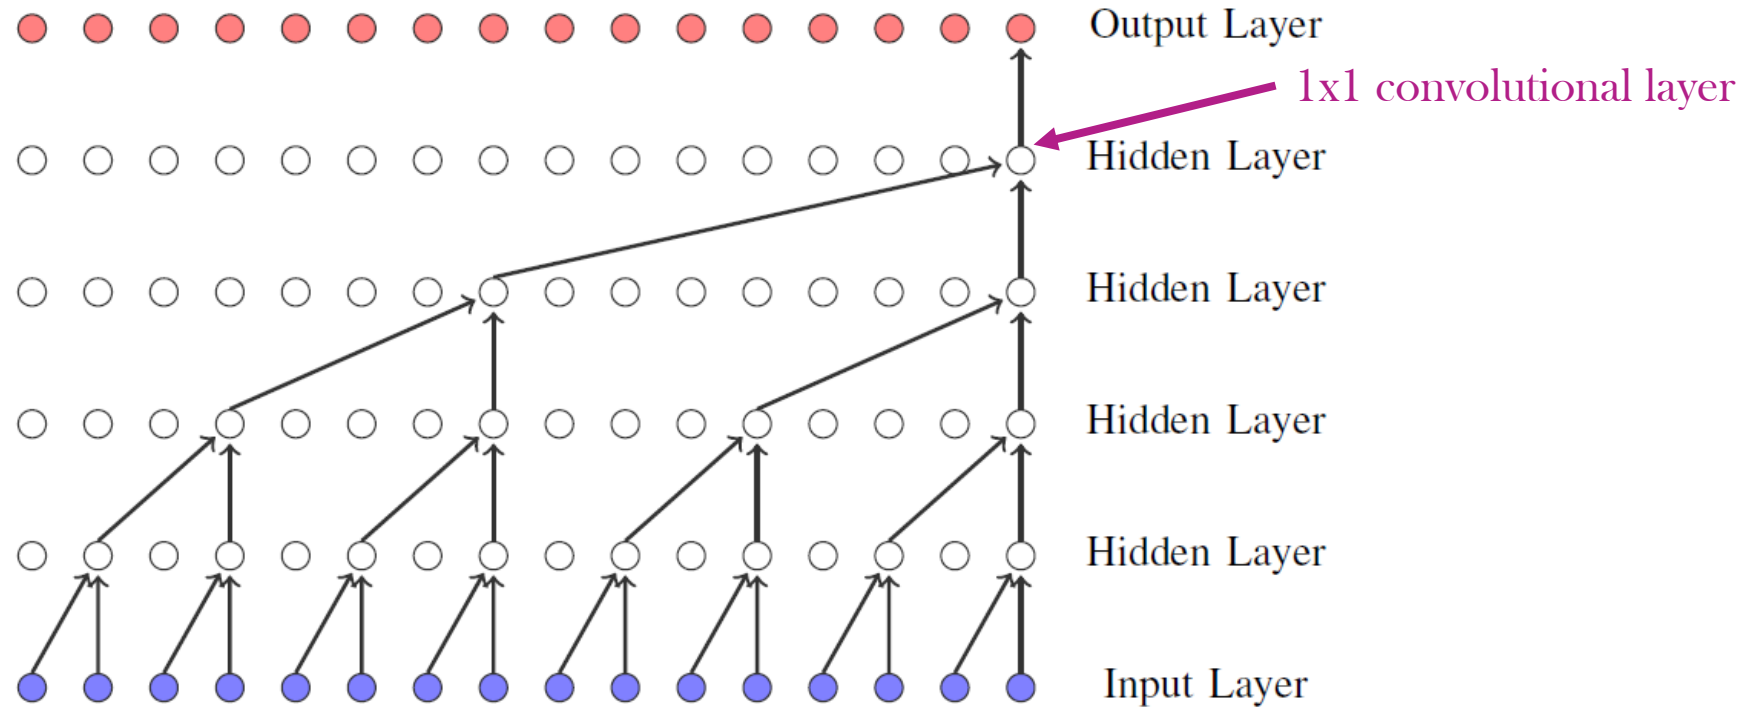
# Block Module



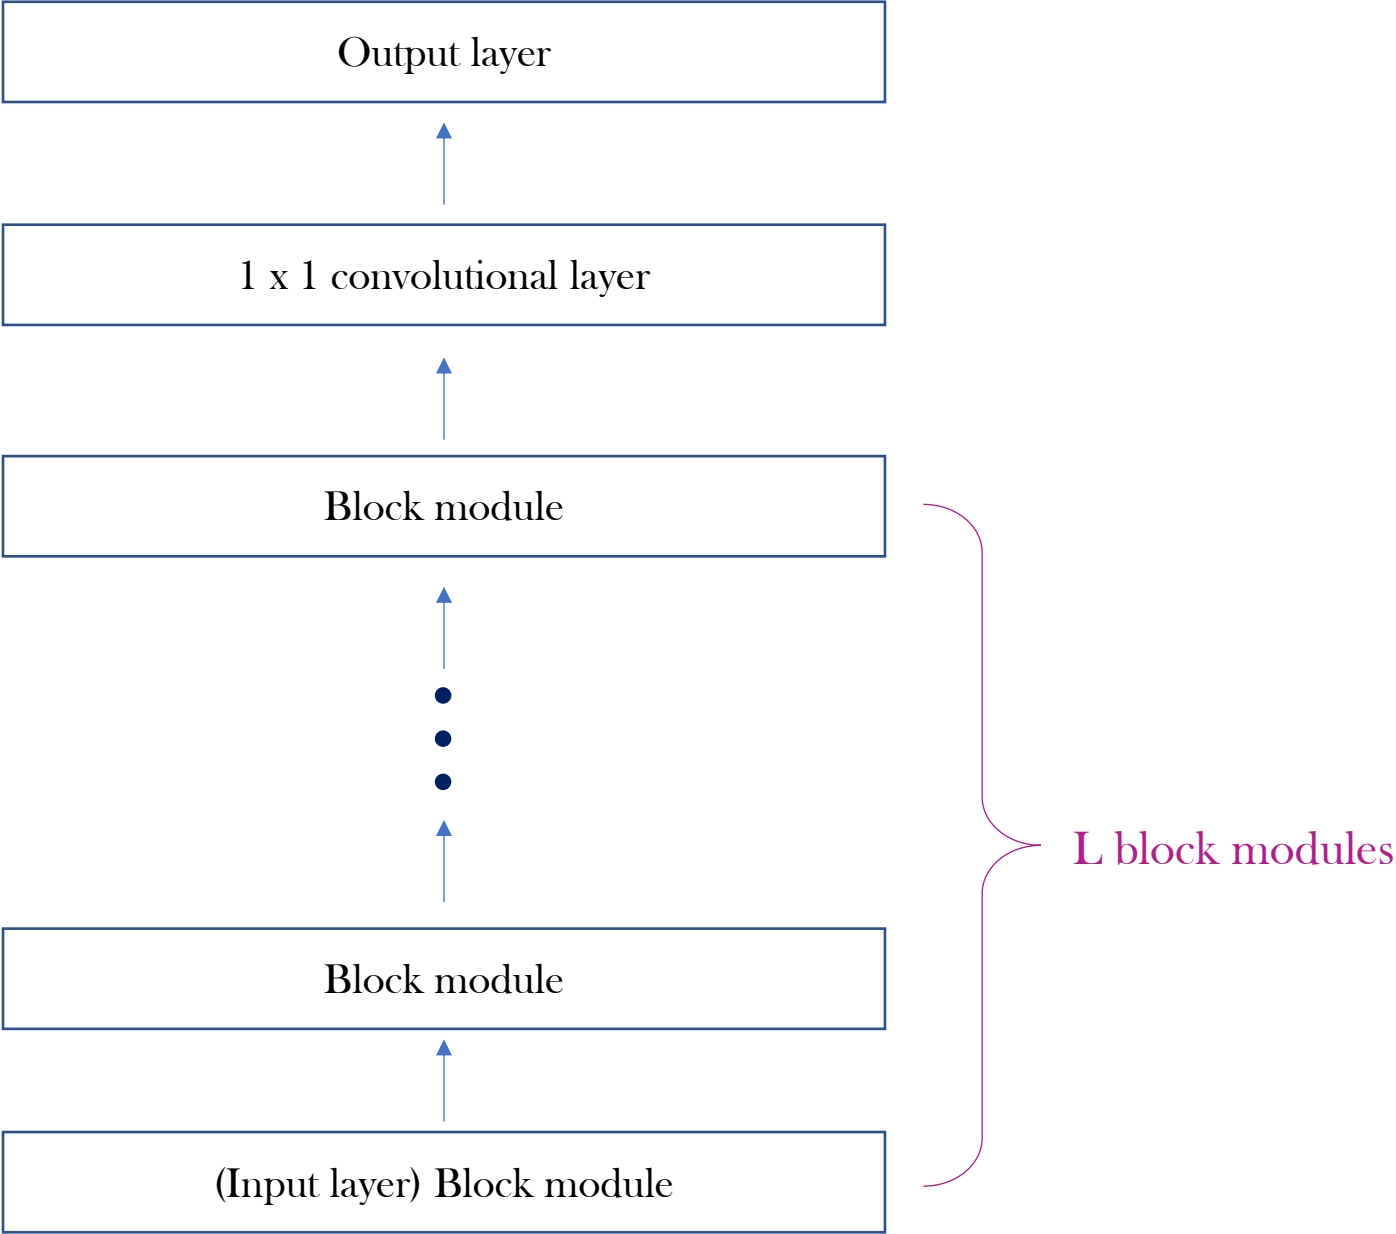
ex)



# 1x1 Convolutional Layer



Temporal Convolutional Network



# Skip Connections

## Skip Connections




Figure 2. Residual learning: a building block.

