NeurIPS 2025

TimePerceiver: An Encoder-Decoder Framework for Generalized Time-Series Forecasting

Jaebin Lee Hankook Lee

Sungkyunkwan University (SKKU)

Preliminaries

✓ Time series forecasting is a task to predict future horizons from past observations.
Note. f_θ is a function that best explains the data.

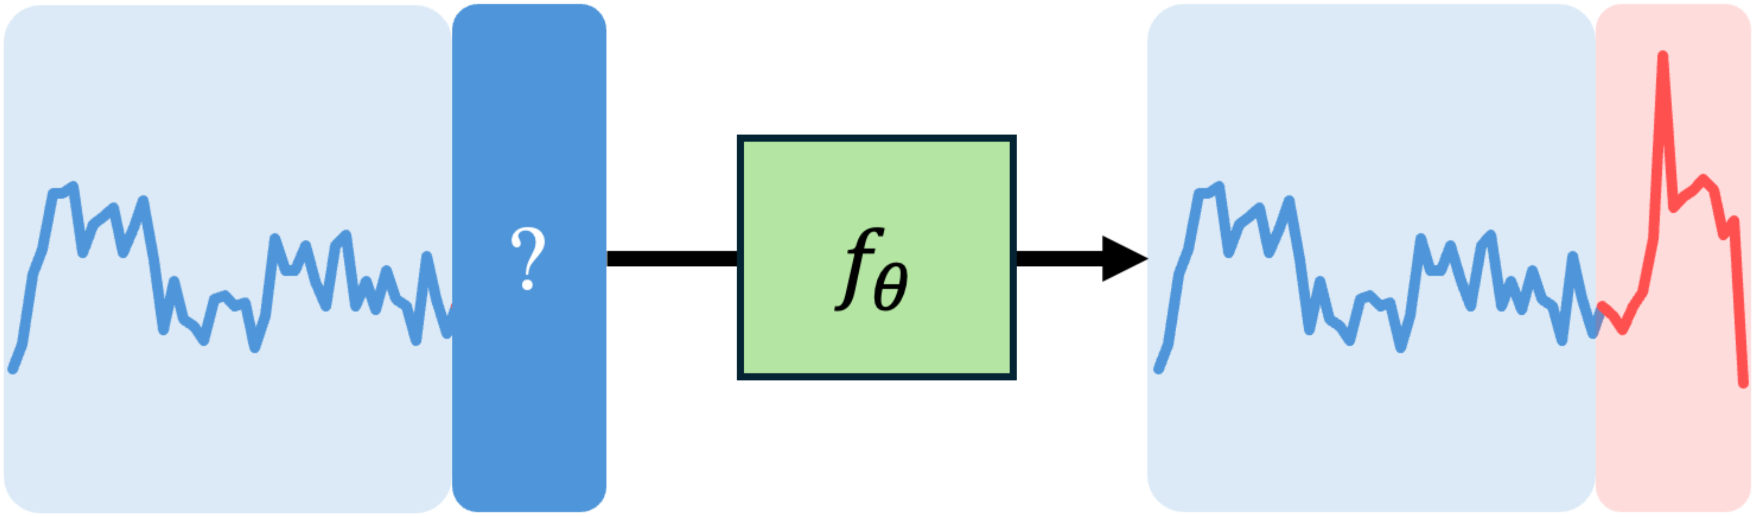

Preliminaries

✓ In general, the standard formulation is aligned with time-series forecasting objective.

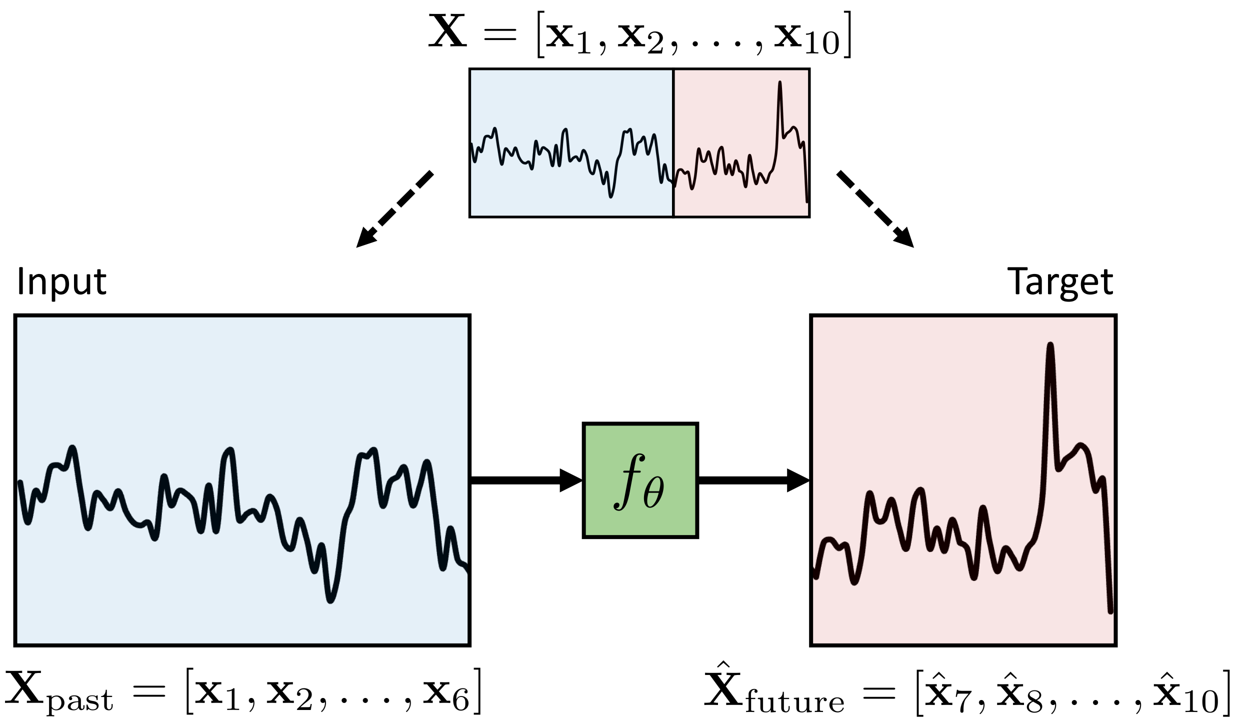

Motivation

✓ A unified approach that can handle extrapolation, interpolation, and imputation may help the model understand time series more deeply.

$$\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{10}]$$

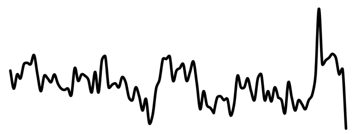

Motivation

✓ A unified approach that can handle extrapolation, interpolation, and imputation may help the model understand time series more deeply.

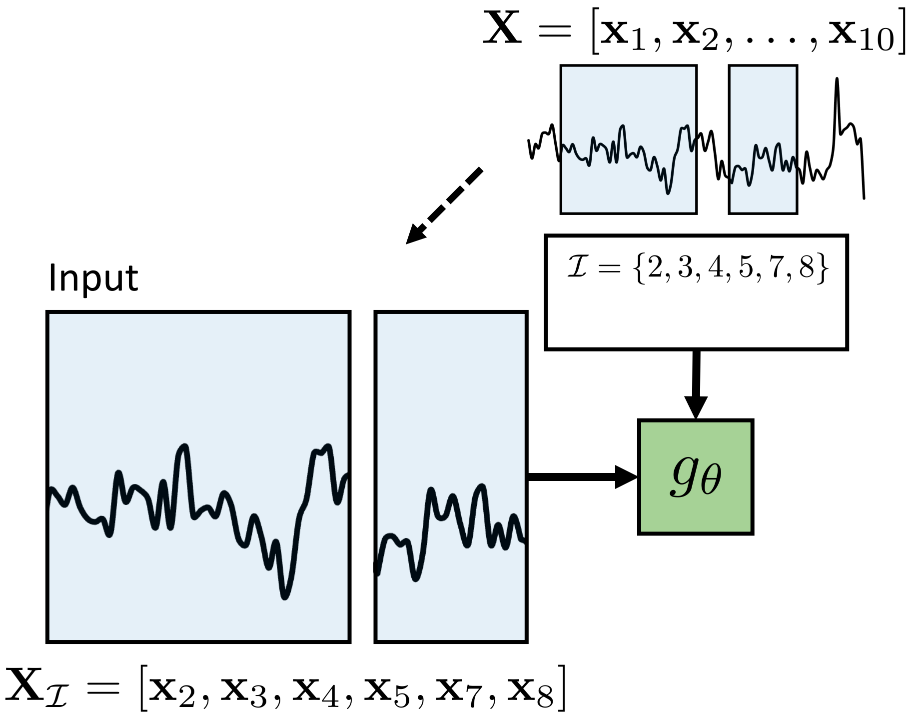

Motivation

✓ A unified approach that can handle extrapolation, interpolation, and imputation may help the model understand time series more deeply.

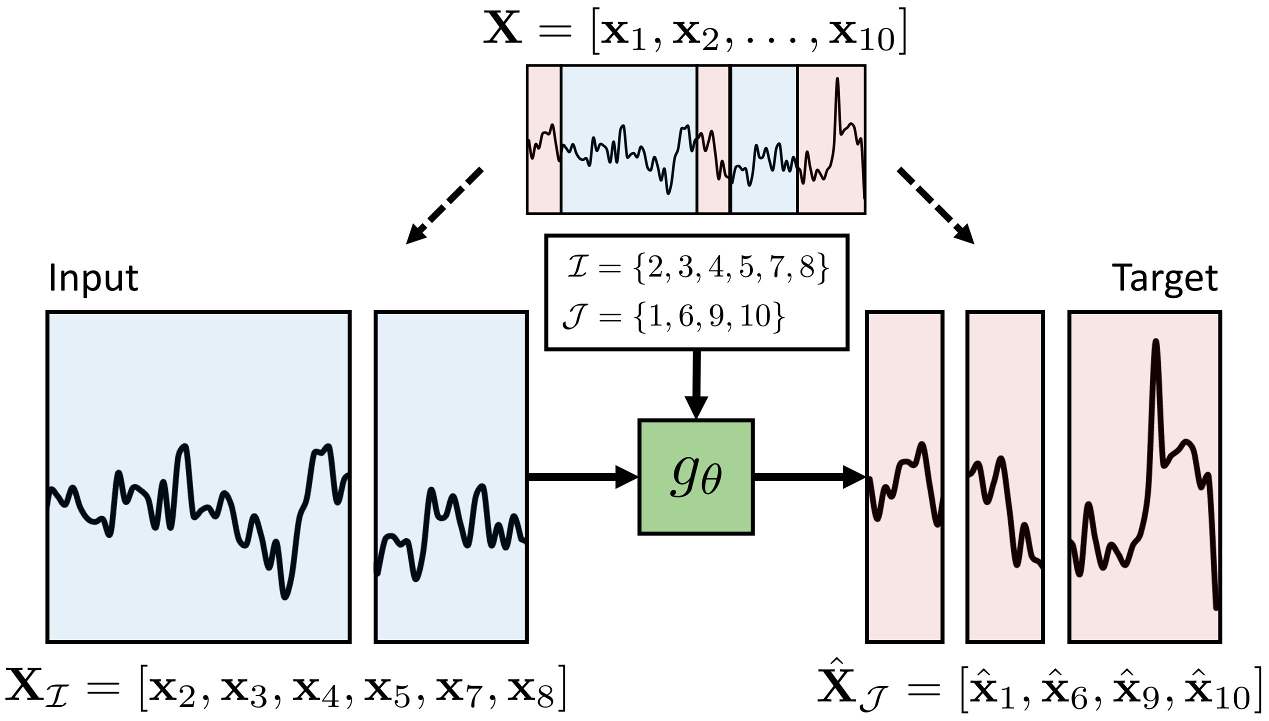

Motivation

✓ A unified approach that can handle extrapolation, interpolation, and imputation may help the model understand time series more deeply.

Note. A time-series forecasting objective is a special case of the generalized formulation.

