Towards Editing Time Series

Baoyu Jing[2*], Shuqi Gu[1*], Tianyu Chen[1], Zhiyu Yang[1],
Dongsheng Li[3], Jingrui He[2], Kan Ren[1†]

[1]ShanghaiTech University
[2]University of Illinois Urbana-Champaign
[3]Microsoft Research
[*]Equal contribution. [†]Corresponding author.

Roadmap

- Background
- Problem Formulation
- Methodology
- Experiments
- Conclusion

Time Series Synthesize

- Time series associated with many attributes.
- Real-world time series are sparse and privacy-sensitive.

Attribute	City	Season
Value	London	Spring

Attribute	City	Season
Value	London	Summer

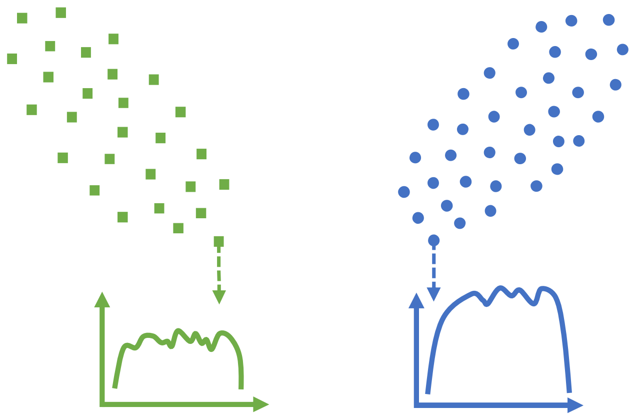

Conditional Time Series Generation

- Synthesize time series based on the condition.
- Do not support sample-level time series manipulation.

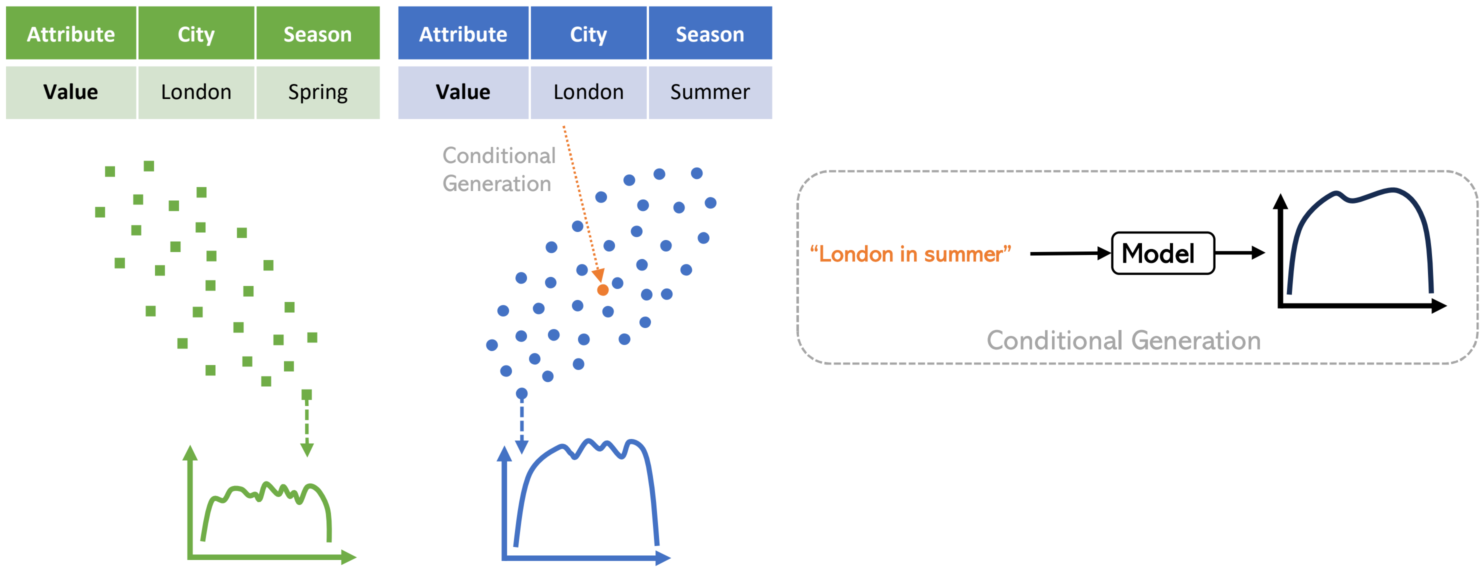

Time Series Editing

- Question: given a time series, what would it become if some of its attributes are modified?

