Bridges-2 Webinar

Recent Advances in Time Series Foundation Modeling

Artur Dubrawski, Ph.D., M. Eng.
Mononito Goswami, Ph.D.

Auton Lab, Carnegie Mellon University

June 11, 2025

Pittsburgh Supercomputing Center

enabling discovery since 1986

The **Pittsburgh Supercomputing Center (PSC)** provides advanced research computing capability, education, and expertise to the national research community.

Since 1986, PSC has provided university, government, and industry researchers with access to some of the most powerful systems available for high-performance computing, enabling discovery across all fields of science.

OUR AREAS OF EXPERTISE

- high-performance and data-intensive computing
- data management technologies
- software architecture, implementation, and optimization
- enabling ground-breaking science, computer science, and engineering
- user support for all phases of research and education
- STEM outreach in data science, bioinformatics, and coding

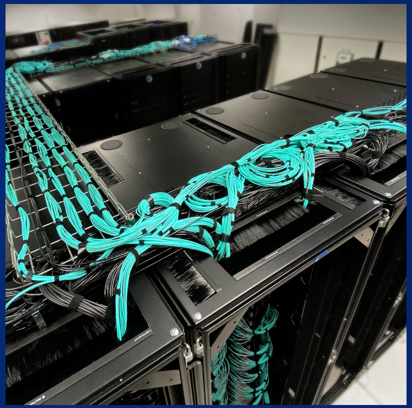

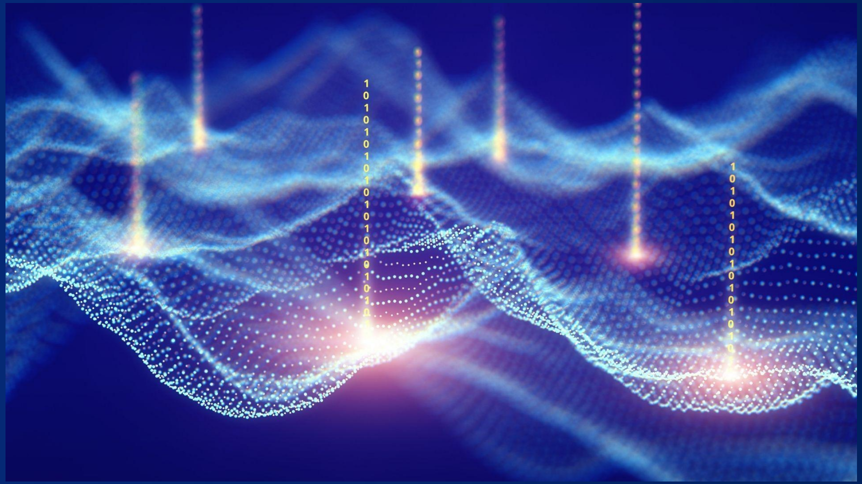

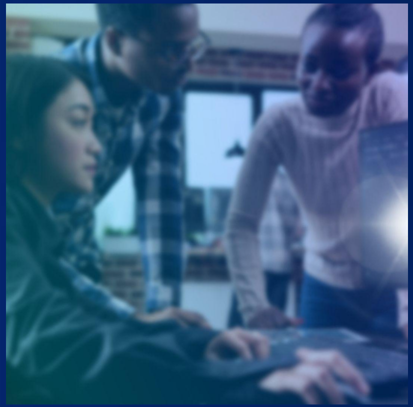

Welcome!

Hewlett Packard Enterprise is delivering *Bridges-2*

Bridges-2 Leadership Team

Sergiu Sanielevici
PI & Dir. Support
for Sci. Apps.

Robin Scibek
Dir. Comms.
co-PI

Paola Buitrago
Dir. AI & Big Data
co-PI

Edward Hanna
Dir. Systems & Ops.
co-PI

Tom Maiden
User Services Mgr.
co-PI

Riaz Khatri
Project Manager

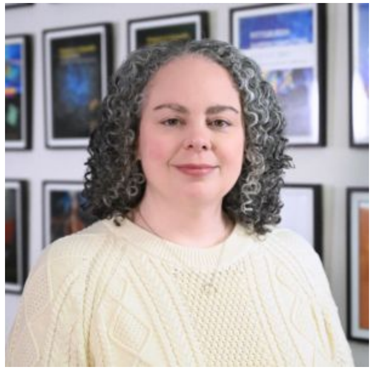

Joanne Peca
Information Security Officer

Bridges-2 Webinars

- A forum for the Bridges-2 community to learn and share ideas and achievements: Bridges-2 Webinar series | PSC
- Topics and speakers of interest to work that is being done, or that may be done in future.
- Please suggest future speakers (including from your own team) and/or topics (including your own)!

Just email: sergiu@psc.edu

Introducing today’s presenters

- **Artur Dubrawski, Ph.D. M.Eng,** is an Alumni Research Professor Chair of Computer Science at Carnegie Mellon University where he directs the Auton Lab, one of the largest applied machine learning and artificial intelligence teams in academia. For more than 3 decades he has been working on the forefront of development of AI serving in technical leadership roles in industry and leading multiple research endeavors in academia.
- **Mononito Goswami** recently graduated with a **Ph.D. in Robotics** from Carnegie Mellon University. He is interested in developing foundational machine learning (ML) techniques for real-world applications. His research tackles the limitations of traditional ML approaches, focusing on scenarios with inaccurate, decentralized, and insufficient data, all in effort to democratize ML. He led the development of one of the first open-source foundation models for time series data.

Q&A Logistics

- **We abide by https://support.access-ci.org/code-of-conduct**
- All of us except our speakers will be muted during their presentation.
- Please type your questions into the Zoom chat.
- After the presentation, our speakers will answer questions live during the final ~10 minutes of this webinar.
- The video recording and slides of this webinar will be linked from https://www.psc.edu/events/bridges-2-webinar-series/time-series-foundation-modeling/ next week.

Recent Advances in Time Series Foundation Modeling

Artur Dubrawski
Alumni Research Professor of Computer Science
awd@cs.cmu.edu

Mononito Goswami
PhD Graduate
mgoswami@cs.cmu.edu

Auton Lab

Carnegie Mellon University

FOUNDATIONS OF MACHINE LEARNING → PROOFS OF CONCEPTS → DEPLOYMENTS


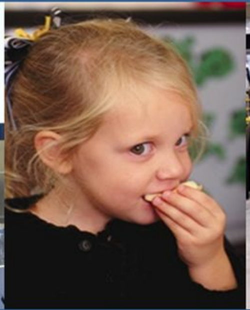

Radiation Safety | Food Safety | Predictive Maintenance | Counter-Trafficking



Predictions in Emergency, Intensive, and Surgical Care; Infection Control

- ✓ Making AI accessible, immersive, and useful in the Real World
- ✓ Making intelligent systems easy to deploy and maintain
- ✓ Learning from limited feedback
- ✓ Learning on various types and abstracts of data

Artur Dubrawski

Auton Lab

CMU Auton Lab

30 Autonomous Systems Laboratory Auton Lab SINCE 1993

Celebrated 30th Anniversary in 2023

How to Realize the Potential of AI in Applications?

With a Systematic Approach: The AI Capability Stack

- There are only a few general principles on how to succeed at developing and deploying AI solutions in practice
  - The real-world application of AI is relatively young; recently we can see a broad interest in it as a potential differentiator
- Many experts claim that one cannot fully capitalize on the potential of Artificial Intelligence unless every layer of the AI Stack is involved (or at least considered) in the AI-driven solution

Stack of Artificial Intelligence Capabilities (AI Stack)

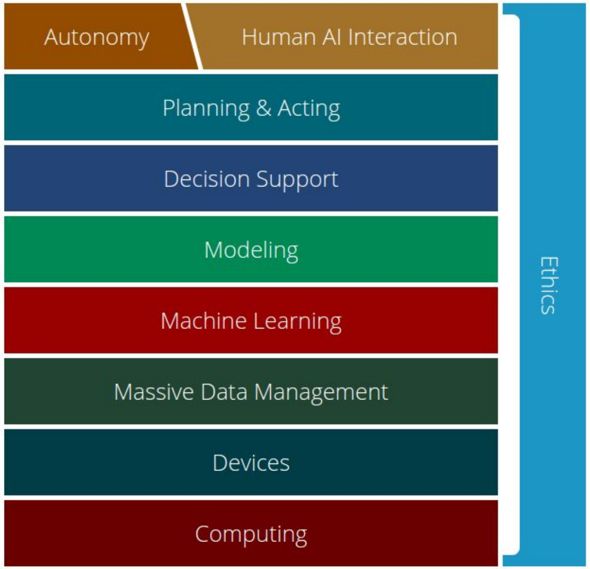

AI Project Cycle Aligns with the AI Stack

- **In the AI Stack and the AI Project Cycle, opportunities downstream depend on the capabilities upstream**

Stack of Artificial Intelligence Capabilities (AI Stack)

Artificial Intelligence Project Cycle

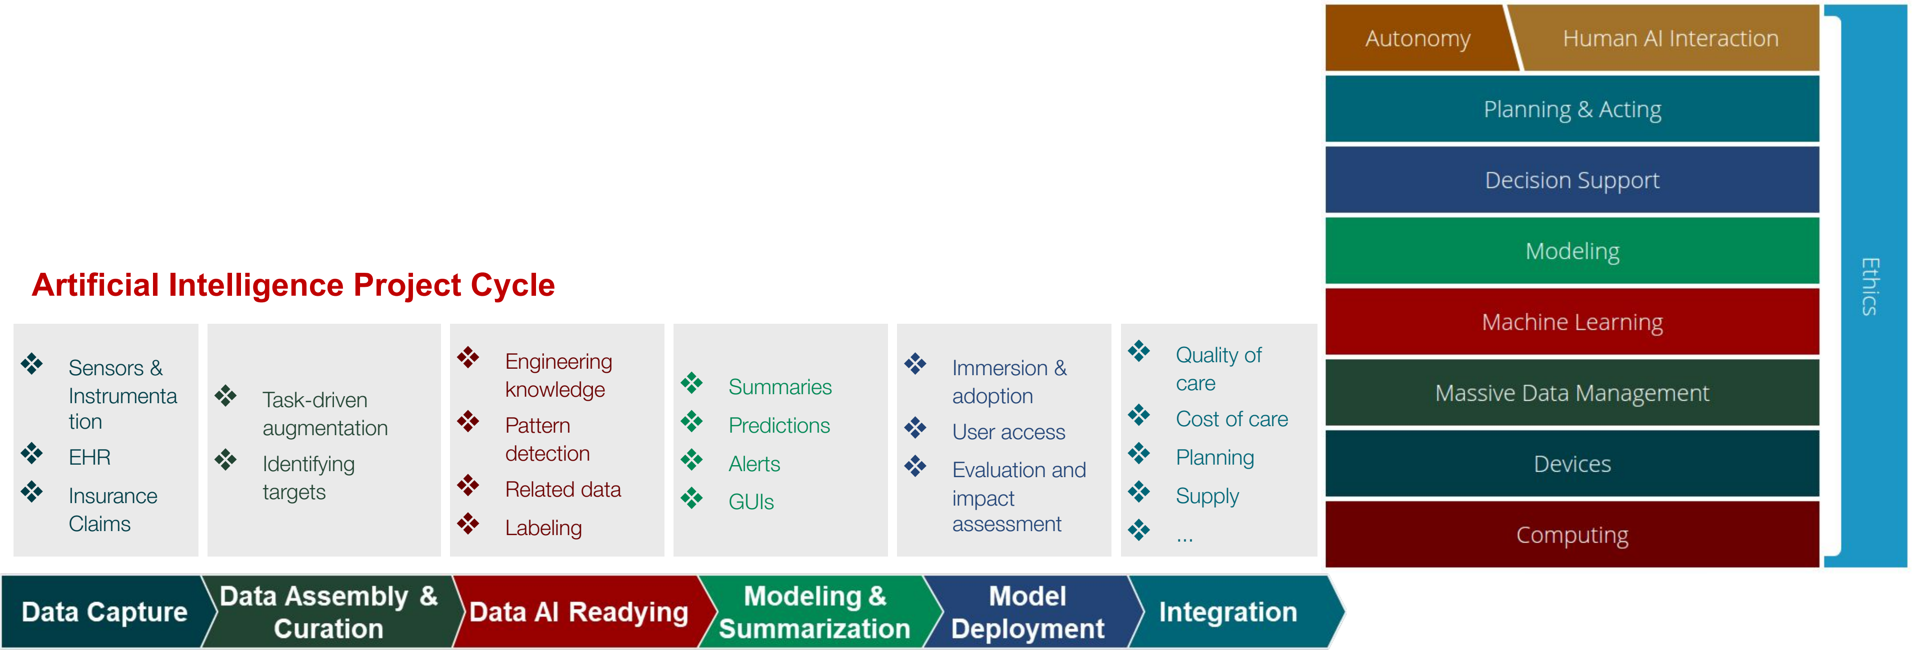

AI Project Cycle Aligns with the AI Stack

- The capability grid, occuring at the intersection, scopes the potential impact of AI

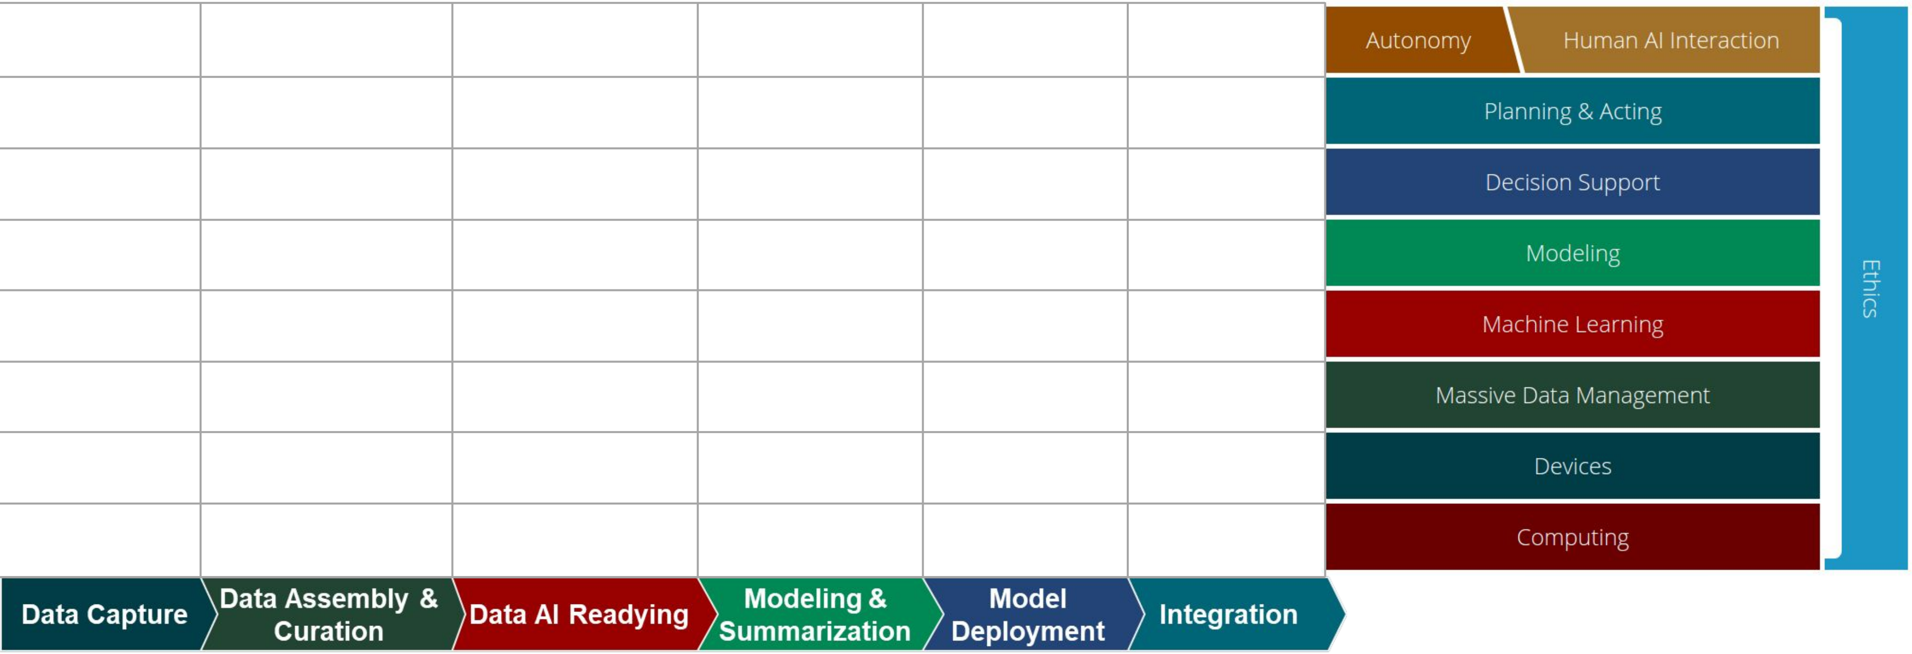

The Capability Grid Scopes Potential Impact of AI

- The capability grid, occuring at the intersection, scopes the potential impact of AI
- It allows us to systematically map the existing capabilities and capability gaps