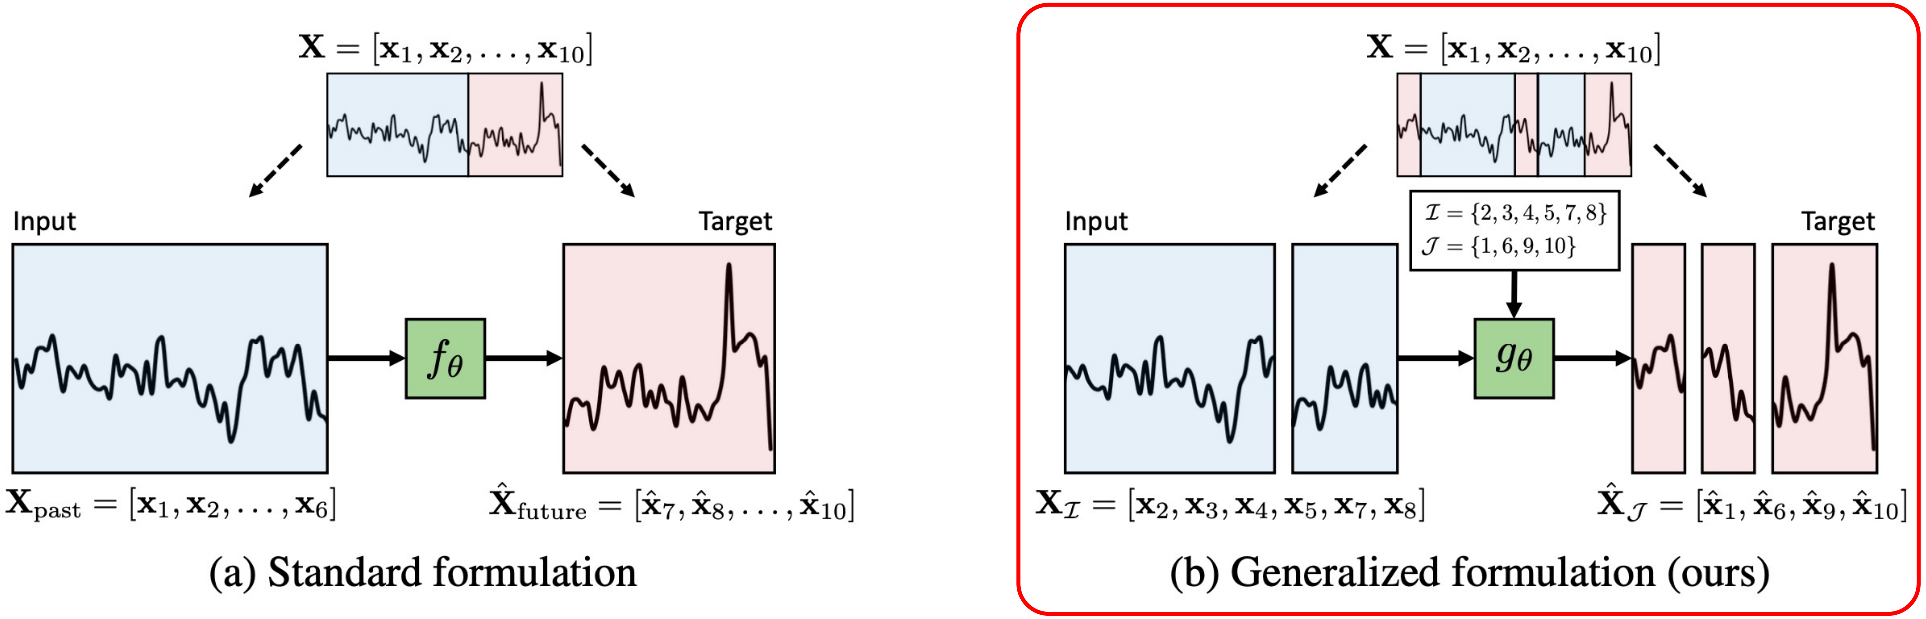

(a) Standard formulation

(b) Generalized formulation (ours)

Architecture

✓ The model need to handle arbitrary segments.

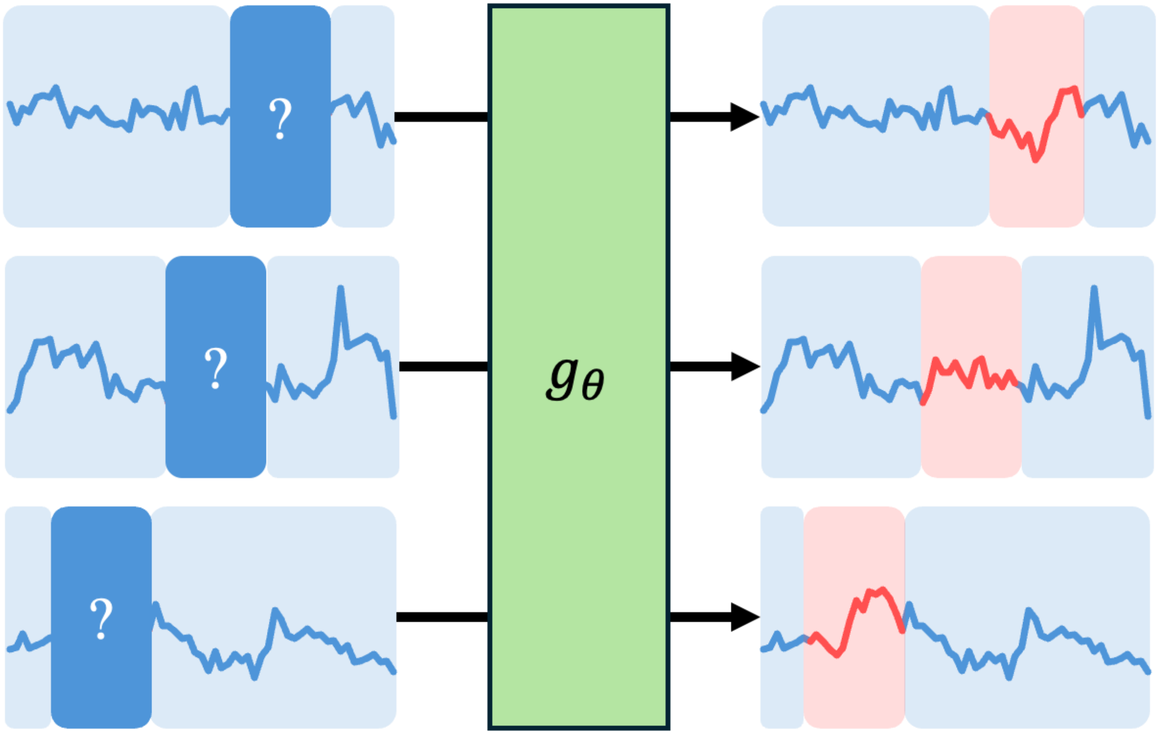

Generalized Formulation (Ours)

Architecture

✓ We propose TimePerceiver, an attention-based encoder-decoder architecture that enables flexible forecasting at arbitrary target positions.

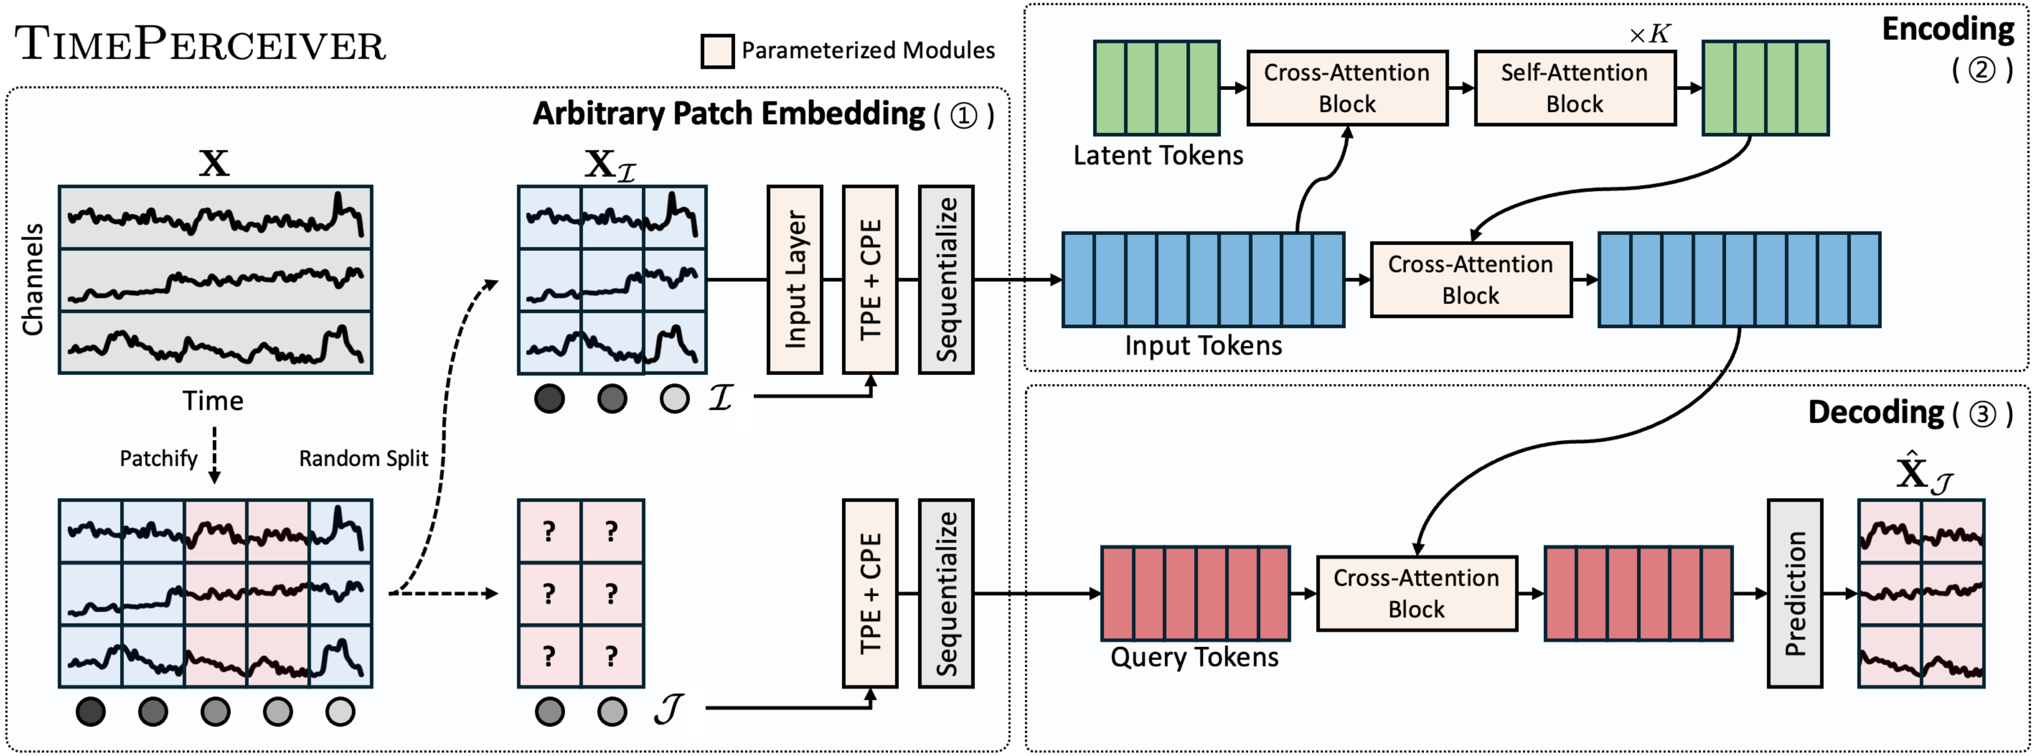

Architecture

✓ Split time series into patches and randomly assign them as input or target patches.

