Dynamical System Perspectives of LLMs

Guannan Qu
Associate Professor, CMU

Joint work with Muhammed Ustaömeroğlu (CMU)

European Control Conference
July 10, 2026

LLMs are incredibly capable

ChatGPT

Claude

Gemini

Grok

deepseek

Qwen3

Meta AI
Muse Spark

LLMs are incredibly capable

Chat

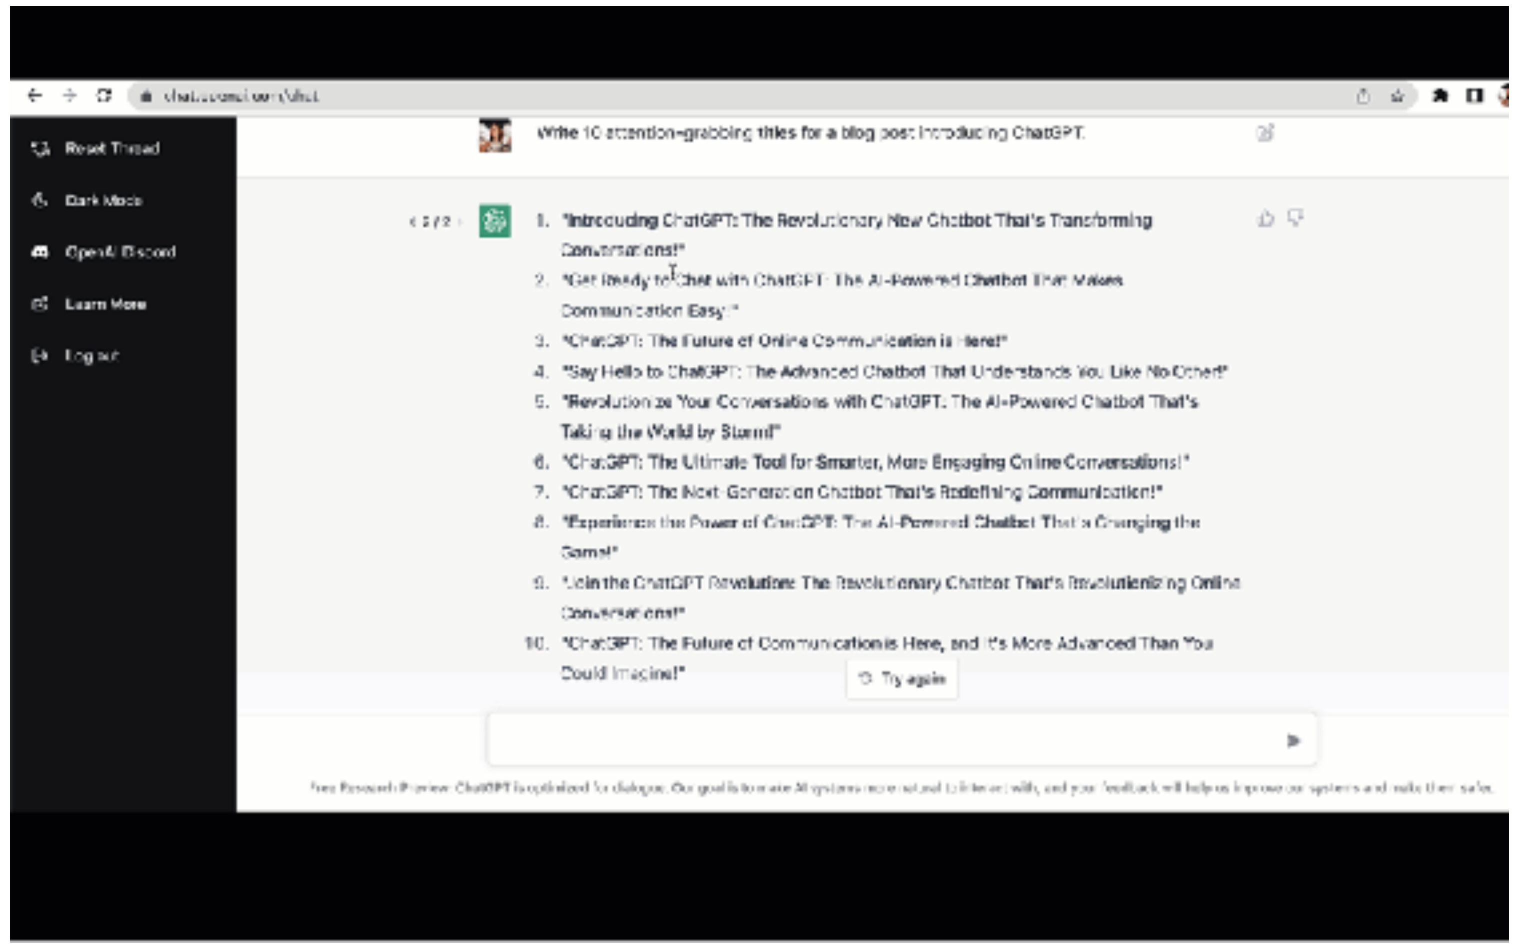

Reason

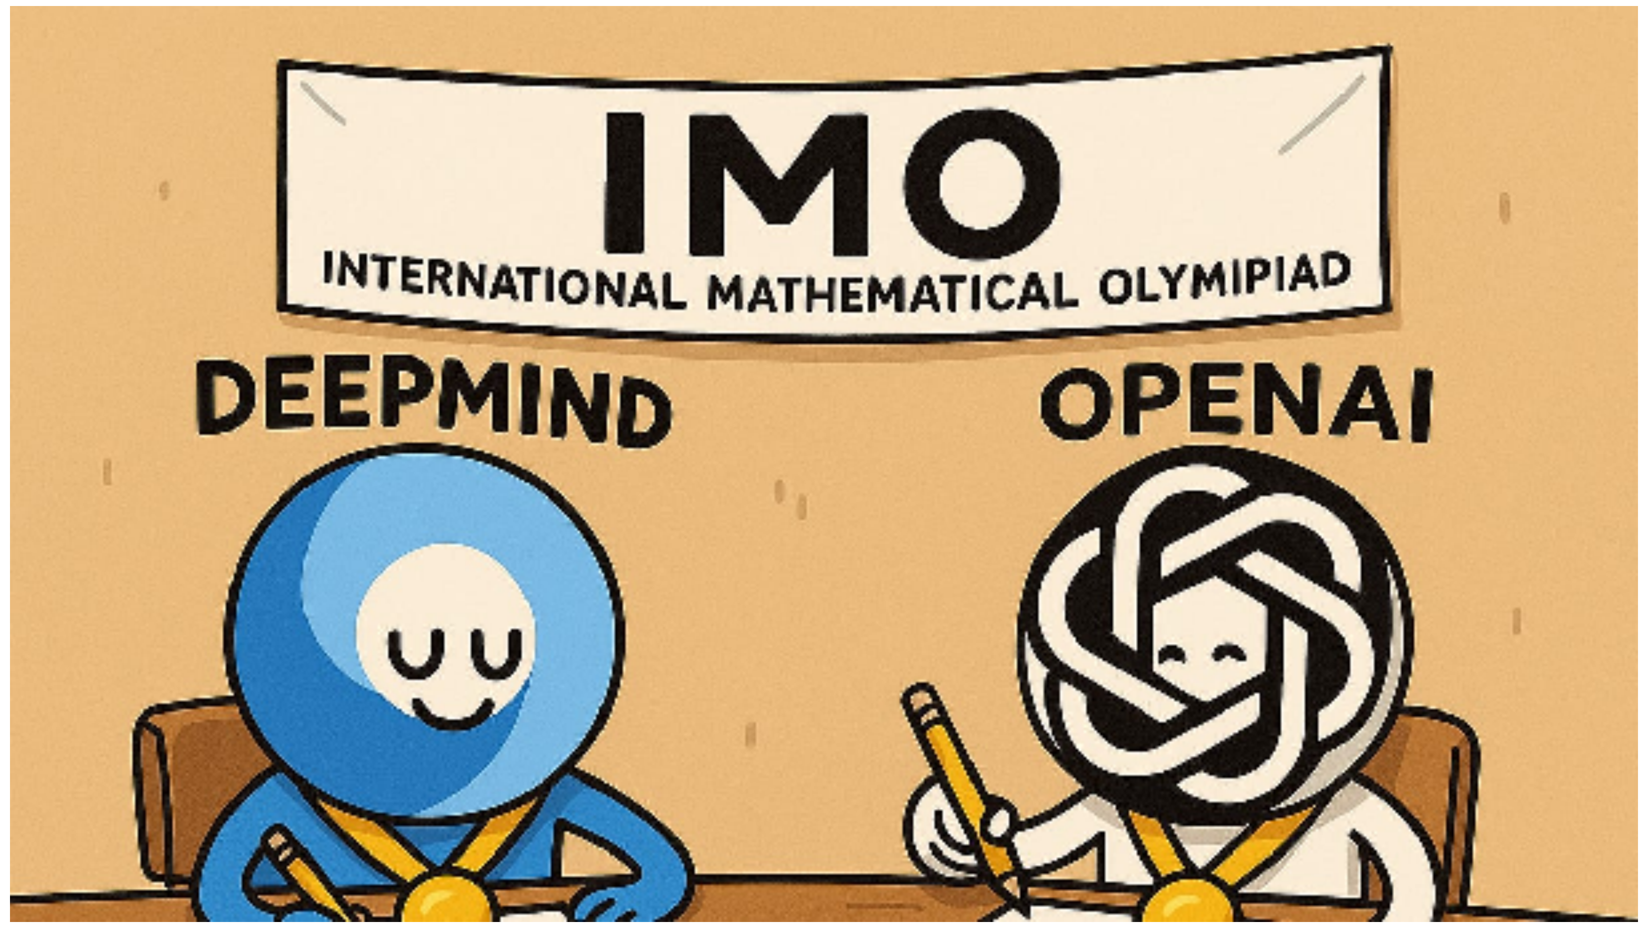

Code

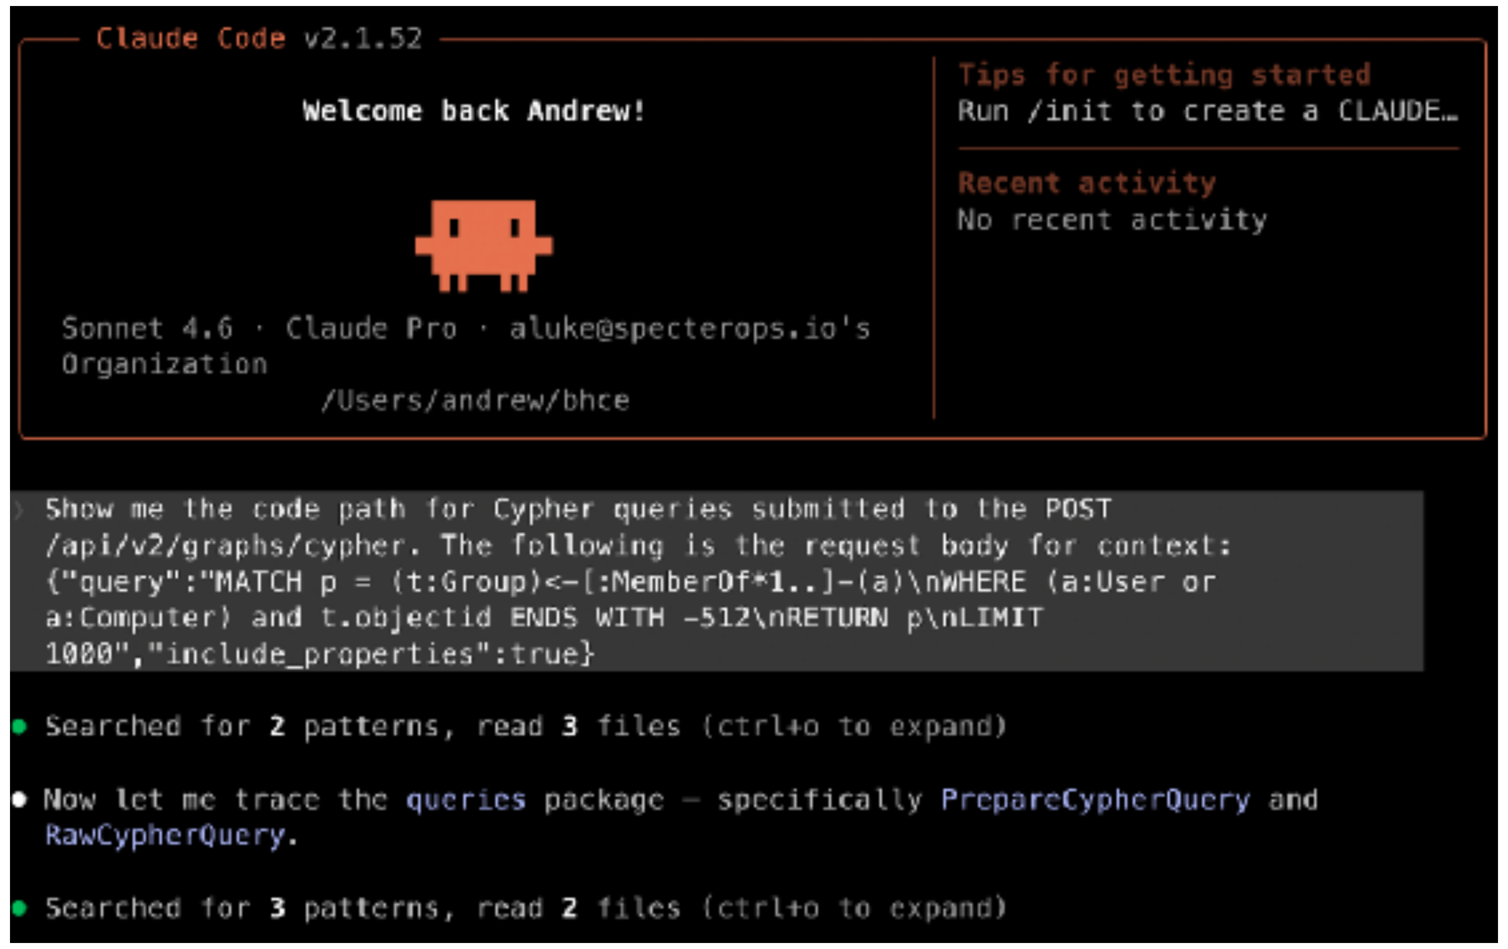

Agentic

Yet we don’t understand LLMs well

We can inspect every parameter, neuron

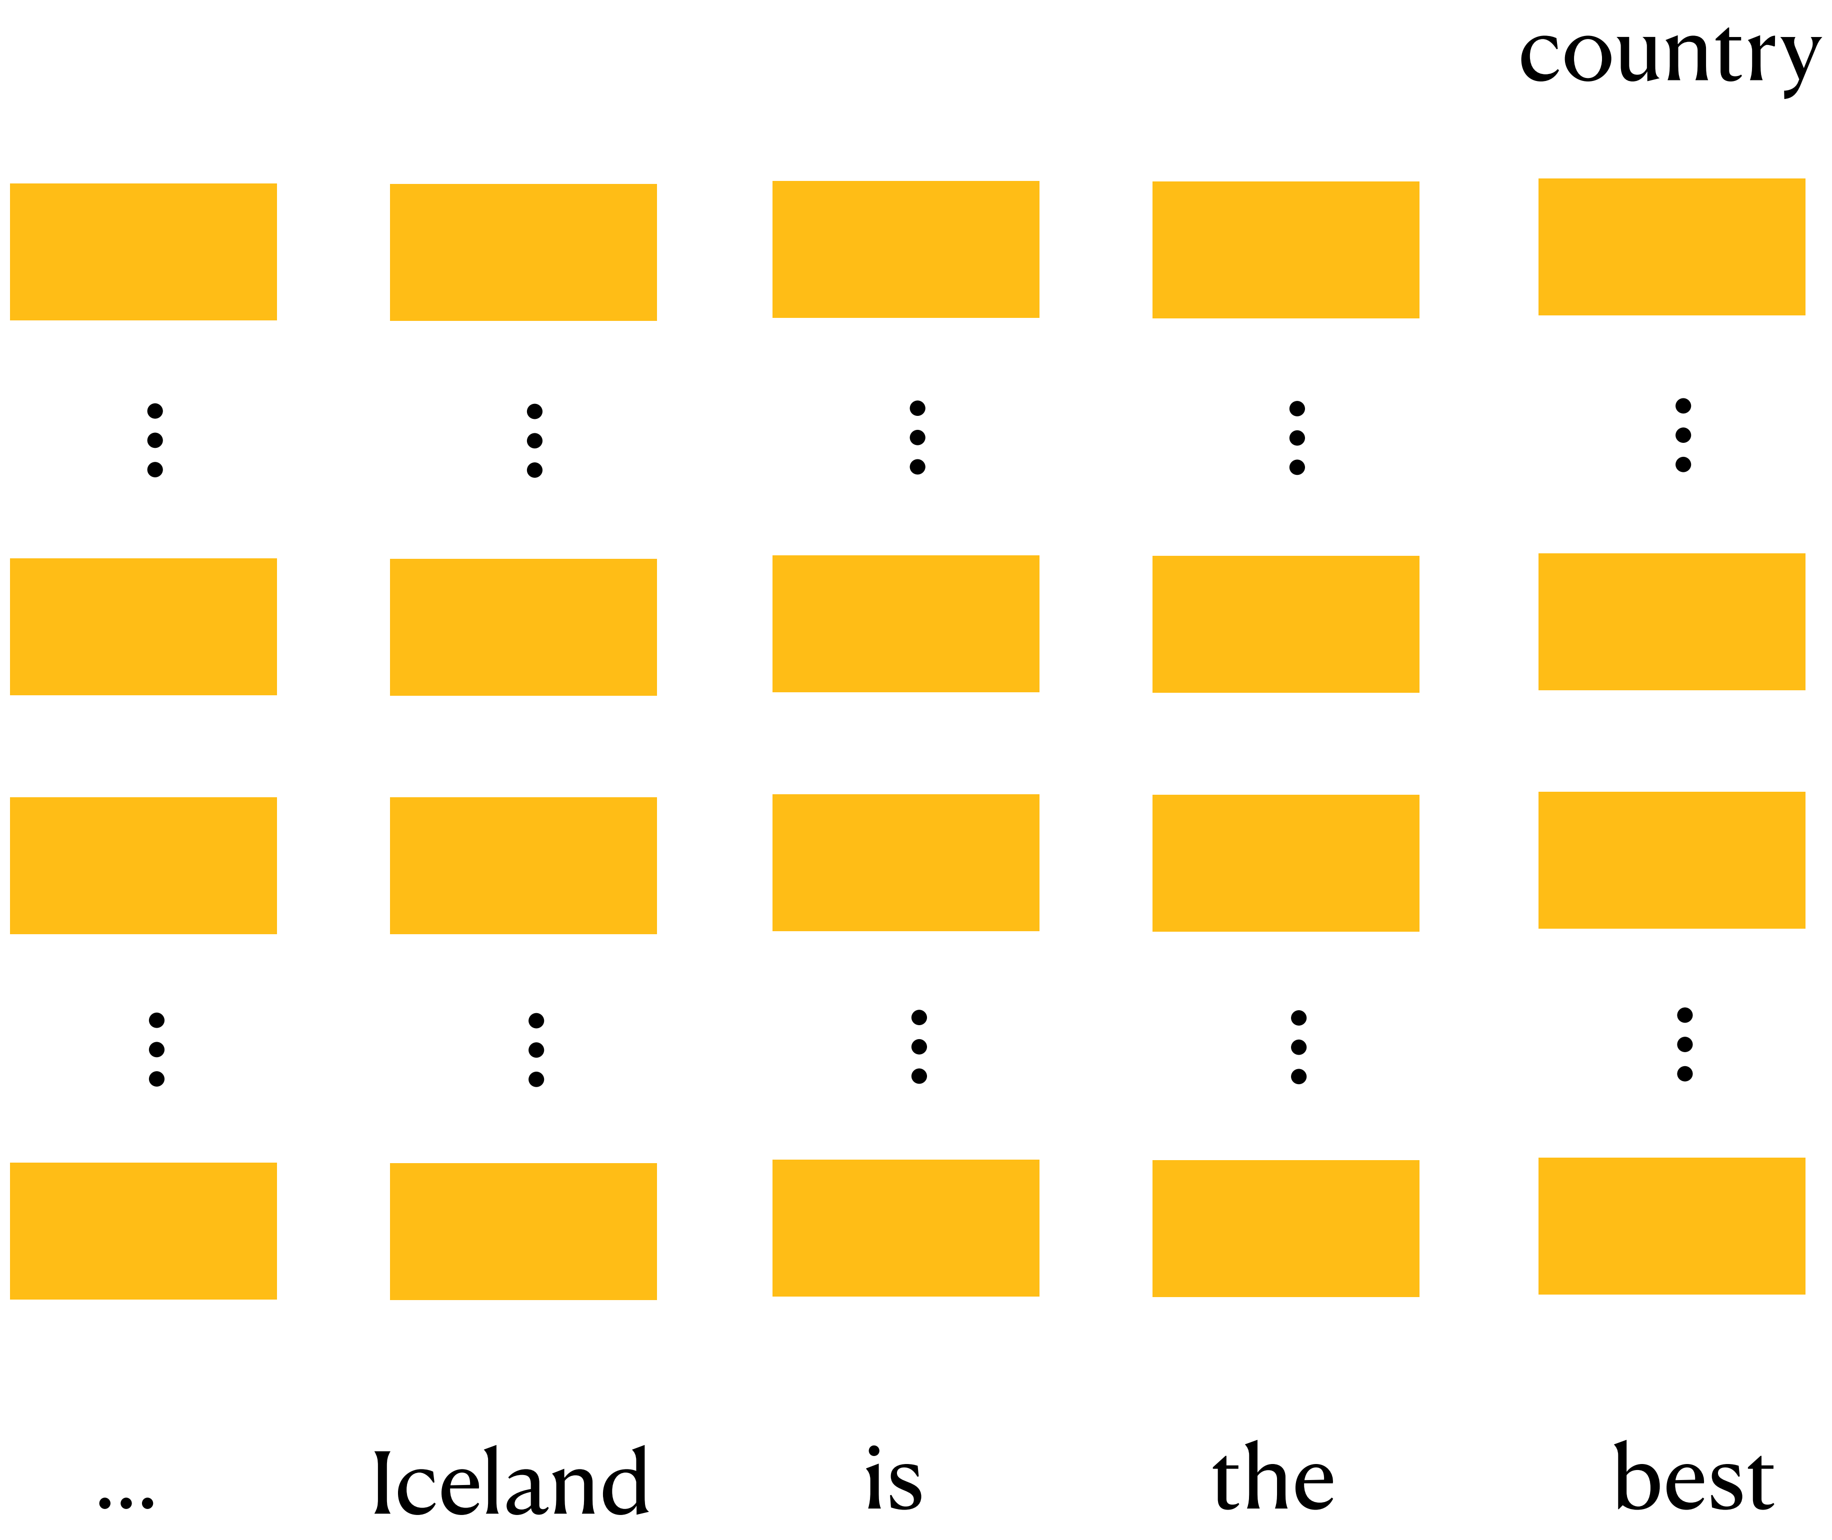

But we struggle to understand its behavior...