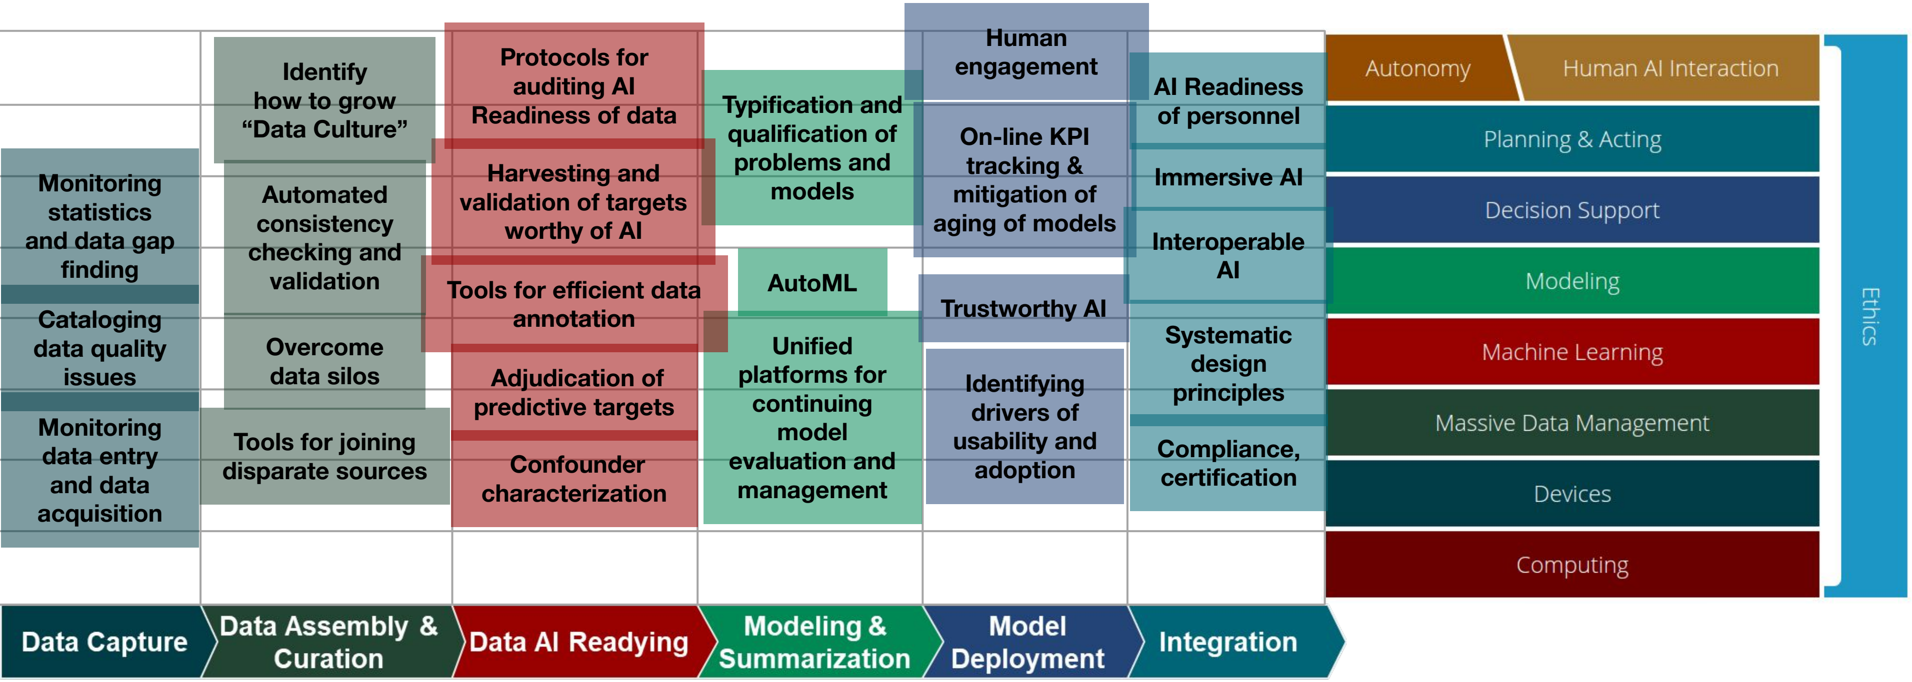

The Capability Grid Scopes Potential Impact of AI

- Self-limiting reduces the potential impact of academic AI initiatives

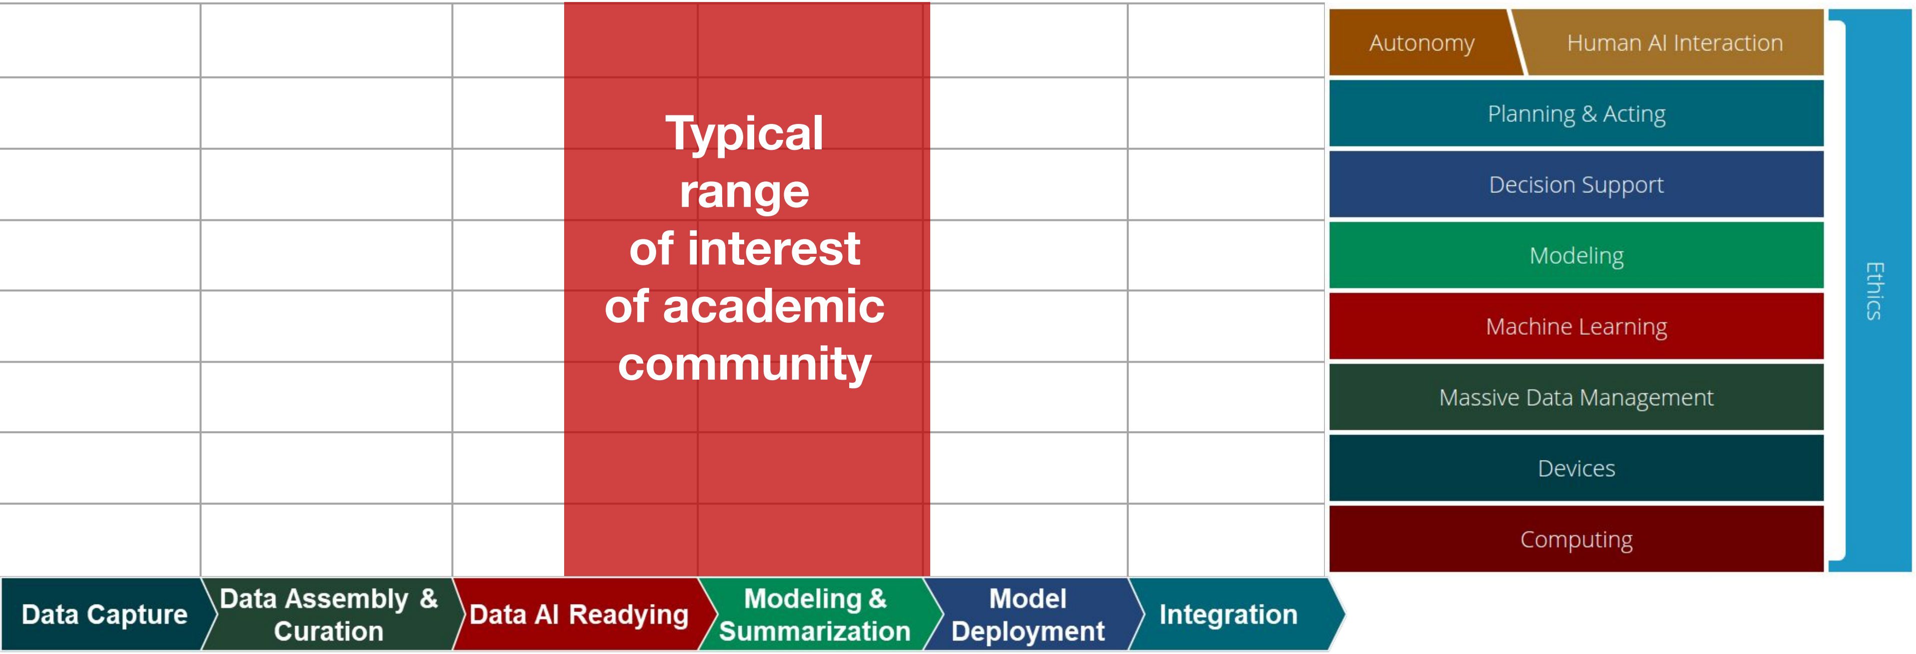

Key Limitations of AI in Practice
are due to the DATA, the MODELS …and HUMANS

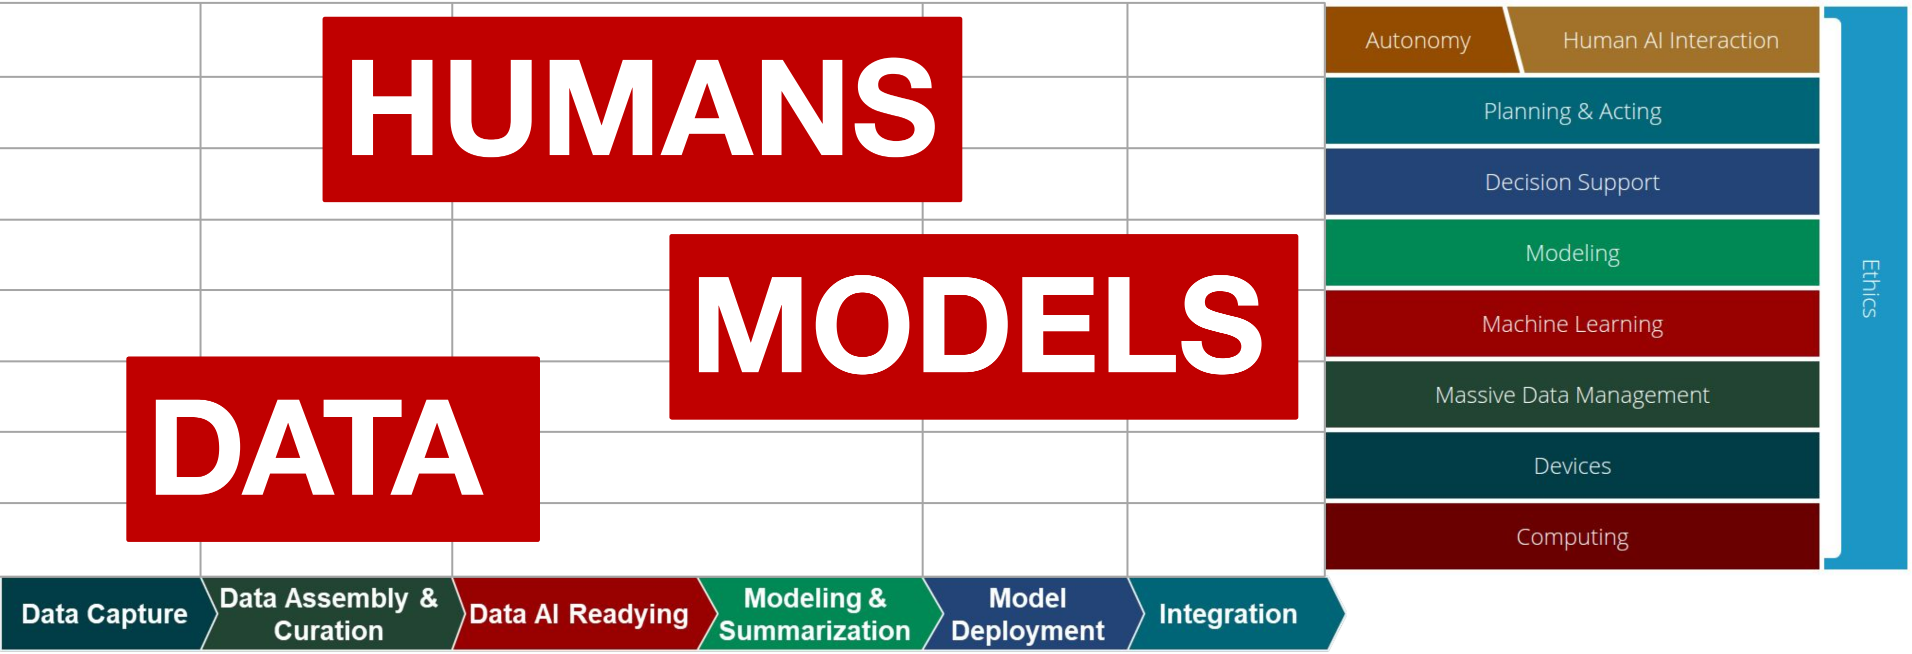

Foundation Models for Language and Vision Data Are Prevalent

(unlike other modalities such as Time Series or Tabular Data)

Forecasting-only

Chronos
(Amazon, Mar'24)

MOIRAI
(Salesforce AI Research Asia, Feb'24)

Lag-Llama
(Morgan Stanley, ServiceNow Research ..., Feb'24)

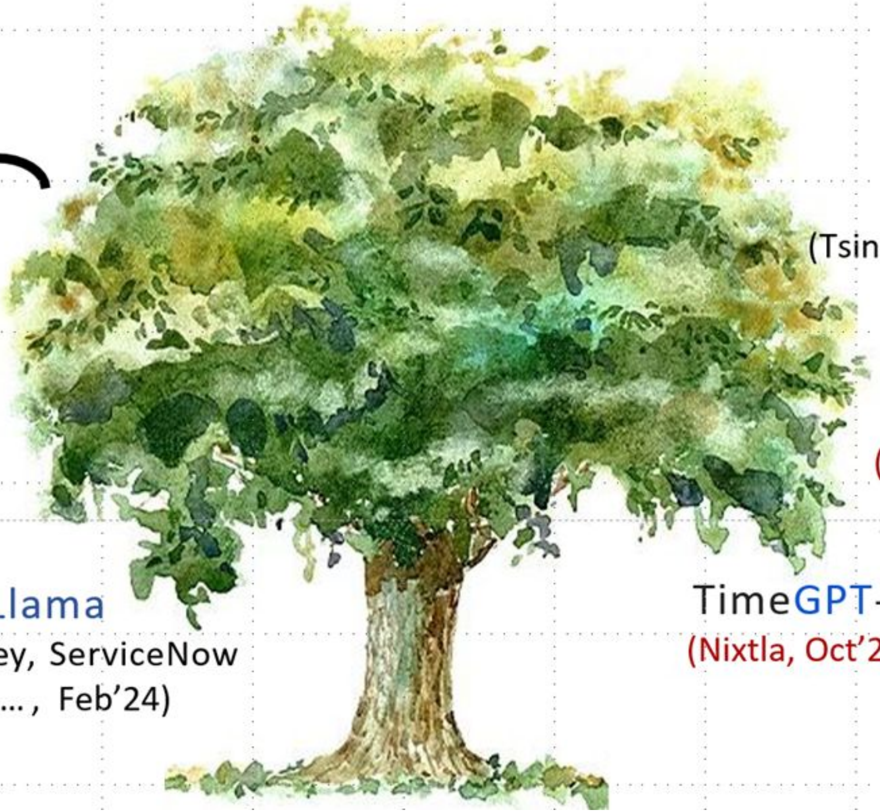

First Open Source, **Multi-task** Model

Timer
(Tsinghua University, Feb'24)

MOMENT
(CMU, Feb'24)

Auton Lab

TimeGPT-1
(Nixtla, Oct'23)

Imagine a Neurologist Reviewing Brain EEG

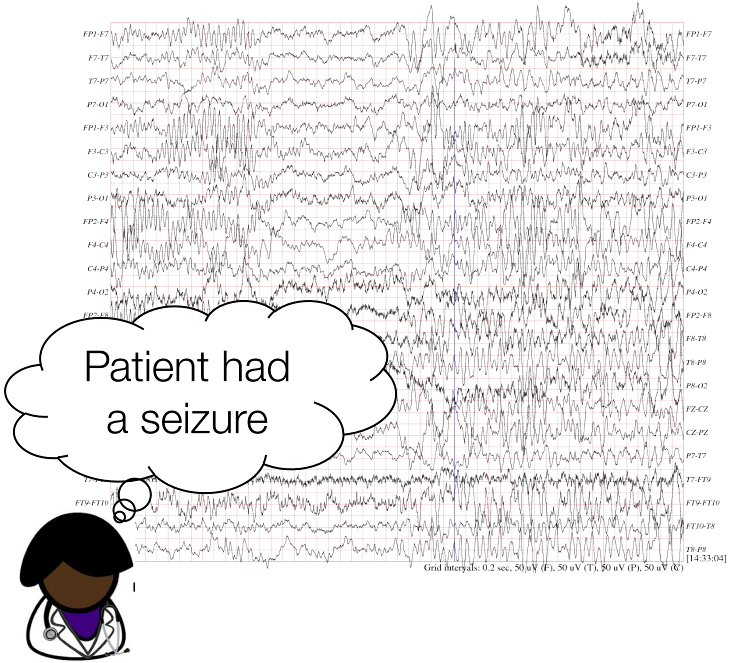

Long, Multivariate Time Series
Hours of data, 240 Hz, 21 leads (features)

Review Complementary Modalities

Biological covariates
(age, sex, …),
past medical history etc.

Text Data

Blood Report	
Haemoglobin	14.0
RBC Count	5.21

Tables

Challenge 1: Domain Experts Must Review Complex Multimodal Information

Building ML Models that Save Lives is Tedious

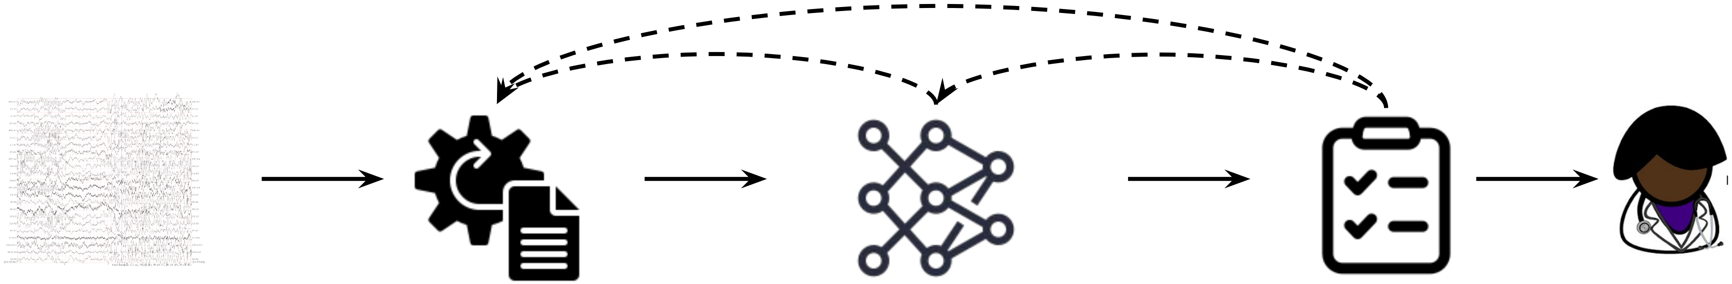

Problem
Seizure vs Normal

Data Preparation
(Preprocessing, Labeling, etc.)