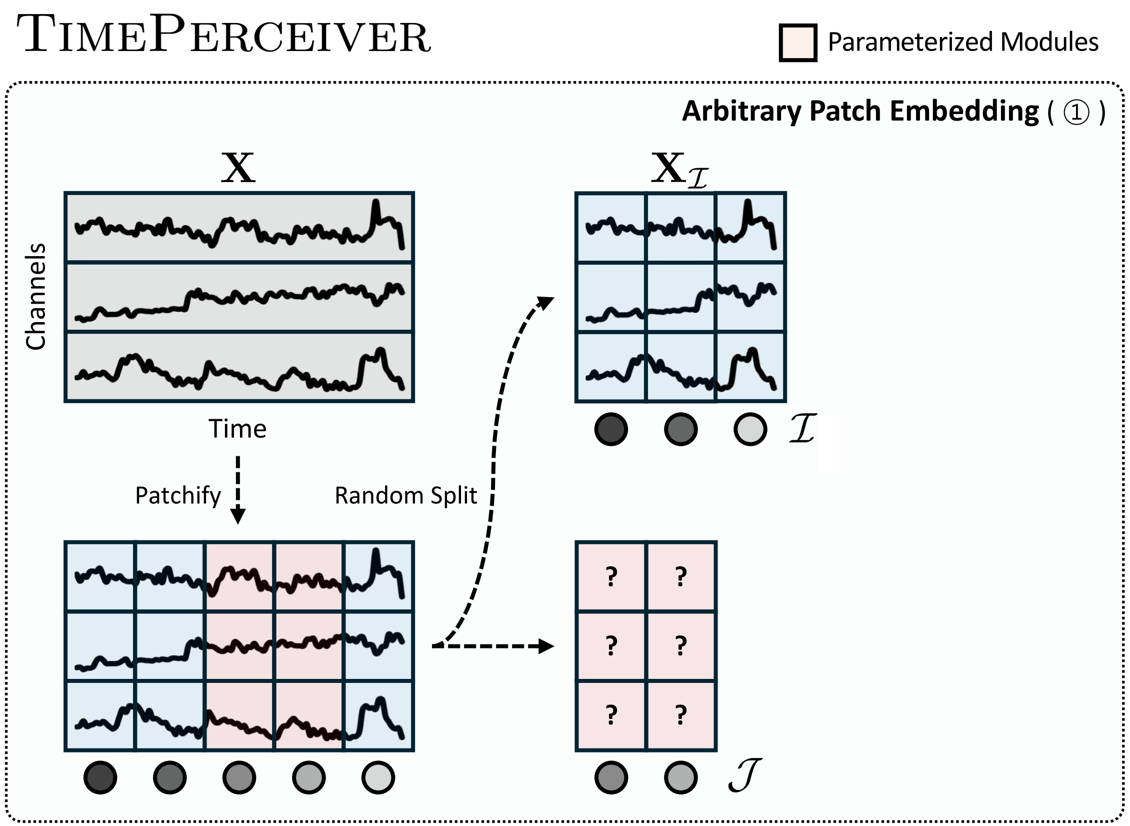

Architecture

✓ Embed the input patches and add positional embeddings to each one.

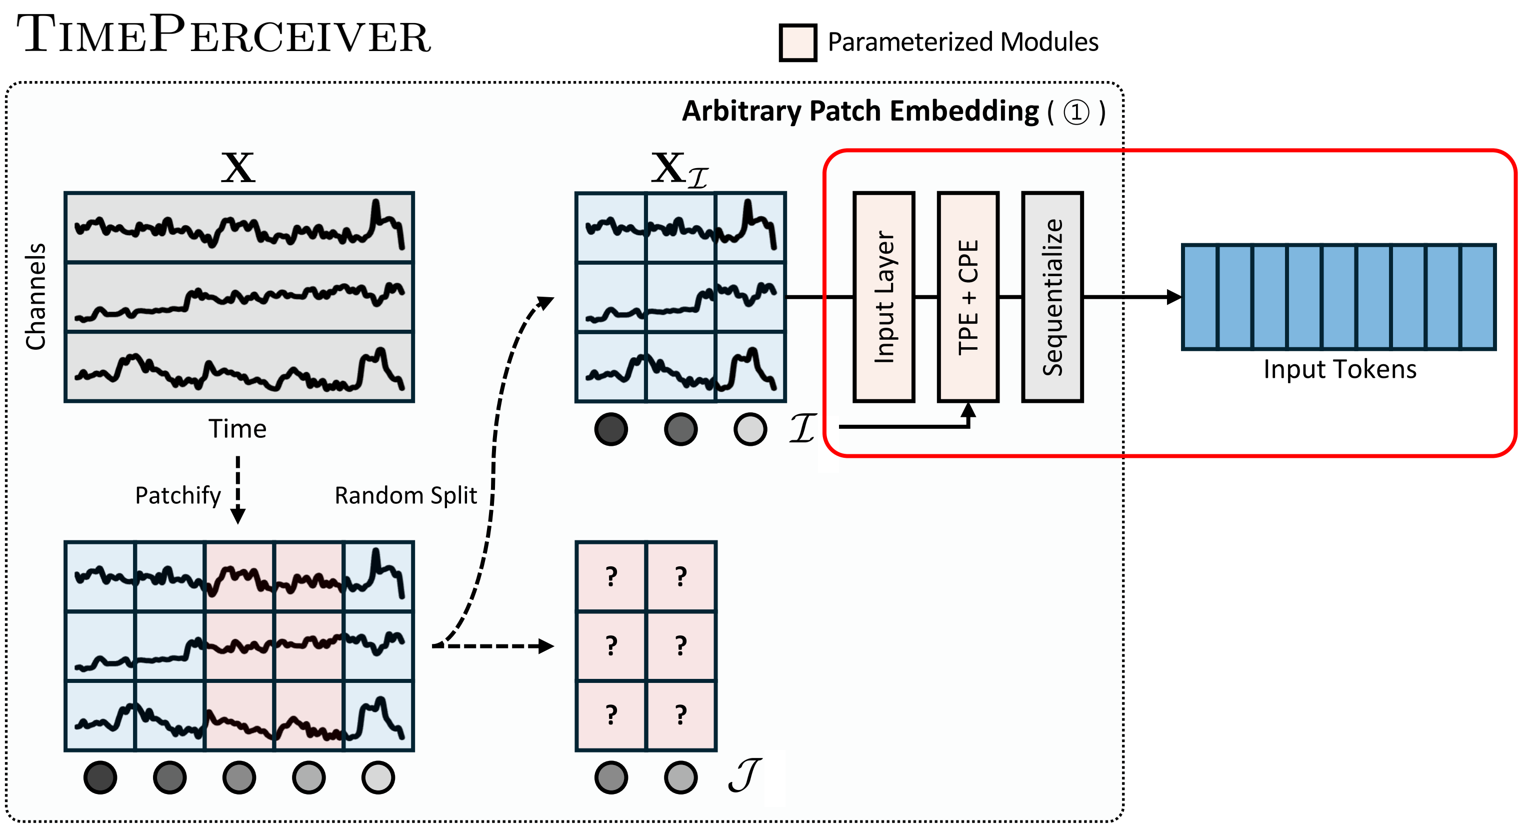

Architecture

✓ Define a **learnable latent bottleneck** to capture temporal and cross-channel dependencies.

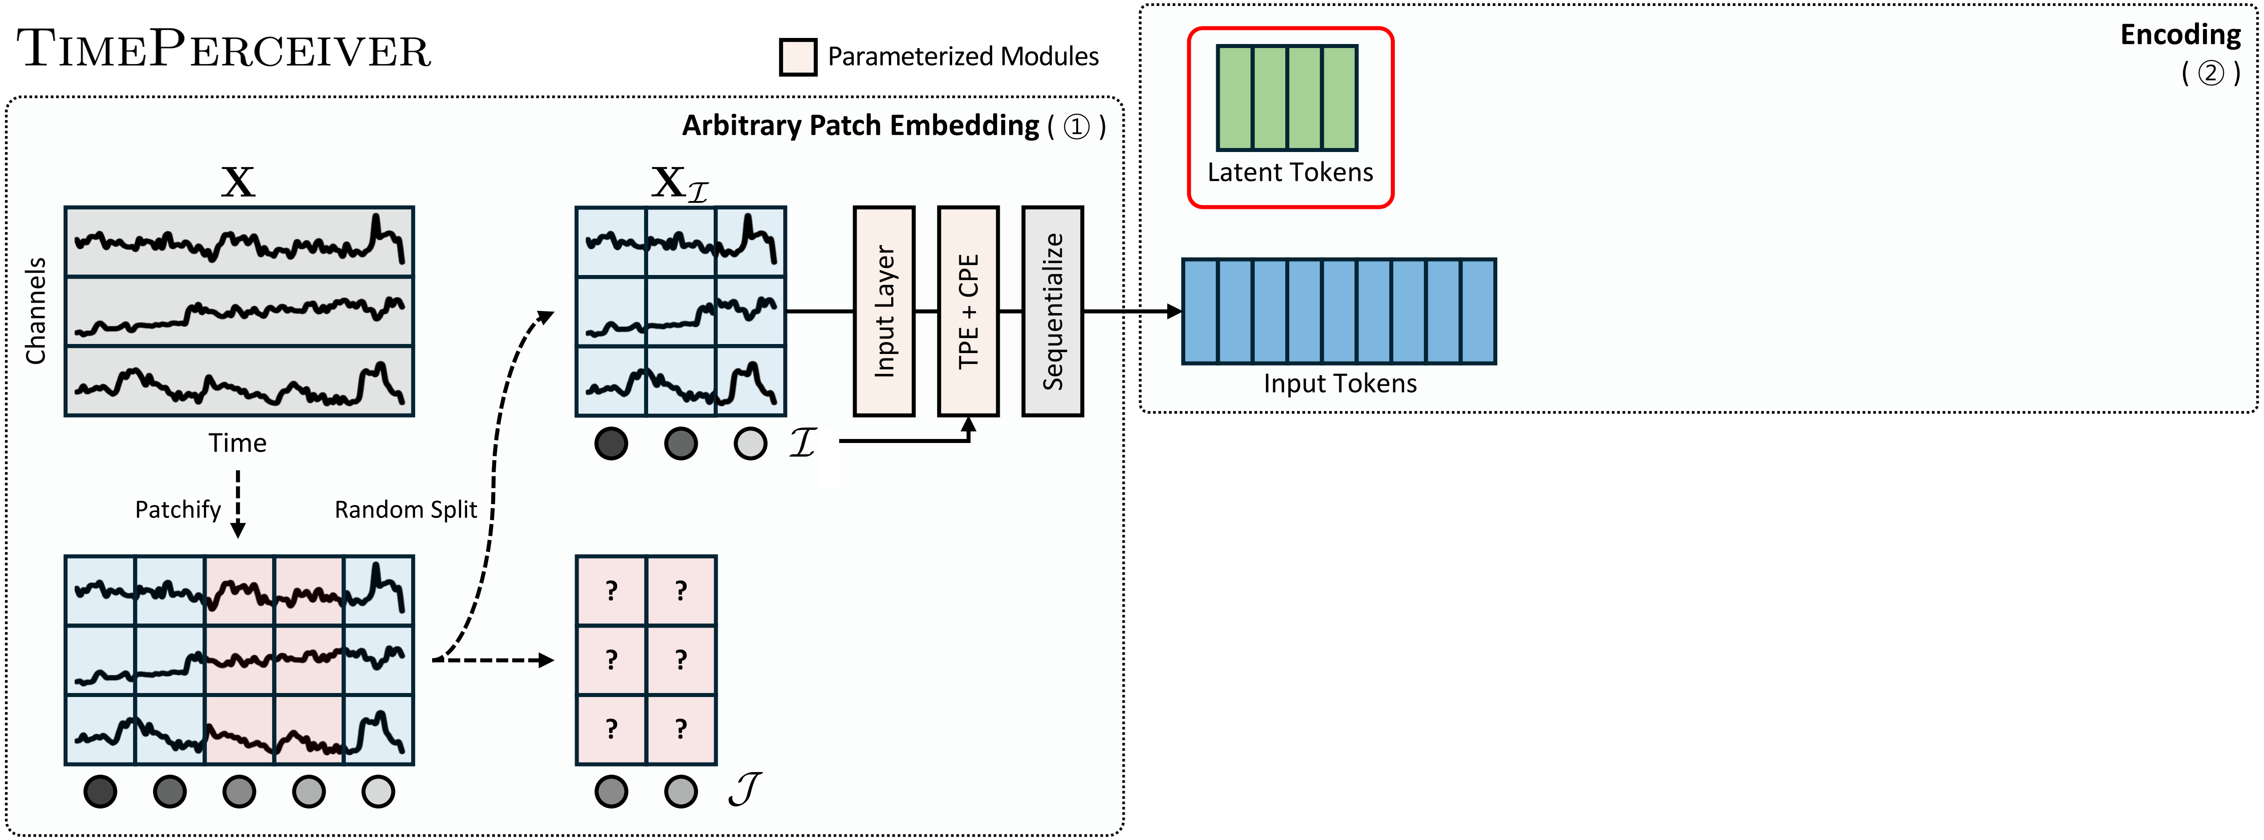

Architecture

✓ **Aggregate** information from the input tokens into the latent tokens via **cross-attention.**

