QuITE: Query-Based Irregular Time Series Embedding

Junghoon Lim · SK Shieldus, Seongnam, Republic of Korea · junghoon9502@gmail.com

1. Motivation

- **Irregularities in IMTS**

1. Misaligned time points across features (i.e., **misalignment**)
2. Inconsistent intervals between observations (i.e., **inconsistency**)

→ Hard to Modeling

- **A new approach : *Input-embedding-based***

Approach	No Artificial Value	Model Flexibility
Architecture-based	✓	×
Data-based	×	✓
Input-embedding-based (Ours)	✓	✓

➢ **Existing approaches & Limits**

1. Architecture-based : design specialized architectures **but limit the reuse of proven Multivariate Time Series (MTS) models.**
2. Data-based : map IMTS onto regular temporal grids through interpolation at either the raw-data or representation level, **which may distort temporal dynamics by introducing artificial values**

"The key bottleneck is not the backbone architecture, but the input embedding layer: conventional MTS embeddings assume regular sampling and thus struggle with IMTS."

By handling irregularity at the embedding stage, SOTA MTS models can be directly adapted to IMTS **without interpolation or architectural redesign.**

2. QuITE: Query-Based Irregular Time Series Embedding

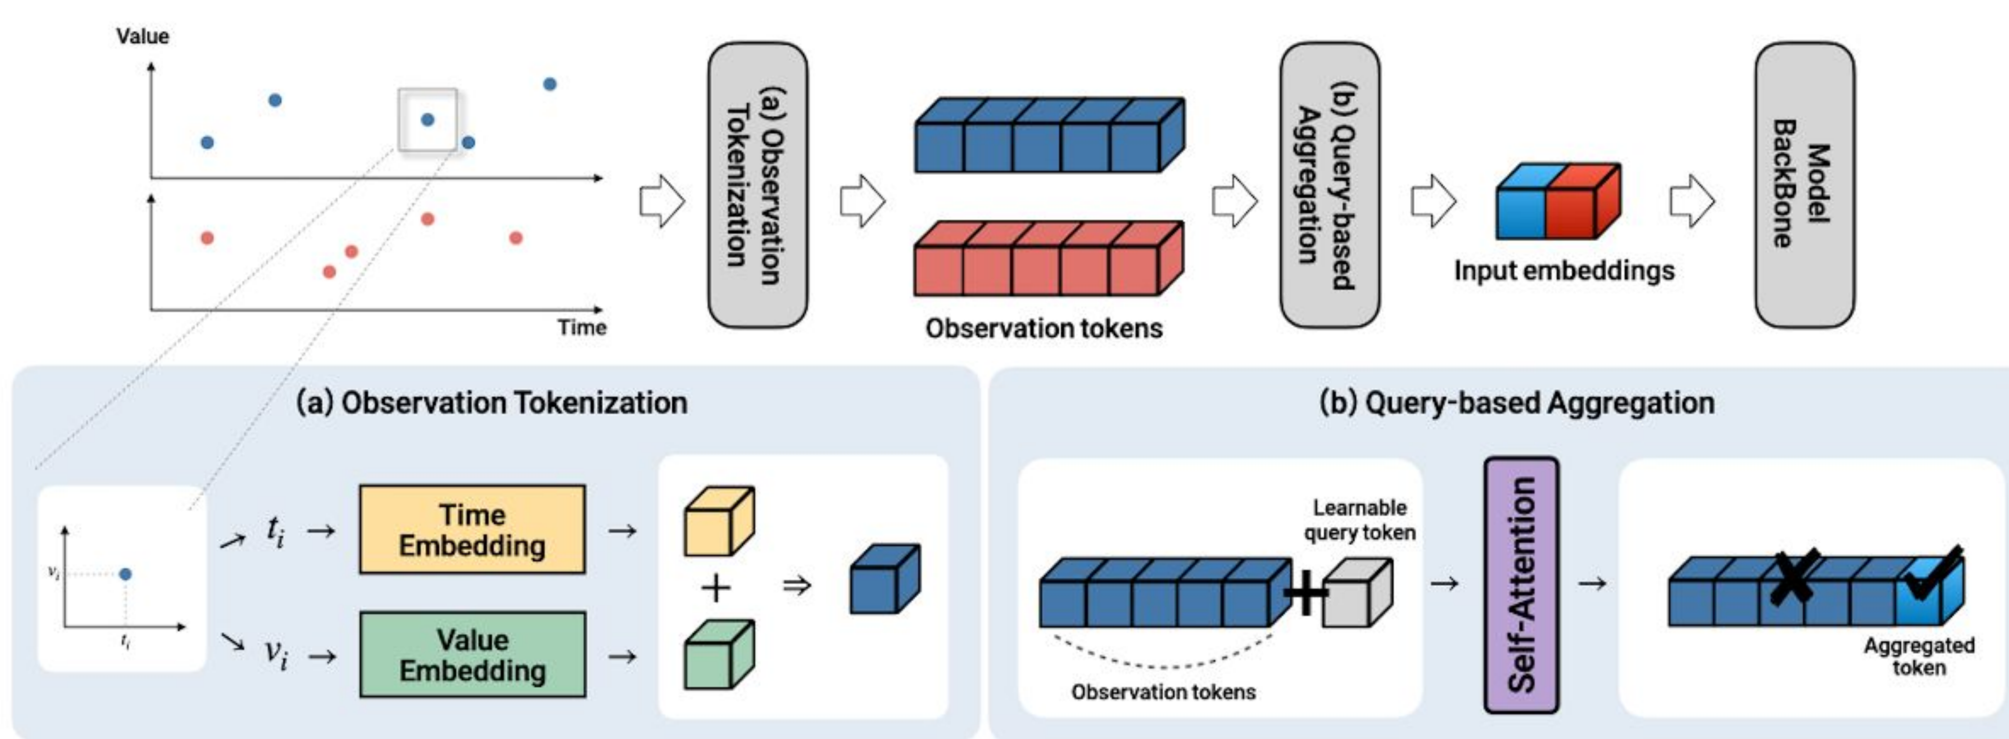

Challenge

Existing input embedding strategies for IMTS are limited

- Simple value-time fusion **assumes regular sampling.**
- Attention mechanisms **require pooling (e.g., [B, N, L, D] → [B, N, D]), which can lose fine-grained temporal information.**

QuITE

"Inspired by BERT's [CLS] token, uses learnable query tokens as IMTS embedding anchors."

- **A plug-and-play IMTS embedding module** for existing MTS models to handle IMTS **without interpolation, pooling, or architectural changes.**
- Converts IMTS into backbone-compatible embeddings via self-attention.

Embedding Process

(A) Observation Tokenization

$$\mathbf{z}_{n,i} = f_{\text{val}}(x_{n,i}) + \phi(t_{n,i})$$

Combines a harmonic time embedding with a learned value projection.

(B) Query-based Aggregation

- Prepends learnable query tokens to irregular observation tokens.
- Aggregates observed entries through a single masked self-attention layer.
- Supports variable- and patch-level aggregation for different MTS backbones.