Attribute	City	Season
Value	London	Spring

Attribute	City	Season
Value	London	Summer

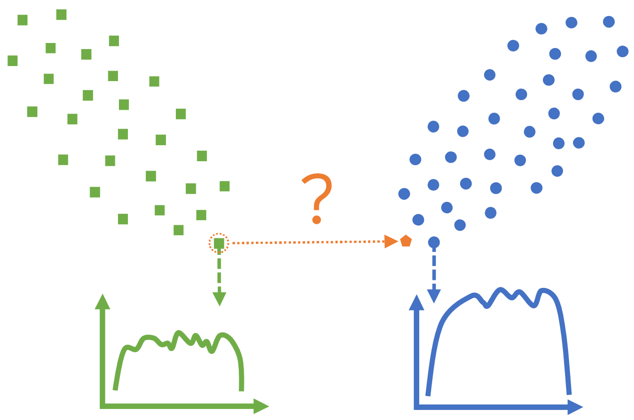

Time Series Editing (Cont.)

- We introduce a novel task – Time Series Editing (TSE).
    - Edit certain attributes of the given time series sample.
    - Preserve other information.

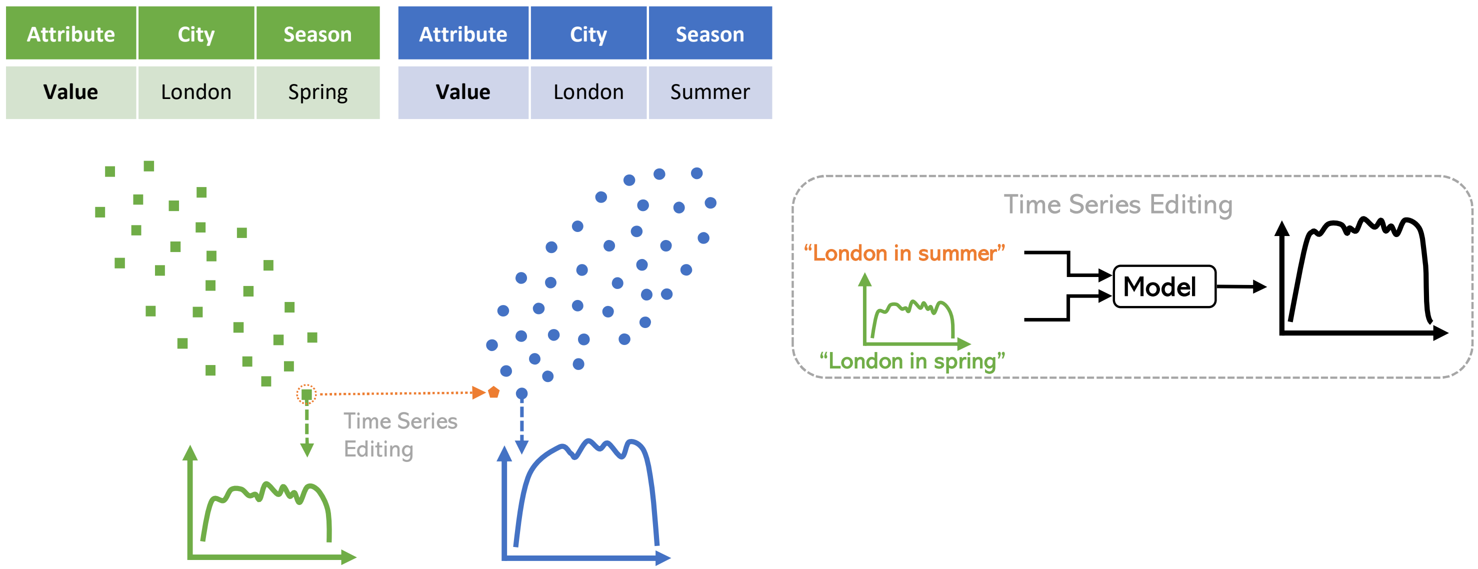

Challenges of Time Series Editing

- The time series data distribution over the attribute space is biased and may not be adequately covered.
  - E.g., in climate data, temperature and humidity are observable while atmospheric pressure variations missing.

