




SD4H



FMSD

Adapt the Data, Not the Model: Input-Space Adaptation for Frozen Time-Series Predictors

Omer Gotfrid Dvir Aran Barak Gahtan Alex M. Bronstein


Technion - Israel Institute of Technology | ICML 2026 Workshops: SD4H (Structured Data for Health) & FMSD (Foundation Models for Structured Data)

A trained **input adapter** lets a **strictly frozen, certified** predictor match or beat models retrained in the new hospital/sensor - **no weight changes, no re-certification, no source data required**.

The Clinical Problem

- **Hospital predictors are frozen.** FDA Software-as-a-Medical-Device, vendor firmware, and closed scoring APIs expose *inference only*; any weight change triggers a new approval cycle.
- **They degrade silently across hospitals** (below) - yet cannot be retrained in place.
- **Source data is usually inaccessible:** privacy law, data-use agreements, and networks like Sentinel / OHDSI / N3C / EHDEN keep patient records at their origin.

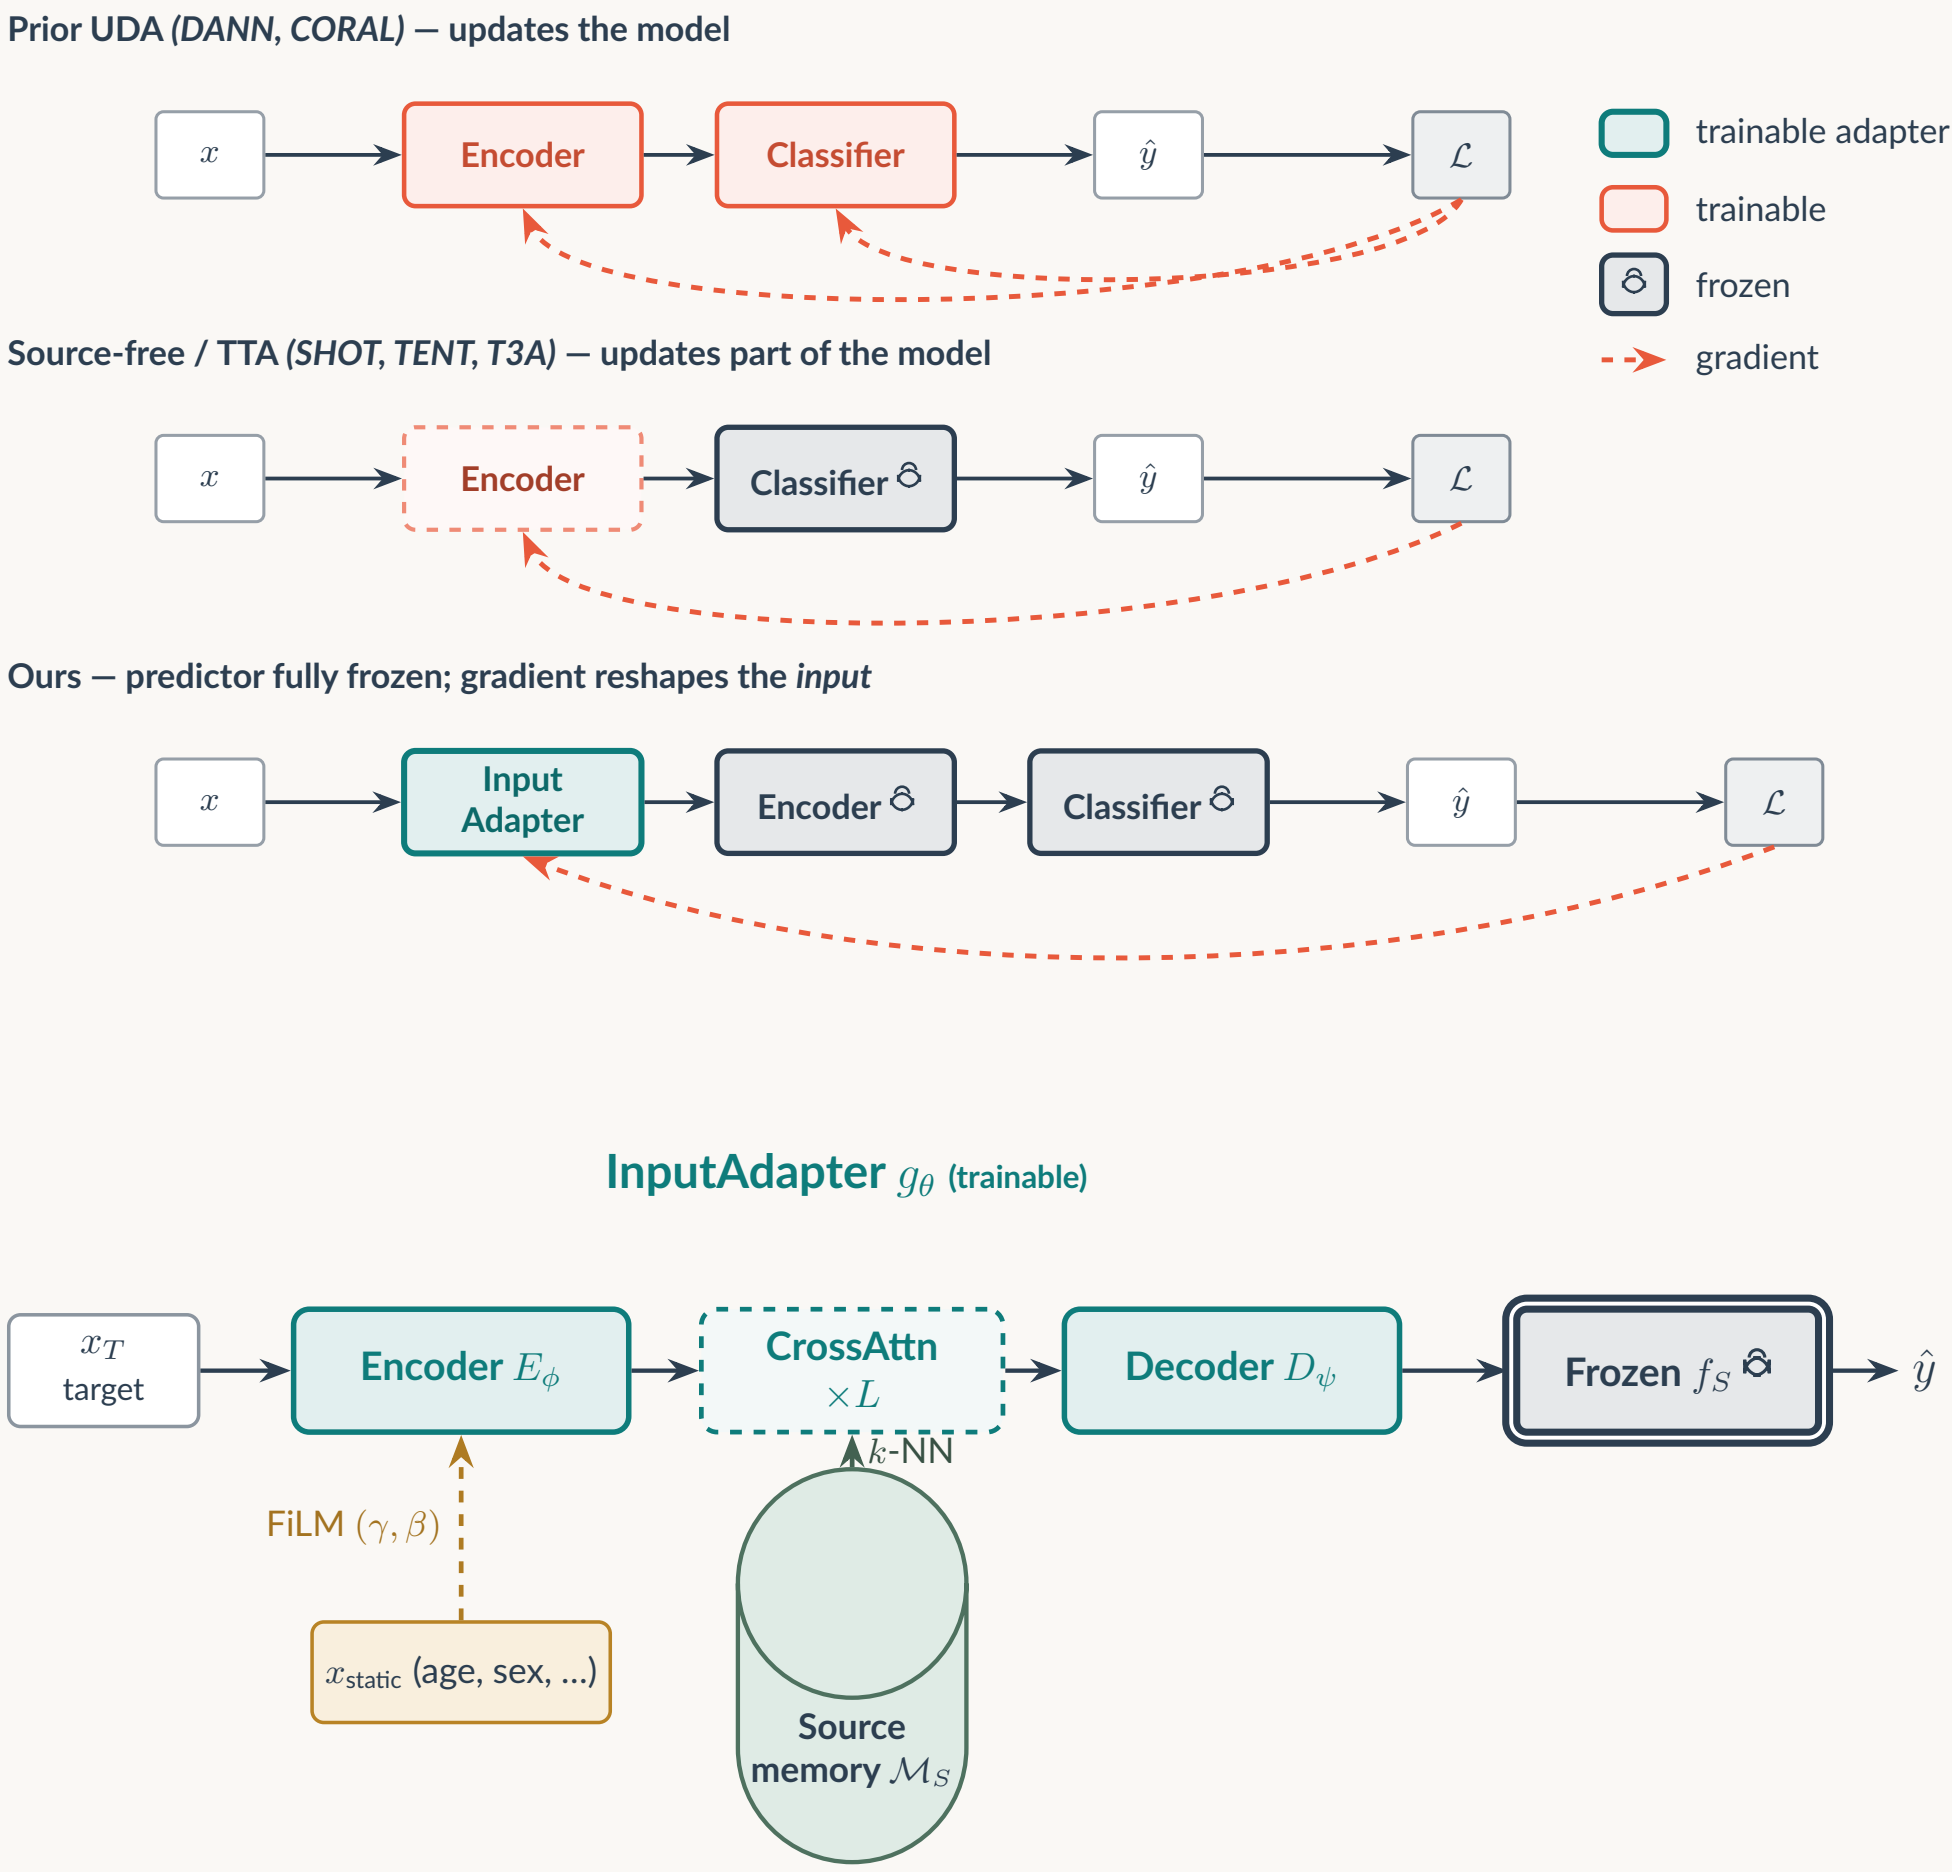
A MIMIC-Trained, FDA-Locked Predictor Degrades Silently Across Hospitals