3. QuITE++

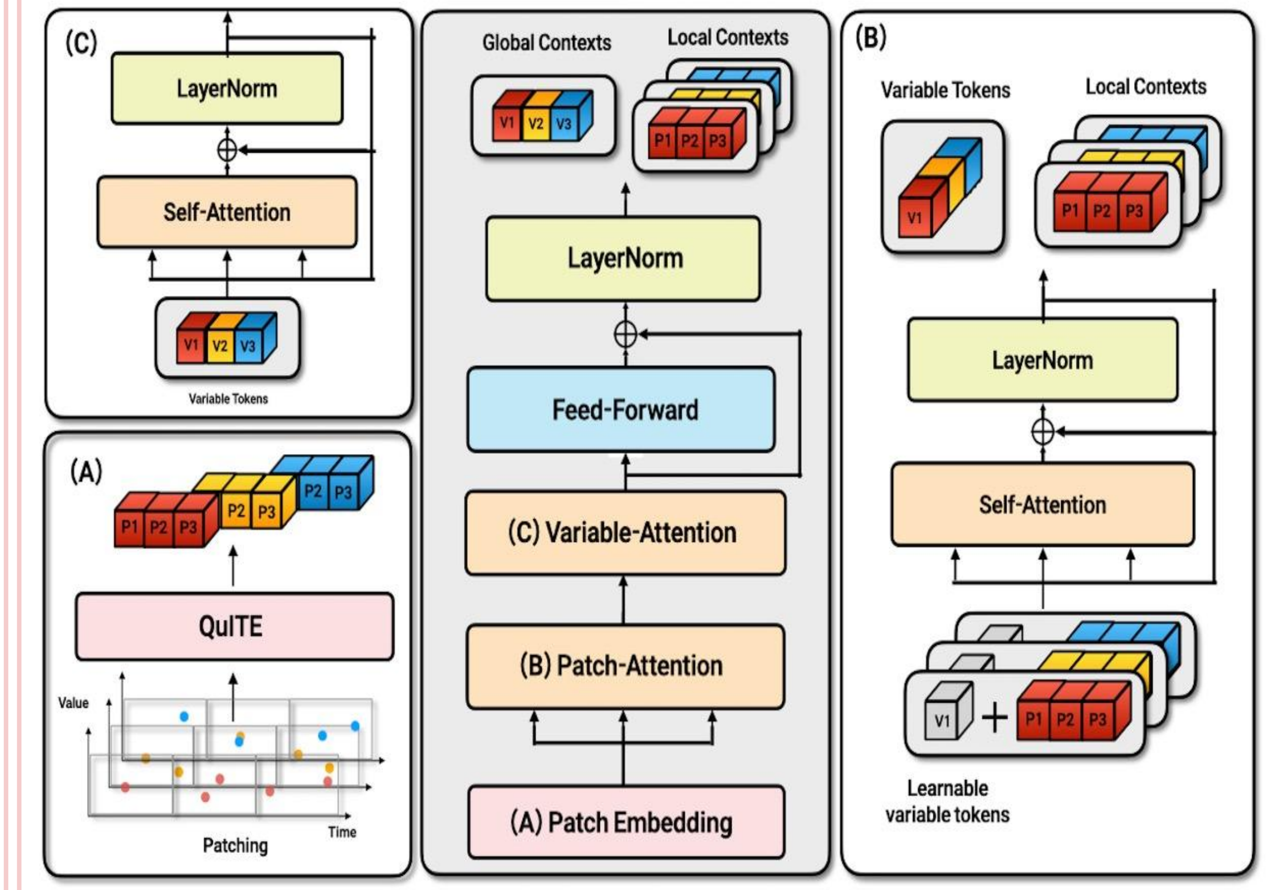

QuITE++

- **Extends the learnable query-token principle** into a forecasting architecture.
- **A hierarchical encoder** that models intra-variable patch-level temporal dependencies and inter-variable interactions.

Patch-level attention

- Aggregates information across temporal patches within each variable.

Variable-level attention

- Models cross-variable dependencies along the variate axis.

4. Experiments

- **Dataset**

(A) Forecasting

	Human Activity	USHCN	PhysioNet	MIMIC-III
# Samples	5400	26736	12000	23457
# Variables	12	5	36	96
# Avg. Length	120	163	74	46
Missing Ratio	75%	77.9%	88.4%	96.7%

(B) Classification

	P19	P12	PAM
# Samples	38803	11988	5333
# Variables	34	36	17
# Classes	2	2	8
Missing Ratio	94.9%	88.4%	60.0%