Modeling
(Build, Train, Tune ML Models)

Evaluation
(Benchmarking, Model Selection)

Challenge 2: Building ML Models for a Specific Problem is Tedious and Repetitive

Most Data is Not Ready for Machine Learning

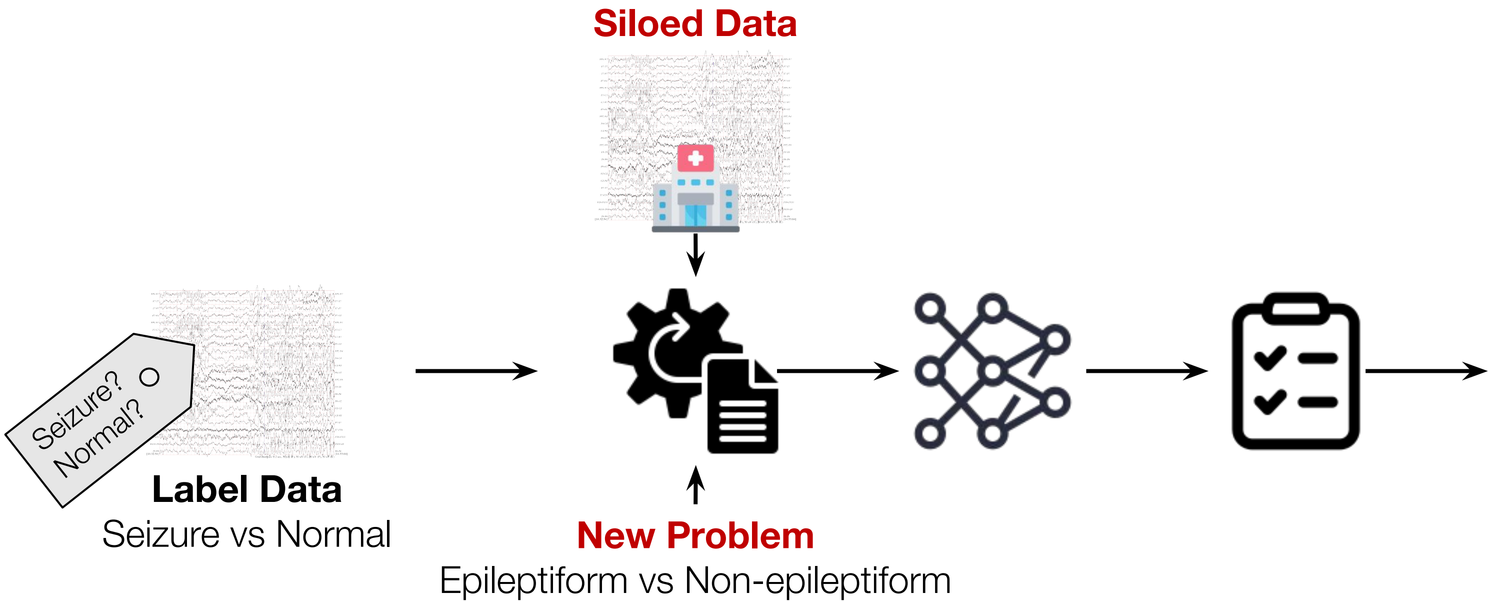

Challenge 3: Preparing Data for Machine Learning
is Time Consuming, Cumbersome, and Error Prone

Challenges Limiting Widespread Adoption of Time Series Machine Learning

Prepare Data Model Evaluate

Building ML Models is
Tedious & Repetitive

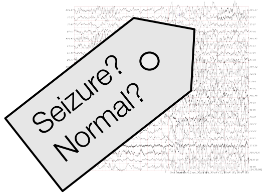

Preparing Data for ML
is Costly, and Error Prone

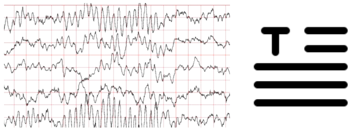

Time Series Text

Large Volumes of Complex
Multimodal Information

Building ML Models is Tedious → Foundation Models

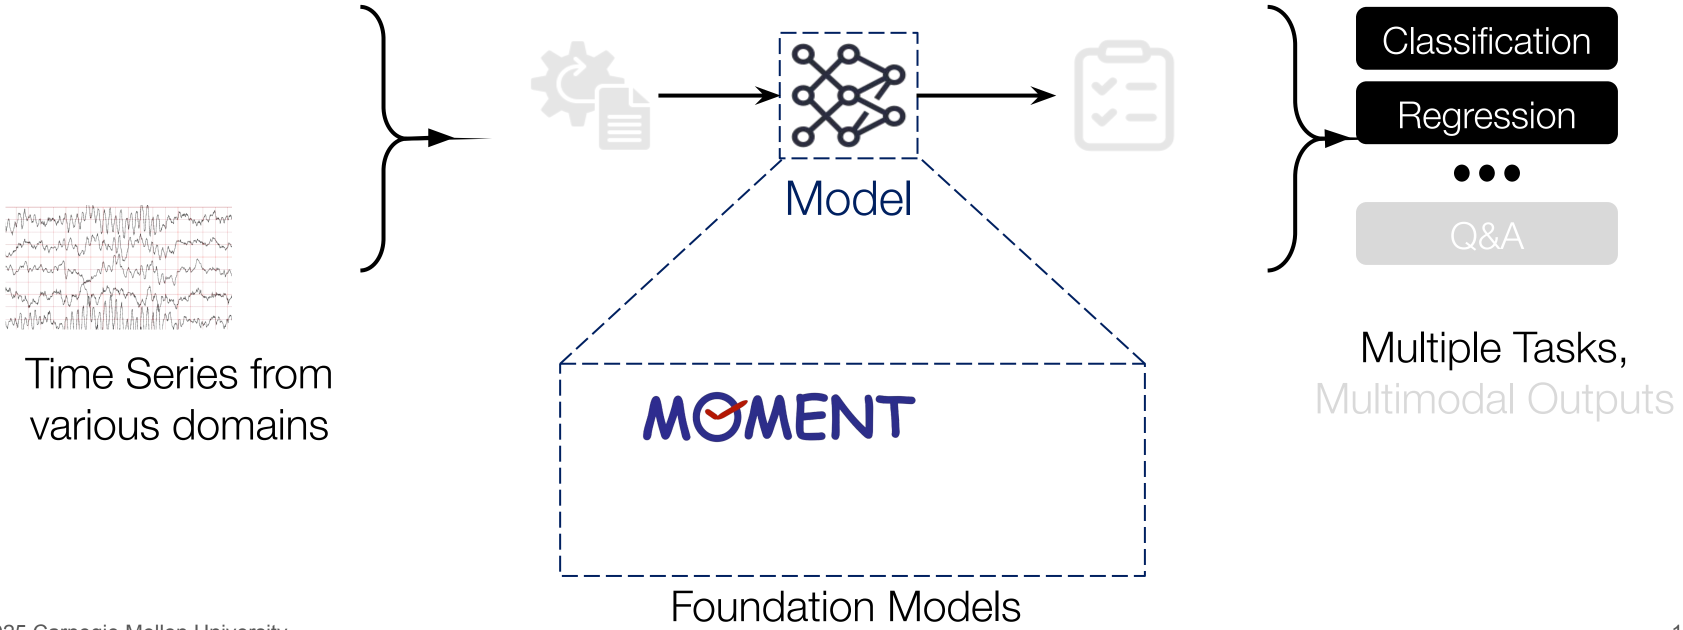

Complex Information →
Long Context, Multimodal Foundation Models

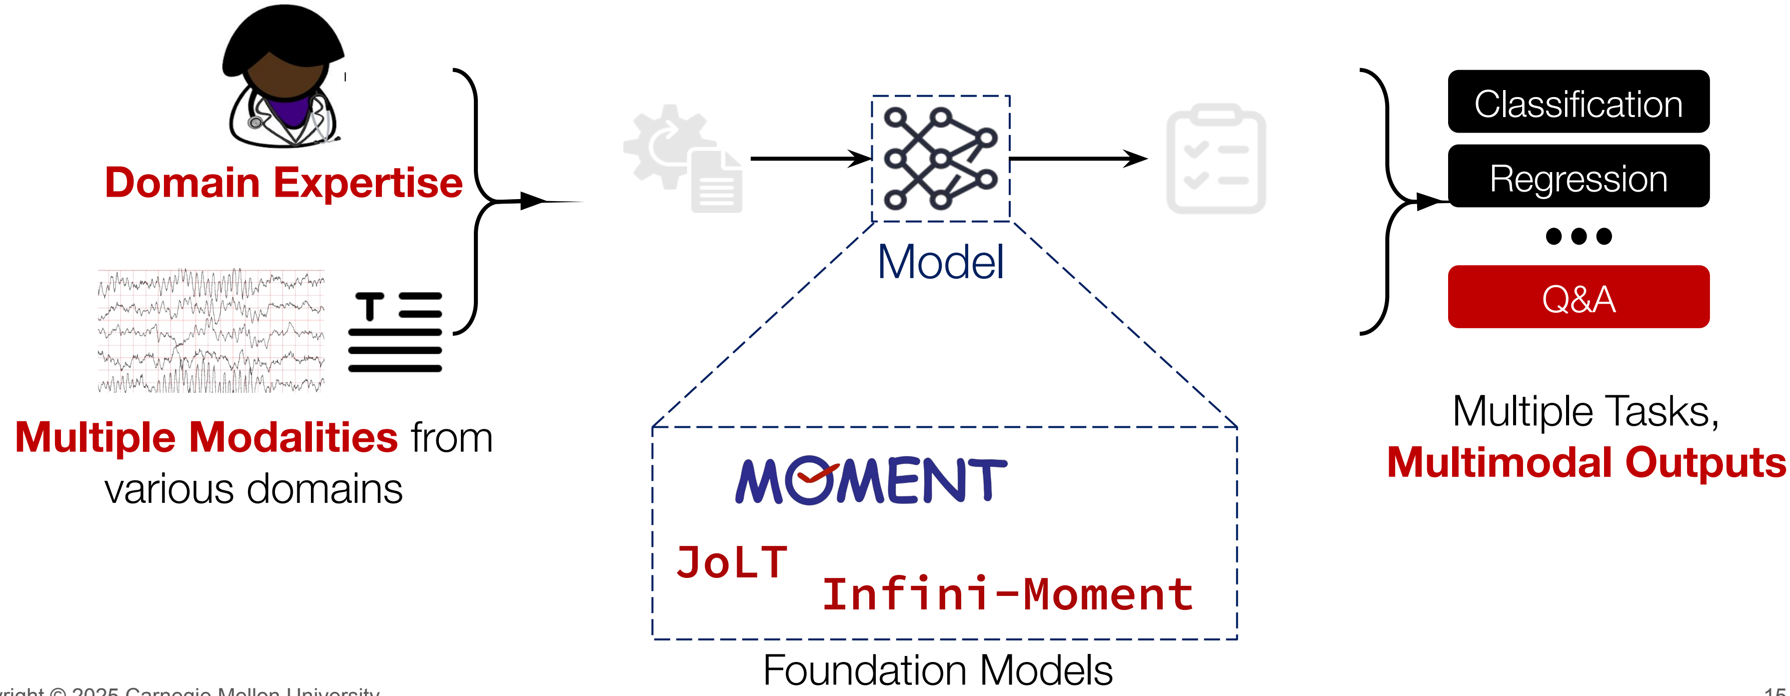

Most Data is Not Ready for Modeling → Tools to Prepare Data for Machine Learning

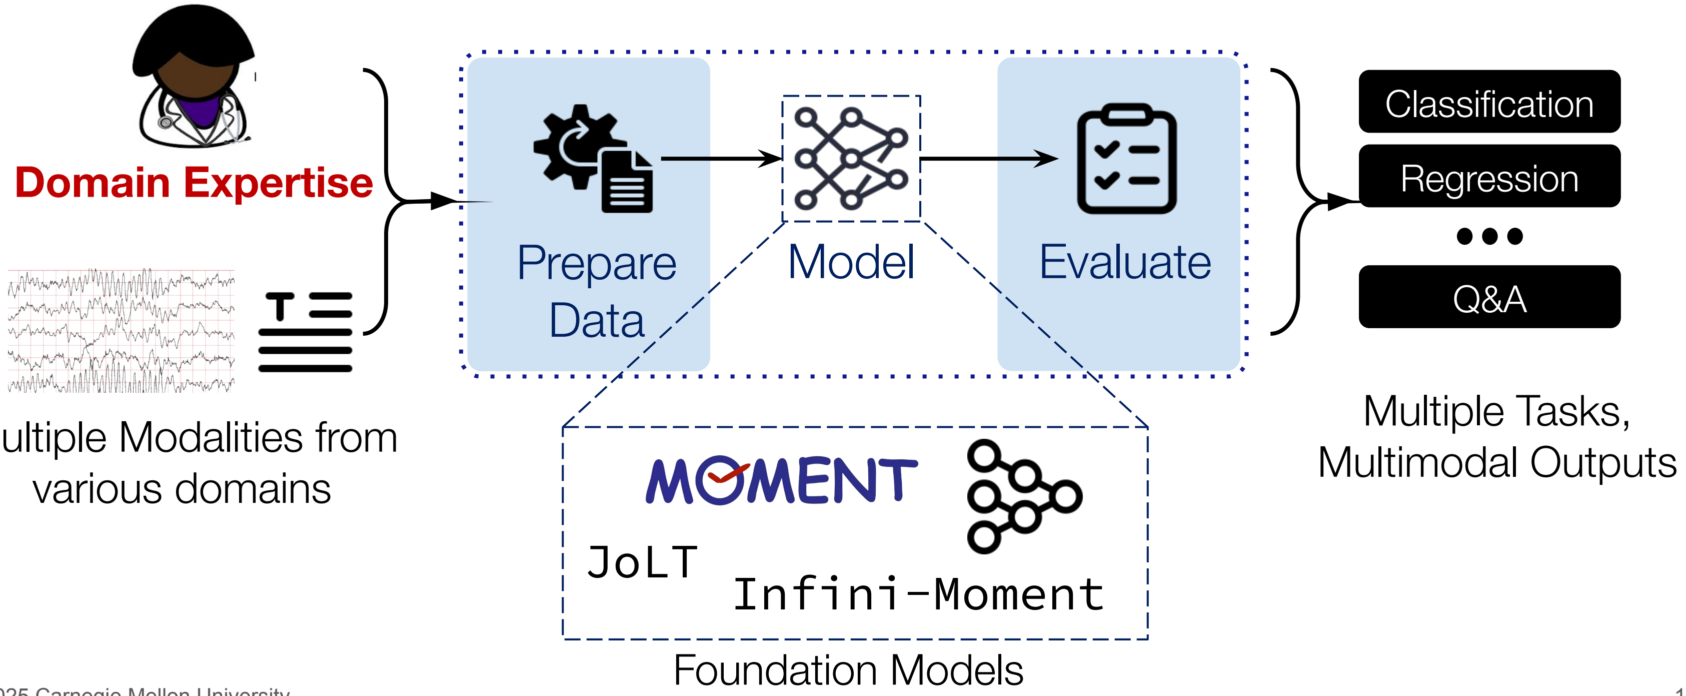

Empower Domain Experts with
Time Series Intelligence at their Fingertips

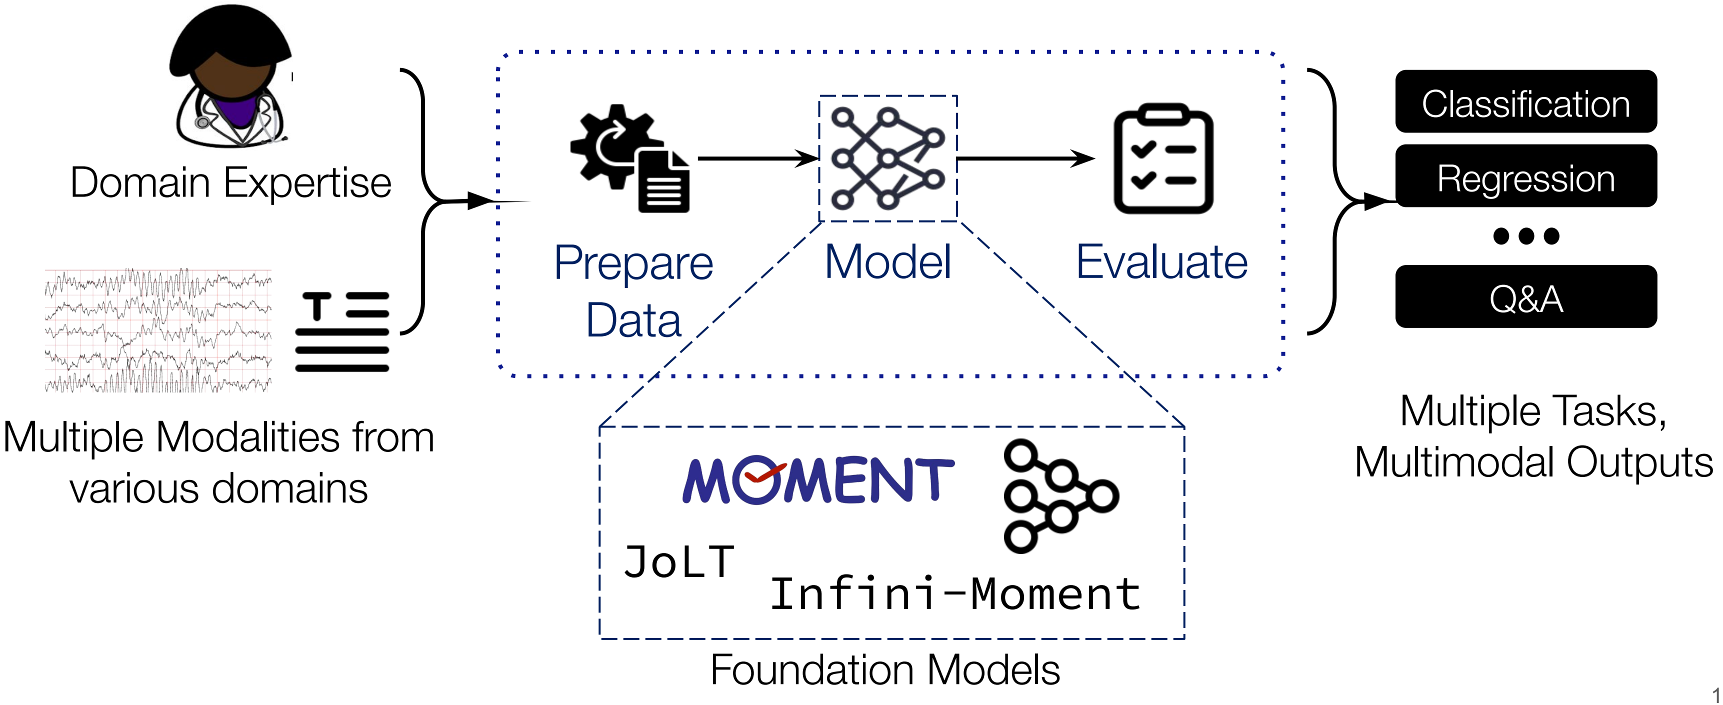

Empower Domain Experts with Time Series Intelligence at their Fingertips

Our hypothesis:

Time series intelligence can be substantially **advanced** by

1. **Building capable foundation models dedicated to time series data**,
2. **Improving our understanding of these models**, and
3. **Addressing practical challenges.**