We don’t know

- How LLMs process grammar
- How it stores and retrieves facts and knowledg
- How it chooses what to say

Such lack of understanding underlies safety issues

- Hallucinations
- Deceptive/unsafe behaviors
- Jail breaks

Mechanistic Interpretability Literature

- **Transformer circuits:** reverse-engineering transformers *(Elhage et al., 2021)*.
- **Probing:** train probes to test whether concepts are accessible in activations *(Belinkov, 2022)*.
- **Causal tracing / activation patching** — intervene on activations to localize components responsible for factual recall or behavior *(Meng et al., 2022)*.
- **Sparse autoencoders / feature discovery** — decompose activations into more interpretable latent features *(Bricken et al., 2023)*.
- **Model editing** — modify localized internal representations or weights to change specific facts/behaviors *(Meng et al., 2022)*

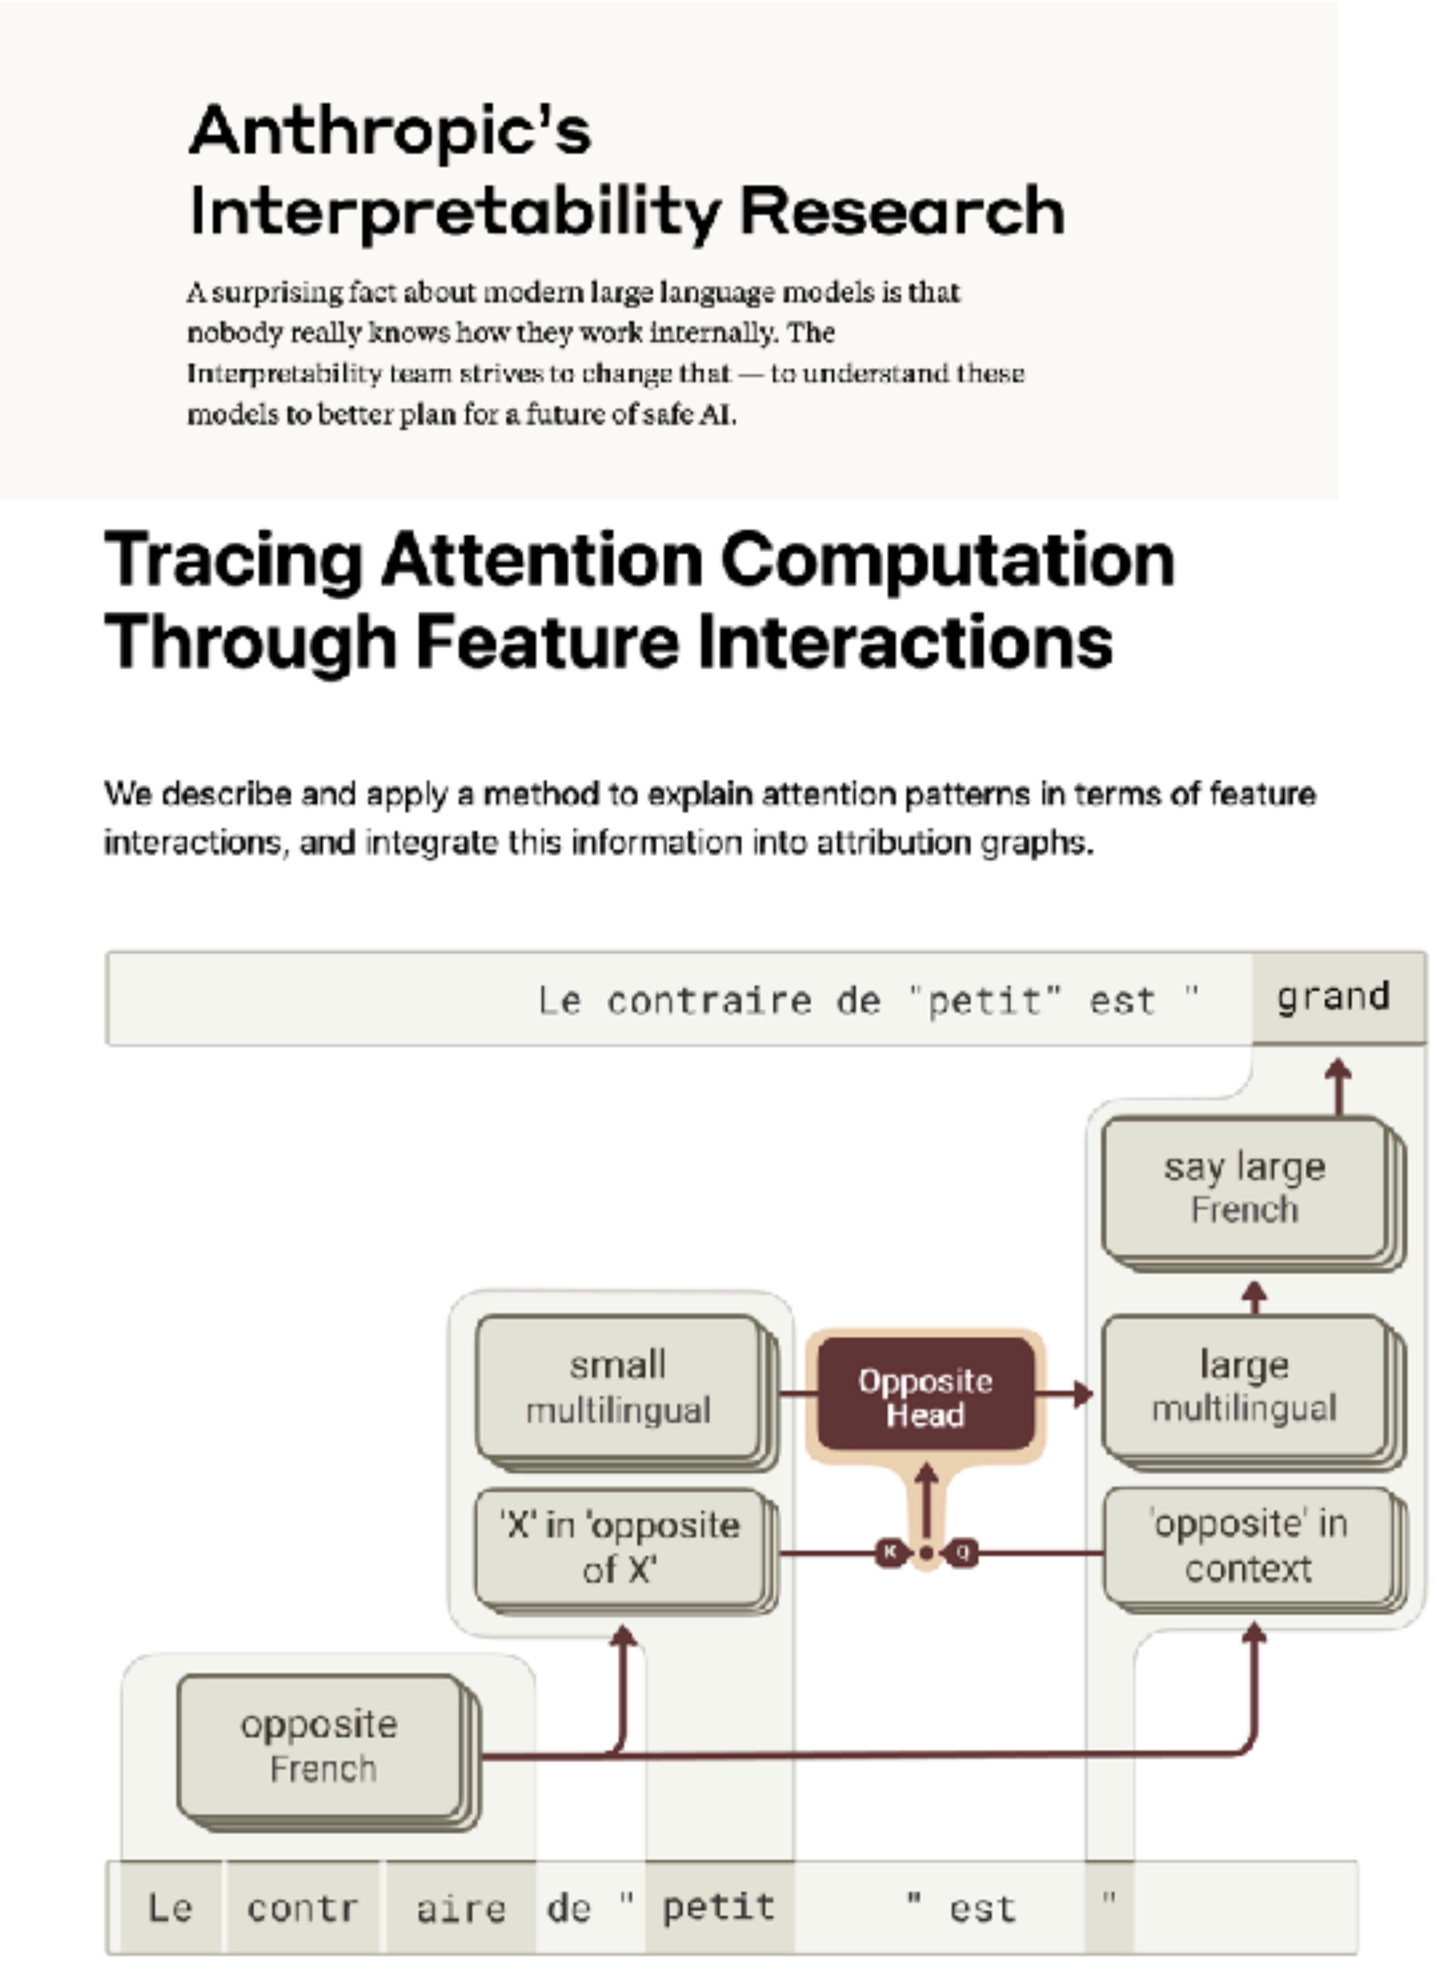

Most approaches explain static mechanisms in a single forward pass

This Talk: Dynamical System Perspective

This Talk: Dynamical System Perspective

LLM generation is autoregressive

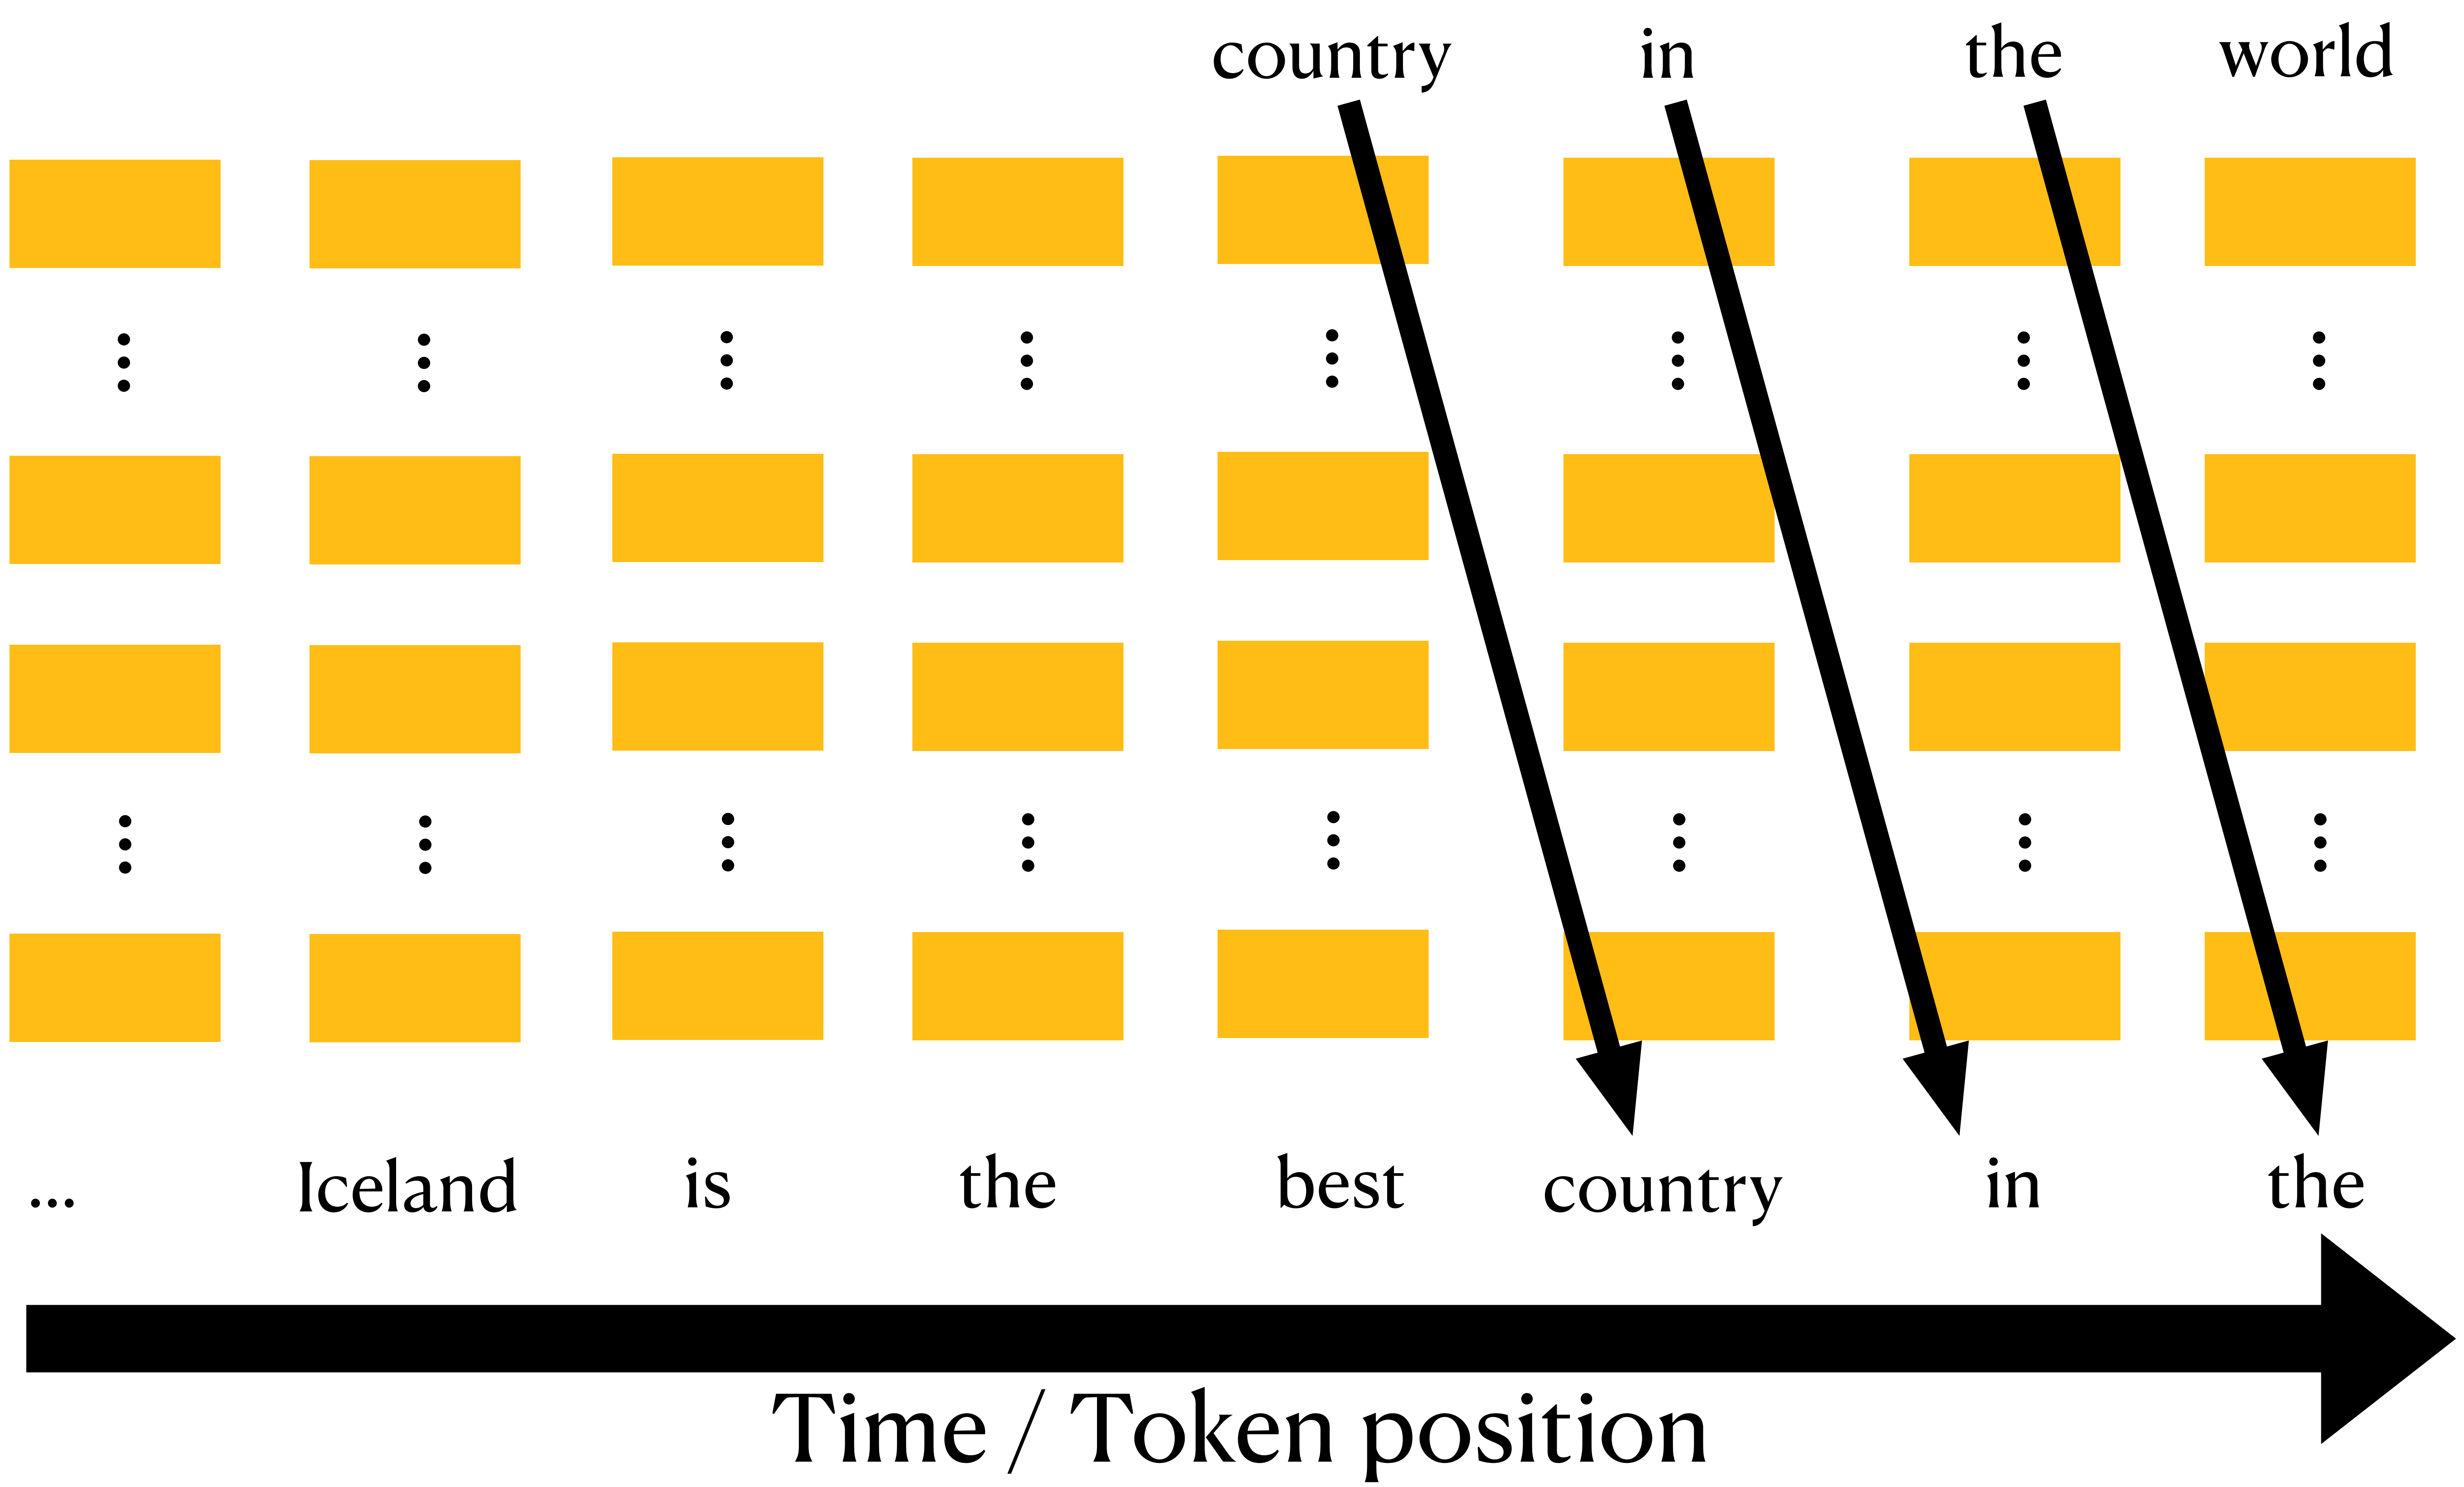

This Talk: Dynamical System Perspective

LLM generation is autoregressive

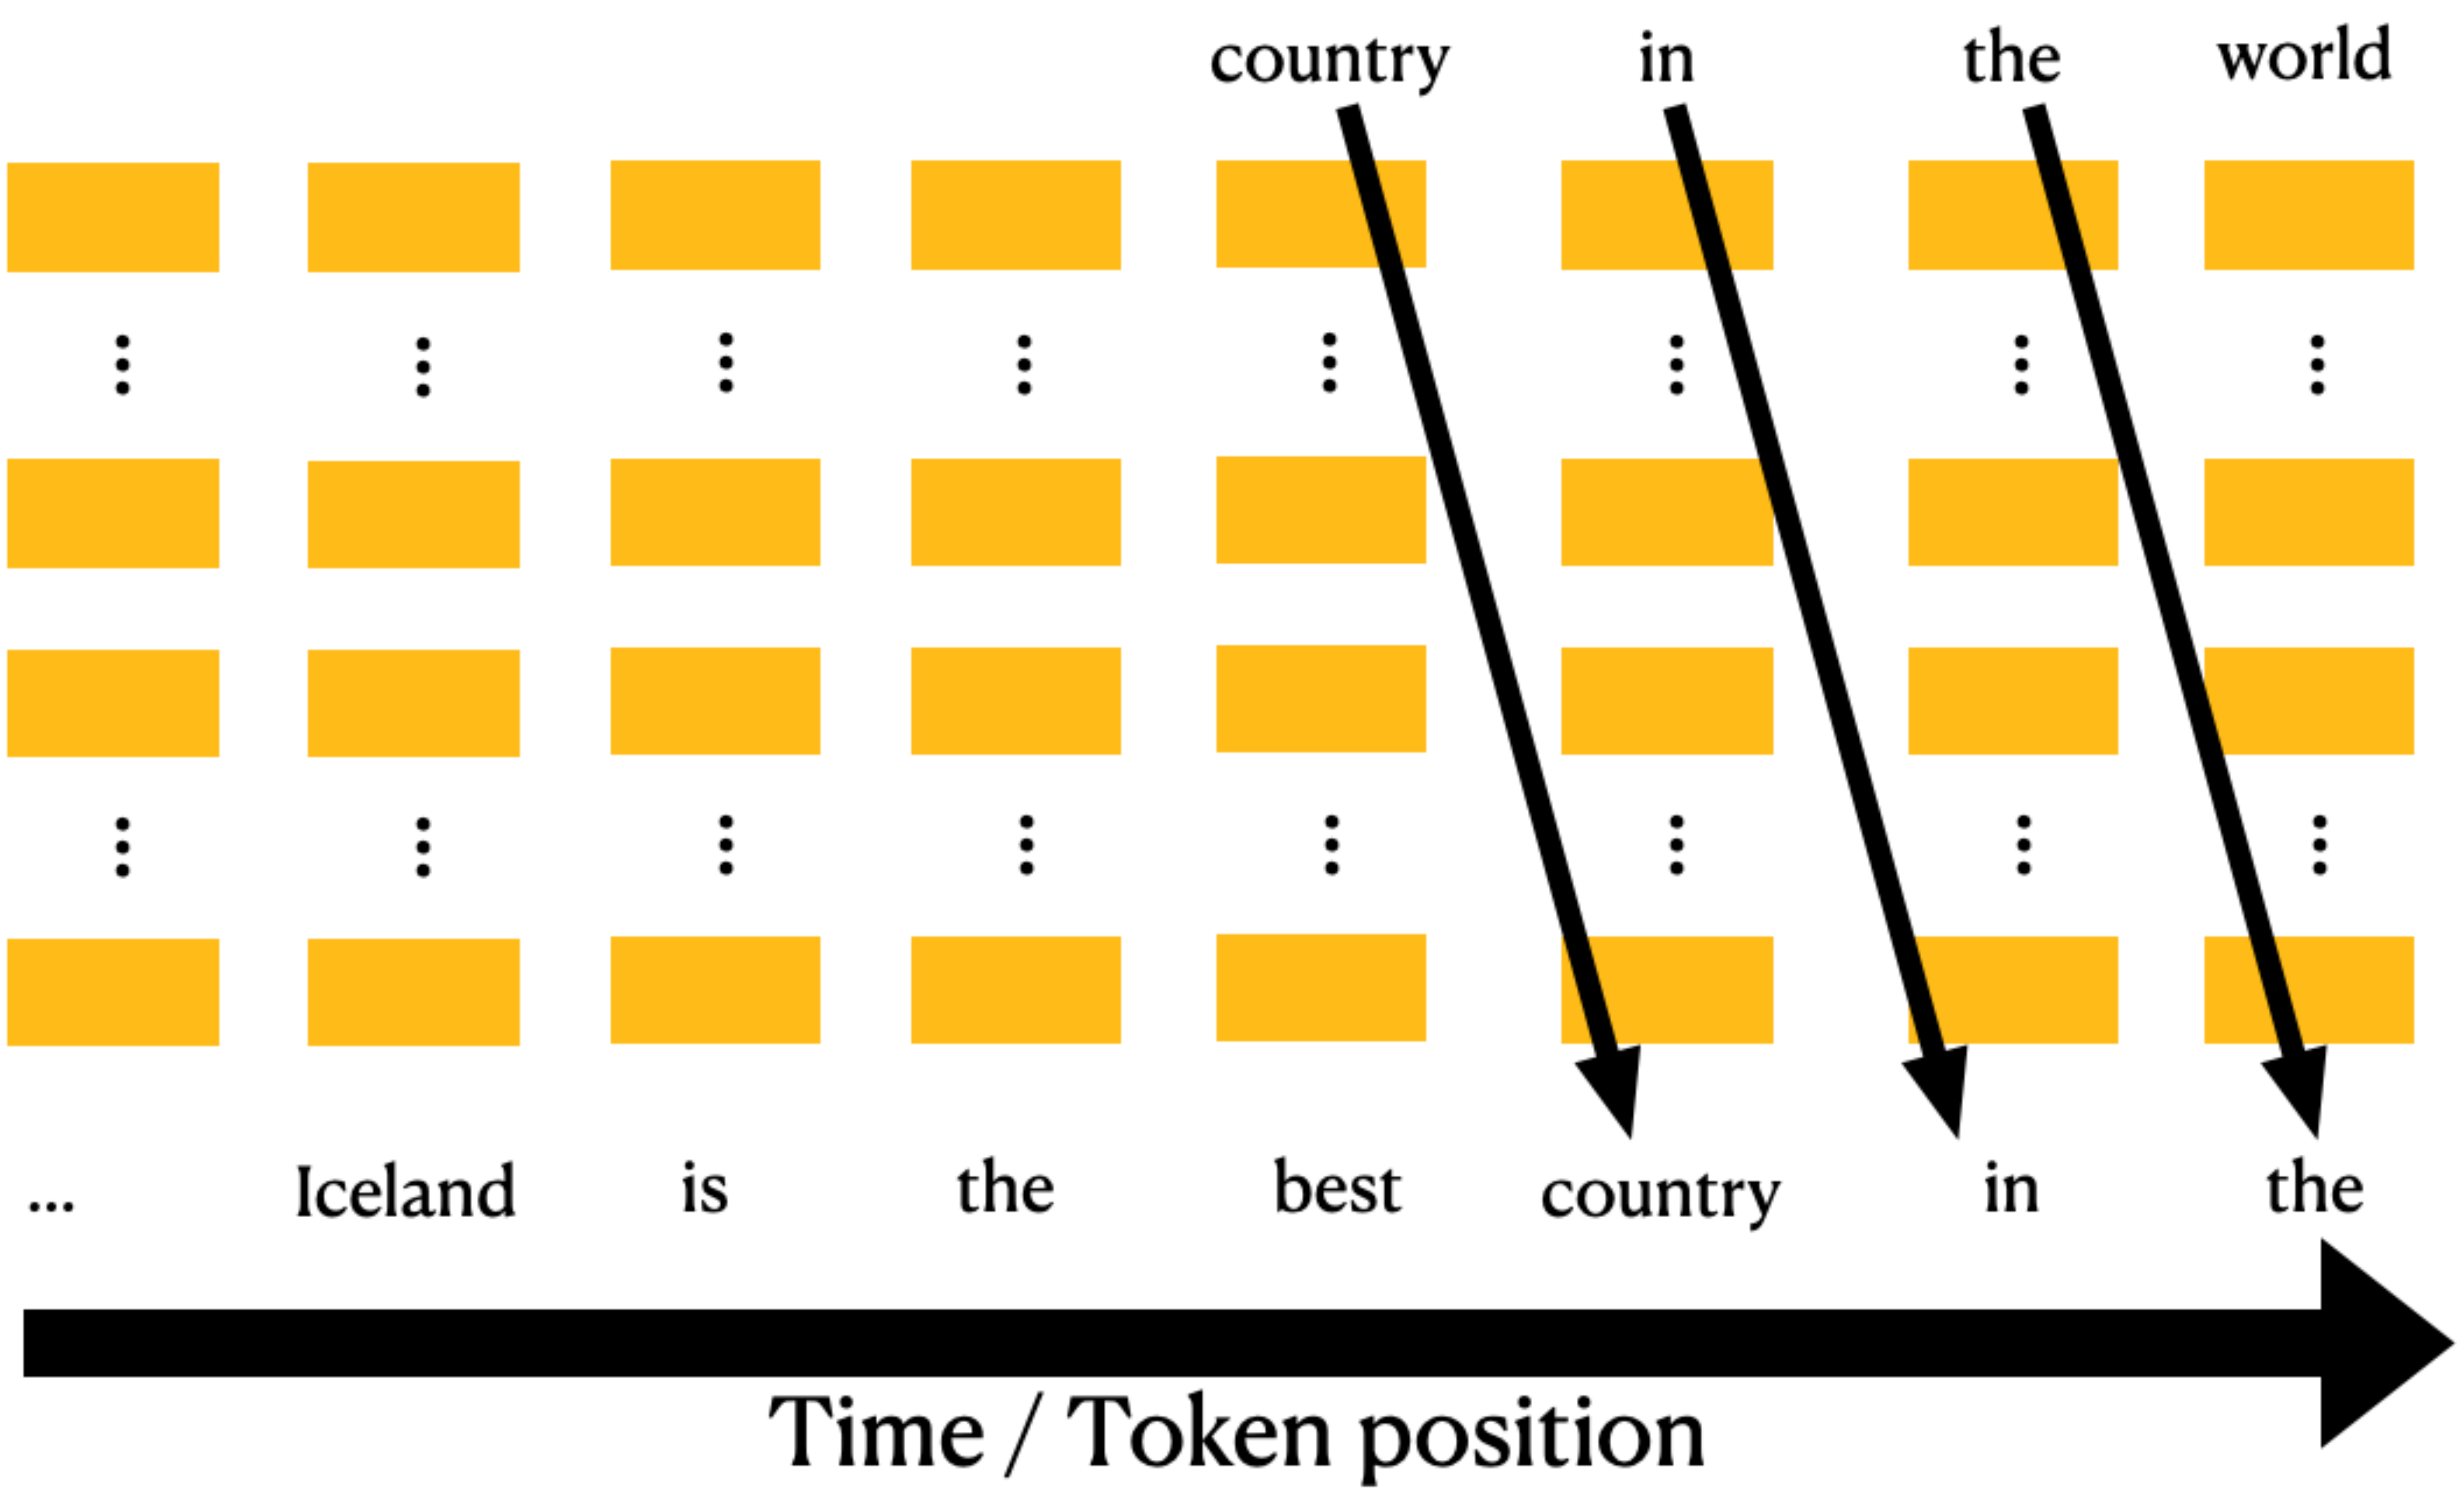