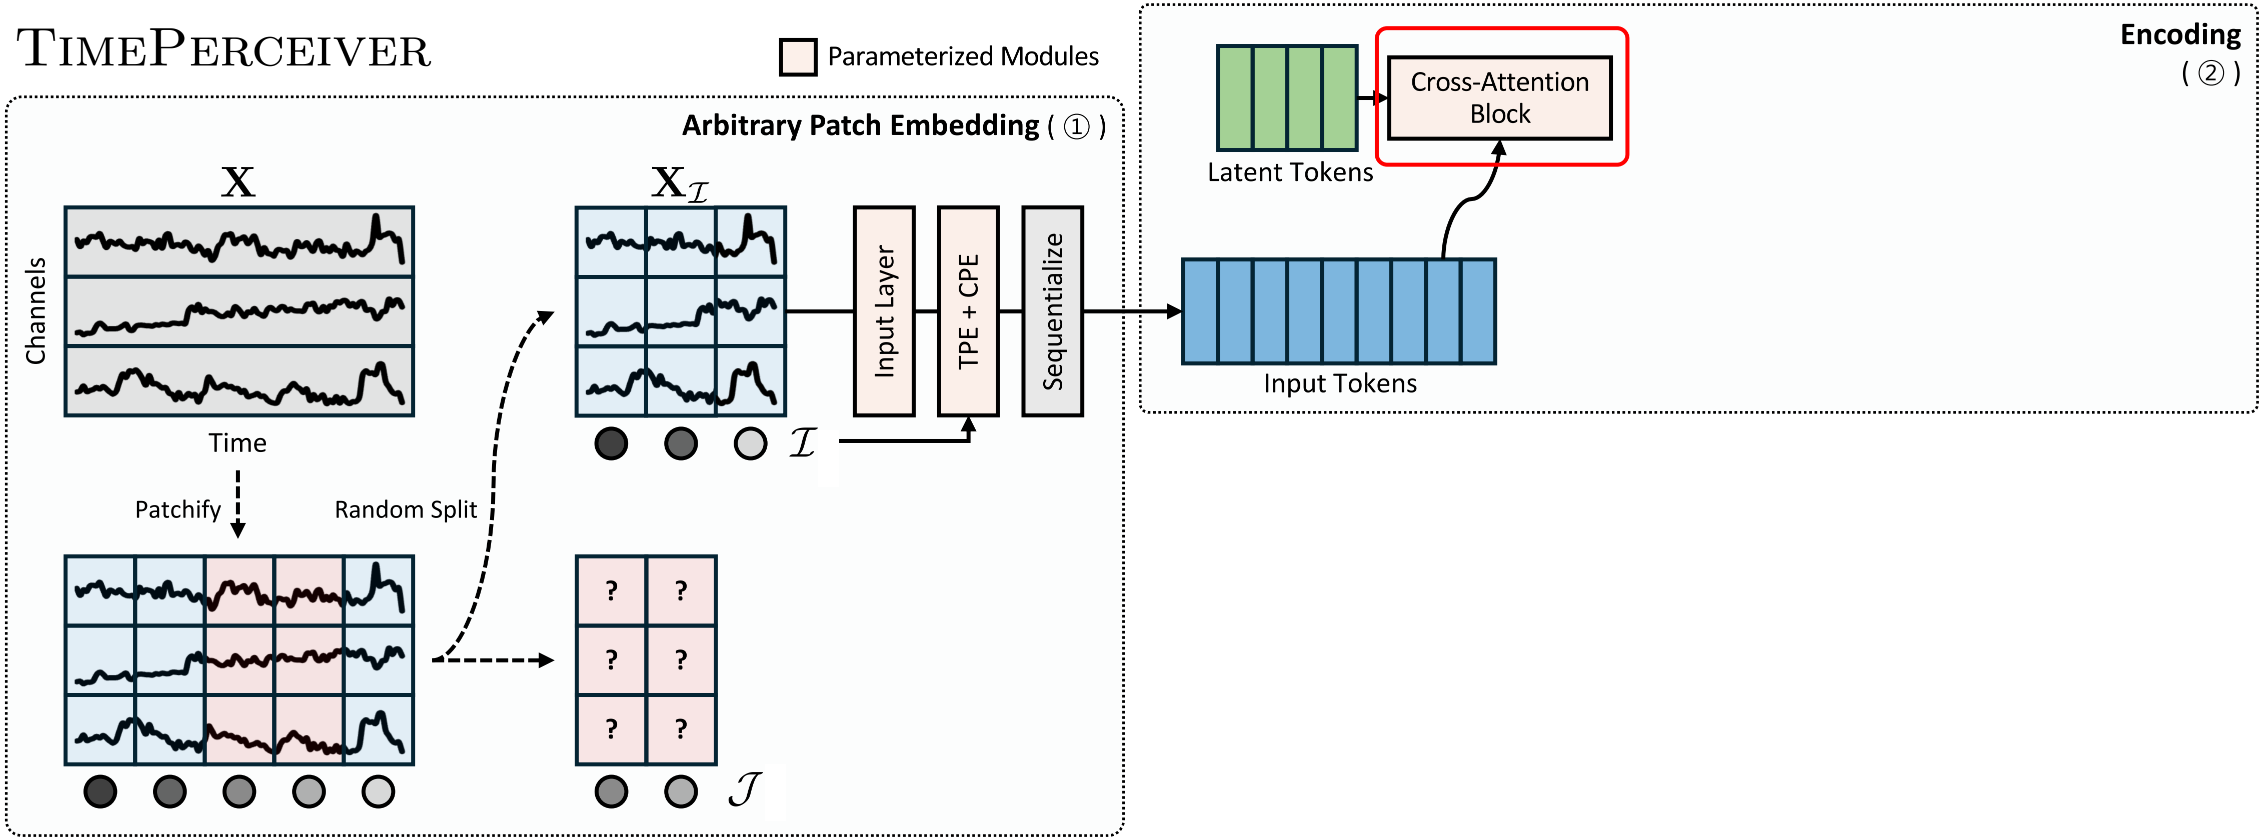

Architecture

✓ **Refine** the aggregated information via **self-attention** within the latent space.

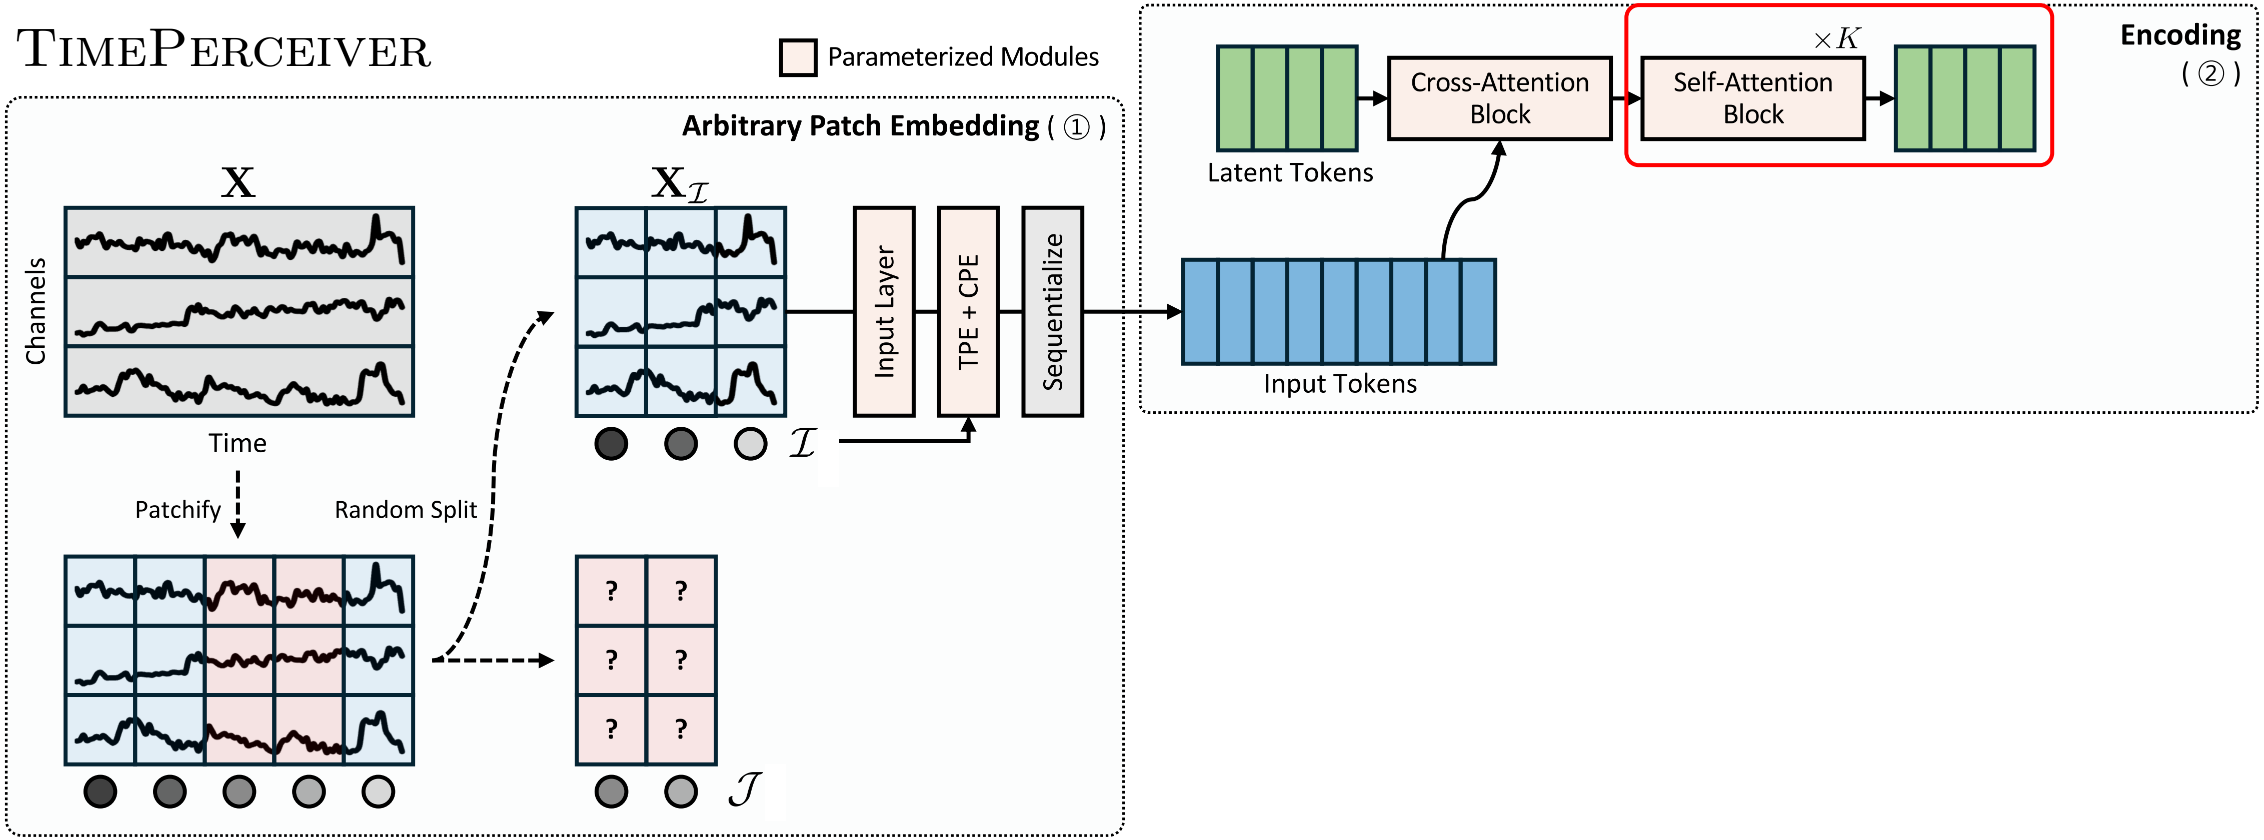

Architecture

✓ **Write** the refined information back to the input space via **cross-attention.**