Foundation Models

Part 1: Building Capable Foundation Models

+ Long-Context

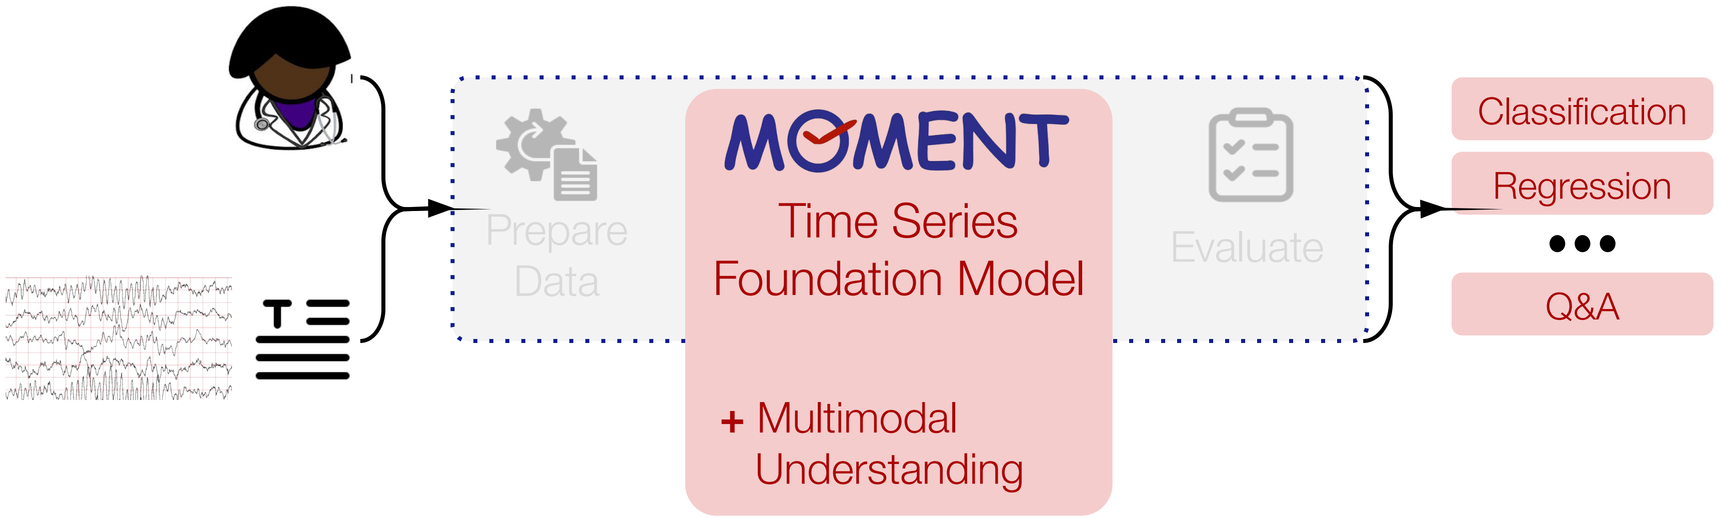

Part 2: Improving Understanding of Foundation Models

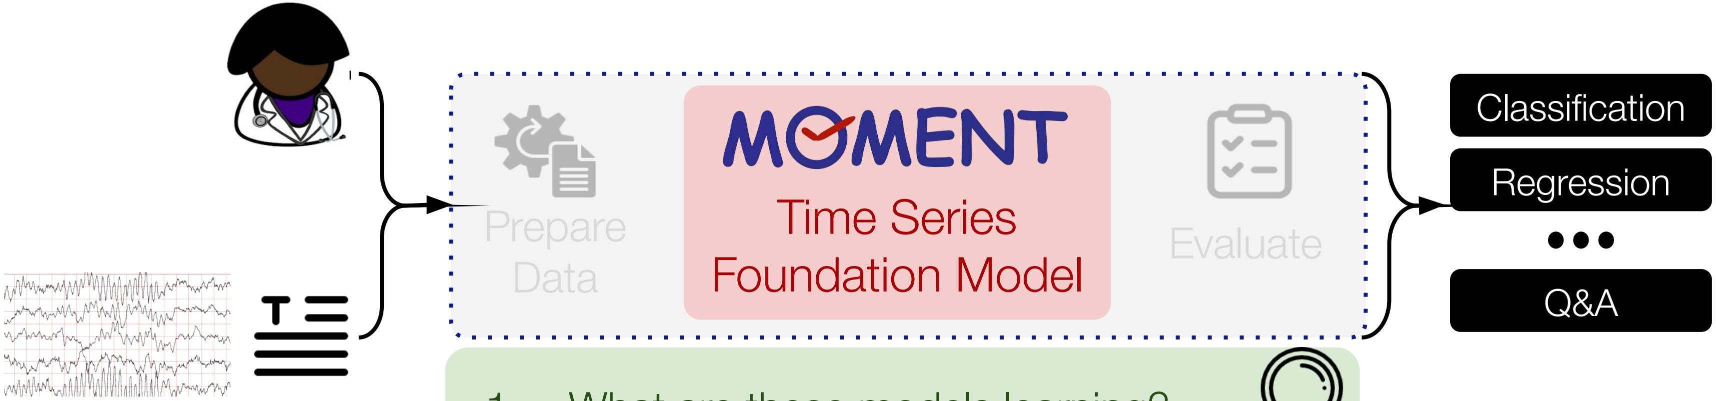

1. What are these models learning?
2. How can we steer these models?
3. Can these models reason?
4. Can they understand basic time series concepts?

Part 3: Addressing Practical Challenges

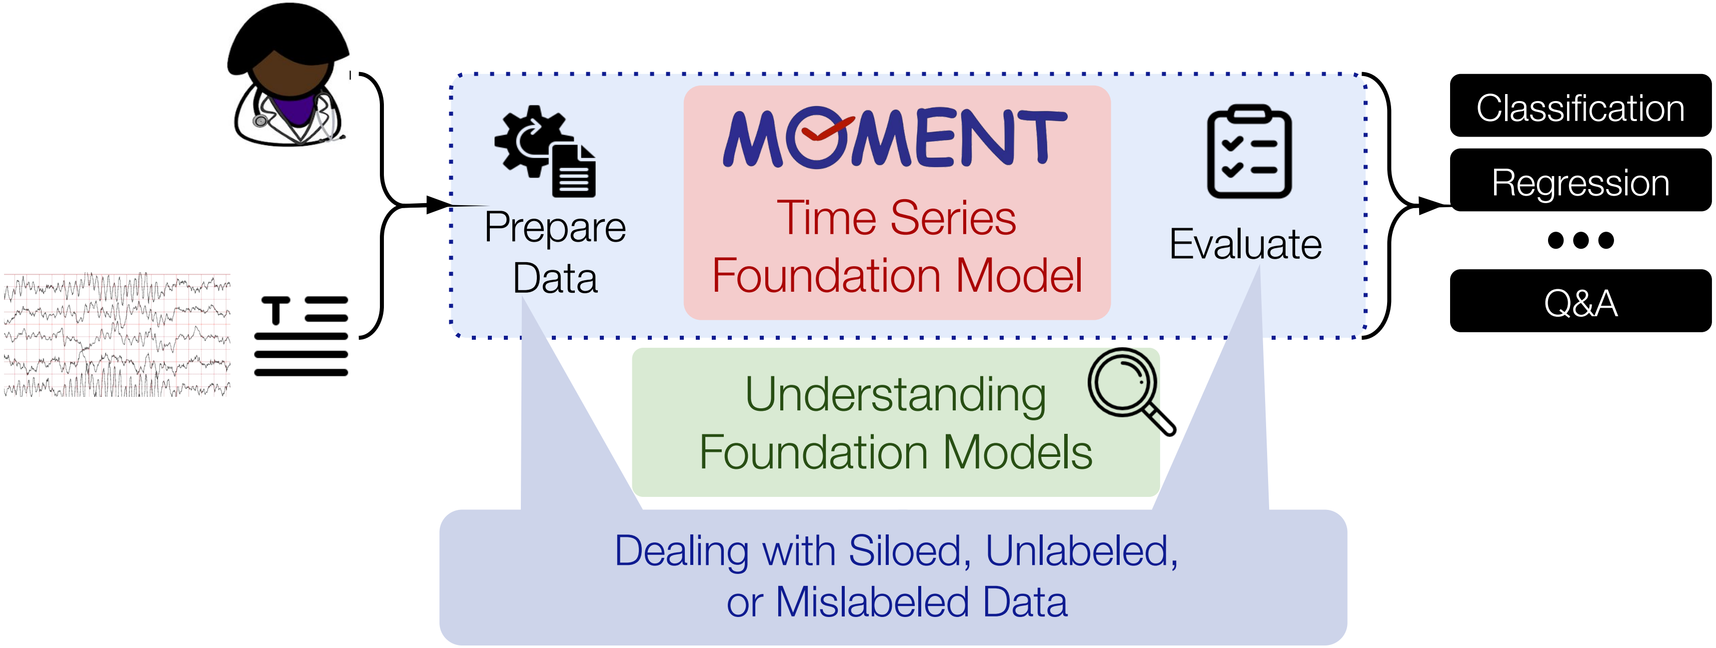

Time Series

Sequence of **numerical observations** of a set of **features** obtained sequentially in **time**

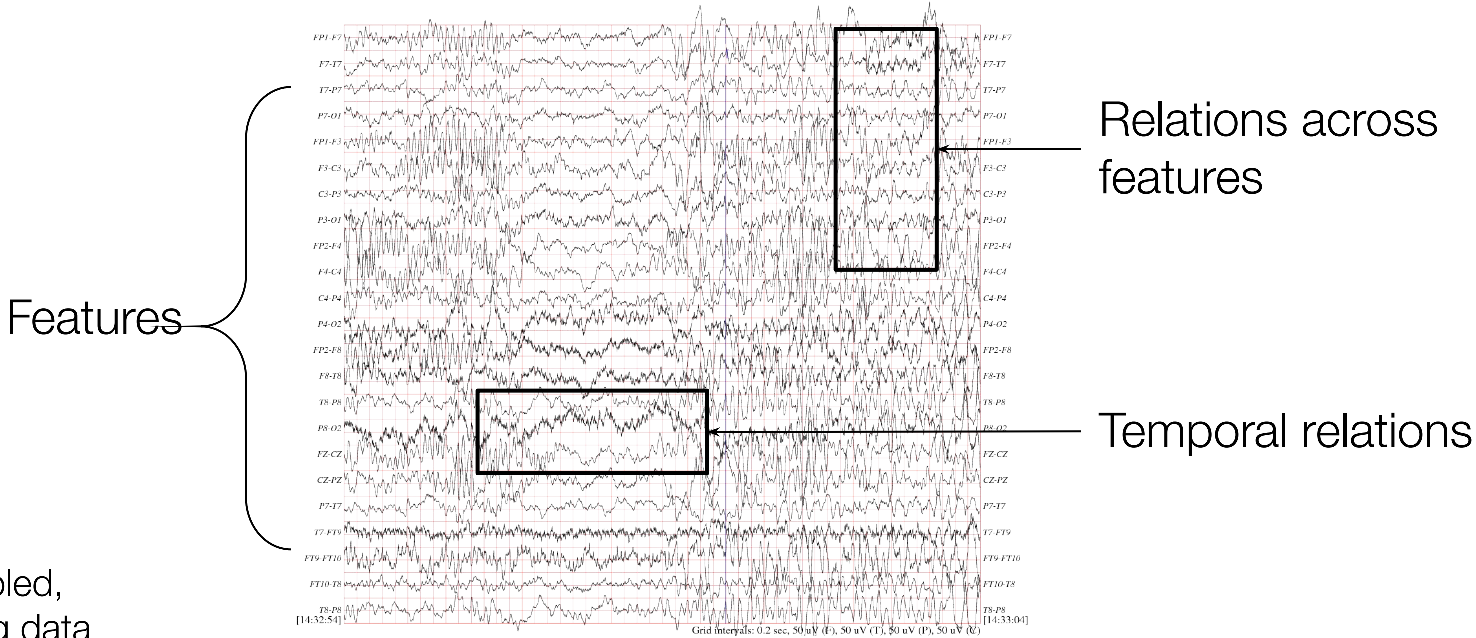

Assumptions:
Uniformly sampled,
minimal missing data

Time Series Data are Prevalent and Modeling it is Impactful

Weather

Temperature

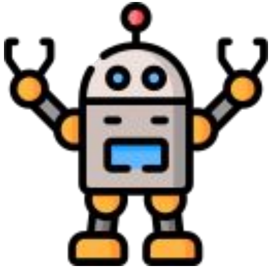

Robotics

Trajectories

Economics

GDP

Healthcare

Electrocardiogram

And many more…

Part 1: Building Capable Foundation Models

+ Long-Context

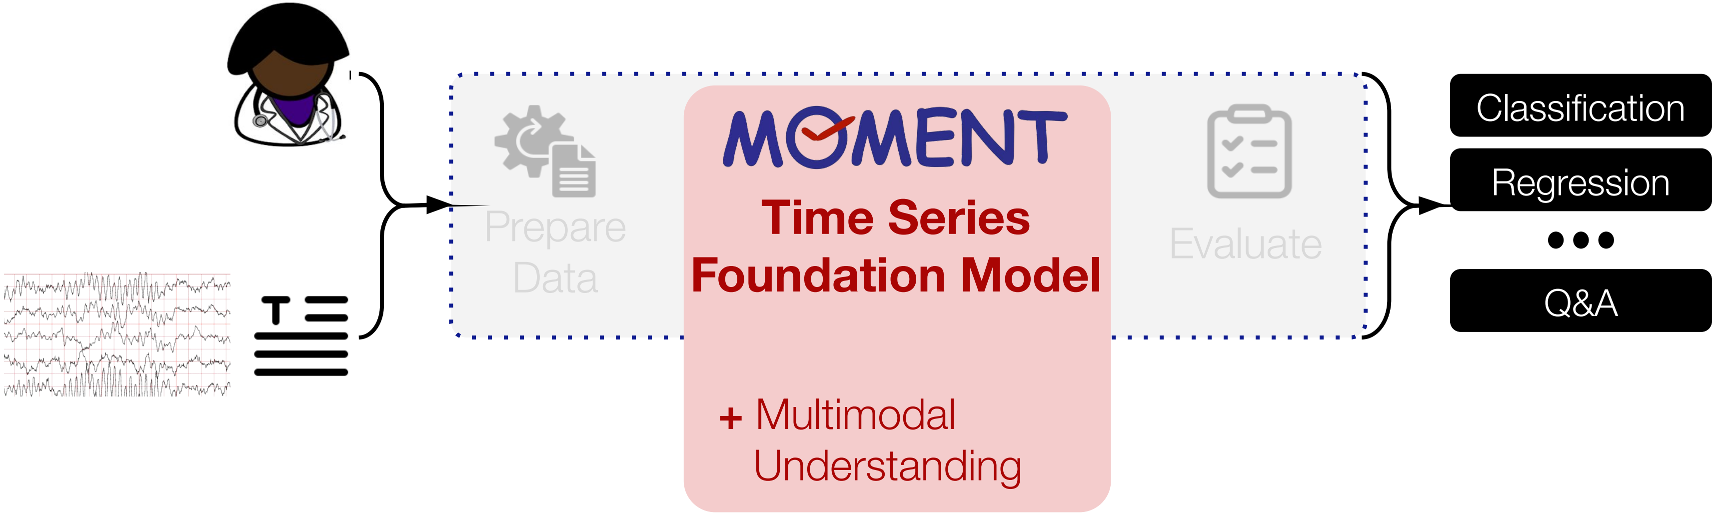

Goswami, M., Szafer, K., Choudhry, A., Cai, Y., Li, S., & Dubrawski, A. (2024, July). MOMENT: A Family of Open Time-series Foundation Models. In International Conference on Machine Learning (pp. 16115-16152). PMLR.

What is a Time Series Foundation Model?

ChatGPT

- Multiple Tasks
- Multimodality
- Fewer Labels
- Cost Efficiency

Classification

Forecasting

Anomaly Detection

Imputation

Building block for diverse time series analysis tasks

Effective out-of-the-box (Zero-shot performance)

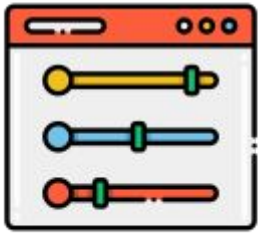

Tunable using in-distribution and task-specific data

Missing Ingredients of a Time Series Foundation Model

as of February 2024

	Data	Modeling (Model & Training Objective)	Benchmarking
LLM	The Pile	Decoder-only Transformer & Next token Prediction	Massive Multitask Language Understanding
Time Series Foundation Model	**Time Series Pile** 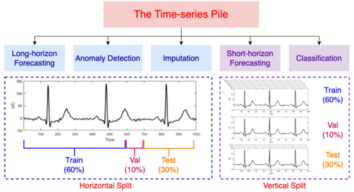	MOMENT 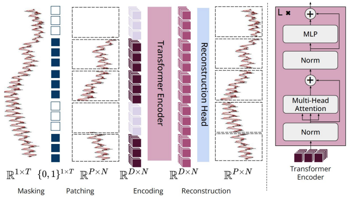	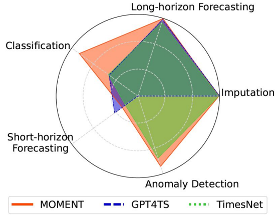
Roadblocks	No Large, Cohesive, Public Dataset	Time series characteristics	Task specific benchmarks, unfit to evaluate FMs

This Talk

Design Space of a Time Series Foundation Model

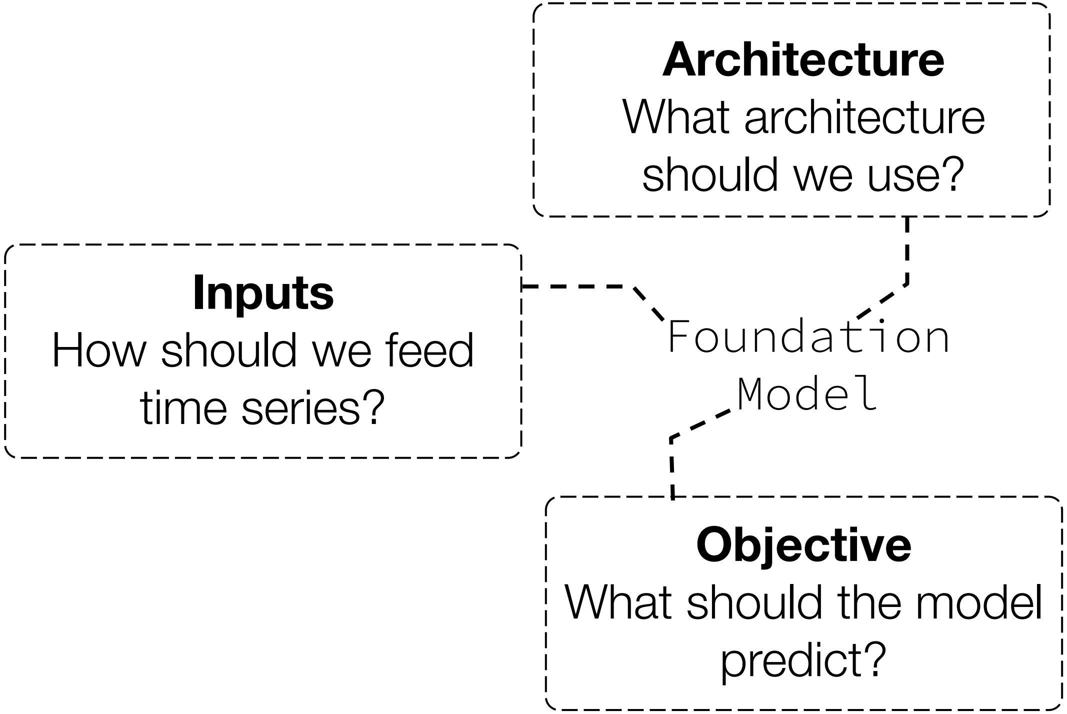

What Type of Architecture Should We Use?

Architecture

Encoder-only Transformer

Inputs

Disjoint time series subsequences (patches)

Foundation Model

Objective

What should the model predict?