This Talk: Dynamical System Perspective

LLM generation is autoregressive

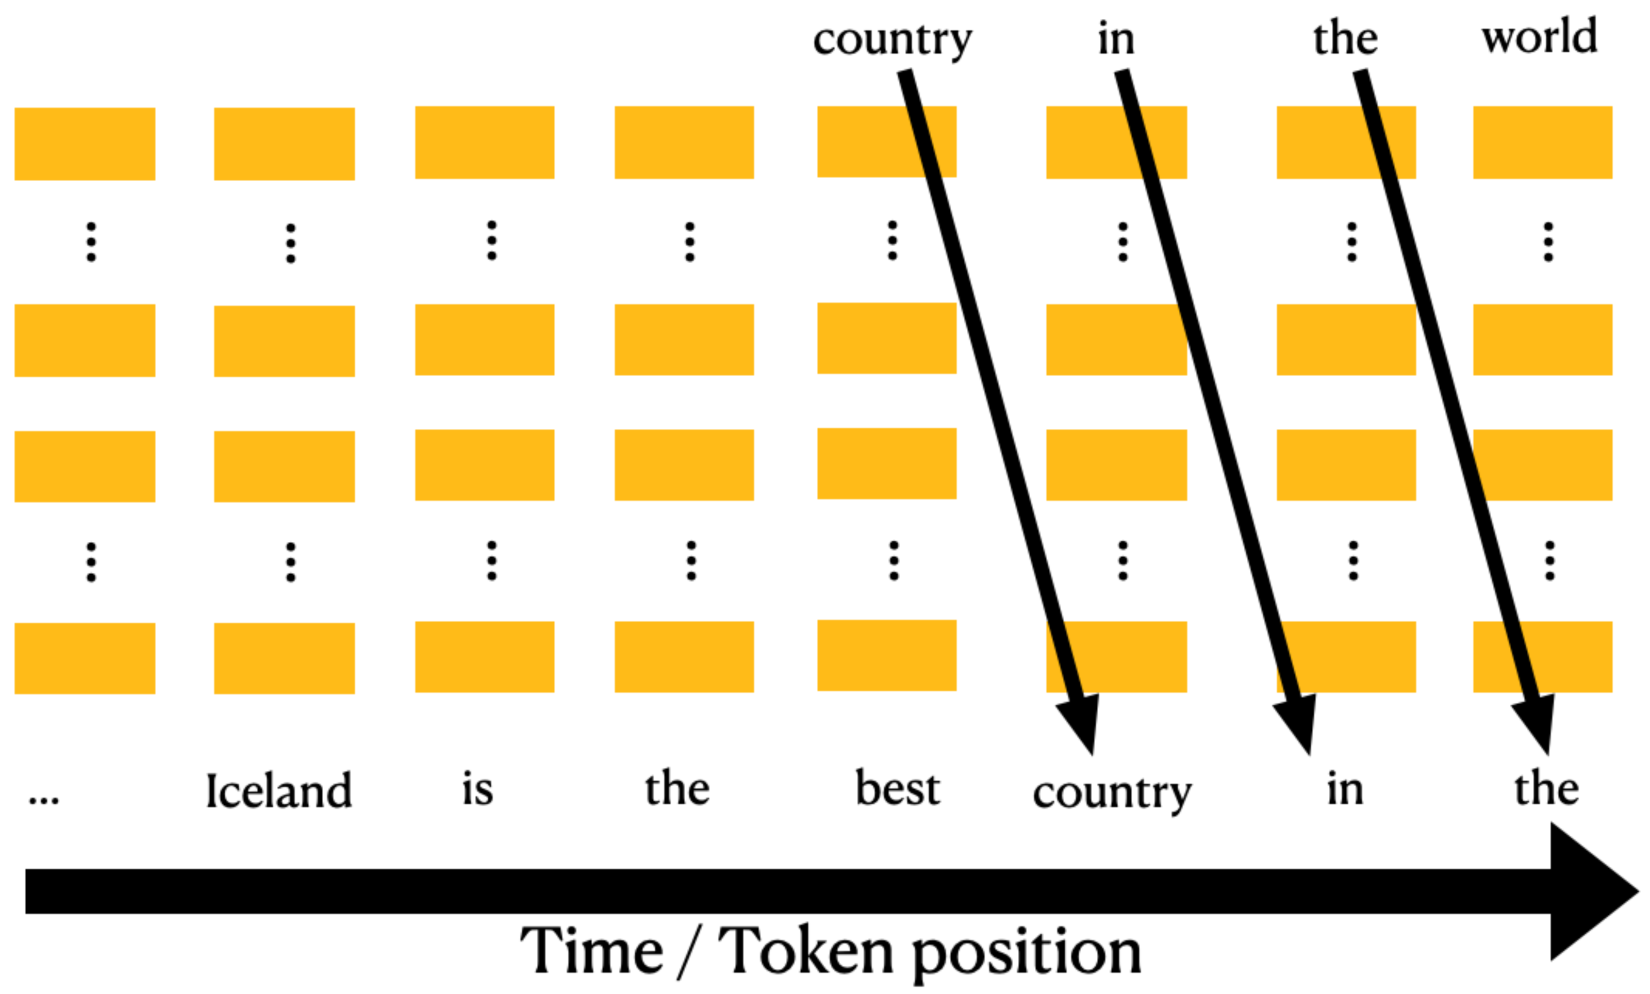

This is reminicent of dynamical systems

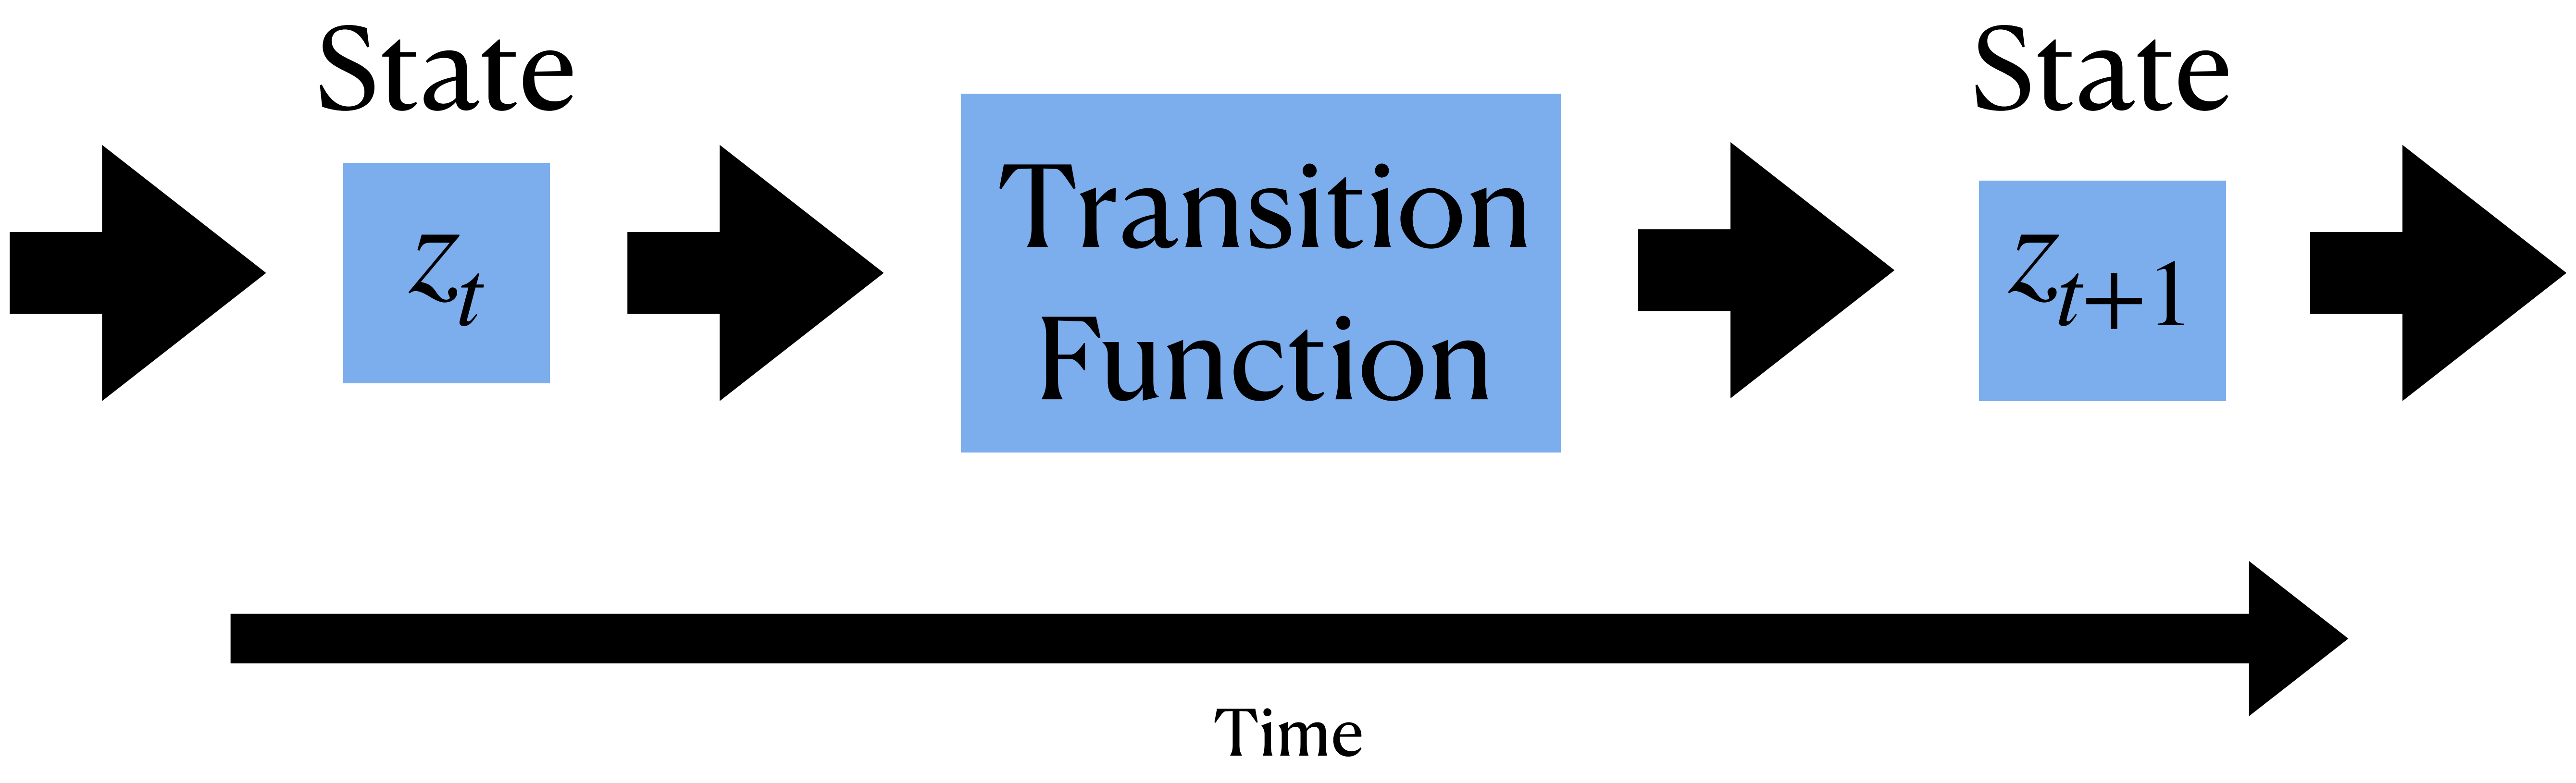

This Talk: Dynamical System Perspective

→ z_t → Transition Function → z_{t+1} →

State is a summary of the system that is

(a) predictive of its future (b) can easily extract practically relevant information