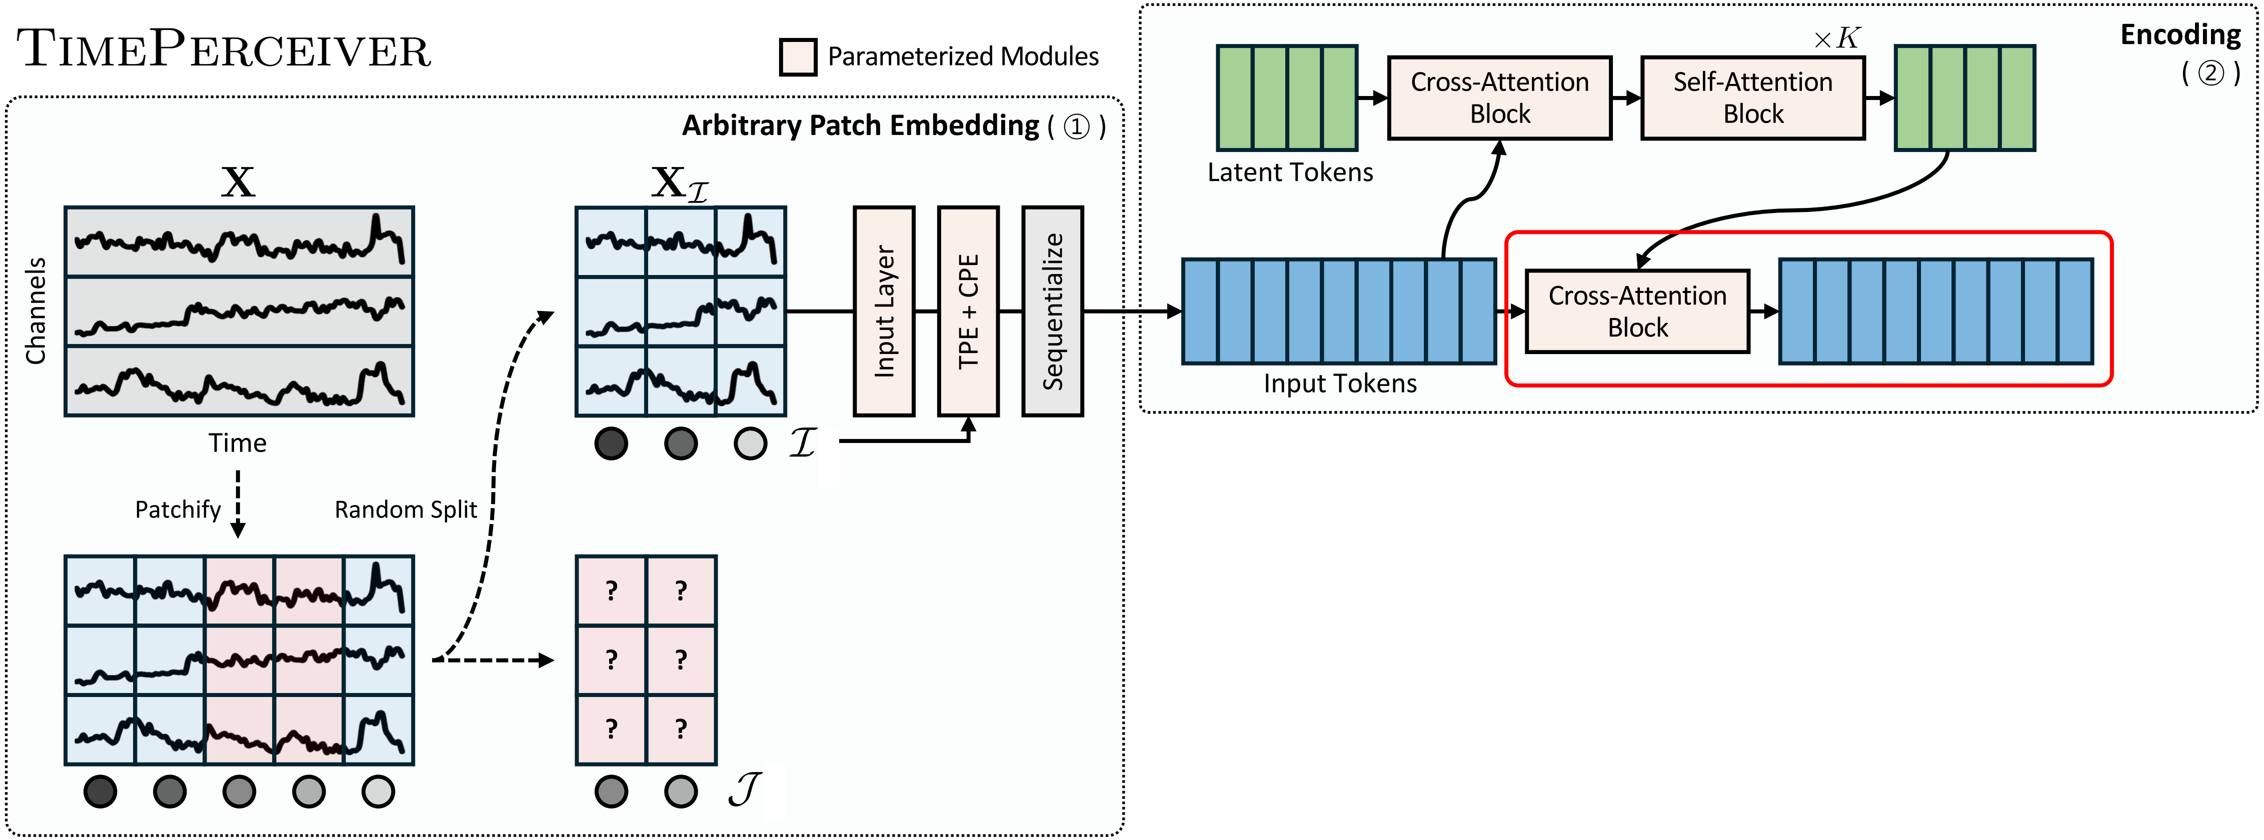

Architecture

✓ Define queries with **positional embeddings to perceive** which target patches to predict.

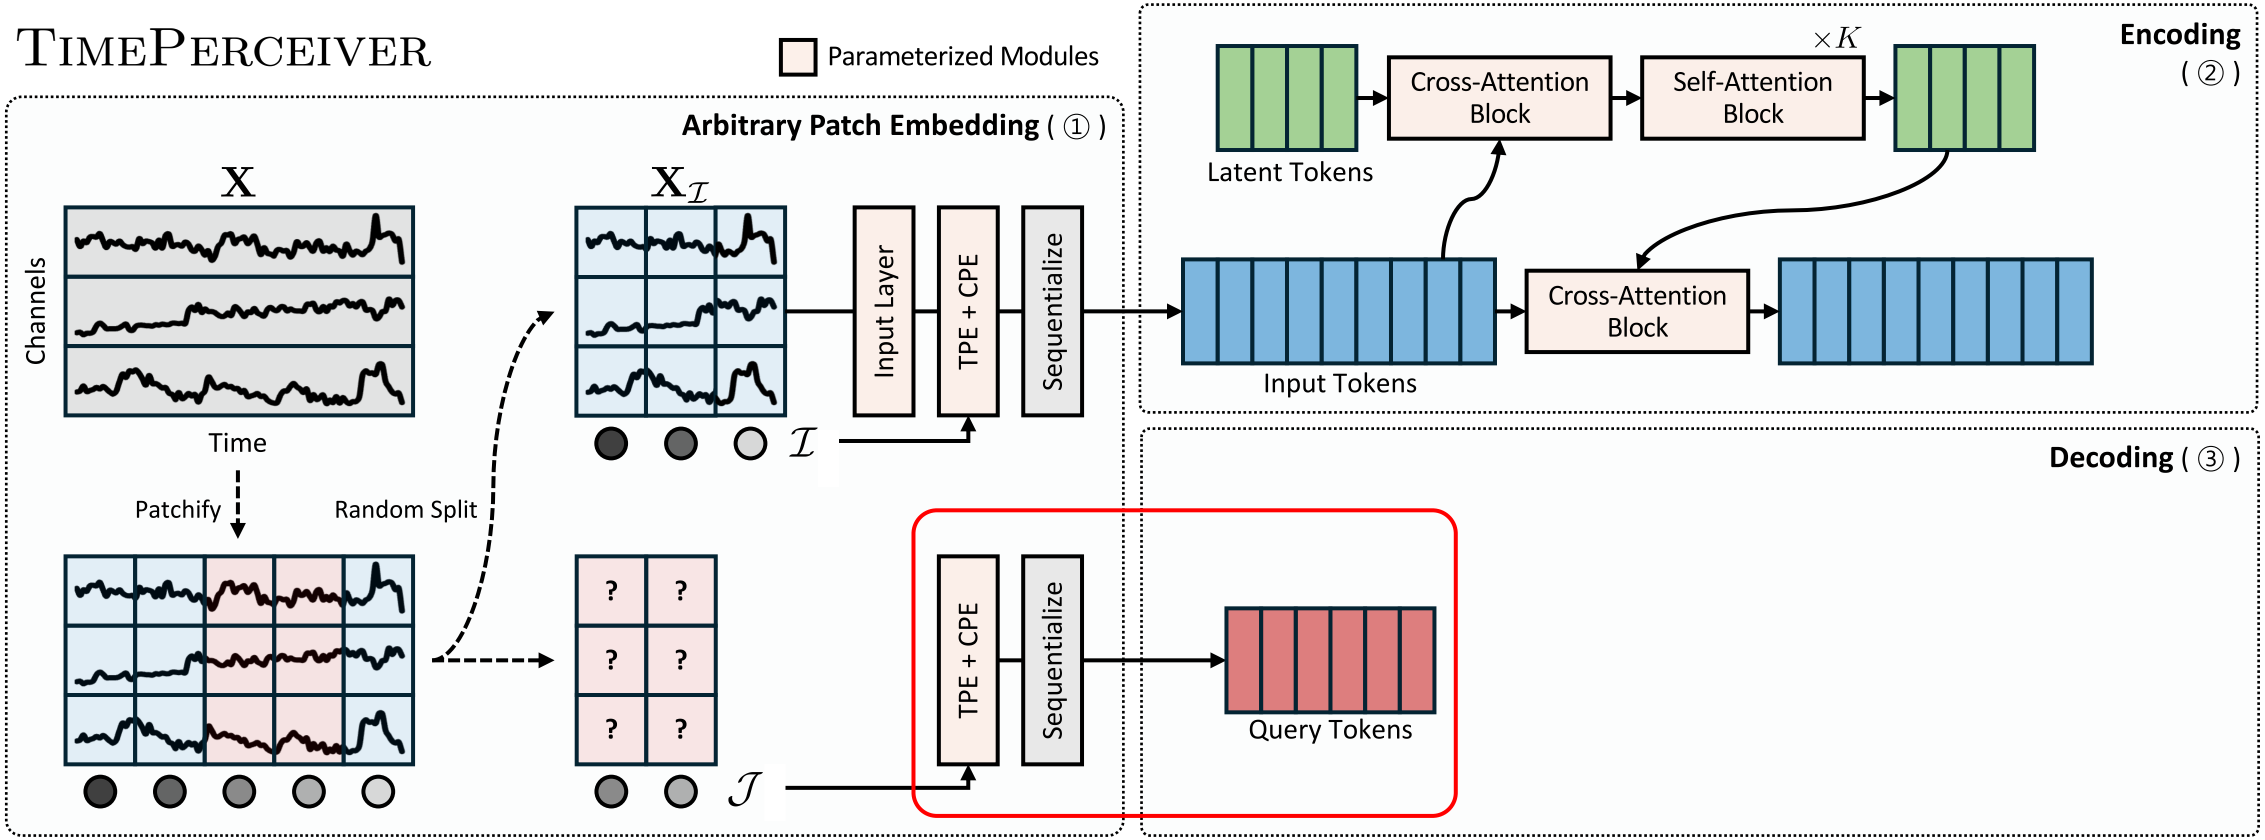

Architecture

✓ Use these queries **to aggregate** information from the input through **cross-attention.**