Intuition

Primary choice for LLMs

Leverage advances in LLMs

Bidirectional attention important for imputation

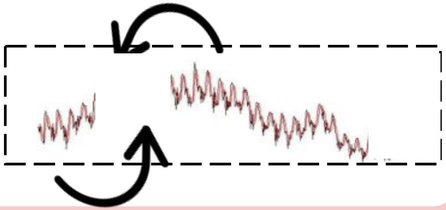

How Should We Feed Time Series Data?

Architecture

What architecture should we use?

Inputs

Disjoint time series subsequences (patches)

Foundation Model

Objective

What should the model predict?

Intuition

Masking patches (vs. timesteps) is efficient & it forces the model to learn good representations

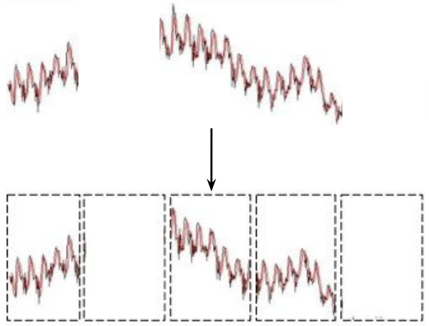

What Should the Model Predict?

Architecture

Encoder-only Transformer

Inputs

Disjoint time series subsequences (patches)

Foundation Model

Objective

Masked Time Series Reconstruction

Intuition

All tasks are instances of masked time series reconstruction

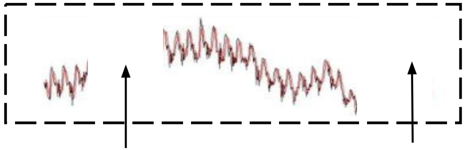

Imputation Forecasting

One Model, Multiple Tasks, Datasets & Domains

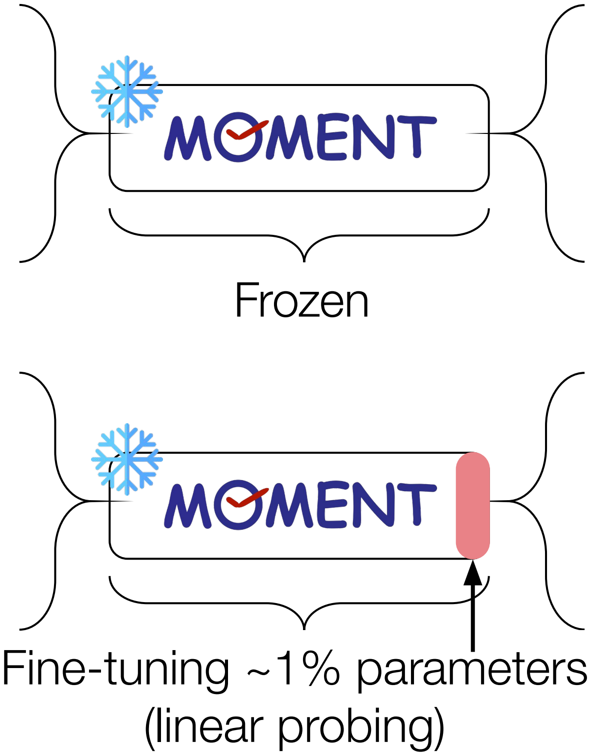

On par with statistical baselines

Large-scale pre-training enables MOMENT to accelerate the development of good time series ML models

Competitive with state-of-the-art

Part 1: Building Capable Foundation Models

+ Long-Context

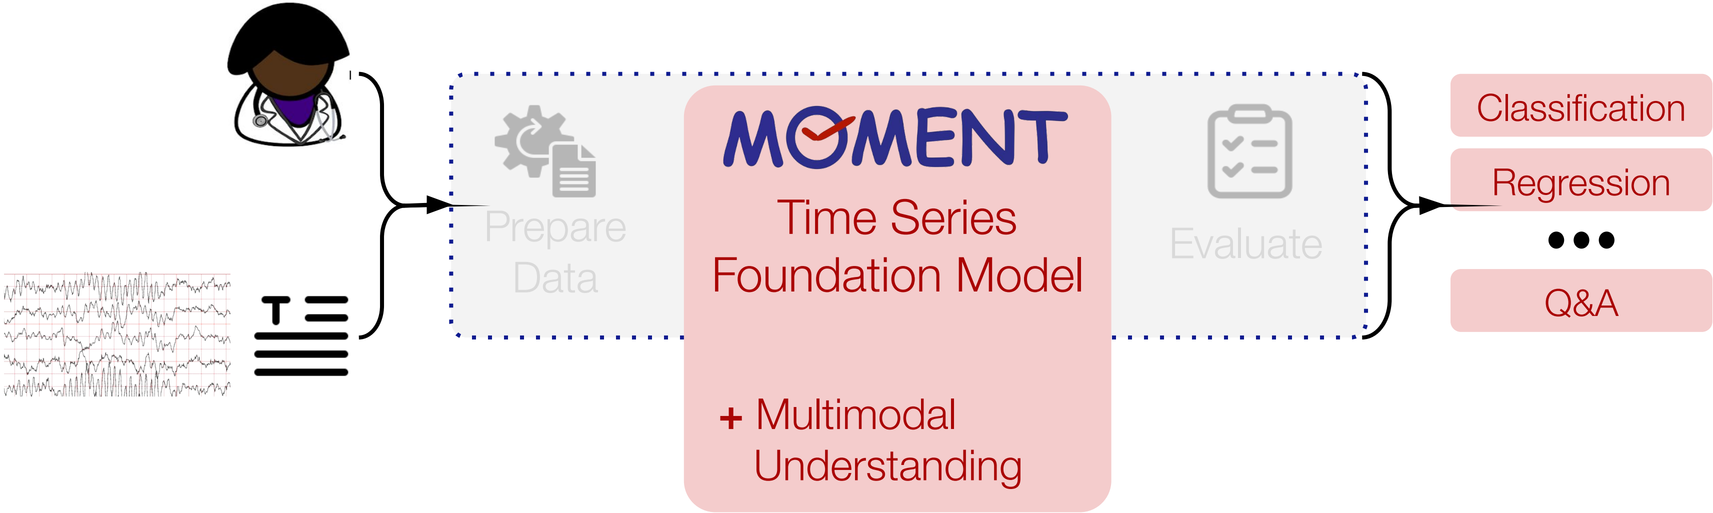

Żukowska, N., Goswami, M., Wiliński, M., Potosnak, W., & Dubrawski, A. Towards Long-Context Time Series Foundation Models With A Handful Of Additional Parameters. In NeurIPS 2024 Workshop on Fine-Tuning in Modern Machine Learning: Principles and Scalability.

Long Context Time Series Foundation Models

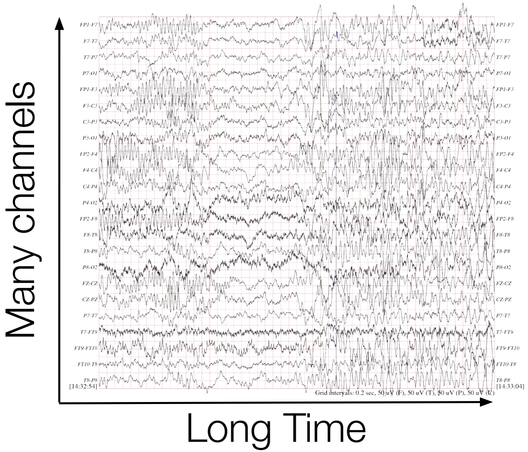

Long Context models should be able to model both long and **multivariate** time series

Focus on Multivariate Time Series Foundation Models

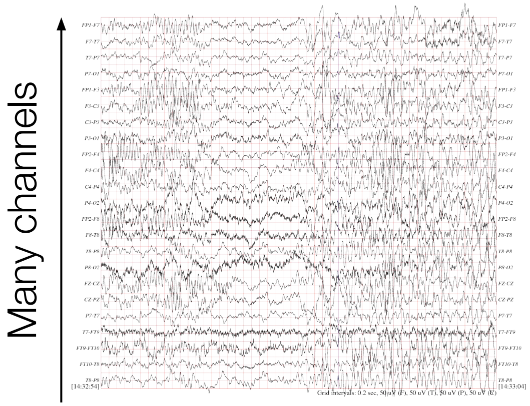

Long Context models should be able to model both long and **multivariate** time series

Multivariate models can be easily extended to model long time series

The Ideal Approach
Minimal changes to the architecture and training procedure

Key Challenge
Standard transformers can only model a univariate sequences of limited lengths

Naive Solution
Training with a larger context length (# of input patches) does not scale due quadratic complexity of attention

Key Idea: Introduce Memory in Transformers

Recurrent Neural Networks

$$\mathrm{RNN}(x_t, h_{t-1})$$

Update hidden state as additional tokens are processed

Stateless Transformers

$$O_s = \mathrm{attention}(X_s)$$

Process a sequence of tokens, but don't maintain state

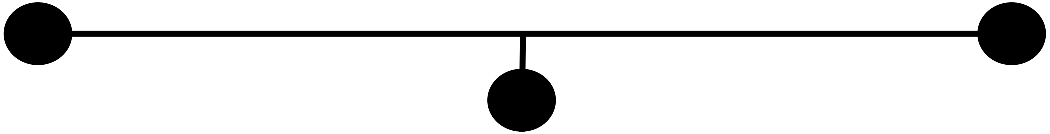

Infini-Attention: **Recurrent Attention Layer**

Attend to sequence of tokens in context, and in the memory

Munkhdalai, T., Faruqui, M., & Gopal, S. (2024). Leave no context behind: Efficient infinite context transformers with infini-attention. arXiv preprint arXiv:2404.07143, 101.

Infini-Attention: Recurrent Attention Layer

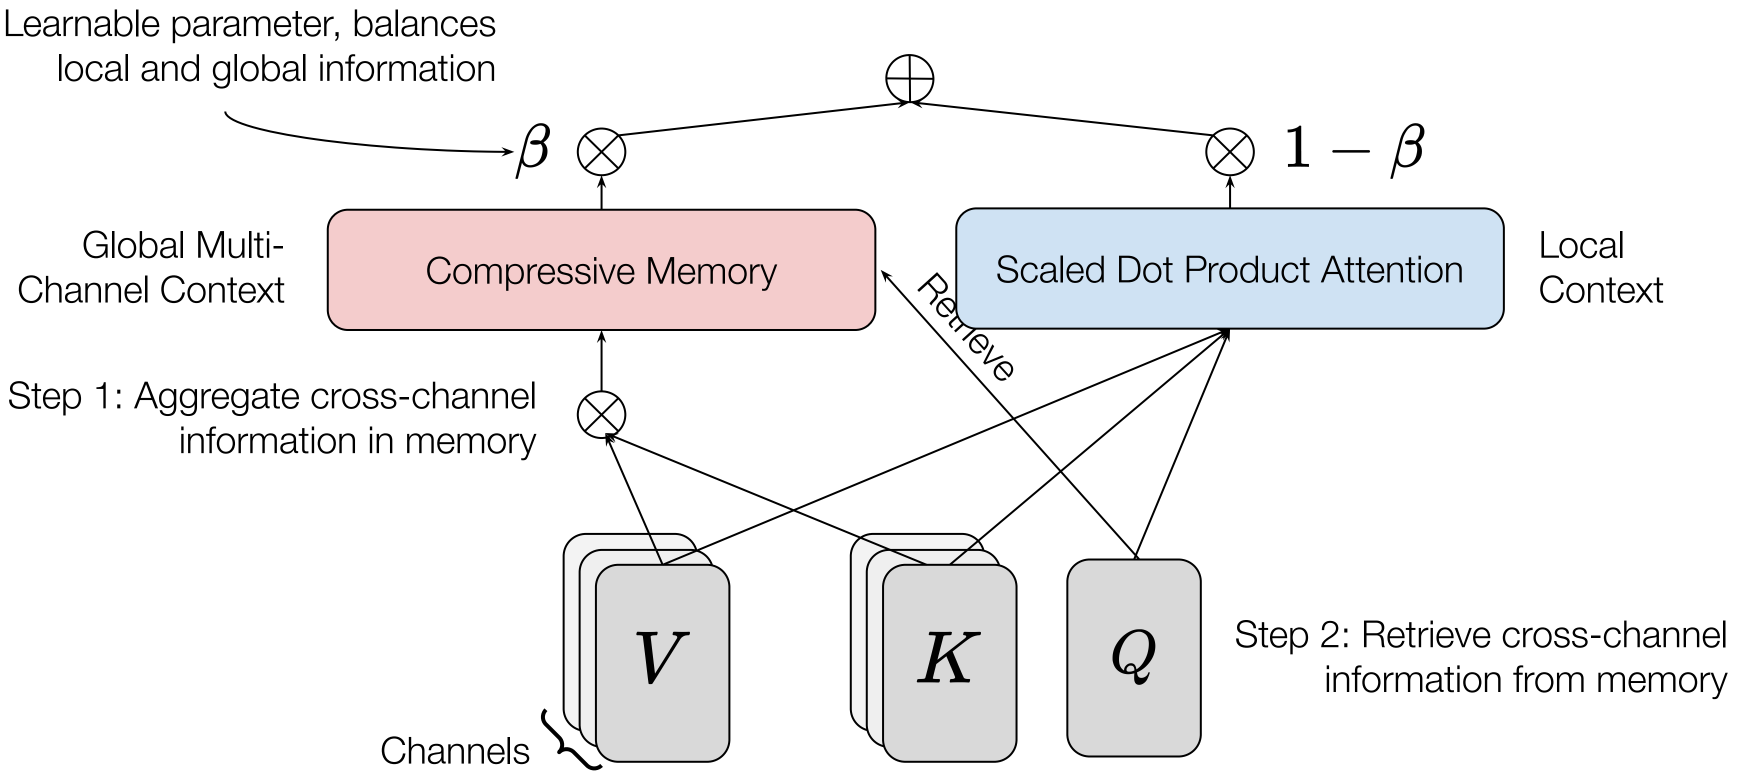

Infini-MOMENT has Nice Properties

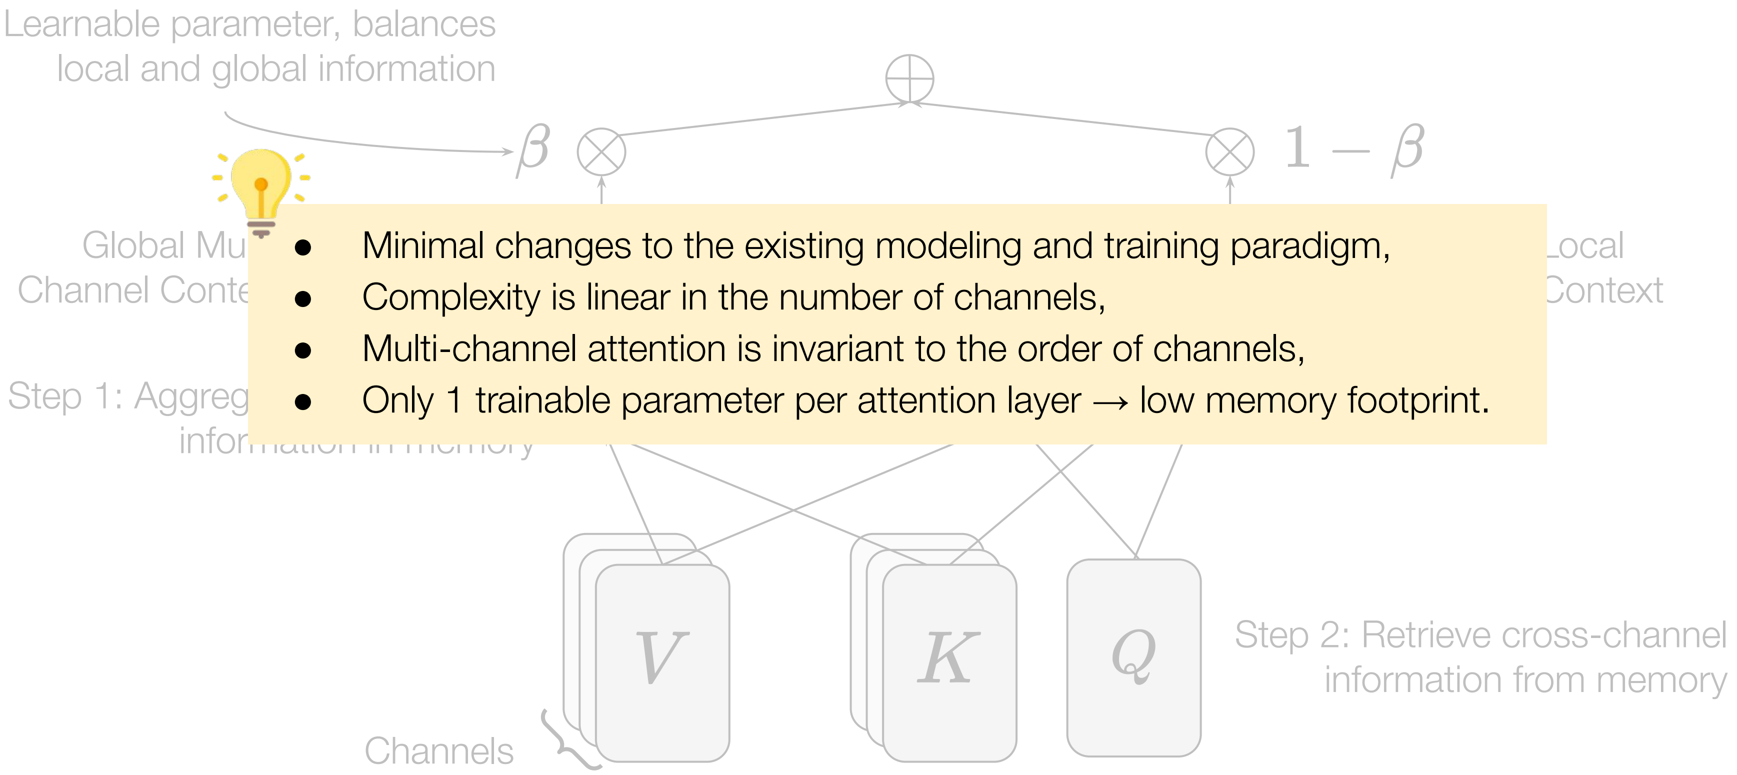

- Minimal changes to the existing modeling and training paradigm,
- Complexity is linear in the number of channels,
- Multi-channel attention is invariant to the order of channels,
- Only 1 trainable parameter per attention layer → low memory footprint.

Infini-MOMENT Outperforms MOMENT on some Multivariate Tasks

MSE (↓) on Multivariate Forecasting Datasets

	Exchange	ETTh1	ETTm1
MOMENT	0.240	0.435	0.340
Infini-MOMENT	**0.232**	**0.416**	**0.333**

Infini-MOMENT outperforms MOMENT with only 20 additional parameters

- Fine-tuning beta parameters improves performance,
- Multivariate modeling doesn't always improve performance.