# TCN with Skip Connections

**Definition 3.15** (TCN with skip connections). Assume the notation from Definition 3.10 and for  $N_{skip} \in \mathbb{N}$  let

$$\gamma_l : \mathbb{R}^{N_{l-l} \times T_{l-1}} \rightarrow \mathbb{R}^{N_l \times T_l} \times \mathbb{R}^{N_{skip} \times T_L} \quad \text{for } l \in \{1, \dots, L\}$$

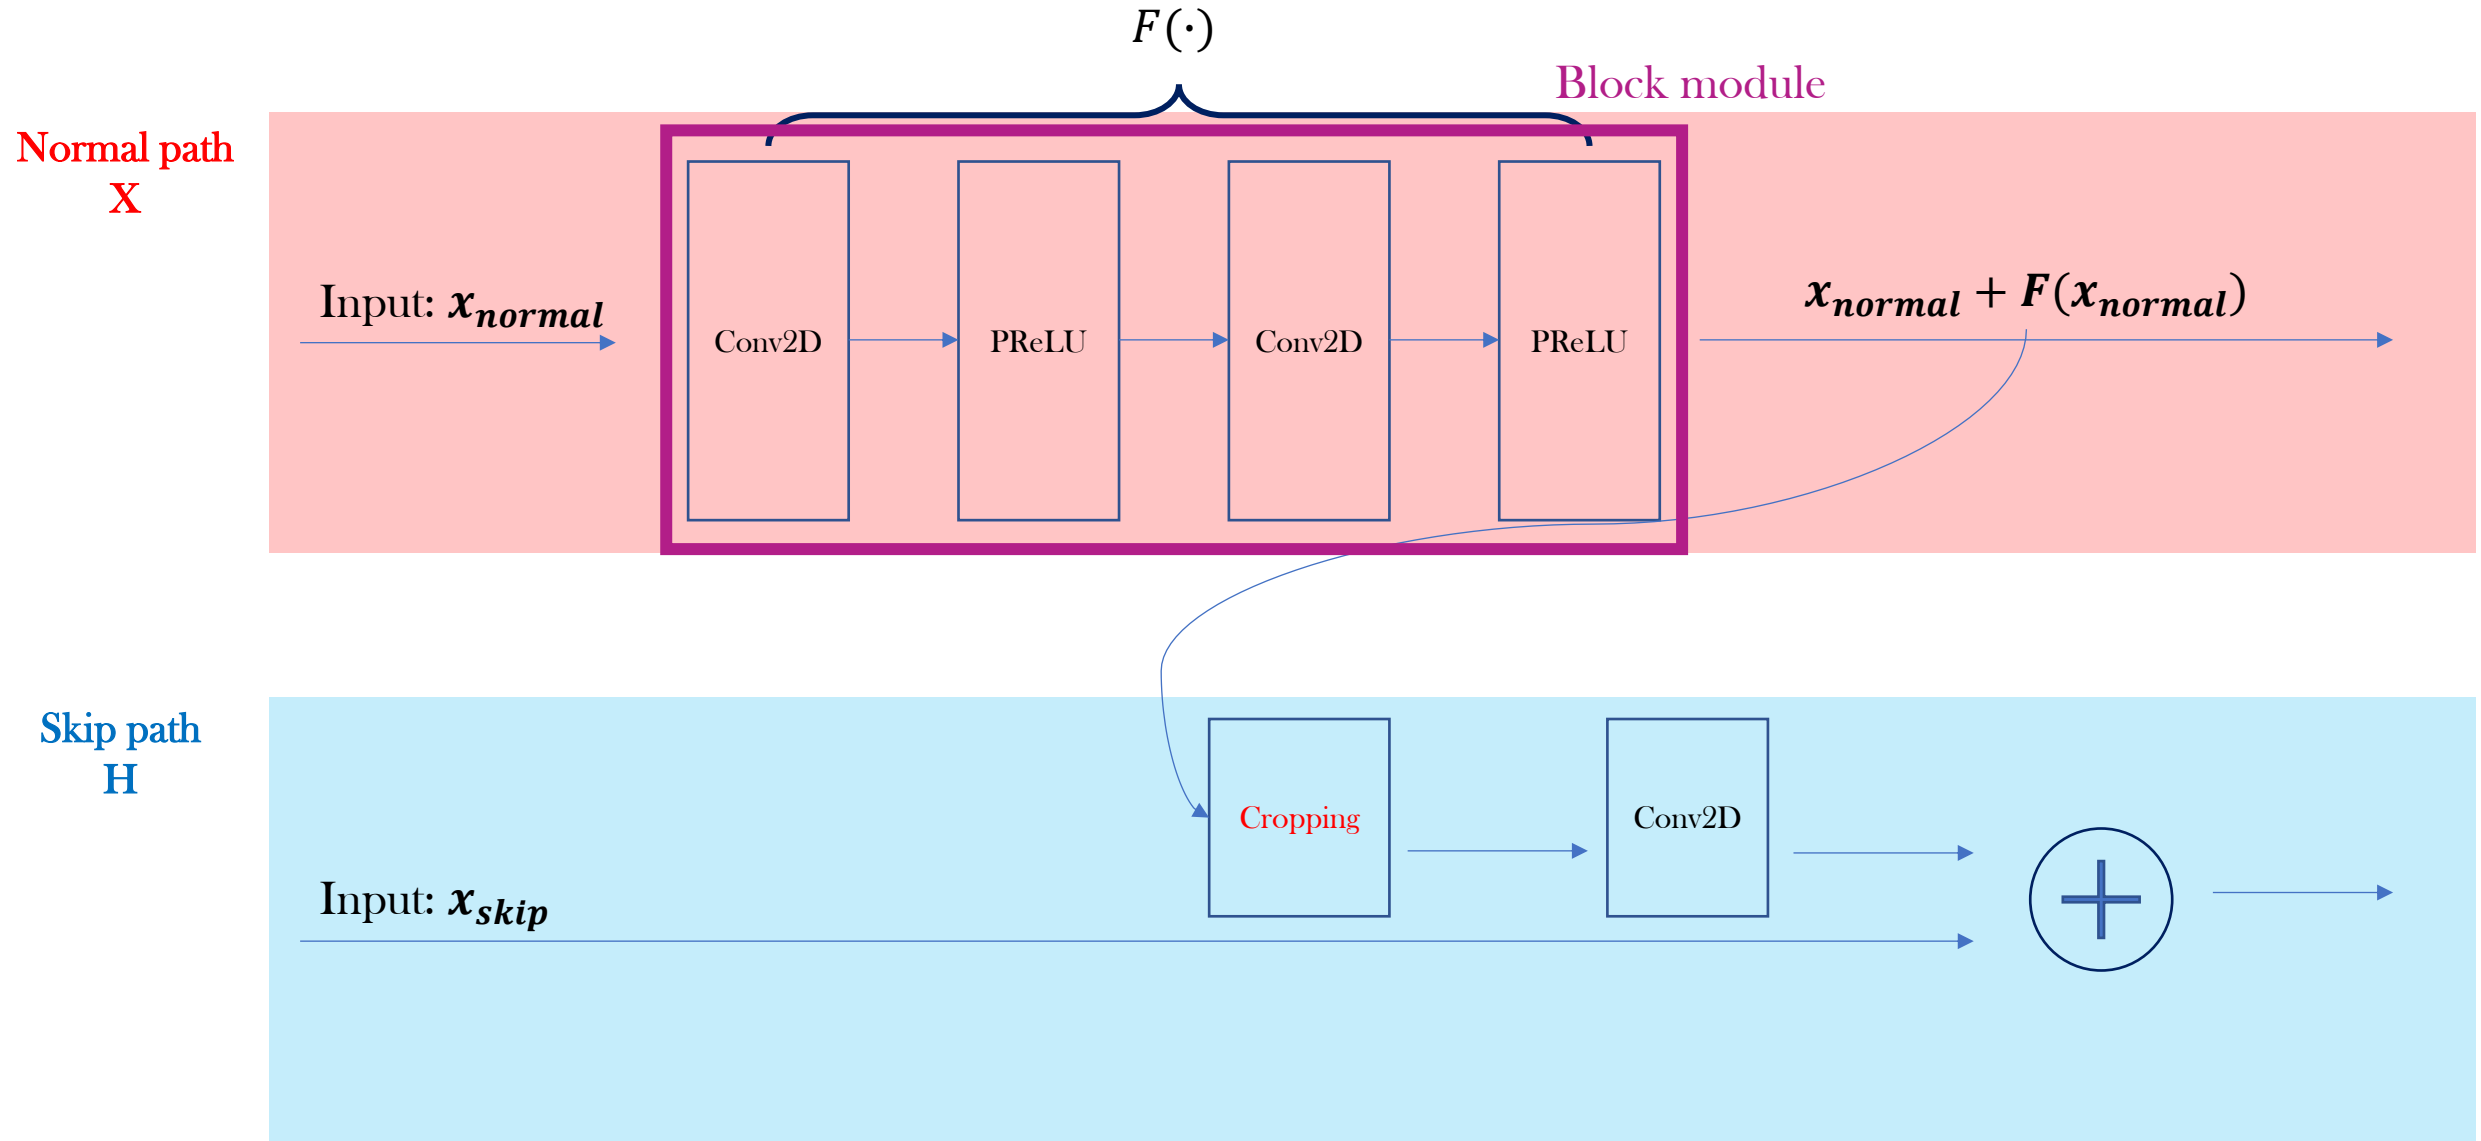
denote block modules. Moreover, let  $\gamma$  be a block module with arguments  $(N_{skip}, N_{L+1}, 0)$ . If the output  $Y \in \mathbb{R}^{N_{L+1} \times T_L}$  of a TCN  $f : \mathbb{R}^{N_0 \times T_0} \times \Theta \rightarrow \mathbb{R}^{N_{L+1} \times T_L}$  is defined recursively by