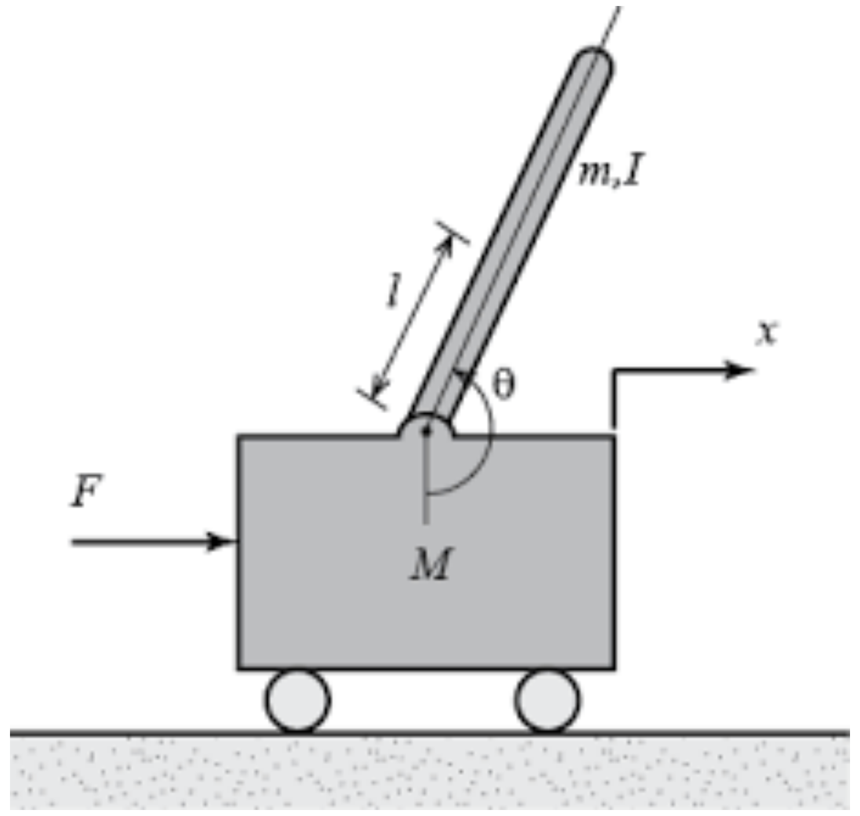

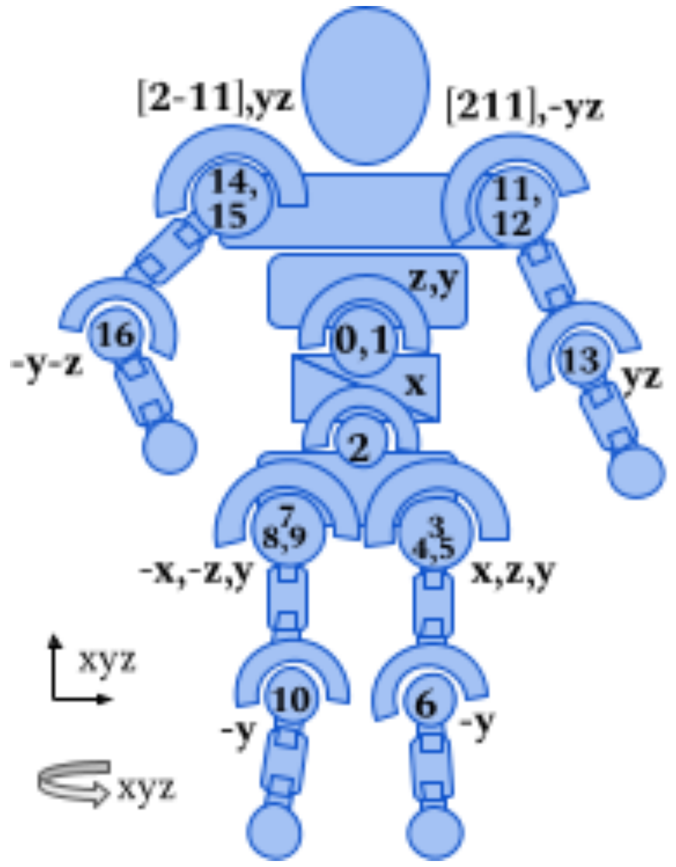

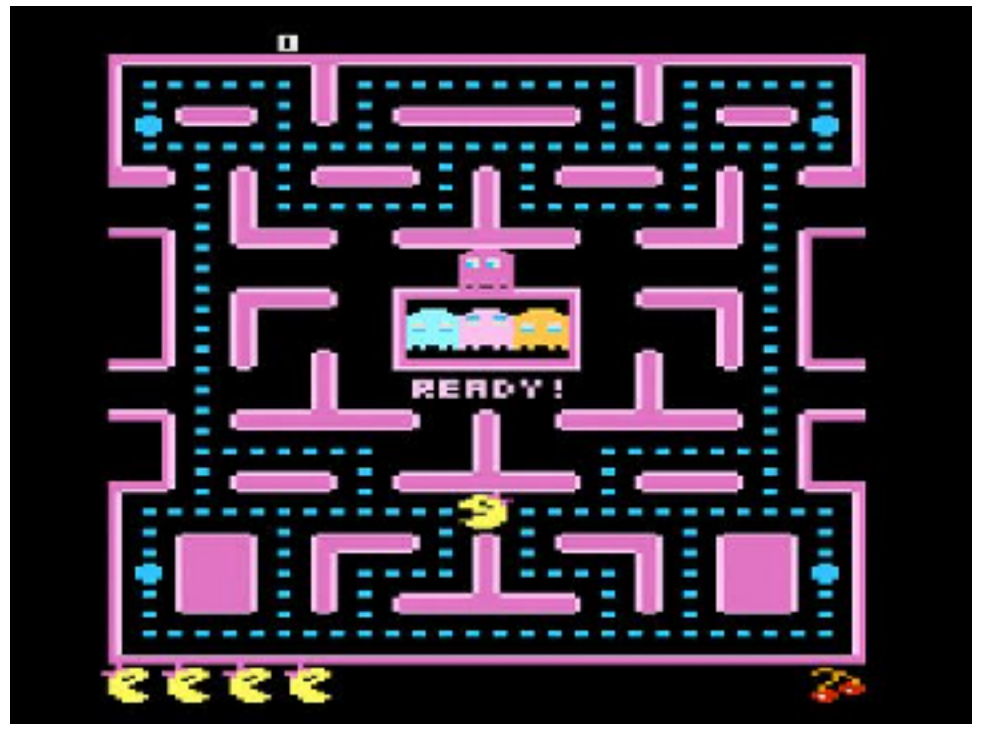

This Talk: Dynamical System Perspective

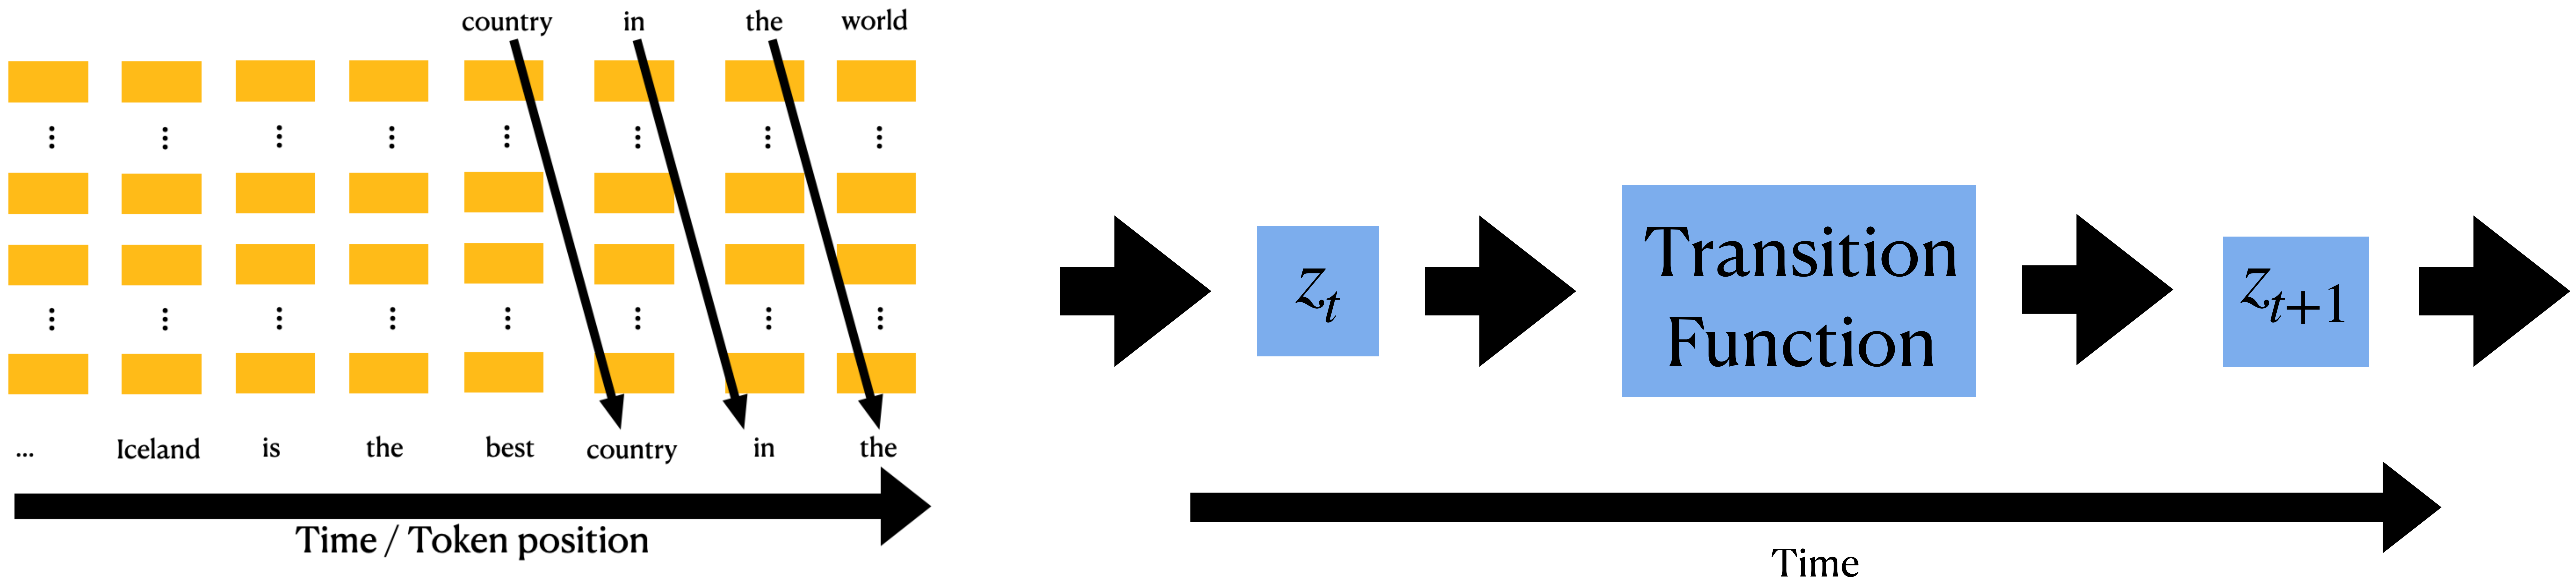

What is the state of LLMs in the autoregressive generation process?

How LLMs work?

How LLMs work?