Part 1: Building Capable Foundation Models

+ Long-Context

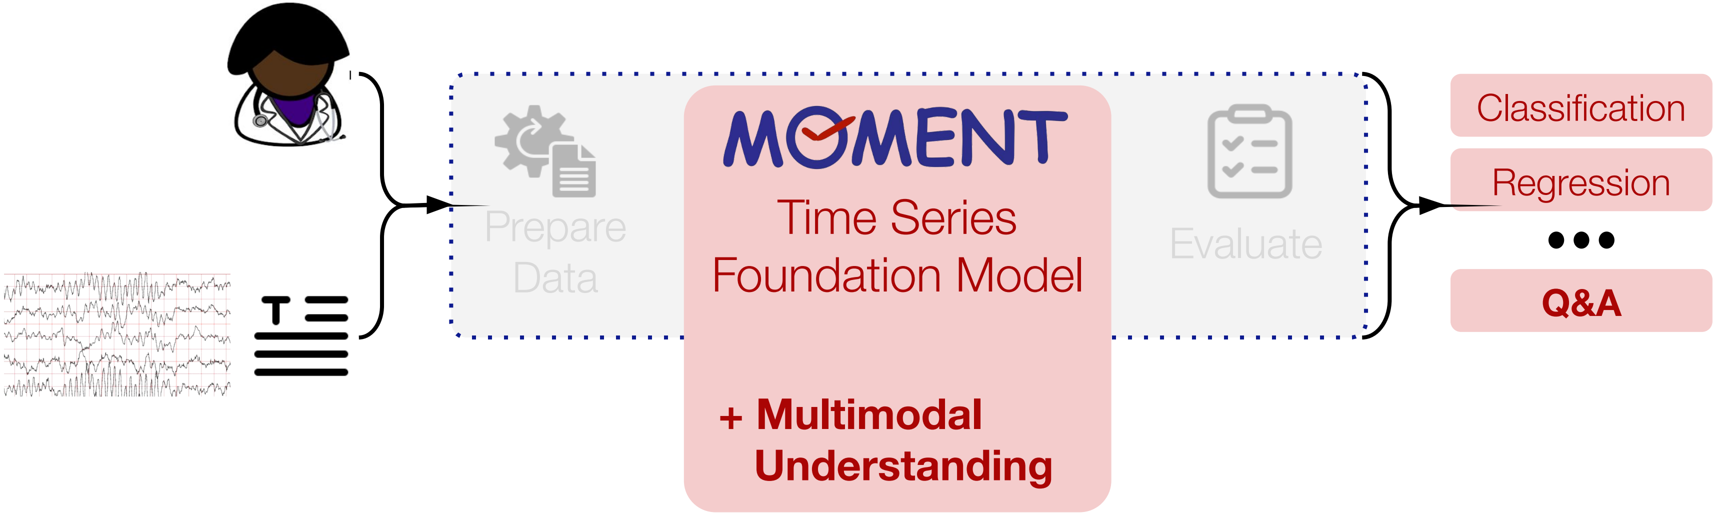

Cai, Y., Goswami, M., Choudhry, A., Srinivasan, A., & Dubrawski, A. (2023). Jolt: Jointly learned representations of language and time-series. In Deep Generative Models for Health Workshop NeurIPS 2023.

Cai, Y., Srinivasan, A., Goswami, M., Choudhry, A., & Dubrawski, A. (2024, March). JoLT: jointly learned representations of language and time-series for clinical time-series interpretation (student abstract). AAAI International Conference on Artificial Intelligence (Vol. 38, No. 21, pp. 23447-23448).

Multimodal Time Series Foundation Models

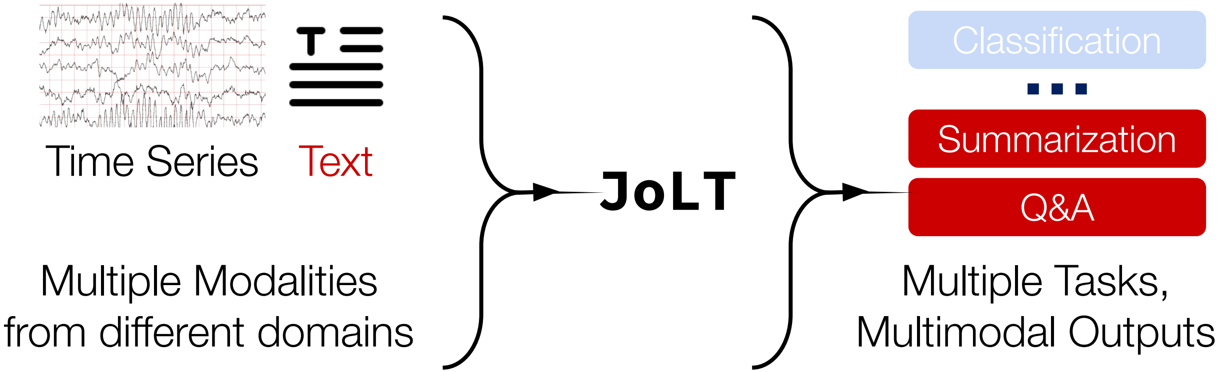

Focus on time series & text-conditioned text generation

Generating Text Conditioned on Time Series is a Useful Task

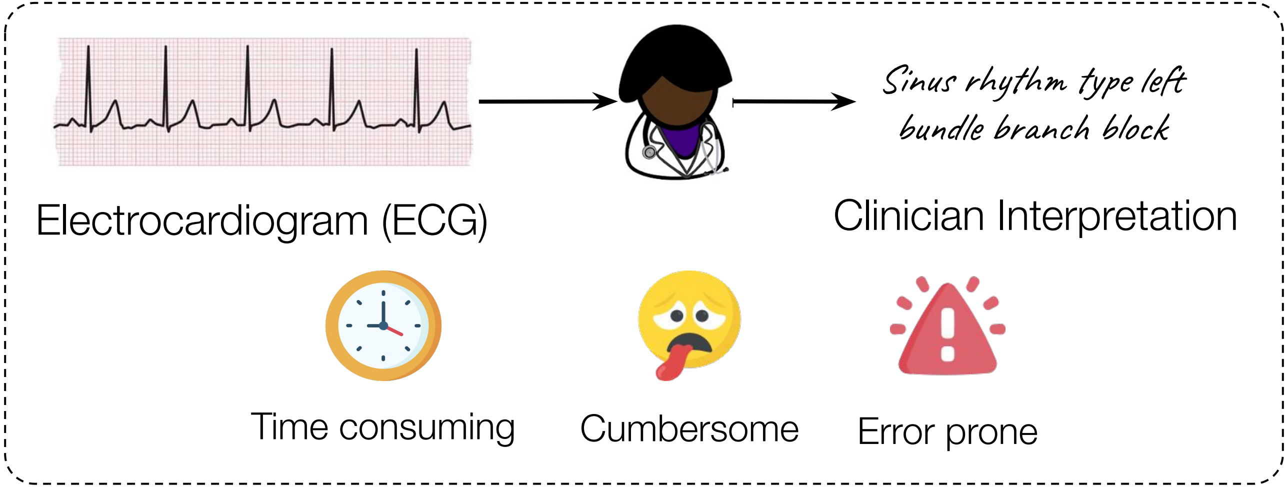

Existing ML Solutions

- ❌ Hand-designed rule-based
- ❌ Model time series as an image or graph

Generating Text Conditioned on Time Series is a Useful Task

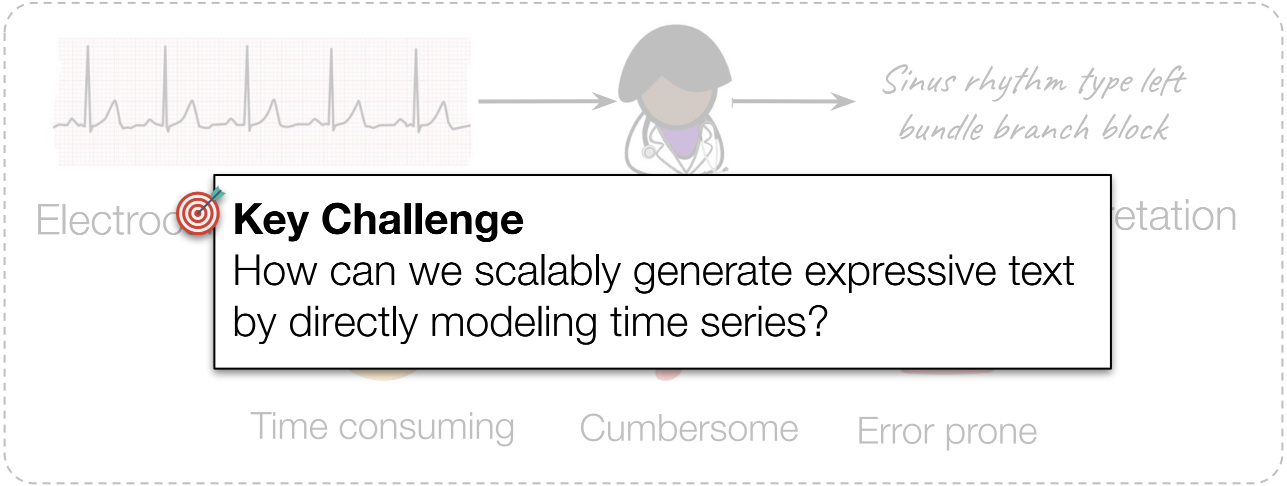

Key Challenge
How can we scalably generate expressive text by directly modeling time series?

Existing ML Solutions

- ✗ Hand-designed rule-based
- ✗ Model time series as an image or graph

JoLT: Align Pre-trained Unimodal Foundation Models

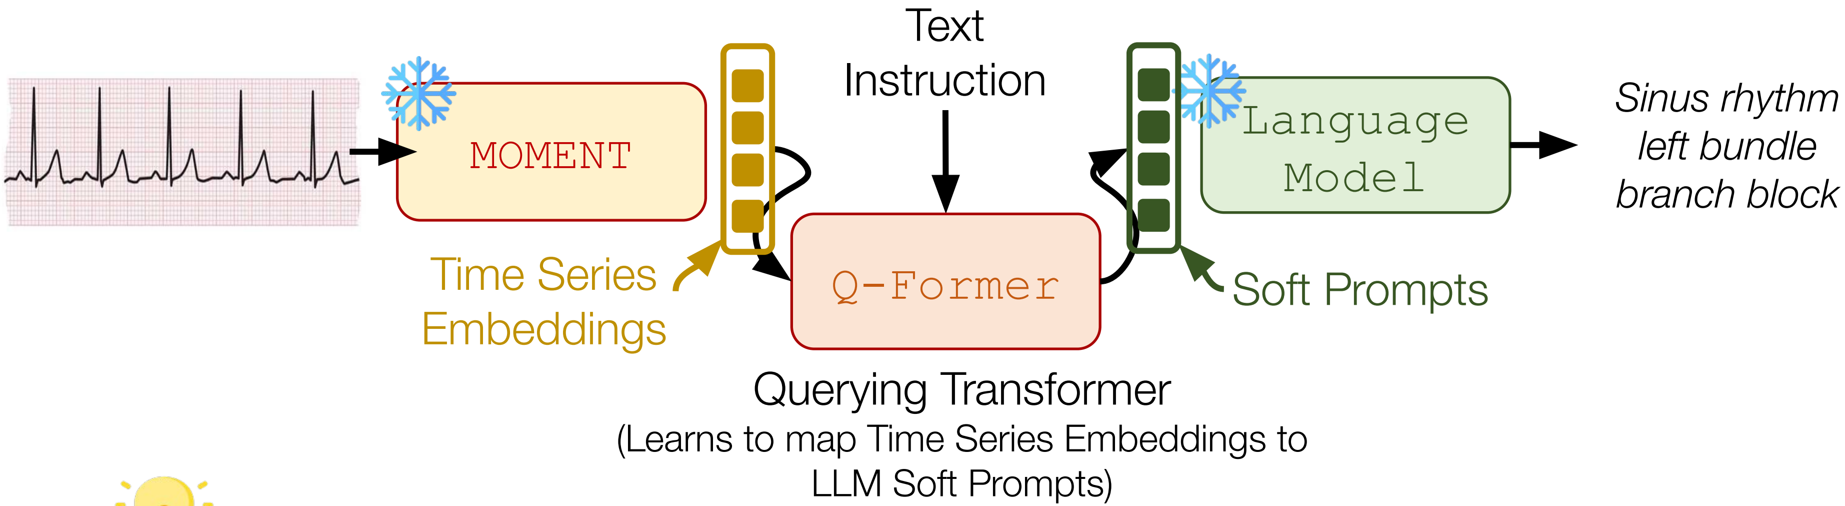

- Leverage frozen pre-trained models, so little training needed,
- Enables conditioning by text (instructions, questions).

❄️ means Frozen

Wenliang, D., et al. "InstructBLIP: Towards General-purpose Vision-Language Models with Instruction Tuning [C]." Advances in Neural Information Processing Systems 36 (2023).

New Time Series to Text Generation Tasks

PTB-XL ECG Arrhythmia Classification Dataset

Paired ECG time series & clinical interpretation

ECG Clinical Summarization **(New Task)**

sinus rhythm position type normal left bundle branch block left hypertrophy possible 4.46 unconfirmed report

Not natural language!
ECG statements in German, translated to English

ECG Question Answering **(New Task)**

Question: Which of these 3 diagnostic classes does this ECG belong to?
Options: (a) Normal ECG, (b) ST/T Change, (c) Hypertrophy
Answer: (a) Normal ECG

Wagner, P., Strodthoff, N., Bousseljot, R. D., Kreiseler, D., Lunze, F. I., Samek, W., & Schaeffter, T. (2020). PTB-XL, a large publicly available electrocardiography dataset. Scientific data, 7(1), 1-15.

Time Series Should be Modeled Explicitly

ECG Clinical Summarization (New Task)

sinus rhythm position type normal left bundle branch block left hypertrophy possible 4.46 unconfirmed report — **Ground Truth**

sinus rhythm left type left bundle branch block left hypertrophy possible 4.46 unconfirmed report — **JoLT (Ours)**

sinus rhythm. normal ecg — **BLIP-2**

BLIP-2 takes an image of ECG as input

- Time series should be modeled explicitly.
- We are limited by the availability of paired data and not technology

Part 2: Improving Understanding of Foundation Models

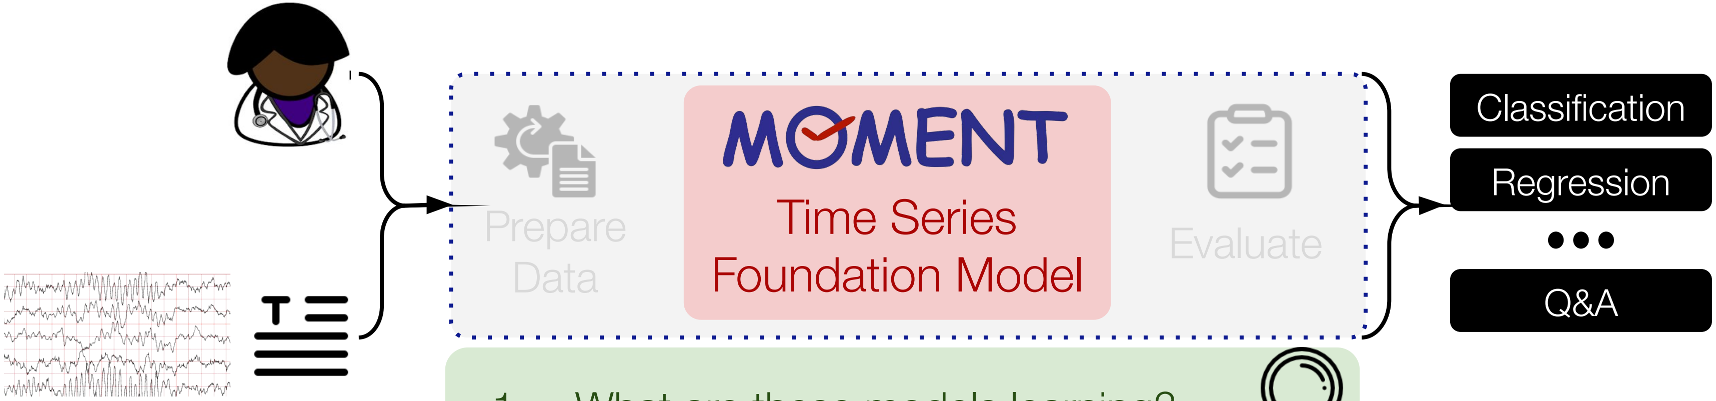

1. What are these models learning?
2. Can we steer these models?
3. Can these models reason?
4. Can they understand basic time series concepts?

1. **What are These Models Learning?** Representational Similarity

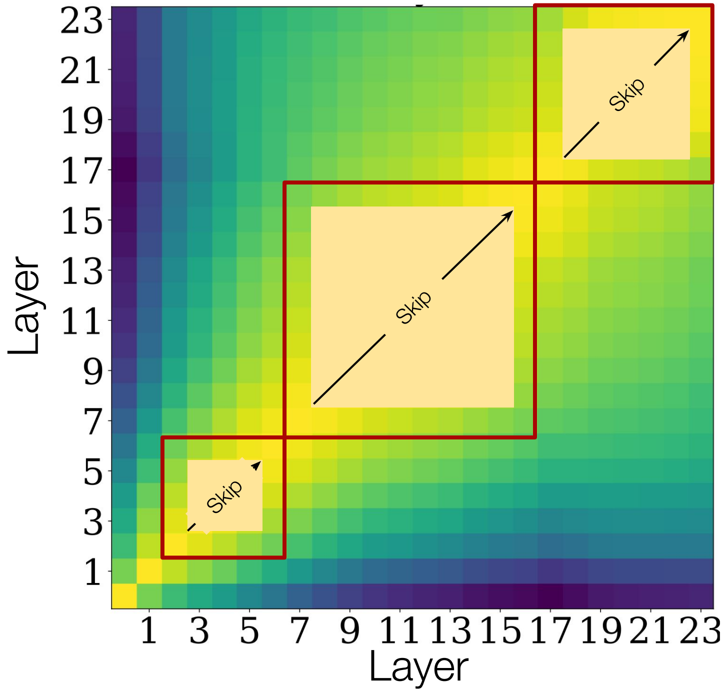

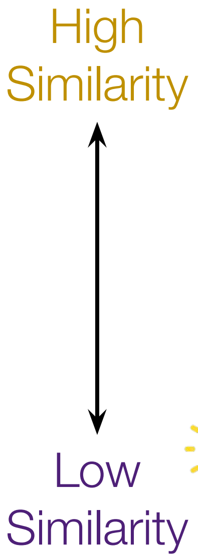

MOMENT pairwise similarity of hidden activations, measured using **Centered Kernel Alignment**

MSE (↓) of Fine-Tuned Vanilla & Pruned Models on Forecasting Datasets

	ETTm2	Illness	Weather
Vanilla	**0.171**	3.260	0.153
Pruned	0.173	**2.981**	**0.152**

Marginal reduction in performance, but **50% faster** inference time

TSFMs learn redundant representations. Reducing redundancy improves speed, while maintaining accuracy.