$$\left( X^{(l)}, H^{(l)} \right) = \gamma_l \left( X^{(l-1)} \right) \quad \text{for } l \in \{1, \dots, L\}$$

$$Y = \gamma \left( \sum_{l=1}^L H^{(l)} \right),$$

where  $X^{(0)} \in \mathbb{R}^{N_0 \times T_0}$ , then  $f$  is called a *temporal convolutional network with skip connections*.

# TCN with Skip Connections



# Vanilla TCN with Skip Connection




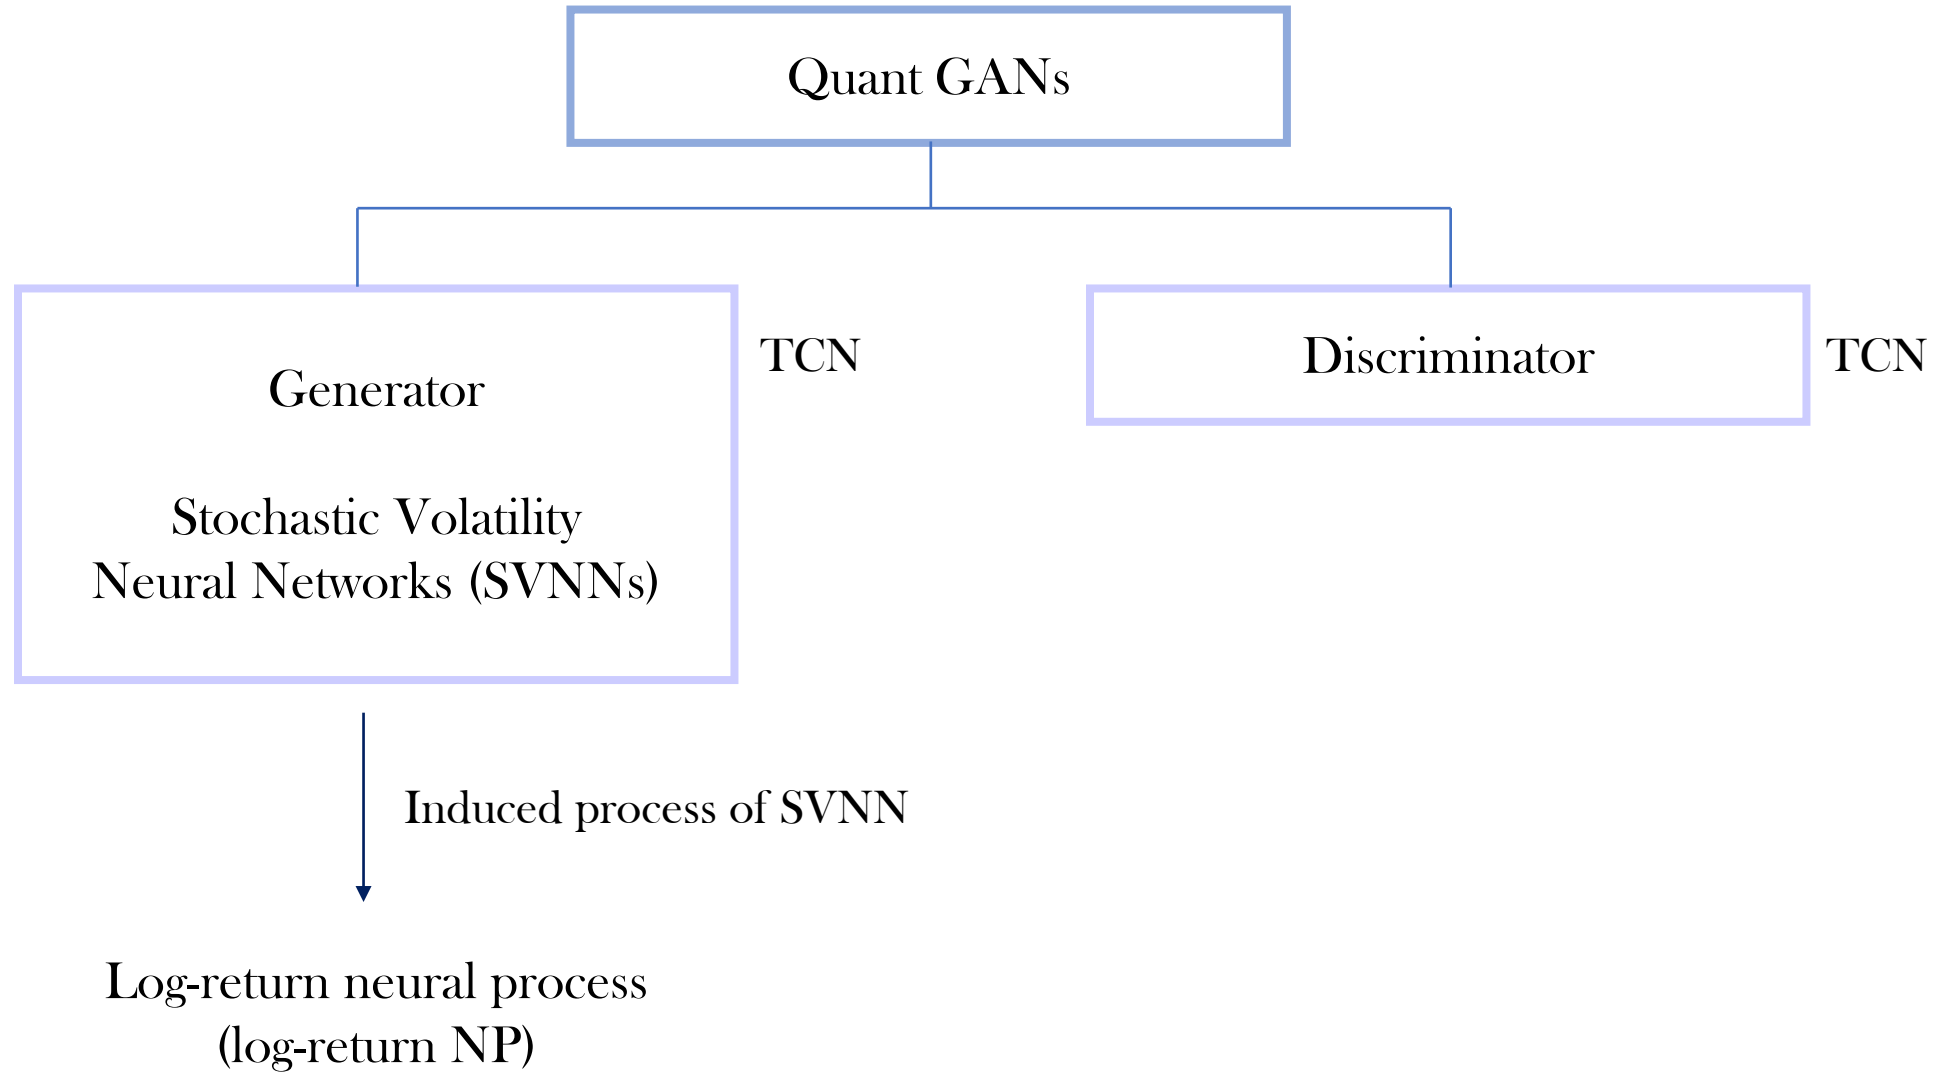
Figure 7: Vanilla TCN with skip connections.