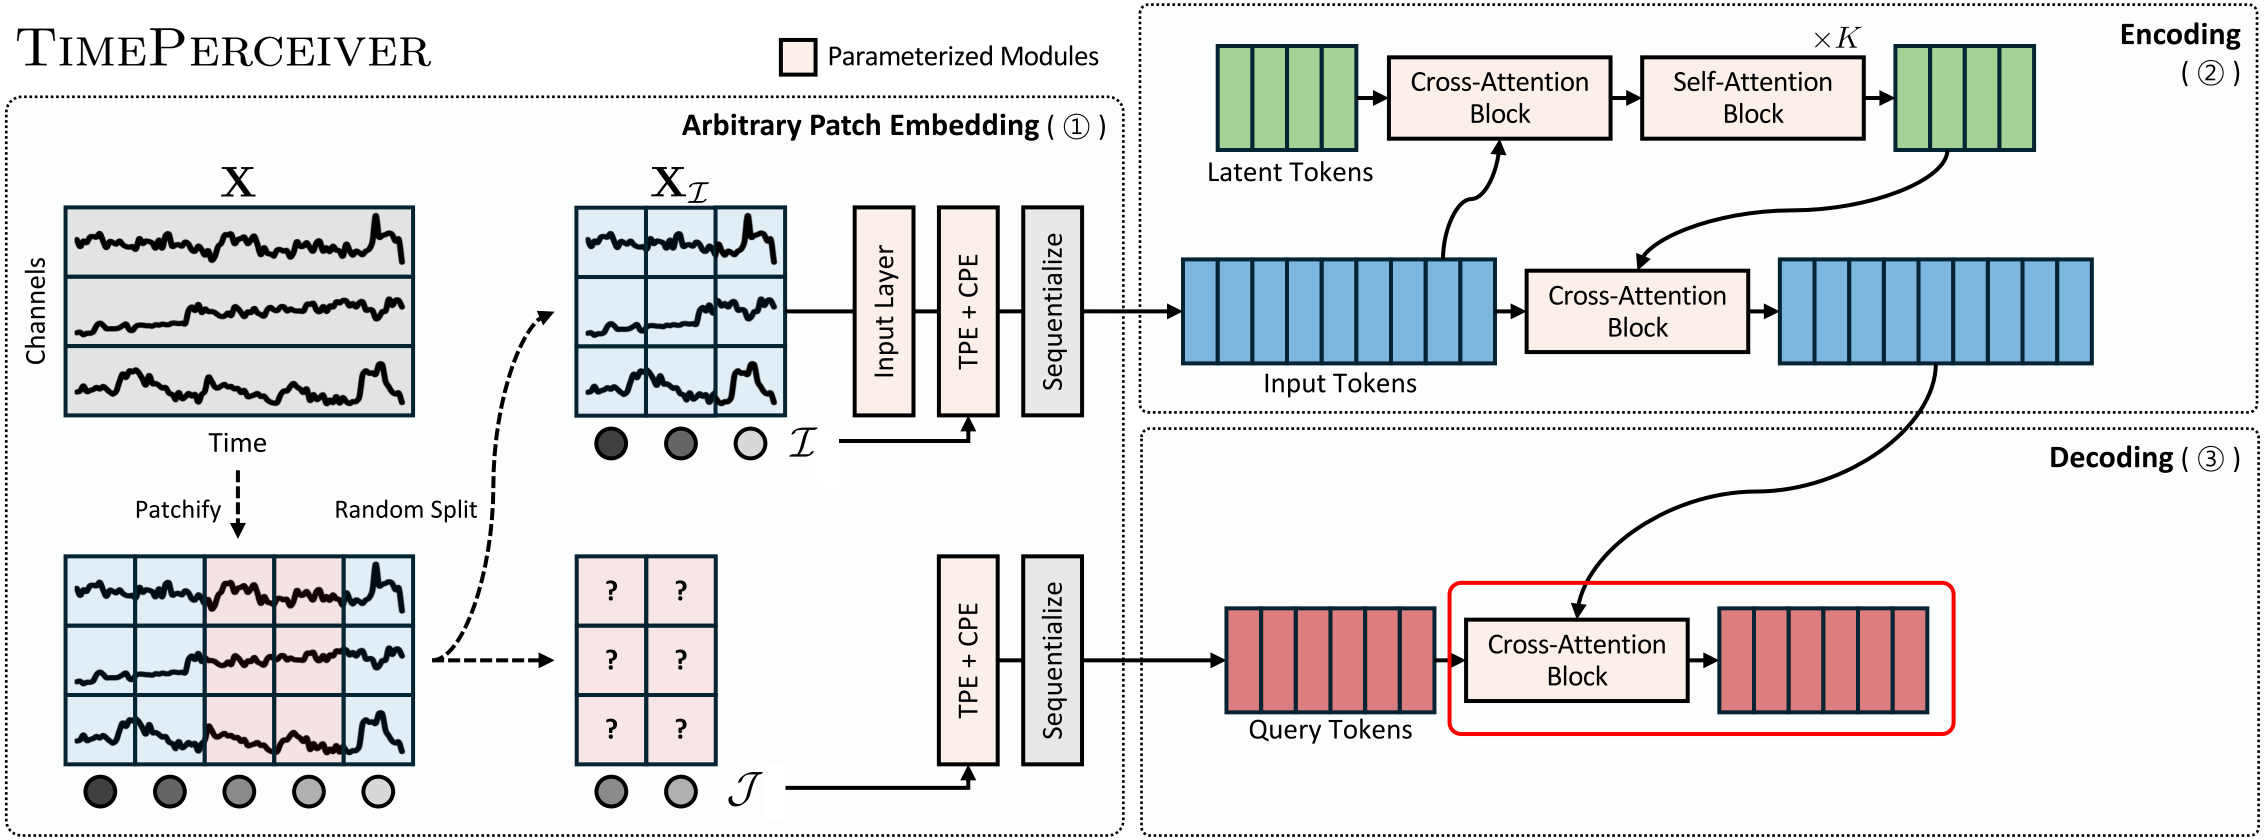

Architecture

✓ **Predict** the output through a **patch-to-patch linear layer.**

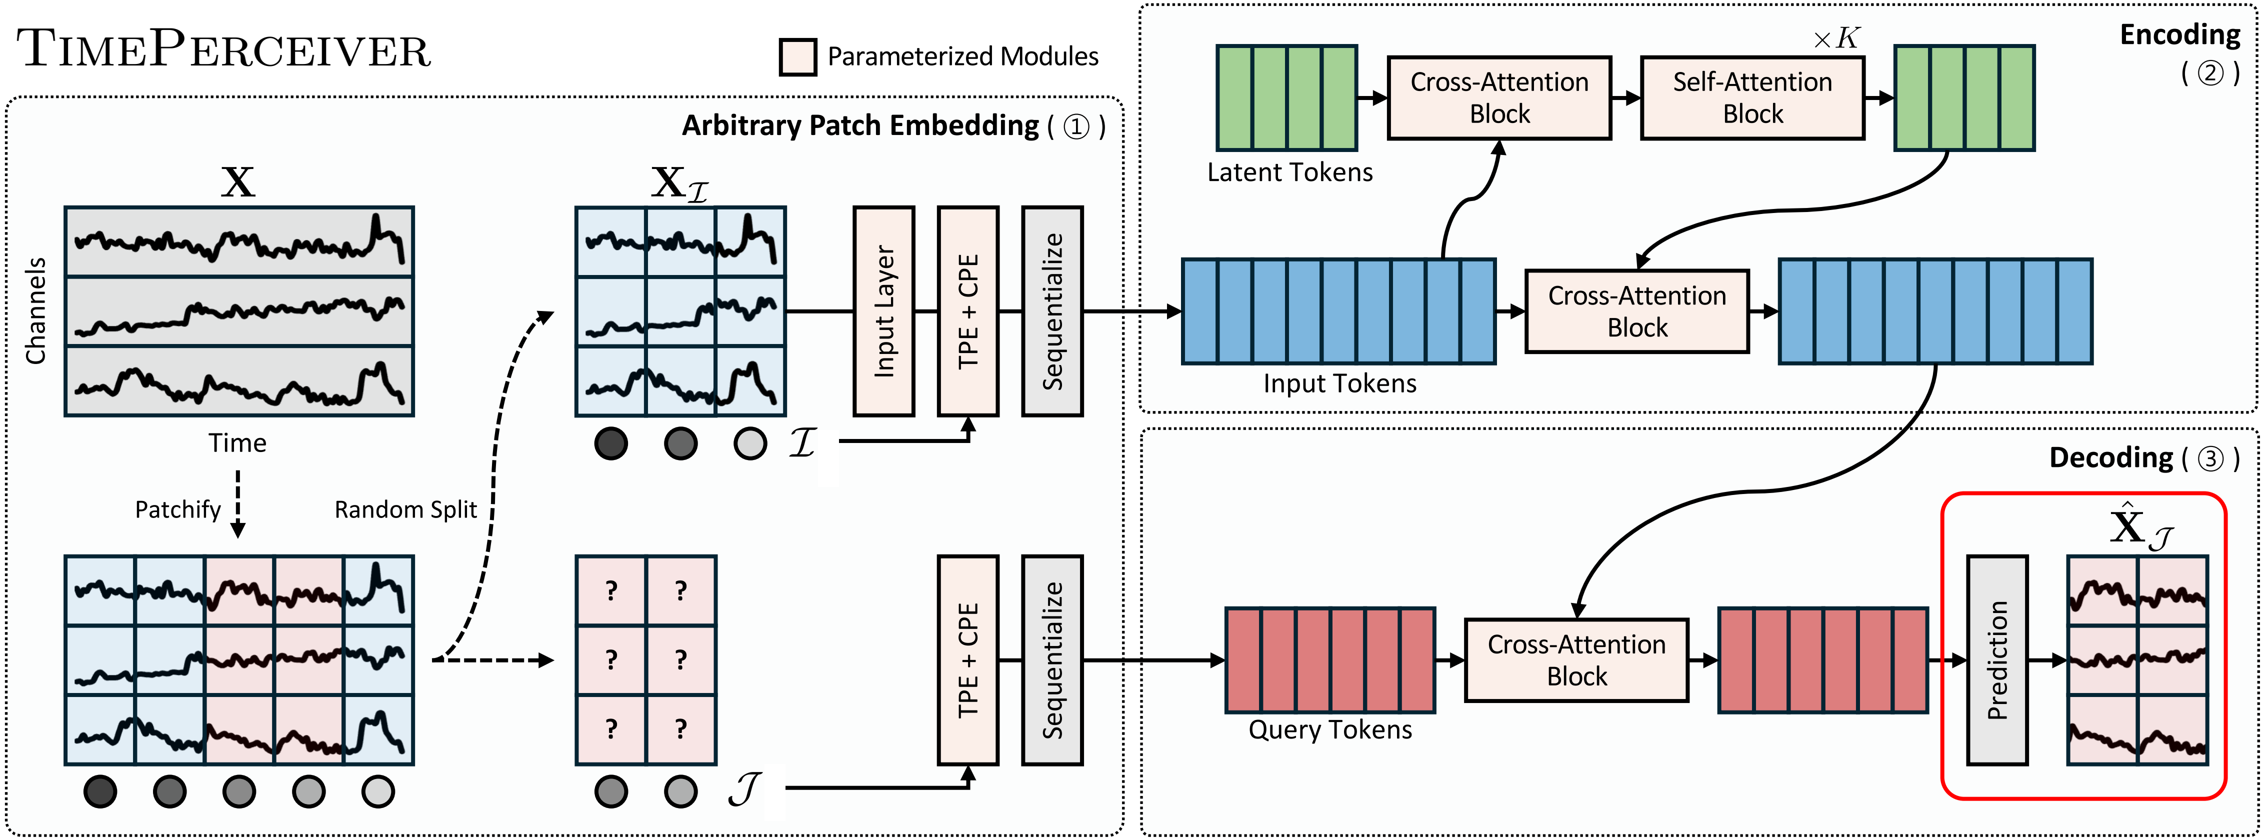

Architecture

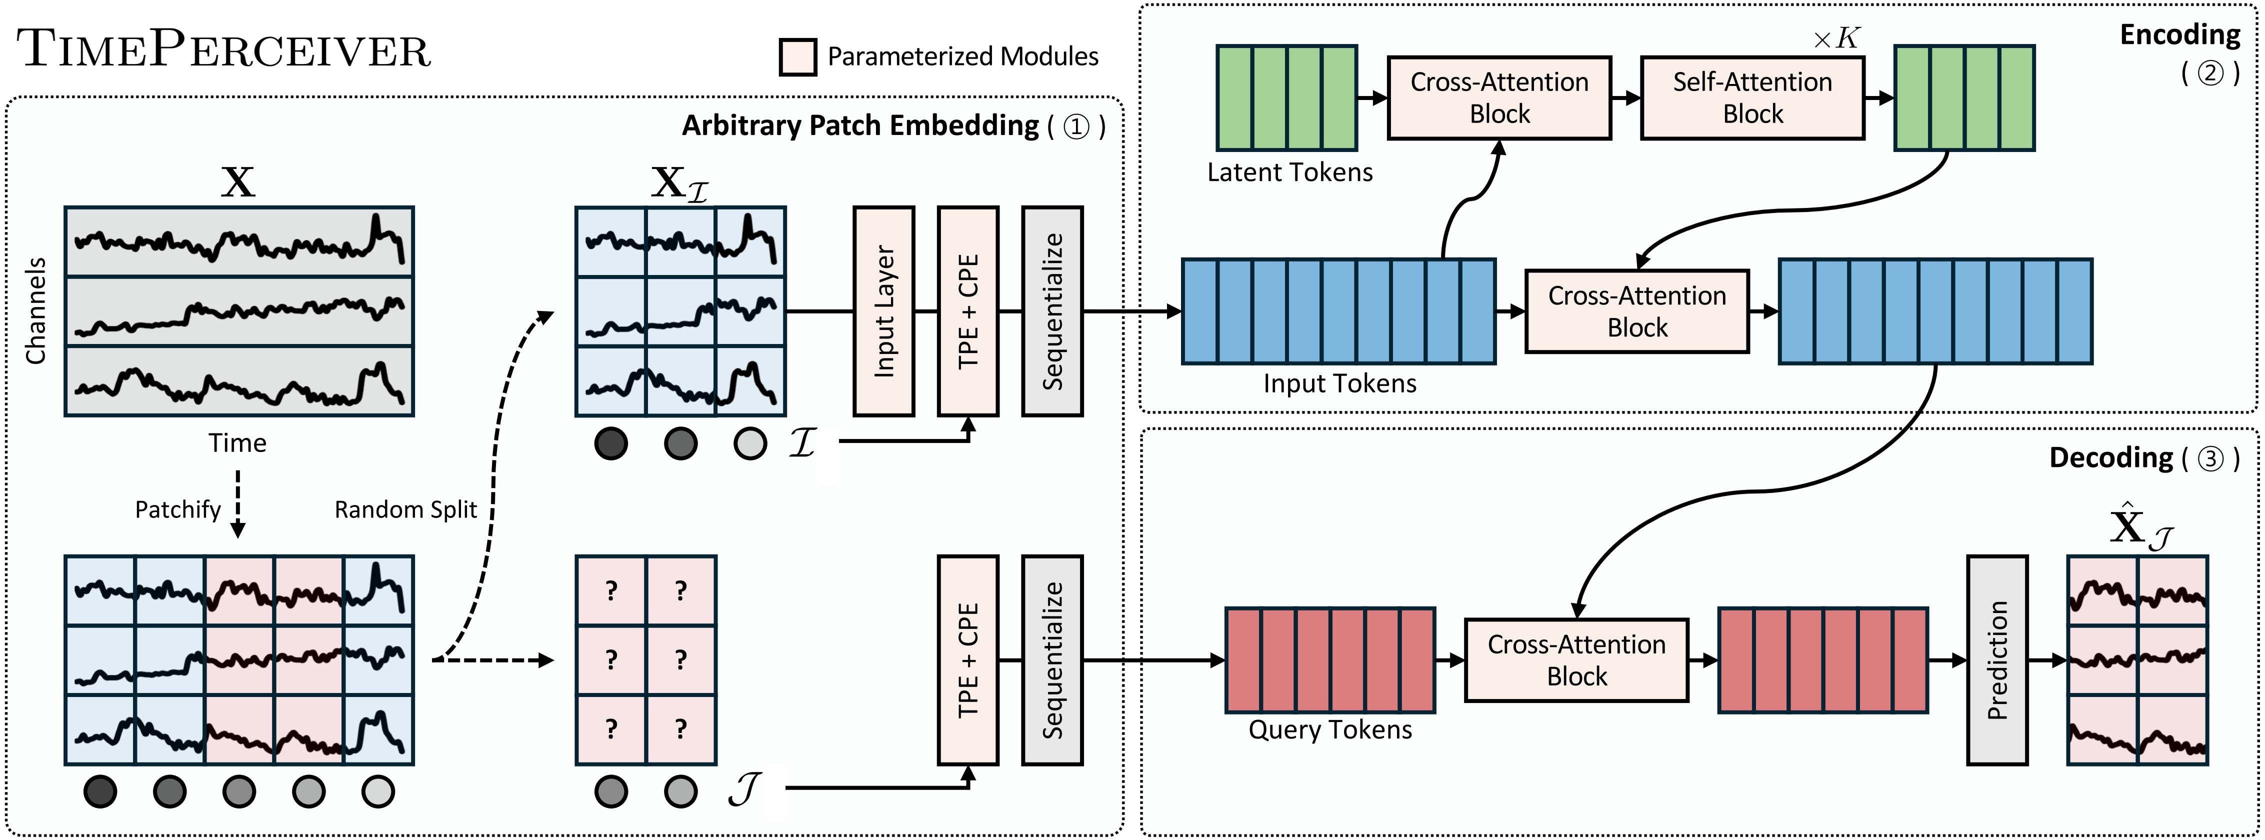

Experiment

✓ In the main results, our model achieves the best performance in 55 out of 80 settings and achieves the second best performance in 17 settings.

Models		TimePerceiver		DeformableTST		CARD		CATS		PatchTST		iTransformer		S-Mamba		DLinear		TimesNet		TiDE	
Metric		MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
Weather	96	**0.147**	**0.199**	0.155	0.203	0.156	0.208	0.150	0.201	0.158	0.207	0.166	0.214	0.162	0.212	0.178	0.238	0.169	0.221	0.183	0.238
	192	**0.193**	**0.242**	0.200	0.244	0.214	0.258	0.196	0.244	0.204	0.248	0.210	0.253	0.209	0.254	0.221	0.278	0.237	0.278	0.225	0.274
	336	**0.246**	**0.283**	0.252	0.284	0.263	0.295	0.249	0.284	0.255	0.286	0.263	0.293	0.266	0.296	0.268	0.319	0.292	0.319	0.270	0.309
	720	0.323	0.338	0.325	**0.335**	0.354	0.351	**0.322**	**0.335**	0.328	0.336	0.335	0.343	0.334	0.342	0.328	0.366	0.385	0.377	0.335	0.357
	Avg	**0.227**	**0.265**	0.233	0.266	0.247	0.278	0.229	0.266	0.236	0.269	0.244	0.276	0.243	0.276	0.249	0.300	0.271	0.299	0.254	0.295
Solar	96	**0.173**	**0.221**	0.177	0.237	0.194	0.257	0.184	0.228	0.207	0.283	0.188	0.242	0.191	0.242	0.232	0.322	0.229	0.294	0.246	0.324
	192	**0.195**	**0.238**	0.199	0.254	0.229	0.279	0.210	0.243	0.231	0.304	0.212	0.263	0.219	0.268	0.258	0.341	0.266	0.318	0.271	0.351
	336	**0.206**	**0.248**	0.208	0.263	0.239	0.289	0.218	0.253	0.249	0.312	0.226	0.275	0.229	0.279	0.280	0.355	0.275	0.337	0.291	0.367
	720	0.217	**0.259**	**0.214**	0.266	0.251	0.301	0.233	0.261	0.250	0.314	0.232	0.282	0.236	0.284	0.291	0.360	0.276	0.329	0.300	0.364
	Avg	**0.198**	**0.241**	0.199	0.255	0.228	0.282	0.211	0.247	0.234	0.303	0.214	0.266	0.219	0.268	0.266	0.345	0.261	0.319	0.277	0.351
Electricity	96	**0.128**	**0.223**	0.141	0.240	0.140	0.236	0.134	0.228	0.148	0.246	0.143	0.240	0.136	0.232	0.164	0.258	0.188	0.294	0.180	0.282
	192	**0.146**	**0.240**	0.155	0.254	0.161	0.257	0.153	0.243	0.163	0.256	0.160	0.256	0.157	0.253	0.173	0.268	0.197	0.303	0.186	0.287
	336	**0.165**	**0.259**	0.172	0.271	0.181	0.279	0.171	0.262	0.180	0.276	0.177	0.274	0.173	0.270	0.188	0.285	0.229	0.327	0.202	0.301
	720	0.204	0.293	0.208	0.303	0.212	0.304	0.204	**0.292**	0.218	0.306	0.220	0.313	**0.198**	0.293	0.223	0.318	0.239	0.337	0.236	0.328
	Avg	**0.161**	**0.254**	0.169	0.267	0.174	0.269	0.166	0.257	0.177	0.271	0.175	0.271	0.166	0.262	0.187	0.282	0.213	0.315	0.201	0.300
Traffic	96	0.381	0.263	0.378	0.265	0.399	0.279	0.379	0.255	0.407	0.280	0.389	0.286	**0.357**	**0.254**	0.493	0.329	0.571	0.310	0.582	0.400
	192	0.394	0.265	0.401	0.272	0.415	0.283	0.400	0.266	0.417	0.283	0.410	0.298	**0.371**	**0.262**	0.482	0.323	0.584	0.316	0.569	0.397
	336	0.409	0.273	0.413	0.282	0.429	0.293	0.413	0.271	0.430	0.289	0.428	0.307	**0.382**	**0.270**	0.493	0.328	0.597	0.317	0.581	0.404
	720	0.445	**0.293**	0.448	0.300	0.461	0.308	0.448	0.294	0.465	0.307	0.468	0.331	**0.438**	0.295	0.530	0.350	0.619	0.330	0.583	0.405