- Different attributes influence timeseries at varying resolutions.
  - E.g., trends have a global impact, while seasonality exerts amore localized influence.

Roadmap

- Background
- Problem Formulation
- Methodology
- Experiments
- Conclusion

Time Series Editing

- Generate a target time series $\hat{\mathbf{x}}^{tgt} = \Phi_\theta(\mathbf{x}^{src}, \mathbf{a}^{src}, \mathbf{a}^{tgt})$, and
- Modifying the edited attributes $\mathcal{A}_{edit}$, while
- Maintaining the preserved attributes $\mathcal{A}_{prsv}$ and other information of $\mathbf{x}^{src}$

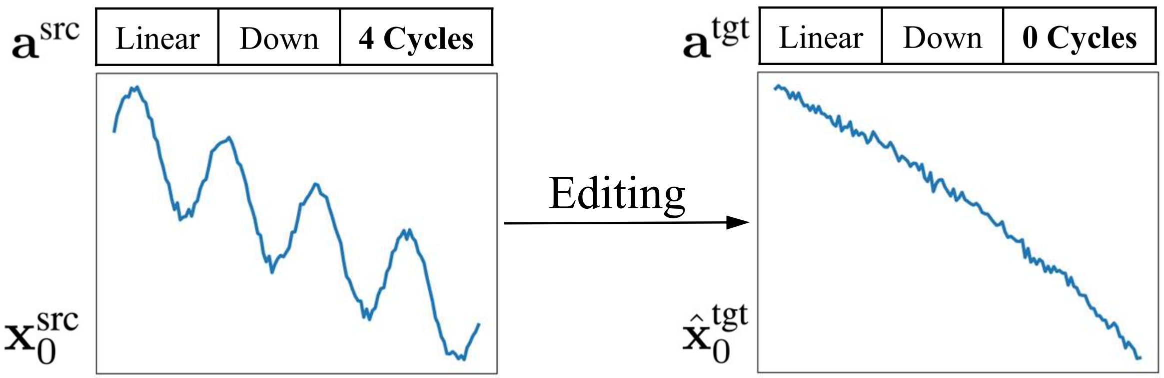

Roadmap

- Background
- Problem Formulation
- Methodology
- Experiments
- Conclusion

Multi-Resolution Noise Estimator

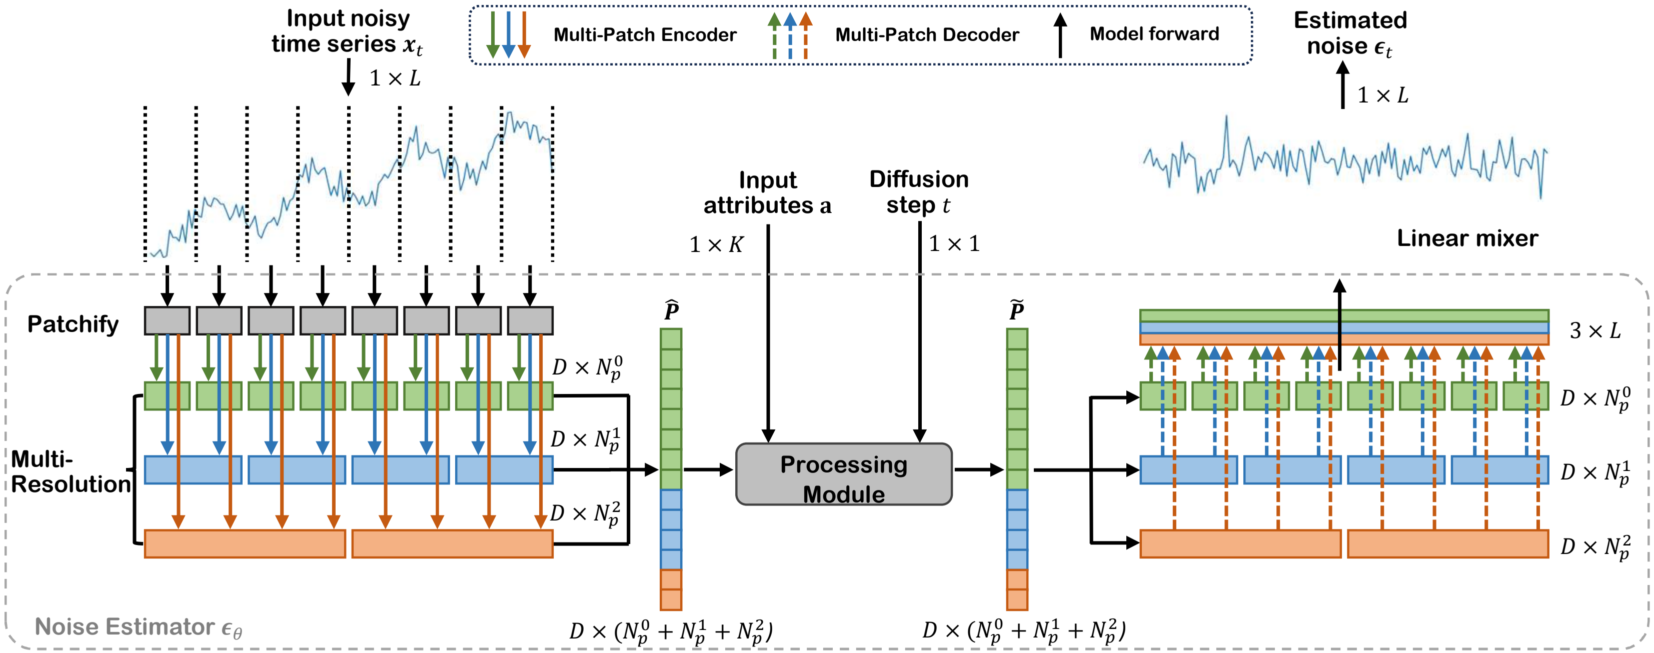

The core component of our proposed diffusion model is the noise estimator $\boldsymbol{\epsilon}_\theta$.
$\boldsymbol{\epsilon}_\theta$ captures the multi-resolution interactions between time series & attributes