Wiliński, M., Goswami, M., Żukowska, N., Potosnak, W., & Dubrawski, A. (2024). Exploring Representations and Interventions in Time Series Foundation Models. In NeurIPS'24 Workshop on Foundation Model Interventions.

1. **What are These Models Learning?** Output Embeddings

Frequency | **Trend**

Embeddings of penultimate layer projected using PCA

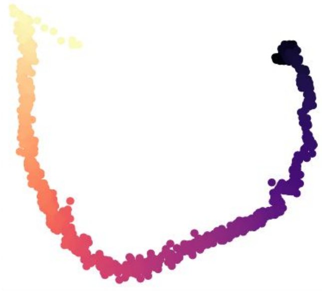

High c

Low c

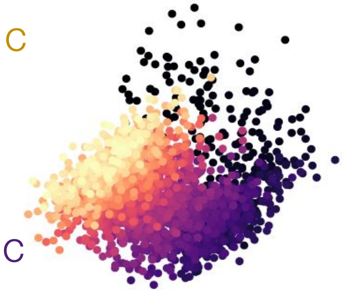

TSFMs capture the language of time series.

Example Time Series

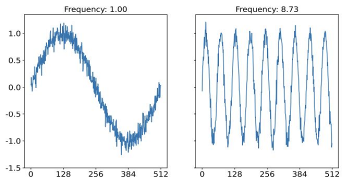

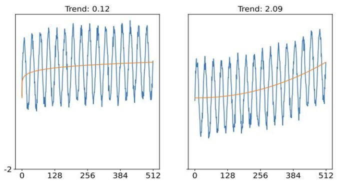

Data Generating Process

$$y = sin(2c\pi x) + \epsilon$$
$$c = [1, 32]$$

$$y = x^c + sin(32\pi x) + \epsilon$$
$$c = [{}^1/_8, 8]$$

Goswami, M., Szafer, K., Choudhry, A., Cai, Y., Li, S., & Dubrawski, A. (2024, July). MOMENT: A Family of Open Time-series Foundation Models. In International Conference on Machine Learning (pp. 16115-16152). PMLR.

1. **What are These Models Learning?** All Embeddings

Frequency

Linear separability between concepts across layers & patches

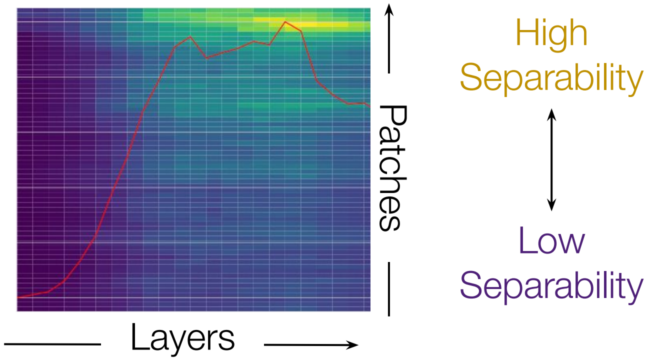

Example Time Series

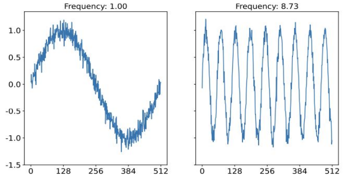

High vs Low Frequency

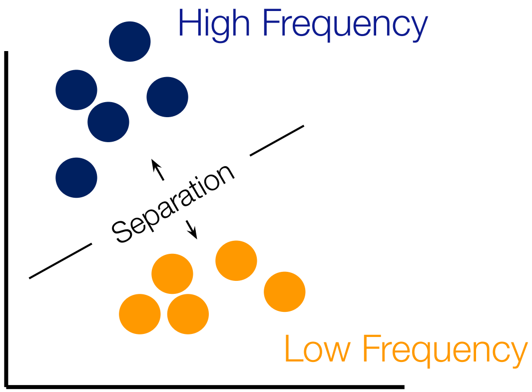

Embeddings at layer *l* & patch *p*

Concepts are learned in the specific parts of the model.

Wiliński, M., Goswami, M., Żukowska, N., Potosnak, W., & Dubrawski, A. (2024). Exploring Representations and Interventions in Time Series Foundation Models. In NeurIPS'24 Workshop on Foundation Model Interventions.

2. Can we Steer These Models?

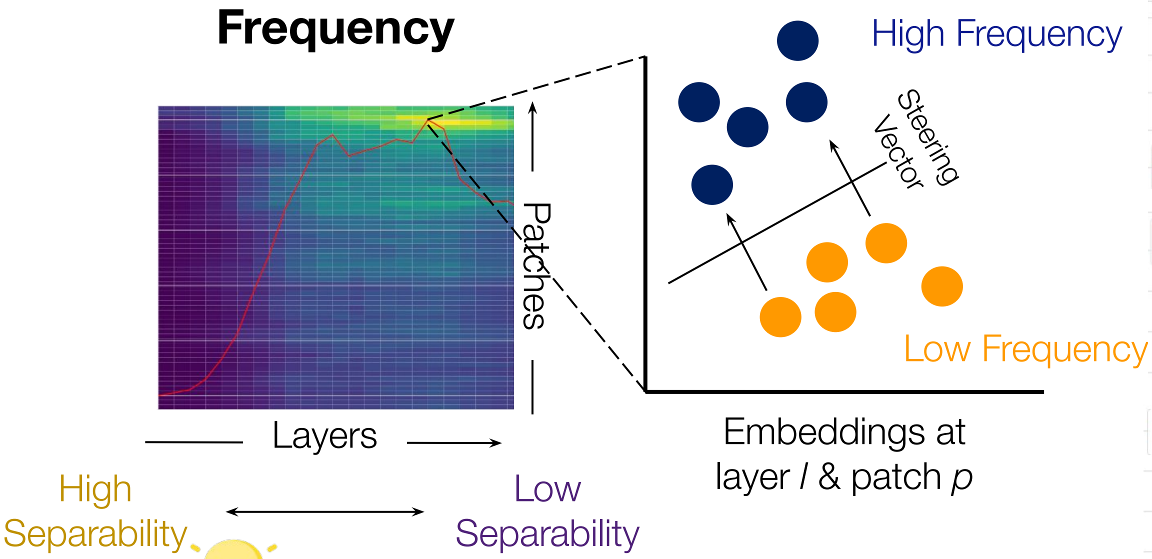

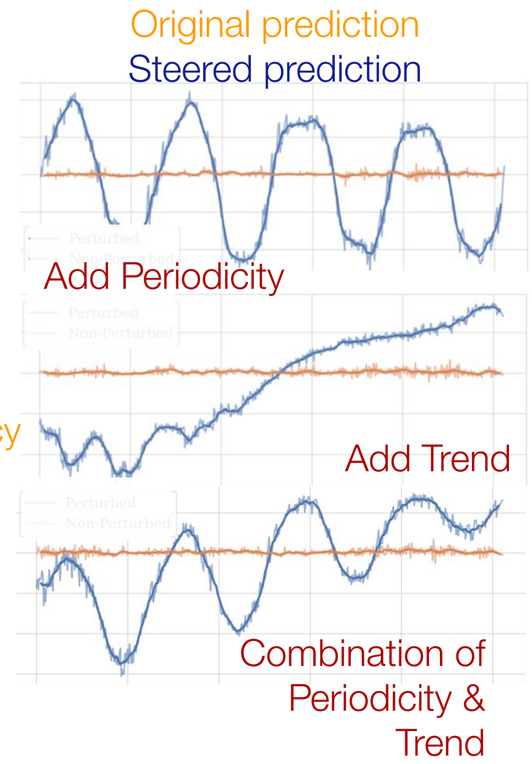

TSFMs predictions can be steered, during inference, without retraining.

Marks et al. "The geometry of truth: Emergent linear structure in large language model representations of true/false datasets." arXiv preprint arXiv:2310.06824 (2023).

Wiliński, M., Goswami, M., Żukowska, N., Potosnak, W., & Dubrawski, A. (2024). Exploring Representations and Interventions in Time Series Foundation Models. In NeurIPS'24 Workshop on Foundation Model Interventions.

Many Practical Applications of Steering

Normal to Abnormal Heartbeat

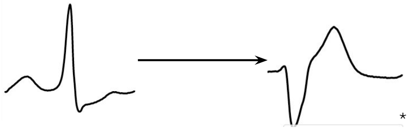

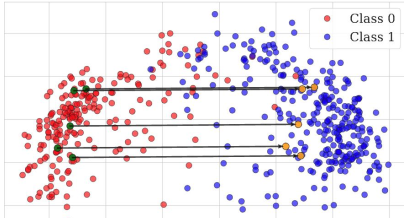

- Introduce Inductive Bias
- Generate Synthetic Data

TSFMs predictions can be steered during inference, improving performance without retraining.

* These images are for illustrative purposes only!

3. Do These Models Reason?

Do Time Series Foundation Models simply memorize training patterns?
Or do they **reason** about patterns?

Compositional Reasoning

Learn simple concepts during training.
Combine learned concepts during inference.

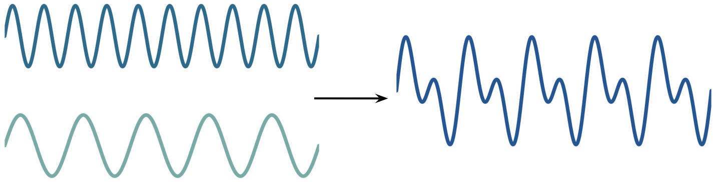

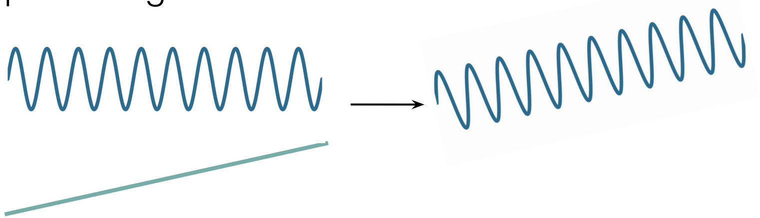

Simple time series (single-frequency)	Complex time series (multi-frequency)	Simple time series with seasonality **or** trend	Complex time series with seasonality **and** trend
Training Set →	Test Set	Training Set →	Test Set

Potosnak, W., Challu, C*., Goswami, M*., Olivares, K. G., Wiliński, M., Żukowska, N., & Dubrawski, A. (2025). Investigating Compositional Reasoning in Time Series Foundation Models. In NeurIPS'24 Workshop on Time Series in the Age of Large Models.

3. Do These Models Reason?

Do Time Series Foundation Models simply memorize training patterns?
Or do they **reason** about patterns?

Compositional Reasoning

Learn simple concepts during training.
Combine learned concepts during inference.

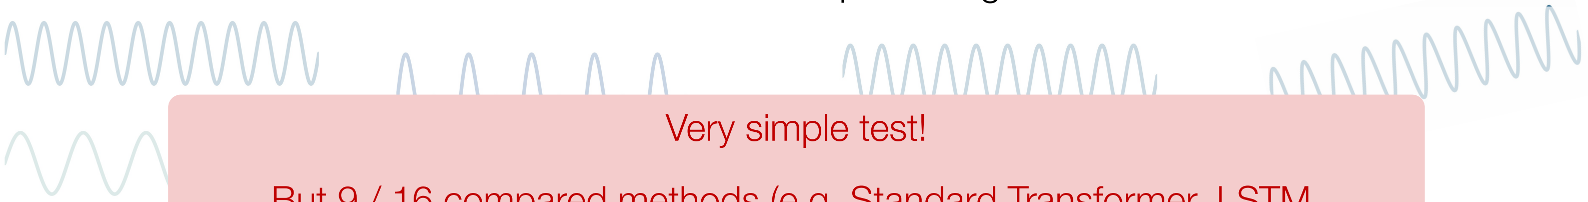

Very simple test!

But 9 / 16 compared methods (e.g. Standard Transformer, LSTM, Linear Models) showed no signs of compositional reasoning!

Simple tim
(single-fre

e series with
and trend

Training Set → Test Set

Training Set → Test Set

3. Do These Models Reason?

Train on **simple** time series with seasonality **or** trend

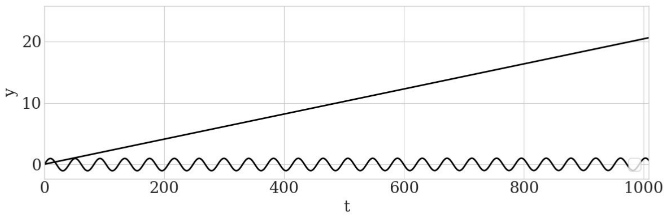

Test on **complex** time series with seasonality **and** trend

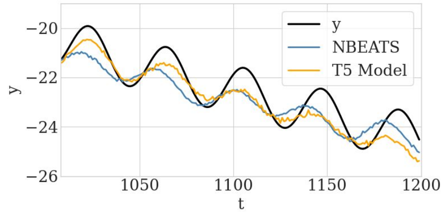

Models never trained on negative trends!

- Patched-transformers (e.g. T5) and MLP-based (e.g NHITs) showed sparks of compositional reasoning
- Input patching unlocks reasoning in Transformer-based TSFMs

4. Can these Models Understand Basic Concepts?

LLMs are used for Time Series Tasks

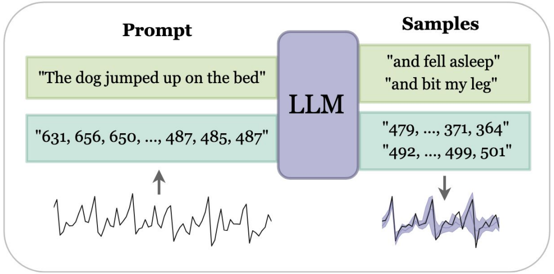

Zero-shot forecasting

`TimeSeriesExam` tests LLMs on 5 core understanding categories

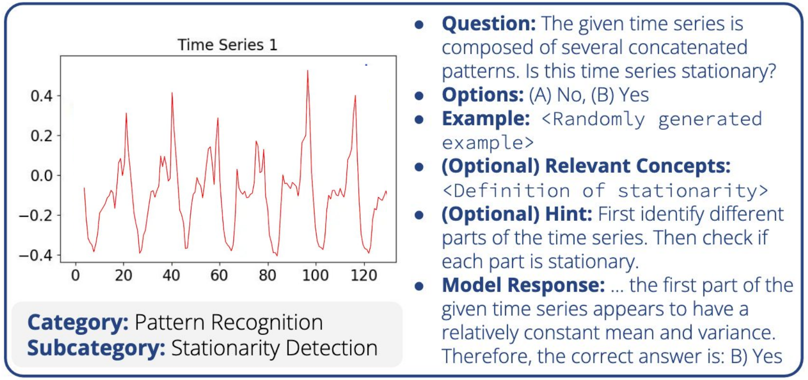

700 questions

Procedurally generated using templates

Refined using Item Response Theory

Gruver, N., Finzi, M., Qiu, S., & Wilson, A. G. (2023). Large language models are zero-shot time series forecasters. Advances in Neural Information Processing Systems, 36, 19622-19635.

Cai, Y., Choudhry, A., Goswami, M., & Dubrawski, A. TimeSeriesExam: A Time Series Understanding Exam. In NeurIPS Workshop on Time Series in the Age of Large Models.

Randomly Generated Questions are Improved using Item Response Theory

Each question can maximally distinguish the abilities of candidate LLMs

2 parameter logistic model:

$$\mathbb{P}(r_{ij} = 1 \mid a_i, b_i, \theta_j) = \frac{1}{1+e^{-a_i(\theta_j - b_i)}}$$

Probability that LLM *j* responds to question *i* correctly

Question *i*'s difficulty and discrimination ability

LLM *j*'s ability

1–3 rounds of iterative improvement

All candidate LLMs take the exam

Parameters fit using MLE based on LLM responses

Questions with low difficulty & discrimination are dropped

Insights into LLMs

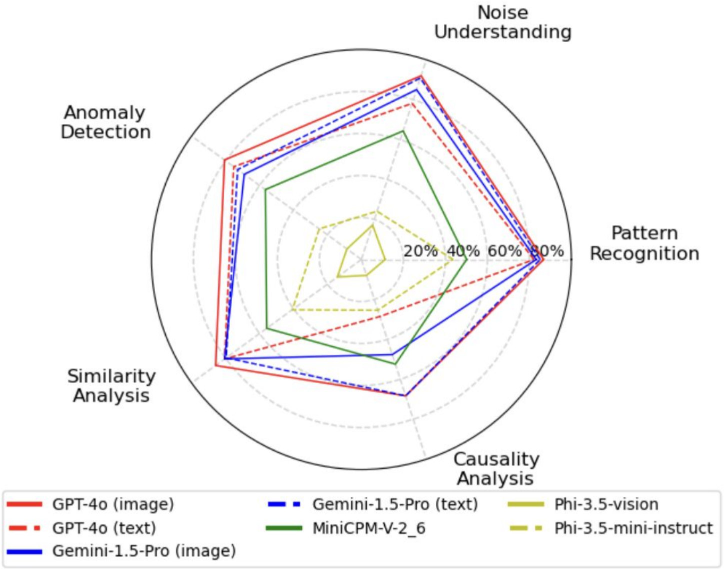

Insights

- LLMs understand basic time series concepts,
- If using LLMs, tokenizing time series as images is better than text.