- **Effectiveness of QuITE**

Forecasting ↓better

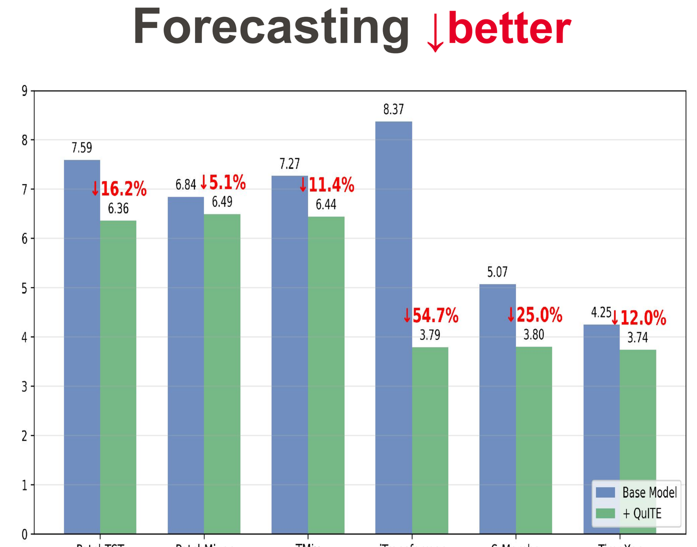

Classification ↑better

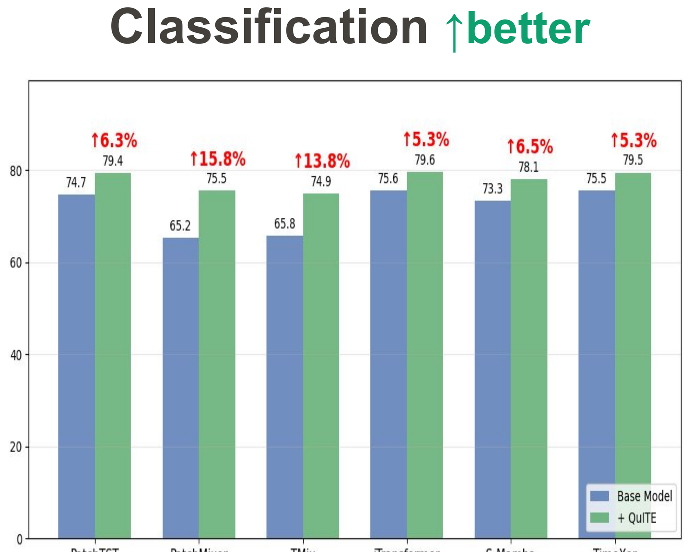

- Replace only the input embedding with QuITE.
- QuITE **consistently improves MTS backbones**, achieving up to **54.7%** forecasting and **15.8%** classification gains.

- **Forecasting**

Dataset	Human Activity (ms)						USHCN (m)					
Horizon	3000 → 1000		2000 → 2000		1000 → 3000		24 → 1		24 → 6		24 → 12	
Metric	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
Warpformer	2.61	3.12	3.60	3.81	4.26	4.26	5.09	3.10	5.12	3.13	5.10	3.13
Raindrop	4.42	4.65	5.57	5.15	5.75	5.37	5.64	3.29	7.01	4.24	7.61	4.61
GRU-D	3.94	4.37	5.93	5.66	6.14	5.75	5.17	3.21	5.29	3.34	5.36	3.25
tPatchGNN	2.79	3.24	3.71	3.89	4.56	4.32	5.00	3.07	5.23	3.24	6.23	3.83
GraFITi	3.03	3.45	4.59	4.45	4.91	4.62	5.07	2.97	5.12	3.09	5.01	3.14
CRU	3.03	3.60	4.12	4.43	4.85	4.86	5.15	3.18	6.77	4.11	6.64	4.08
mTAND	3.14	3.71	4.38	4.59	5.29	5.12	5.03	3.00	5.16	3.10	5.07	3.09
NeuralFlow	4.29	4.61	5.47	5.35	6.01	5.66	5.41	3.35	5.52	3.46	5.48	3.56
Latent-ODE	3.32	3.91	5.04	5.11	5.48	5.33	5.16	3.21	5.18	3.36	5.23	3.35
HyperIMTS	2.49	3.02	3.15	3.58	4.00	4.13	4.96	3.00	4.99	3.10	4.97	3.08
Hi-Patch	2.56	3.12	3.26	3.67	4.20	4.22	5.00	3.03	5.13	3.05	5.04	3.03
PatchTST + QuITE	2.76	3.14	3.62	3.74	4.69	4.29	5.07	3.01	5.06	3.17	5.04	3.00
PatchMixer + QuITE	2.78	3.13	3.67	3.75	4.71	4.31	5.11	3.06	5.02	3.04	5.00	3.07
TMix + QuITE	2.77	3.15	3.66	3.74	4.75	4.38	5.18	3.10	5.52	3.39	5.03	3.05
iTransformer + QuITE	2.58	3.12	3.25	3.65	4.10	4.18	5.06	3.06	4.86	2.96	4.94	3.01
S-Mamba + QuITE	2.72	3.24	3.37	3.72	4.20	4.22	5.04	3.06	4.93	3.04	4.93	3.02
TimeXer + QuITE	2.53	3.04	3.19	3.57	4.04	**4.03**	4.97	2.97	4.98	3.05	4.93	2.97
QuITE^{++}	**2.46**	**2.92**	**3.11**	**3.49**	**3.96**	4.04	**4.84**	**2.92**	**4.81**	**2.94**	**4.81**	**2.93**