Bootstrap Learning Algorithm

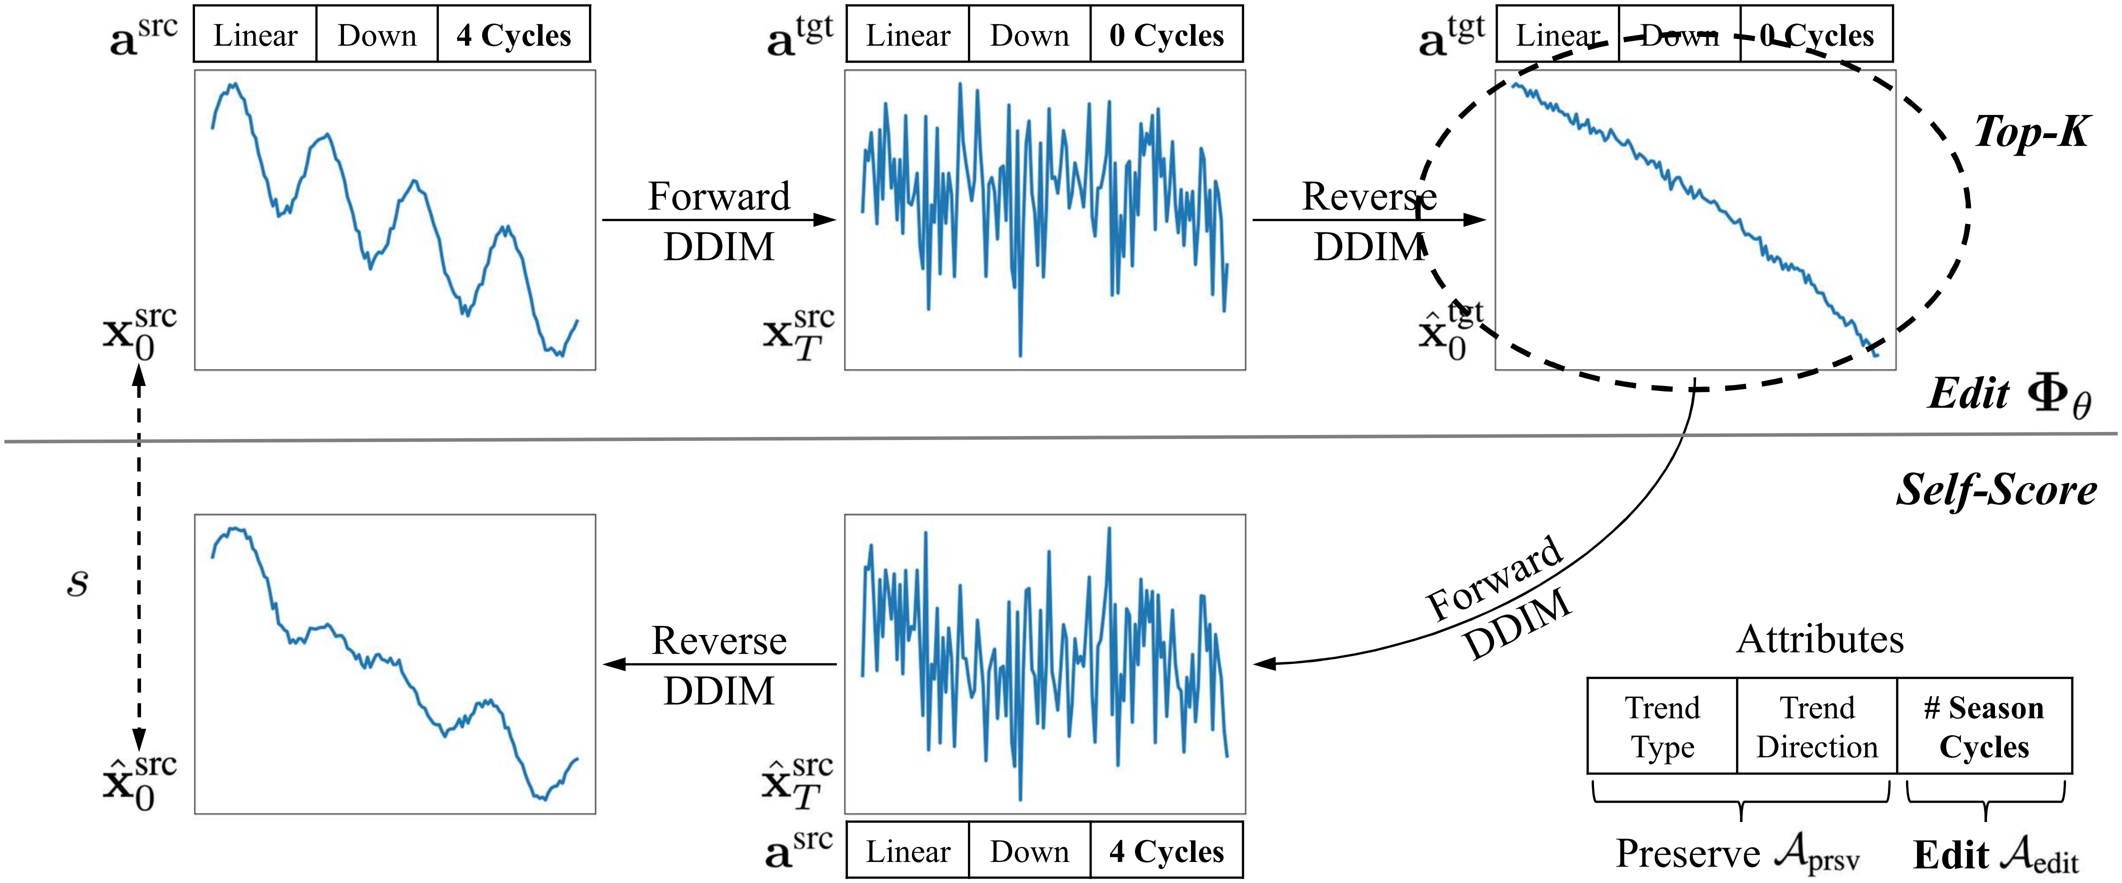

Roadmap

- Background
- Problem Formulation
- Methodology
- Experiments
- Conclusion

Evaluation Metrics

- Editability: for edited attributes: **CTAP** & **RaTS**
- Preservability: for preserved attributes: **CTAP** & **|RaTS|**
- For data with ground truth: **MSE** & **MAE**