# Table of Contents

- I. Introduction
- II. Generative Adversarial Networks (GANs)
- III. Temporal Convolutional Networks (TCNs)
- IV. Log-return Neural Process**
- V. Numerical Results



# Overview



# Log-return Neural Process

**Notation 4.4.** Consider a stochastic process  $(X_t)_{t \in \mathbb{Z}}$  parametrized by some  $\theta \in \Theta$ . For  $s, t \in \mathbb{Z}$ ,  $s \leq t$ , we write

$$X_{s:t,\theta} := (X_{s,\theta}, \dots, X_{t,\theta})$$

and for an  $\omega$ -realization

$$X_{s:t,\theta}(\omega) := (X_{s,\theta}(\omega), \dots, X_{t,\theta}(\omega)) \in \mathbb{R}^{N_X \times (t-s+1)}.$$

We can now introduce the concept of neural (stochastic) processes.

# Log-return Neural Process

**Definition 4.5** (Neural process). Let  $(Z_t)_{t \in \mathbb{Z}}$  be an i.i.d. noise process with values in  $\mathbb{R}^{N_Z}$  and  $g : \mathbb{R}^{N_Z \times T^{(g)}} \times \Theta^{(g)} \rightarrow \mathbb{R}^{N_X}$  a TCN with RFS  $T^{(g)}$  and parameters  $\theta \in \Theta^{(g)}$ . A stochastic process  $\tilde{X}$ , defined by

$$\begin{aligned}\tilde{X} : \Omega \times \mathbb{Z} \times \Theta^{(g)} &\rightarrow \mathbb{R}^{N_X} \\ (\omega, t, \theta) &\mapsto g_\theta(Z_{t-(T^{(g)}-1):t}(\omega))\end{aligned}$$

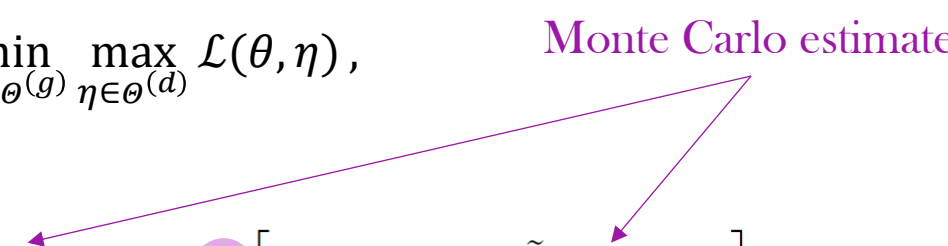
such that  $\tilde{X}_{t,\theta} : \Omega \rightarrow \mathbb{R}^{N_X}$  is a  $\mathcal{F} - \mathcal{B}(\mathbb{R}^{N_X})$ -measurable mapping for all  $t \in \mathbb{Z}$  and  $\theta \in \Theta^{(g)}$ , is called *neural process* and will be denoted by  $\tilde{X}_\theta := (\tilde{X}_{t,\theta})_{t \in \mathbb{Z}}$ .

# Log-return Neural Process

The GAN objective for stochastic processes can be formulated as

$$\min_{\theta \in \Theta(\mathcal{G})} \max_{\eta \in \Theta(\mathcal{D})} \mathcal{L}(\theta, \eta),$$

where

$$\mathcal{L}(\theta, \eta) := \mathbb{E}[\log(d_{\eta}(X_{1:T^{(d)}}))] + \mathbb{E}[\log(1 - d_{\eta}(\tilde{X}_{1:T^{(d)},\theta}))]$$


A diagram with the text "Monte Carlo estimate" in purple. Two purple arrows originate from this text. One arrow points to the first expectation term  $\mathbb{E}[\log(d_{\eta}(X_{1:T^{(d)}}))]$  in the equation below. The other arrow points to the second expectation term  $\mathbb{E}[\log(1 - d_{\eta}(\tilde{X}_{1:T^{(d)},\theta}))]$  in the same equation.

and  $X_{1:T^{(d)}}$  and  $\tilde{X}_{1:T^{(d)},\theta}$  denote the real and the generated process, respectively.