	Avg	0.407	0.273	0.410	0.280	0.426	0.291	0.411	0.272	0.430	0.290	0.424	0.306	**0.387**	**0.270**	0.499	0.333	0.593	0.318	0.579	0.401
ETTh1	96	**0.364**	**0.392**	0.370	0.399	0.384	0.406	0.375	0.402	0.390	0.410	0.400	0.420	0.392	0.416	0.377	0.400	0.436	0.438	0.459	0.459
	192	**0.402**	**0.417**	0.415	0.425	0.421	0.426	0.419	0.432	0.429	0.433	0.441	0.446	0.437	0.444	0.418	0.426	0.481	0.467	0.497	0.481
	336	0.425	0.431	**0.411**	**0.426**	0.451	0.445	0.443	0.445	0.458	0.452	0.472	0.466	0.470	0.467	0.481	0.451	0.561	0.503	0.525	0.499
	720	**0.450**	**0.461**	0.456	0.469	0.463	0.473	0.476	0.470	0.473	0.481	0.533	0.520	0.514	0.508	0.495	0.508	0.615	0.533	0.541	0.529
	Avg	**0.410**	**0.426**	0.413	0.430	0.430	0.438	0.428	0.438	0.438	0.444	0.461	0.463	0.445	0.459	0.436	0.447	0.524	0.485	0.505	0.492
ETTh2	96	0.278	0.336	**0.274**	**0.334**	0.283	0.339	0.282	0.342	0.284	0.341	0.303	0.346	0.298	0.354	0.321	0.380	0.373	0.405	0.378	0.433
	192	0.347	0.388	**0.334**	**0.374**	0.355	0.384	0.348	0.385	0.354	0.384	0.396	0.415	0.371	0.400	0.436	0.451	0.416	0.438	0.468	0.485
	336	0.354	0.407	**0.328**	**0.377**	0.382	0.409	0.372	0.405	0.378	0.408	0.427	0.438	0.403	0.425	0.533	0.508	0.469	0.469	0.524	0.522
	720	**0.395**	0.437	0.406	0.437	0.401	**0.431**	0.418	0.448	0.409	0.439	0.432	0.454	0.423	0.448	0.847	0.656	0.484	0.482	0.624	0.579
	Avg	0.344	0.392	**0.336**	**0.381**	0.355	0.391	0.355	0.395	0.356	0.394	0.390	0.417	0.374	0.407	0.535	0.499	0.436	0.450	0.499	0.505
ETTm1	96	**0.291**	**0.339**	0.300	0.350	0.306	0.349	0.296	0.342	0.305	0.354	0.321	0.368	0.317	0.365	0.318	0.356	0.351	0.383	0.331	0.367
	192	**0.329**	**0.365**	0.338	0.374	0.350	0.374	0.333	0.369	0.345	0.377	0.359	0.389	0.357	0.387	0.351	0.375	0.384	0.396	0.362	0.383
	336	**0.357**	**0.383**	0.368	0.396	0.381	0.394	0.365	0.393	0.377	0.400	0.399	0.412	0.388	0.407	0.382	0.394	0.414	0.415	0.393	0.402
	720	**0.412**	**0.415**	0.426	0.426	0.435	0.424	0.421	0.424	0.432	0.433	0.463	0.450	0.448	0.443	0.436	0.428	0.501	0.463	0.447	0.434
	Avg	**0.347**	**0.375**	0.358	0.386	0.368	0.386	0.354	0.382	0.365	0.391	0.386	0.405	0.377	0.401	0.372	0.388	0.412	0.414	0.383	0.397
ETTm2	96	**0.166**	**0.254**	0.172	0.259	0.168	0.257	0.173	0.261	0.174	0.262	0.181	0.271	0.179	0.268	0.173	0.268	0.216	0.290	0.194	0.296
	192	**0.225**	**0.290**	0.233	0.298	0.234	0.301	0.237	0.307	0.239	0.305	0.244	0.313	0.239	0.309	0.248	0.328	0.262	0.323	0.262	0.343
	336	**0.283**	**0.327**	0.289	0.336	**0.283**	0.333	0.288	0.339	0.296	0.343	0.301	0.349	0.294	0.345	0.325	0.386	0.320	0.360	0.330	0.389
	720	**0.371**	**0.384**	**0.371**	0.389	0.386	0.395	0.377	0.395	0.382	0.396	0.397	0.407	0.380	0.396	0.453	0.457	0.446	0.431	0.454	0.463
	Avg	**0.261**	**0.314**	0.267	0.321	0.268	0.321	0.269	0.326	0.273	0.327	0.281	0.335	0.273	0.330	0.300	0.360	0.311	0.351	0.310	0.373
Rank		**1.375**	**1.550**	2.525	2.800	4.975	4.375	2.850	2.700	5.450	5.225	6.475	6.650	4.950	5.175	7.575	7.750	9.300	8.825	9.175	9.225

Ablation Study

✓ We compare different design choices for each component of our model.