Log Ratio of Target-to-Source (**RaTS**↑ and **|RaTS|**↓)

$$\mathrm{RaTS} = \log \frac{P(a_k^{tgt} | \hat{x}_0^{tgt})}{P(a_k^{tgt} | x_0^{src})}$$

Contrastive Time series-Attribute Pretraining (**CTAP**↑)

$$\mathrm{CTAP} = \mathrm{Sim}(\hat{x}_0^{tgt}, a_k^{tgt})$$

Results

	Synthetic						Air				Motor			
	Overall		Edited		Preserved		Edited		Preserved		Edited		Preserved	
	↓MSE	↓MAE	↑RaTS	↑CTAP	↓|RaTS|	↑CTAP	↑RaTS	↑CTAP	↓|RaTS|	↑CTAP	↑RaTS	↑CTAP	↓|RaTS|	↑CTAP
CSDI	0.1789	0.3221	0.7540	0.5405	0.1439	0.7898	0.7452	0.1581	0.1705	0.6311	0.0939	0.4203	0.1597	0.6617
Time Weaver	0.1454	0.2898	0.9030	0.6943	0.1169	0.8292	0.8956	0.3266	0.1866	0.6299	0.0979	0.4168	**0.1520**	**0.6691**
TEdit-CSDI	**0.1235**	**0.2606**	0.9257	0.7109	0.1021	0.8553	0.8022	0.2179	**0.1614**	**0.6529**	0.1016	0.4186	0.1580	0.6654
TEdit-TW	0.1315	0.2722	**1.0121**	**0.7957**	**0.0995**	**0.8622**	**0.9661**	**0.3930**	0.1916	0.6274	**0.1212**	**0.4348**	0.1571	0.6621

Table 1: Averaged performance over all finetuning sets for Synthetic (left), Air (middle), and Motor (right). "Edited" and "Preserved" are the average results of all edited and preserved attributes.

Our method perform better than baselines both on editing and preserving the attributes