# Log-return Neural Process

**Definition 5.1** (Log return neural process). Let  $Z = (Z_t)_{t \in \mathbb{Z}}$  be  $\mathbb{R}^{N_Z}$ -valued i.i.d. Gaussian noise,  $g^{(\text{TCN})} : \mathbb{R}^{N_Z \times T^{(g)}} \times \Theta^{(\text{TCN})} \rightarrow \mathbb{R}^{2N_X}$  a TCN with RFS  $T^{(g)}$  and  $g^{(\epsilon)} : \mathbb{R}^{N_Z} \times \Theta^{(\epsilon)} \rightarrow \mathbb{R}^{N_X}$  be a network. Furthermore, let  $\alpha \in \Theta^{(\text{TCN})}$  and  $\beta \in \Theta^{(\epsilon)}$  denote some parameters. A stochastic process  $R$ , defined by


$$R : \Omega \times \mathbb{Z} \times \Theta^{(\text{TCN})} \times \Theta^{(\epsilon)} \rightarrow \mathbb{R}^{N_X} \\ (\omega, t, \alpha, \beta) \mapsto [\sigma_{t,\alpha} \odot \epsilon_{t,\beta} + \mu_{t,\alpha}] (\omega) ,$$

where  $\odot$  denotes the Hadamard product and

$$\begin{aligned} h_t &:= g_{\alpha}^{(\text{TCN})} (Z_{t-T^{(g)}:(t-1)}) \\ \text{Volatility TCN} \quad \sigma_{t,\alpha} &:= |h_{t,1:N_X}| \\ \text{Drift TCN} \quad \mu_{t,\alpha} &:= h_{t,(N_X+1):2N_X} \\ \text{Innovation NN} \quad \epsilon_{t,\beta} &:= g_{\beta}^{(\epsilon)} (Z_t) , \end{aligned}$$

is called *log return neural process*. The generator architecture defining the log return NP is called *stochastic volatility neural network (SVNN)*. The NPs  $\sigma_{\alpha} := (\sigma_{t,\alpha})_{t \in \mathbb{Z}}$ ,  $\mu_{\alpha} := (\mu_{t,\alpha})_{t \in \mathbb{Z}}$  and  $\epsilon_{\beta} := (\epsilon_{t,\beta})_{t \in \mathbb{Z}}$  are called *volatility*, *drift* and *innovation NP*, respectively.

# Log-return Neural Process



Structure of the SVNN architecture. The volatility and drift component are generated by inferring the latent process  $Z_{t-8:t-1}$  through the TCN, whereas the innovation is generated by inferring  $Z_t$ .

# Table of Contents

- I. Introduction
- II. Generative Adversarial Networks (GANs)
- III. Temporal Convolutional Networks (TCNs)
- IV. Log-return Neural Process
- V. Numerical Results**

# Numerical Results

## QuantGANs Using Pure TCN




Figure A1. Five generated driftless log paths.