Regression hit just as hard: up to ~ 11 h worse MAE on HiRID length-of-stay.


Method: INPUTADAPTER

Prior domain adaptation *updates the model*; we don't. We send the **task gradient** through the frozen predictor into an upstream adapter that reshapes the input.




Results - Frozen Predictor + INPUTADAPTER Matches or Beats Target-Native Models on All 5 Clinical Tasks and 5 Sensor Benchmarks

Clinical Classification + Regression



Sensor Classification



Baselines


Beats every comparable baseline on both suites - 15 (clinical - DA and TTA baselines) and 13 (AdaTime benchmark) - and matches or beats end-to-end fine-tuning, while the certified predictor stays bit-identical.

Multiple Targets & Modalities

Generalizes across targets and modalities: a second hospital (HiRID - AKI +7.5 AUROC, beating native by +1.7) and five wearable / EEG / industrial sensor benchmarks (+2.6 mean Macro-F1, all five won).

Deploys Without Source Data


Source-Free = Source-Assisted



The strict **source-free** variant (no source sample touches the gradient) matches the source-assisted regime on every task - and still beats the target-native LSTM on 4/5. Regression unchanged: LoS -5.4 h (both), KF -0.116 / -0.117 mg/dL.

Architecture-Agnostic

Adapter Trained on LSTM, Reused Zero-Shot



An adapter trained on a frozen LSTM transfers **zero-shot** to frozen GRU and TCN predictors, retaining **53-103%** of the gain - all 16 task $\times$ architecture cells positive.