Benchmarks can be generated at scale

Understanding of Foundation Models Can Help Improve their Performance

Questions	Actionable Insight
What are these models learning?	Redundant representations can be pruned to improve efficiency
Can we steer these models?	Steering TSFM predictions can imbue missing domain knowledge
Can these models reason?	Input patching can unlock compositional reasoning
Can they understand basic time series concepts?	Targeted model interrogation can reveal capability gaps

Research Impact: Empower Domain Experts with Time Series Intelligence on their Fingertips

Before

Different models for different tasks, trained from scratch

Now

One model, minimal task-specific adaptation

MOMENT

Near Future

LLM agent uses TSFMs & learns from feedback

Empower Domain Experts with Time Series Intelligence at their Fingertips

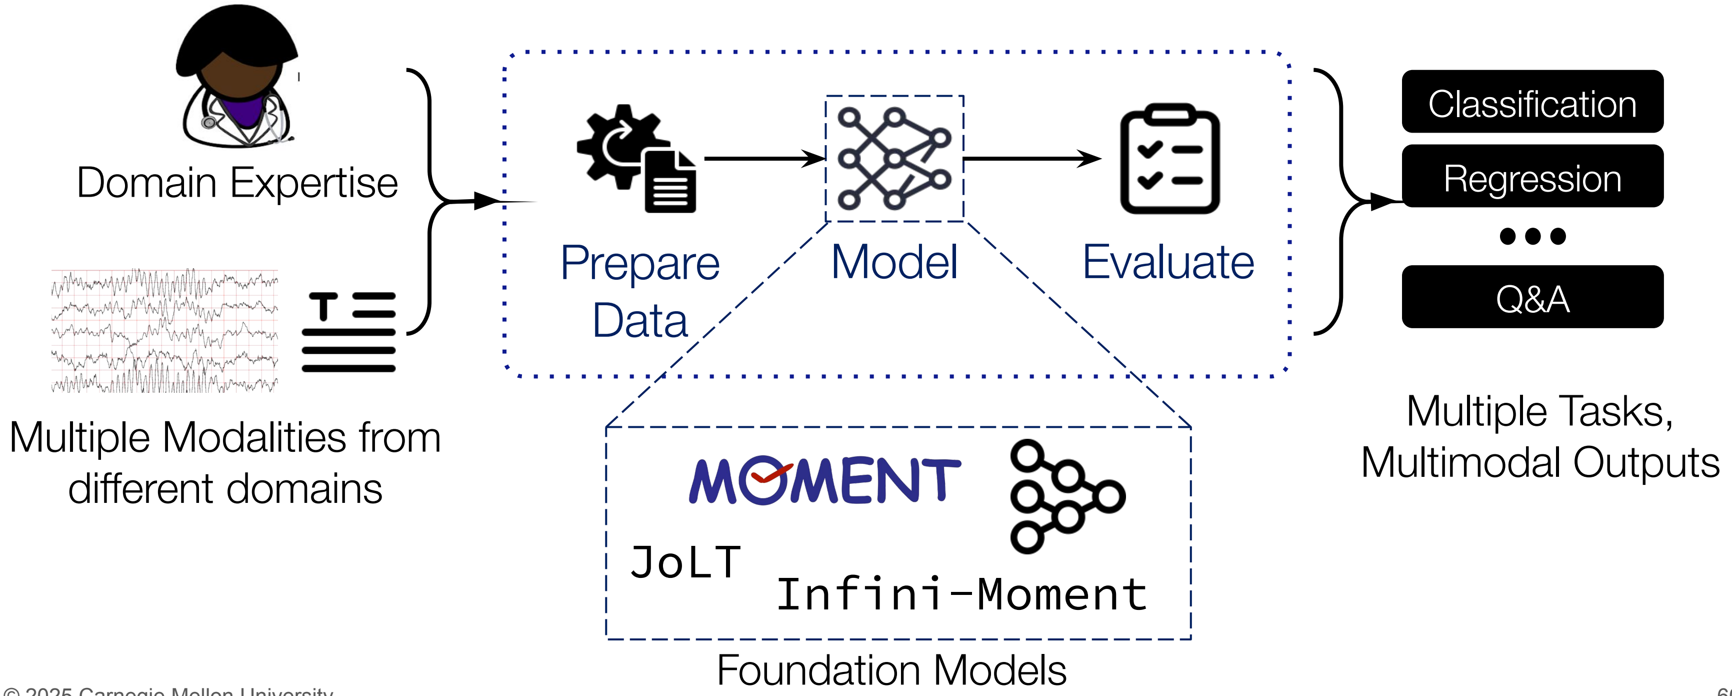

Empower Domain Experts with
Time Series Intelligence at their Fingertips

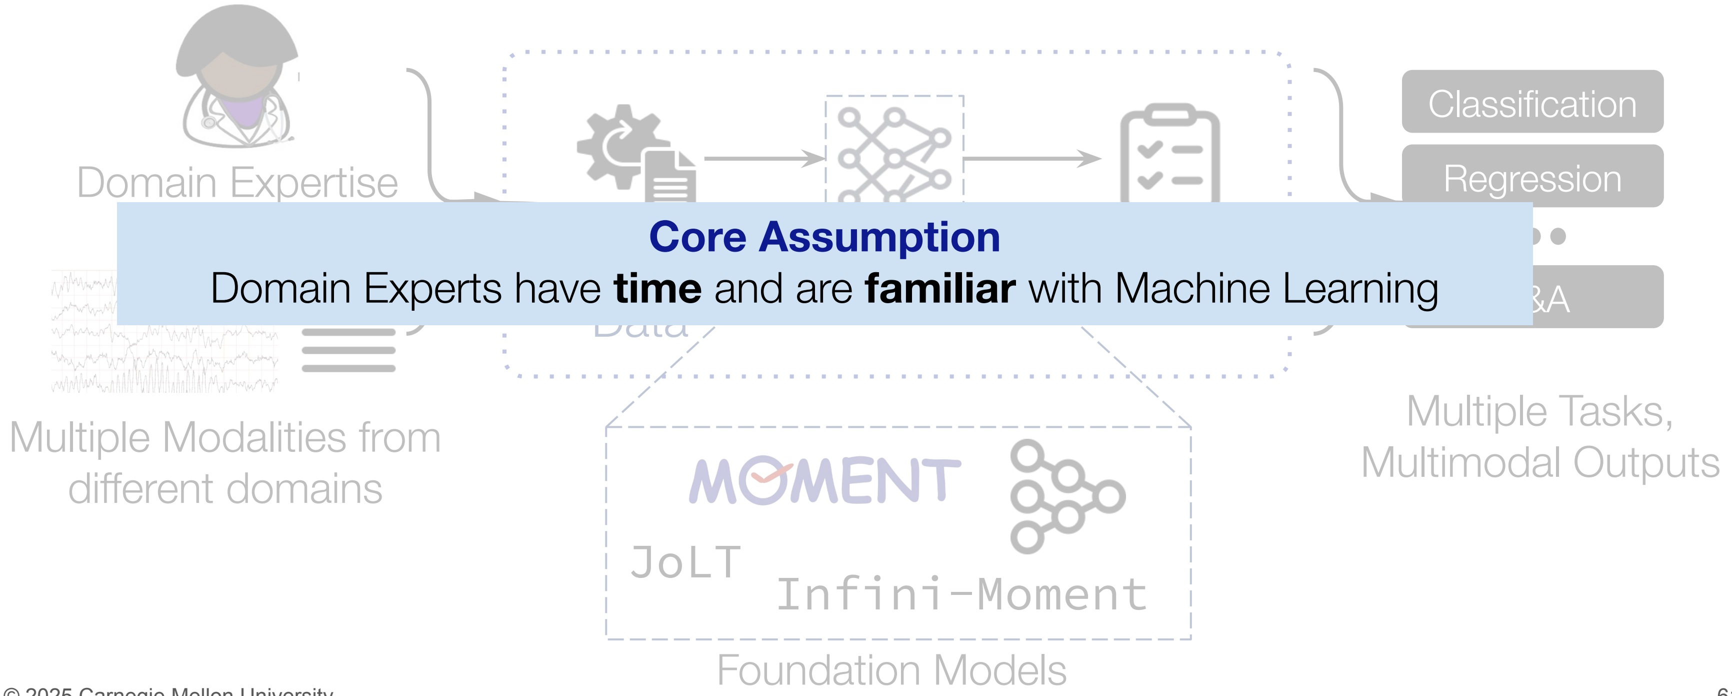

Core Assumption

Domain Experts have **time** and are **familiar** with Machine Learning

LLM Agents can Automate ML Workflows Using Our Tools

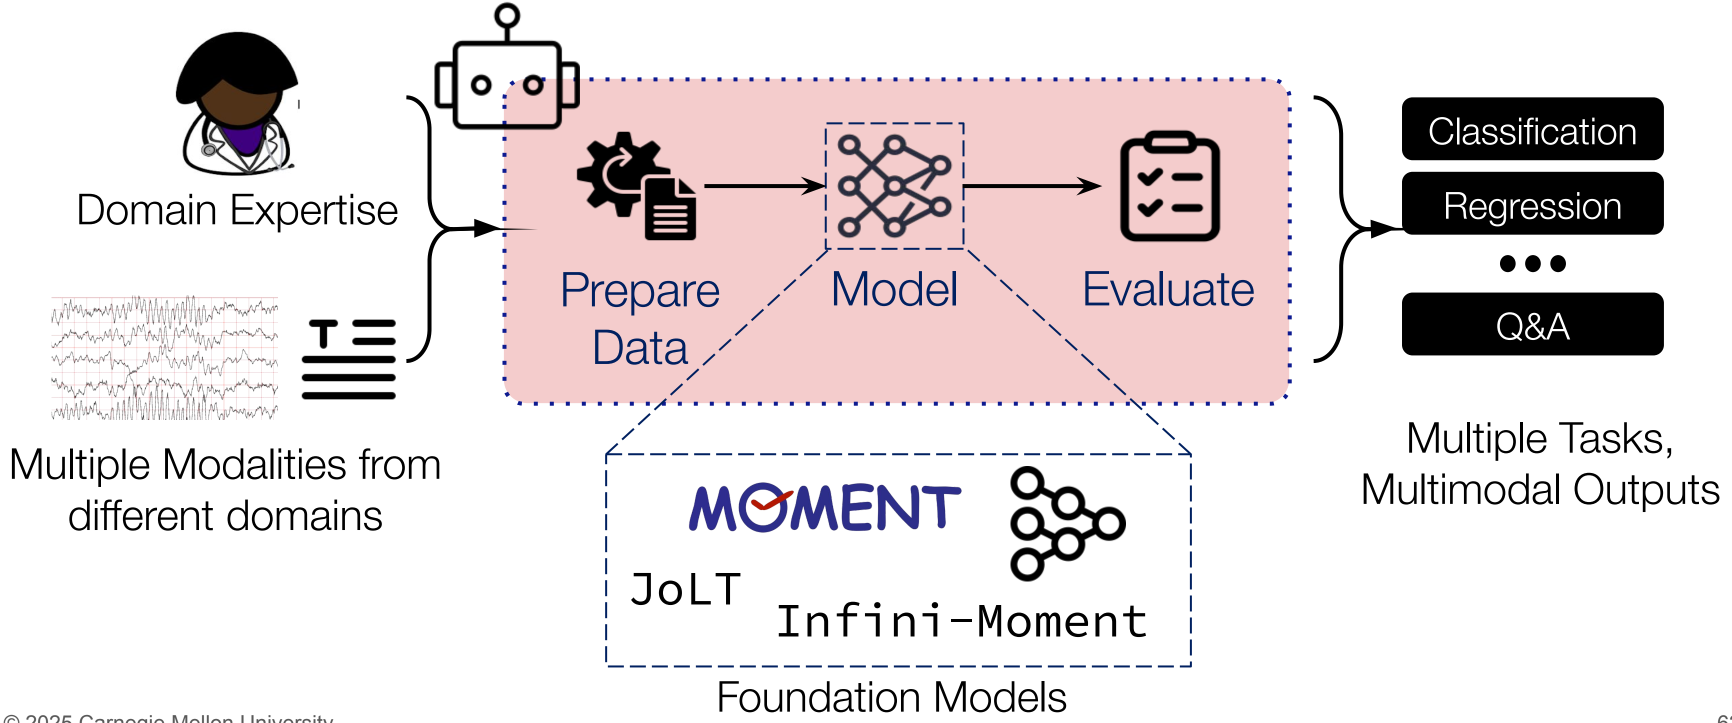

AI Agents are Not Ready Yet for Practical Use

Machine Learning Agents

- Research Agent (Huang, 2024)
- AIDE (Jiang, 2025)
- OpenHands (Wang, 2024)

kaggle

- Minimal ambiguity
- Well structured input/output
- Few time series tasks

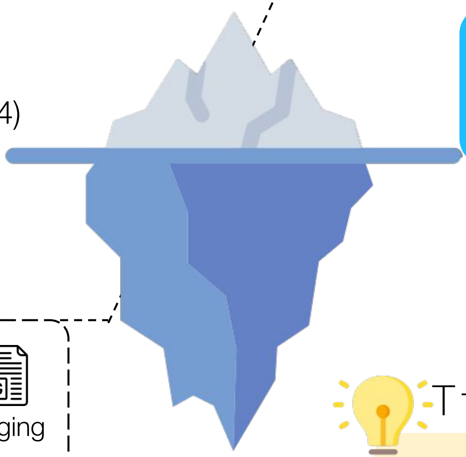

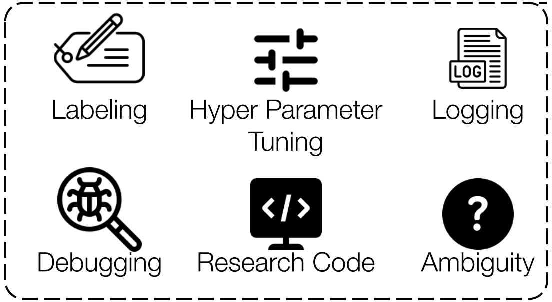

Skills of a Seasoned ML Engineers

TimeSeriesGym

Building Time Series ML Engineering agents will require thinking beyond Kaggle

Cai, Y., Li, X., Goswami, M., Wiliński, M., Welter, G., & Dubrawski, A. (2025). TimeSeriesGym: A Scalable Benchmark for (Time Series) Machine Learning Engineering Agents. arXiv preprint arXiv:2505.13291.

Research on MOMENT Has Been Impactful

Datasets

`Time Series Pile`	Pre-train & Evaluate Multi-task Time Series Foundation Models
`TimeSeriesExam-1`	Evaluate LLMs (and Agents) on Time Series Understanding

10K+ downloads

Research on MOMENT Has Been Impactful

First Open Source, **Multi-task** Model

Models

1.8M+ downloads

500+ stars

Attracted $2M+ in government & industrial funding

Stellar Flare Forecasting
(Zhu et la., 2025)

Machine Fault Diagnosis
(Eldele et la., 2025)

Human Activity Recognition
(Chen et al., 2025)

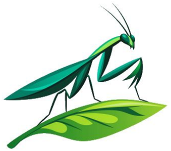

Classification
(Feofanov et la., 2025)

Anomaly Detection
(Liu et la., 2025)

Building Predictive Analytics
(Mulayim et la., 2024;
Dumitru et la., 2024)

Research on MOMENT Has Been Impactful

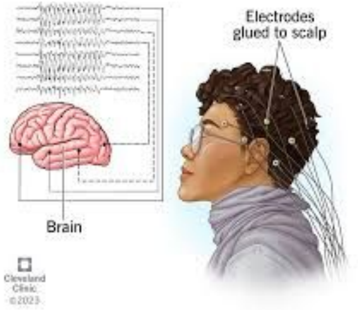

Electroencephalogram (EEG) (Yuan et al., 2025)

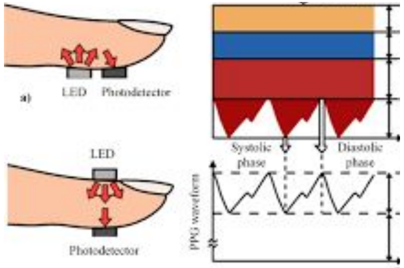

Photoplethysmography (PPG) (Chen et al., 2024)

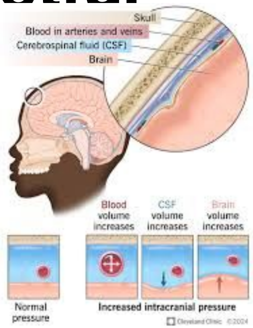

Intracranial Pressure (ICP) (Leeuwen et al., 2025)

MOMENT (Generic)

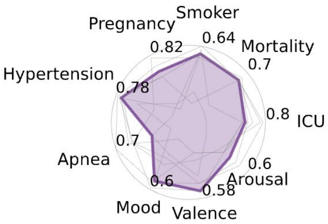

PaPaGei (PPG-specific; Pillai, 2024)

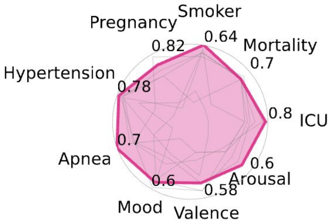

Lower Area → Better

MOMENT without substantial domain-specific training performs well against custom-built domain-specific models.

It Takes a Village to Build a Foundation Model

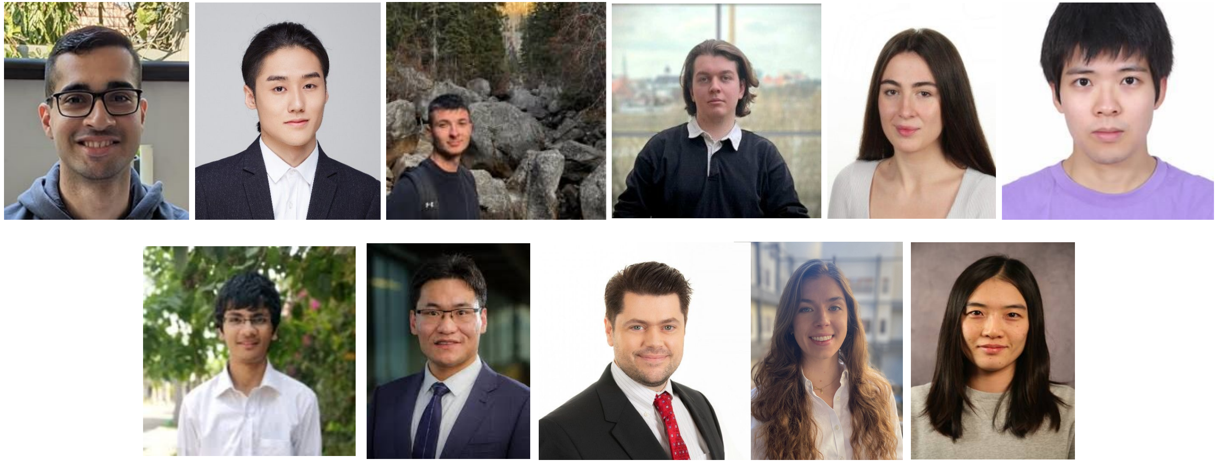

Running MOMENT on Bridges-2

https://www.psc.edu/resources/bridges-2/user-guide/#moment

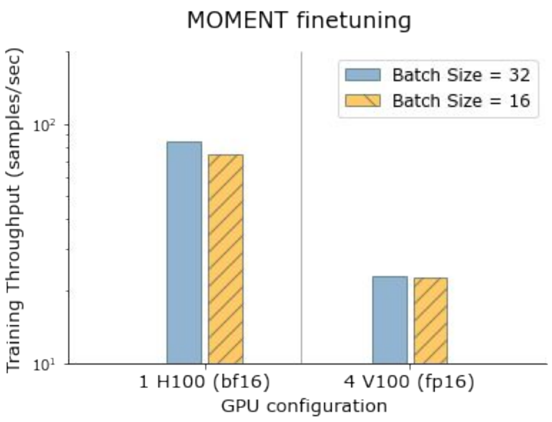

- **It is straightforward to run MOMENT on Bridges-2 with pre-configured environment!**
- **Scan the QR code or click on the link for documentation showing the steps to run through all the tutorials in the MOMENT Github repo. Examples include:**

  Anomaly Detection

  Classification/PTBXL dataset classification

  Forecasting

  Imputation

  Representation learning

  Model Finetuning

Thank You!

Time series intelligence can be substantially **advanced** by