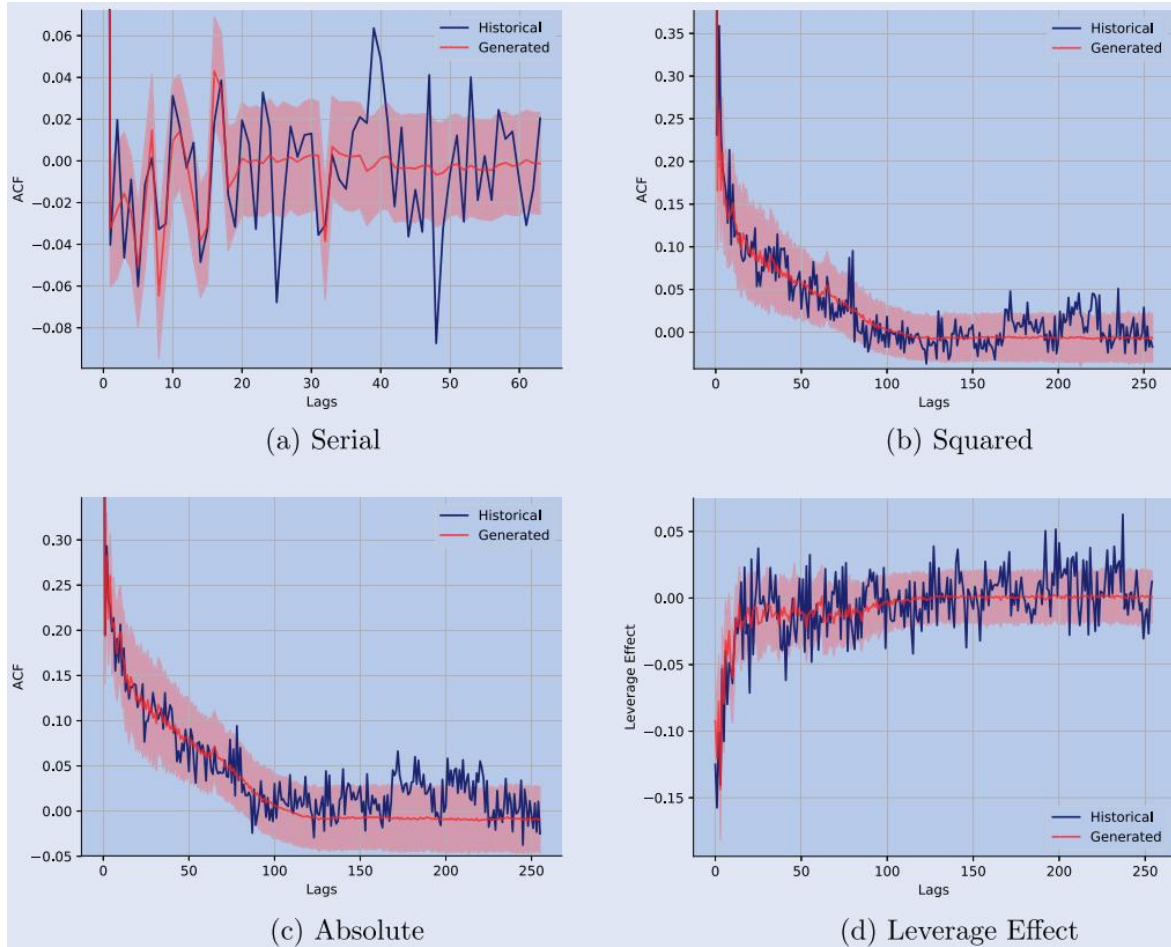


Figure A2. Fifty generated driftless log paths.

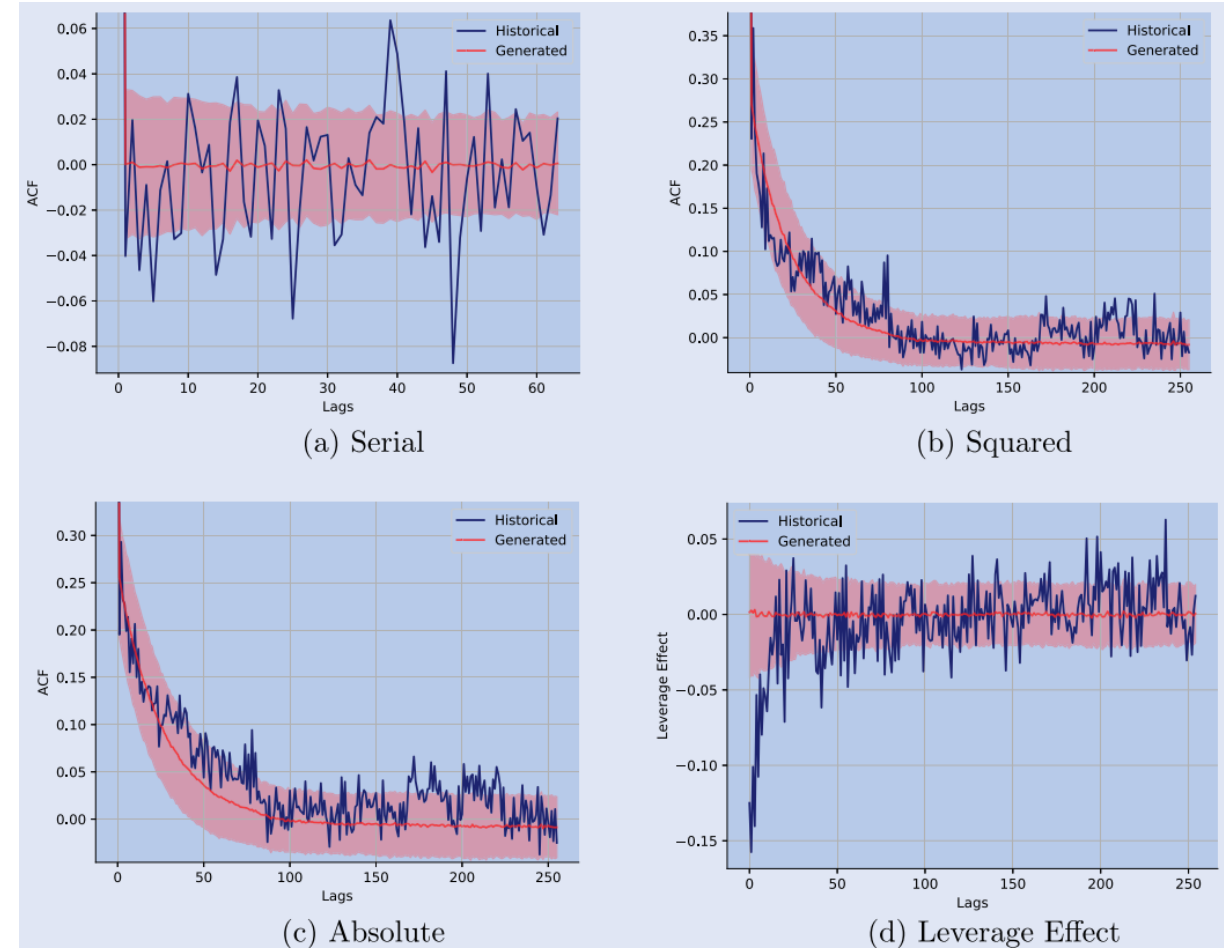


# QuantGANs VS GARCH(1,1) (old model)

## QuantGANs



## GARCH(1,1) (old model)



# Reference

Wiese, Magnus, et al. "Quant GANs: deep generation of financial time series." *Quantitative Finance* 20.9 (2020): 1419-1440.

Thank you for listening