Dataset	PhysioNet (h)						MIMIC-III (h)					
Horizon	12 → 36		24 → 24		36 → 12		12 → 36		24 → 24		36 → 12	
Metric	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
Warpformer	6.51	4.24	5.04	3.72	4.17	3.38	2.32	8.14	1.76	7.27	1.45	6.74
Raindrop	10.24	5.83	10.63	6.02	10.67	5.87	2.36	8.63	2.31	8.61	2.21	9.17
GRU-D	7.80	5.13	5.76	4.53	6.85	4.88	2.39	8.43	2.35	8.34	2.03	8.14
tPatchGNN	6.45	4.24	5.06	3.75	4.22	3.38	2.35	8.23	1.97	7.76	**1.44**	6.78
GraFITi	6.30	4.38	5.11	3.96	4.58	3.65	2.22	8.13	1.76	7.28	1.61	7.16
CRU	7.66	4.97	6.43	4.51	6.74	4.82	2.34	8.32	2.23	7.99	2.00	8.16
mTAND	7.46	4.85	6.18	4.44	5.61	4.15	2.29	8.38	2.15	8.00	2.01	8.13
NeuralFlow	7.98	5.08	7.68	4.84	8.87	5.43	2.26	8.29	2.34	8.09	1.97	8.39
Latent-ODE	7.28	4.83	6.85	4.77	6.99	4.74	2.38	8.35	2.11	7.76	1.90	7.92
HyperIMTS	6.11	4.23	**4.65**	**3.56**	3.99	3.21	1.85	7.71	1.68	6.92	1.52	6.68
Hi-Patch	6.39	4.10	5.07	3.63	4.27	3.30	1.88	7.95	1.70	7.18	1.56	6.74
PatchTST + QuITE	17.47	7.15	10.62	5.17	8.87	4.43	5.37	17.01	3.90	12.35	3.82	11.84
PatchMixer + QuITE	17.52	7.22	11.88	5.62	9.06	4.55	5.37	16.96	3.94	12.34	3.87	11.51
TMix + QuITE	17.48	7.21	10.72	5.22	8.96	4.50	5.41	16.94	3.98	12.32	3.87	11.63
iTransformer + QuITE	6.32	4.15	4.99	3.65	4.33	3.34	1.83	7.64	1.67	6.93	1.56	6.78
S-Mamba + QuITE	6.26	4.11	5.11	3.67	4.11	3.27	1.82	7.56	1.64	6.90	1.52	6.64
TimeXer + QuITE	6.18	4.08	4.91	3.64	4.06	3.27	1.84	7.67	1.68	7.12	1.55	6.73
QuITE^{++}	**6.08**	**3.99**	4.99	3.62	**3.81**	**3.18**	**1.80**	**7.54**	**1.63**	**6.83**	1.48	**6.56**

- QuITE++ achieves the **best performance in 20/24 settings.**
- QuITE **effectively adapts MTS backbones to IMTS** via input-embedding replacement.