Formulation	Encoder Attention	PE Sharing	ETTh1		ETTm1		Solar		ECL	
			MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
Standard	Latent Bottleneck	Sharing	0.420	0.437	0.355	0.382	0.194	0.243	0.169	0.265
Generalized	Latent Bottleneck	Sharing	**0.404**	**0.420**	**0.338**	**0.370**	**0.182**	**0.233**	**0.157**	**0.249**
Generalized	Full Self-Attention	Sharing	0.425	0.434	0.353	0.375	0.192	0.237	0.161	0.254
Generalized	Decoupled Self-Attention	Sharing	0.423	0.433	0.356	0.379	0.189	**0.233**	0.158	0.250
Generalized	Latent Bottleneck	Sharing	**0.404**	**0.420**	**0.338**	**0.370**	**0.182**	**0.233**	**0.157**	**0.249**
Generalized	Latent Bottleneck	Not Sharing	0.423	0.432	0.342	0.374	0.193	**0.232**	0.163	0.254
Generalized	Latent Bottleneck	Sharing	**0.404**	**0.420**	**0.338**	**0.370**	**0.182**	0.233	**0.157**	**0.249**

Ablation Study

✓ To better understand how our model processes information throughout encoding and decoding, we visualize the cross-attention maps.

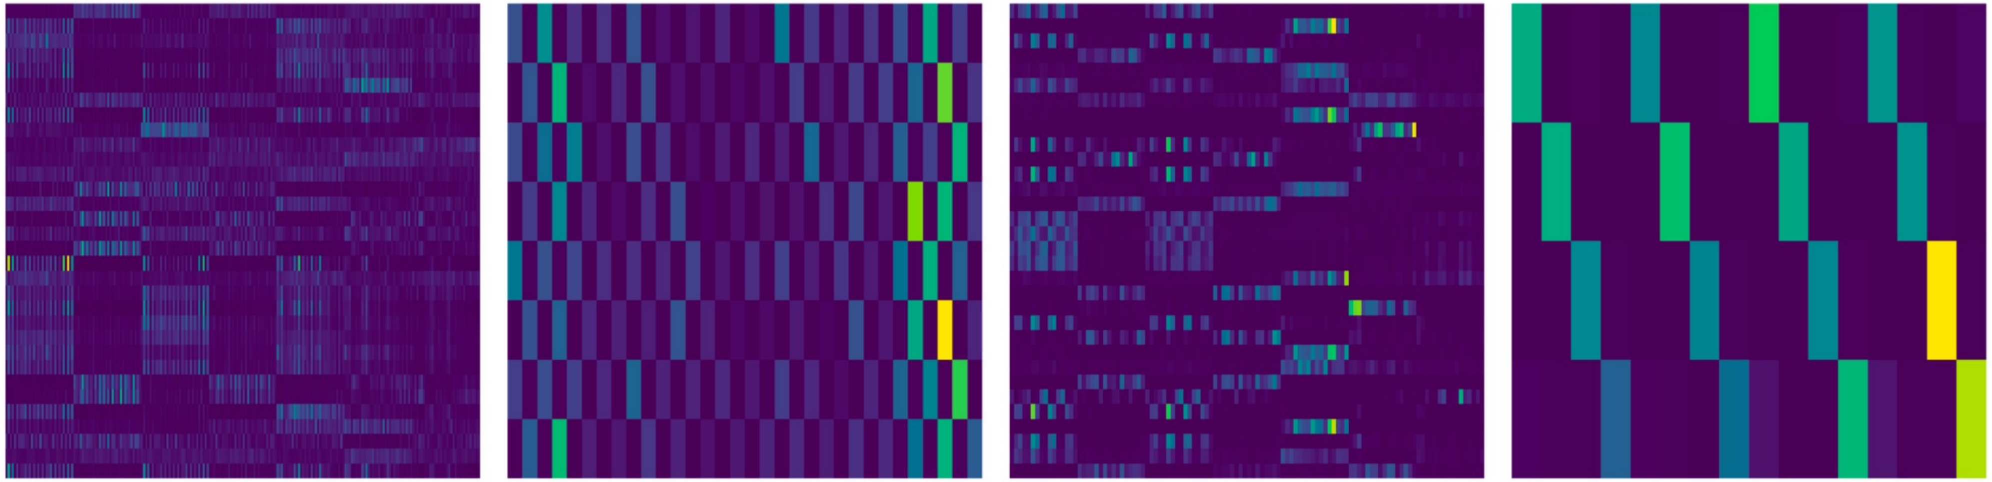

(a) Encoder and decoder attention on **ETTh1** (b) Encoder and decoder attention on **ETTm1**

Ablation Study

✓ To better understand how our model processes information throughout encoding and decoding, we visualize the cross-attention maps.

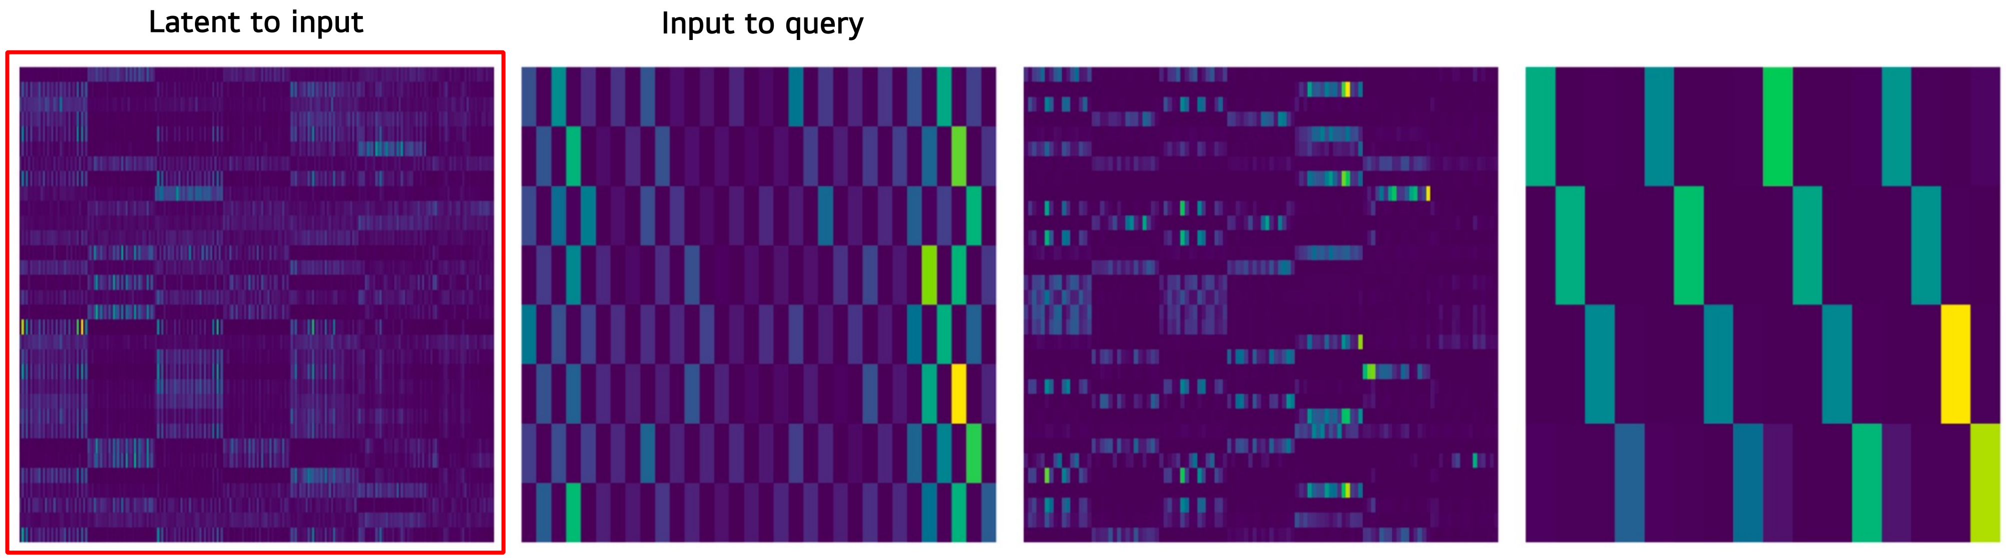

(a) Encoder and decoder attention on **ETTh1**   (b) Encoder and decoder attention on **ETTm1**

Ablation Study

✓ To better understand how our model processes information throughout encoding and decoding, we visualize the cross-attention maps.

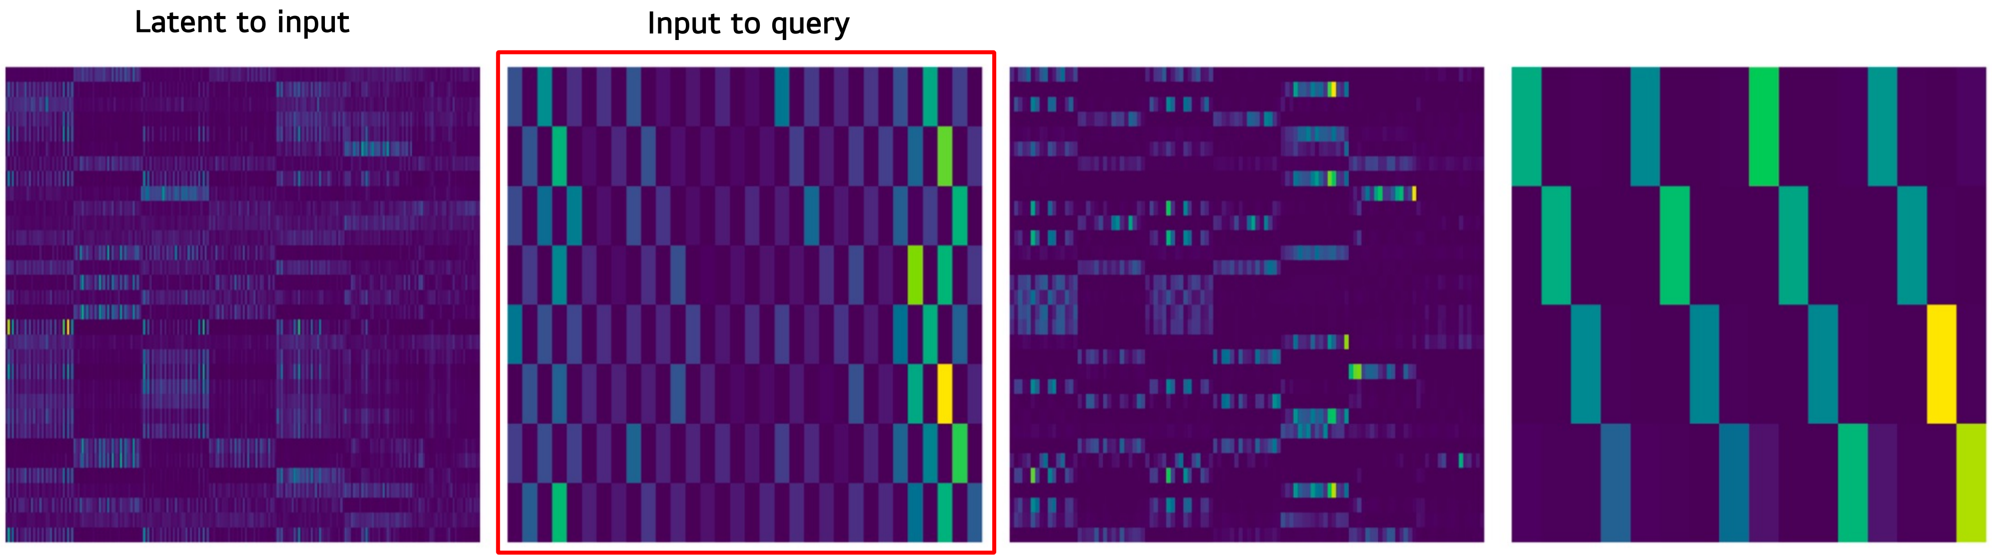

(a) Encoder and decoder attention on **ETTh1** (b) Encoder and decoder attention on **ETTm1**

Conclusion

- ✓ We propose a unified framework that tightly integrates an **attention-based architecture** with a **generalized training strategy** that enables flexible forecasting at arbitrary target positions.
    - Extensive experiments demonstrate that our framework consistently and significantly **outperforms prior state-of-the-art** baselines across a wide range of benchmark datasets.
    - Through ablation studies on **how the model is designed and operates**, we demonstrate the validity of our architectural choices.
- ✓ We hope our work encourages the community to **explore holistic modeling approaches** that align architectural choices more closely with **the underlying forecasting objectives.**

Thank you.

Poster: Session 2, Dec 3(Wed), 16:30

San Diego Convention Center