Results

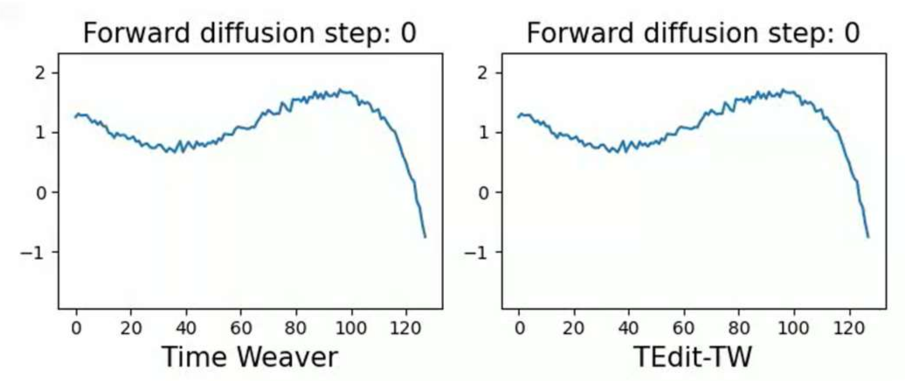

Edit the trend type from **exponential** to **logarithm**

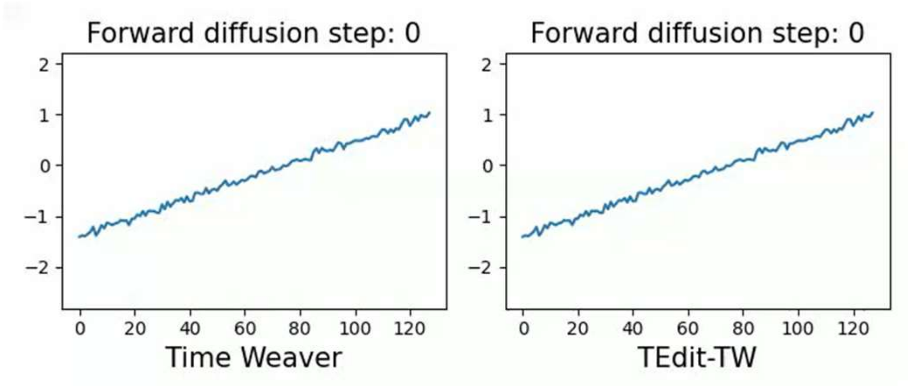

Edit the trend direction from **up** to **down**

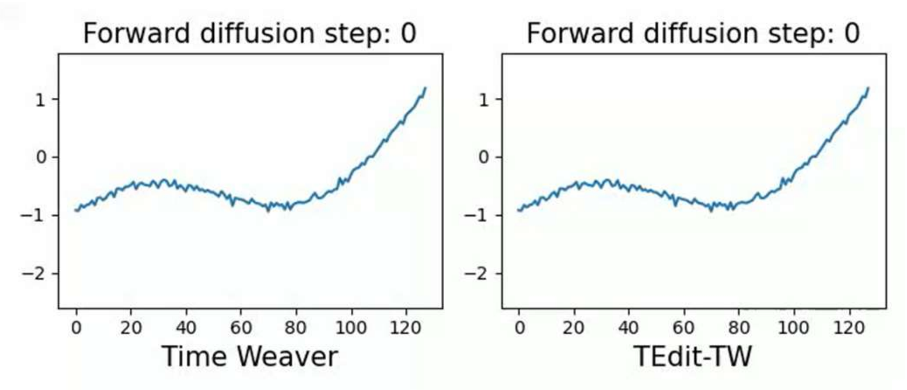

Edit the season cycle from **1** to **4**

Our method perform better than baselines both on editing and preserving the attributes

Ablation study

	Synthetic						Air				Motor			
	Overall		Edited		Preserved		Edited		Preserved		Edited		Preserved	
	↓MSE	↓MAE	↑RaTS	↑CTAP	↓|RaTS|	↑CTAP	↑RaTS	↑CTAP	↓|RaTS|	↑CTAP	↑RaTS	↑CTAP	↓|RaTS|	↑CTAP
TEdit-TW w GT	**0.1233**	**0.2622**	**1.0165**	**0.7991**	0.0984	**0.8650**	-	-	-	-	-	-	-	-
TEdit-TW	0.1315	0.2722	1.0121	0.7957	0.0995	0.8622	**0.9661**	**0.3930**	0.1916	0.6274	**0.1212**	**0.4348**	0.1571	0.6621
w/o BS	0.1376	0.2793	1.0127	0.7952	**0.0962**	0.8632	0.9524	0.3792	**0.1839**	**0.6418**	0.1113	0.4289	0.1571	0.6658
w/o BS & MR	0.1454	0.2898	0.9030	0.6943	0.1169	0.8292	0.8956	0.3266	0.1866	0.6299	0.0978	0.4168	**0.1520**	**0.6691**

Table 2: Ablation studies on the Synthetic, Air and Motor datasets. GT, BS and MR refer to Ground Truth source and target pairs, BootStrap and Multi-Resolution. Results are averaged over all finetuning sets. "Edited" and "Preserved" are the average results of all edited and preserved attributes.

Both multi-resolution architecture and bootstrapped training algorithm improve the performance of editing and preserving attributes.

Roadmap

- Background
- Problem Formulation
- Methodology
- Experiments
- Conclusion

Conclusion

- Contribution

  1. We introduce a novel task called **Time Series Editing (TSE)** for **sample-level time series manipulation.**

  2. We introduce a novel **diffusion-based** method: **TEdit**, which is equipped with:

     - **A bootstrap learning algorithm** for the problem of data coverage.
     - **A multi-resolution noise estimator** for the multi-scale interaction between time series and attributes.

- Application
  - Climate monitoring, healthcare, and urban management

Thanks :)