5. Analysis

- **Comparison of Different Embedding Methods**

Model		PatchTST				iTransformer				QuITE^{++}			
Method	Metric	Activity	USHCN	PhysioNet	MIMIC-III	Activity	USHCN	PhysioNet	MIMIC-III	Activity	USHCN	PhysioNet	MIMIC-III
Add	MSE	4.00	5.23	13.79	4.71	4.98	6.26	18.26	6.34	3.44	5.05	5.34	1.71
	MAE	4.03	3.17	6.54	14.74	4.84	3.80	8.01	19.73	3.76	3.04	3.86	7.31
Concat	MSE	3.90	5.21	12.97	4.52	5.77	6.10	18.27	6.48	3.35	4.99	5.43	1.75
	MAE	3.91	3.18	5.96	14.13	5.21	3.67	8.01	20.28	3.71	3.02	3.90	7.28
mTAND	MSE	3.74	5.21	13.38	4.39	3.50	5.23	13.11	4.34	3.34	5.01	**4.96**	1.71
	MAE	3.76	3.20	6.03	13.80	**3.65**	3.15	5.96	13.72	3.64	3.07	**3.60**	7.17
Mean Pooling	MSE	3.75	5.14	12.84	4.43	3.59	5.08	12.15	4.33	3.31	4.94	5.26	1.69
	MAE	3.77	3.19	5.95	13.84	3.72	3.11	5.70	13.61	3.64	3.01	3.89	7.23
QuITE	MSE	**3.69**	**5.06**	**12.32**	**4.36**	**3.31**	**4.95**	**5.21**	**1.69**	**3.18**	**4.82**	**4.96**	**1.64**
	MAE	**3.72**	**3.06**	**5.58**	**13.73**	**3.65**	**3.01**	**3.71**	**7.12**	**3.48**	**2.93**	**3.60**	**6.98**

- **Quality of Embedding Representation**

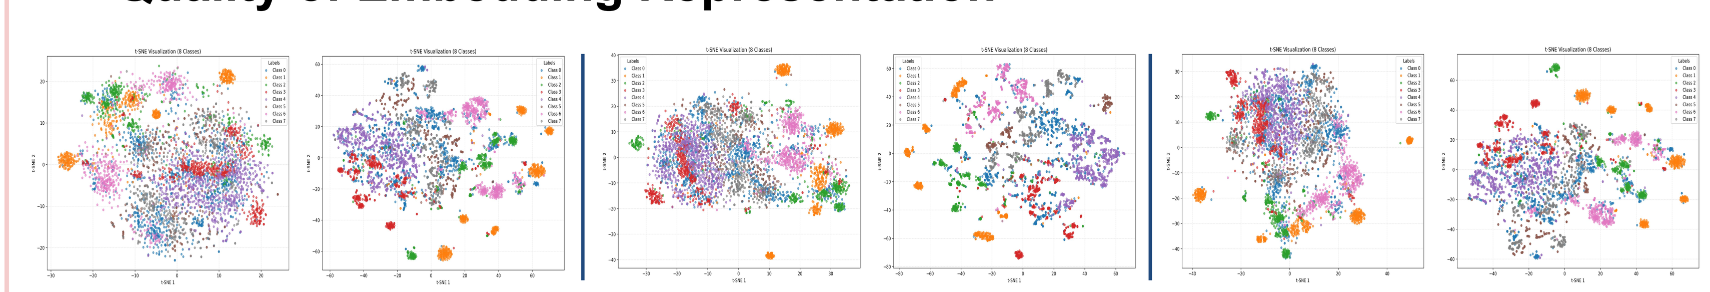

PatchTST · PatchTST w/ QuITE · iTransformer · iTransformer w/ QuITE · TimeXer · TimeXer w/ QuITE

- **Query initialization**

Dataset	Metric	Xavier	Uniform	Zero	Random
Human Activity (3000ms → 1000ms)	MSE	2.46	2.45	2.46	2.46
	MAE	3.00	2.99	3.01	2.92
USHCN (24m → 12m)	MSE	4.83	4.86	4.87	4.81
	MAE	2.99	2.95	2.98	2.93

- **Observation Sparsity**

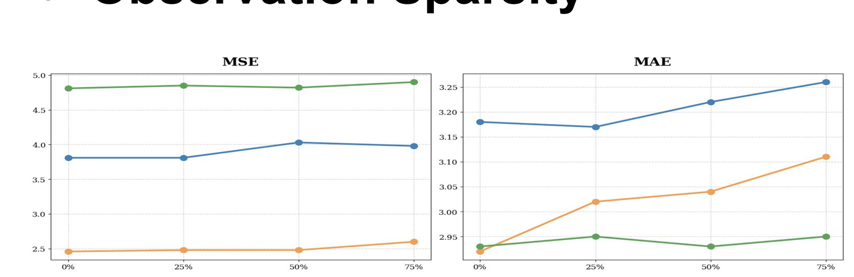

QuITE *is a simple and effective plug-and-play input-embedding module that enables strong MTS backbones to handle irregular multivariate time series.*

Linkedin · Paper · Code