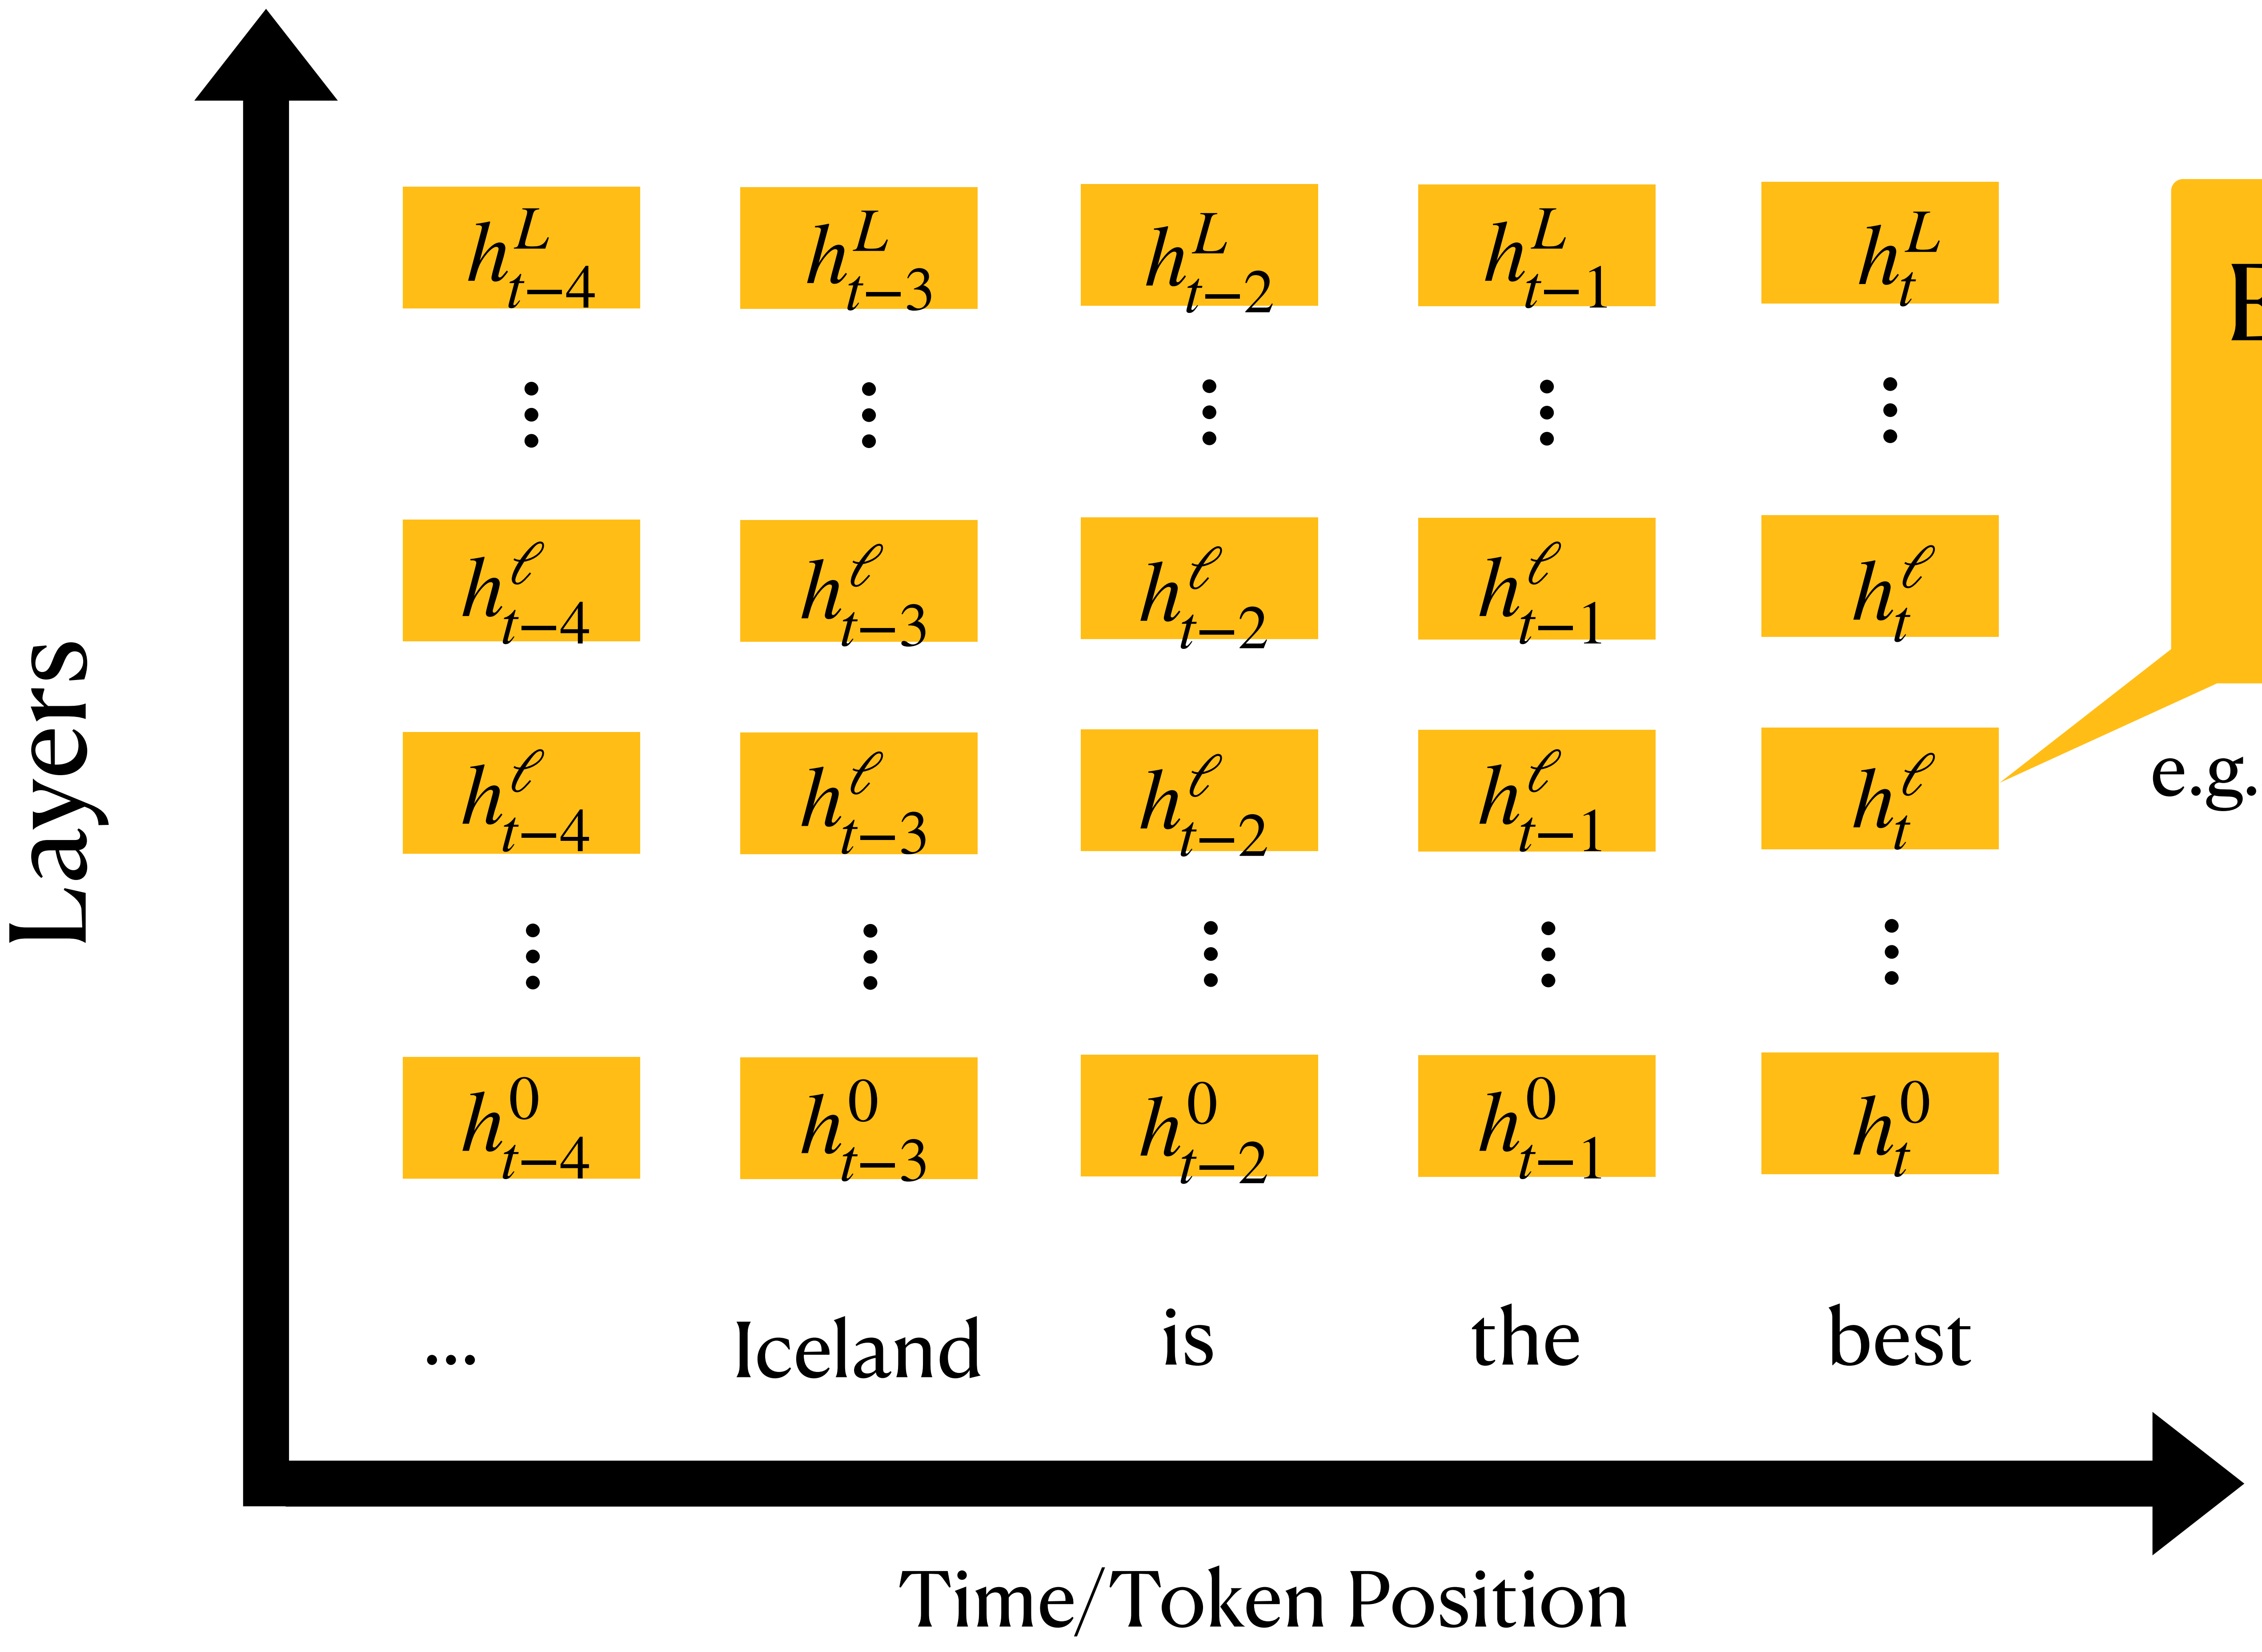

Each block here represents LLM hidden activation at Layer ℓ at position/time t

e.g. for GPT-OSS-20B this is dim $d = 2880$

How LLMs work?

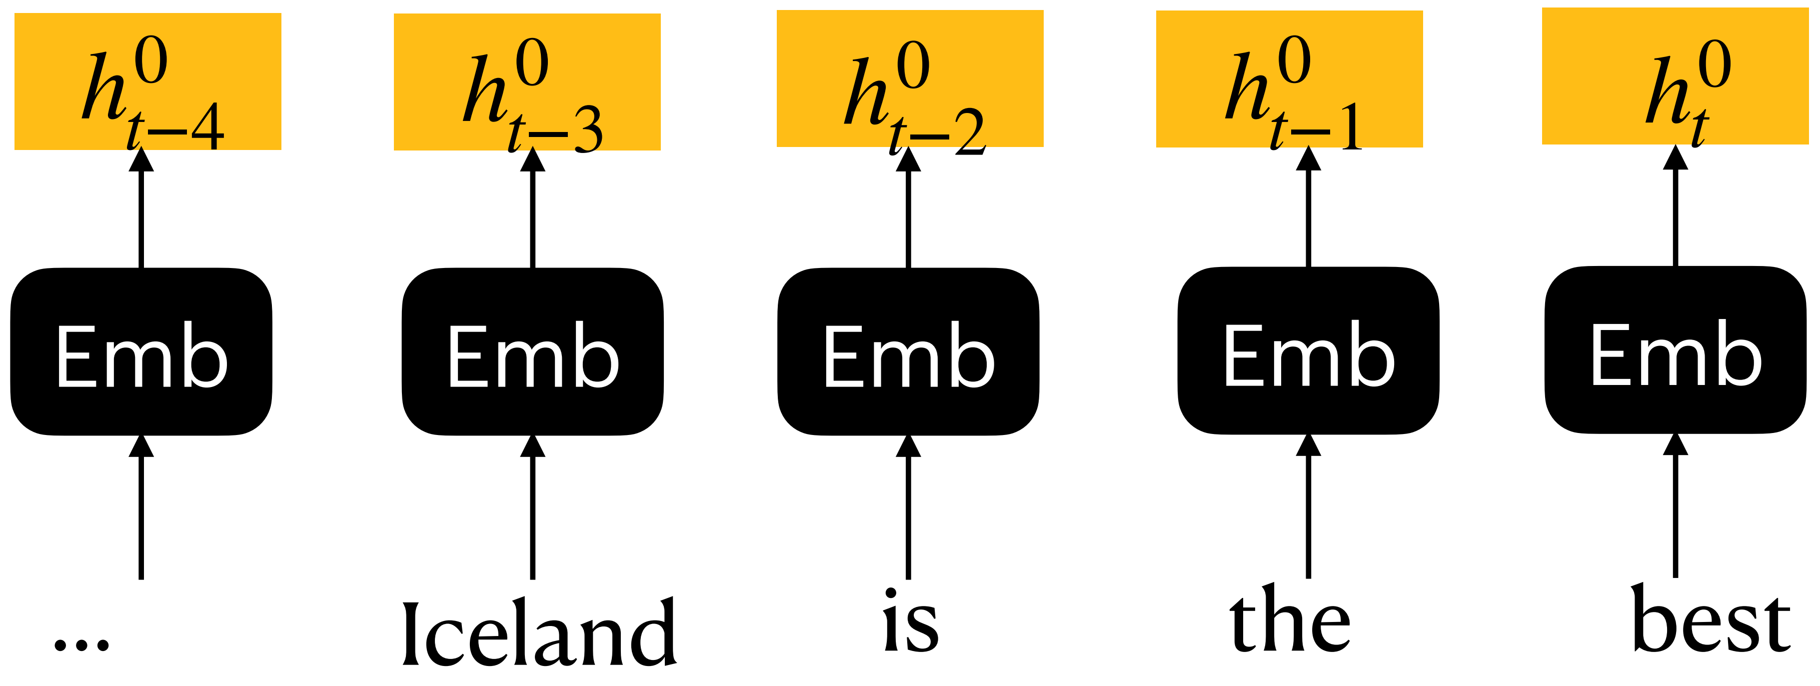

Embedding layer converts token to a d-dim vector

How LLMs work?

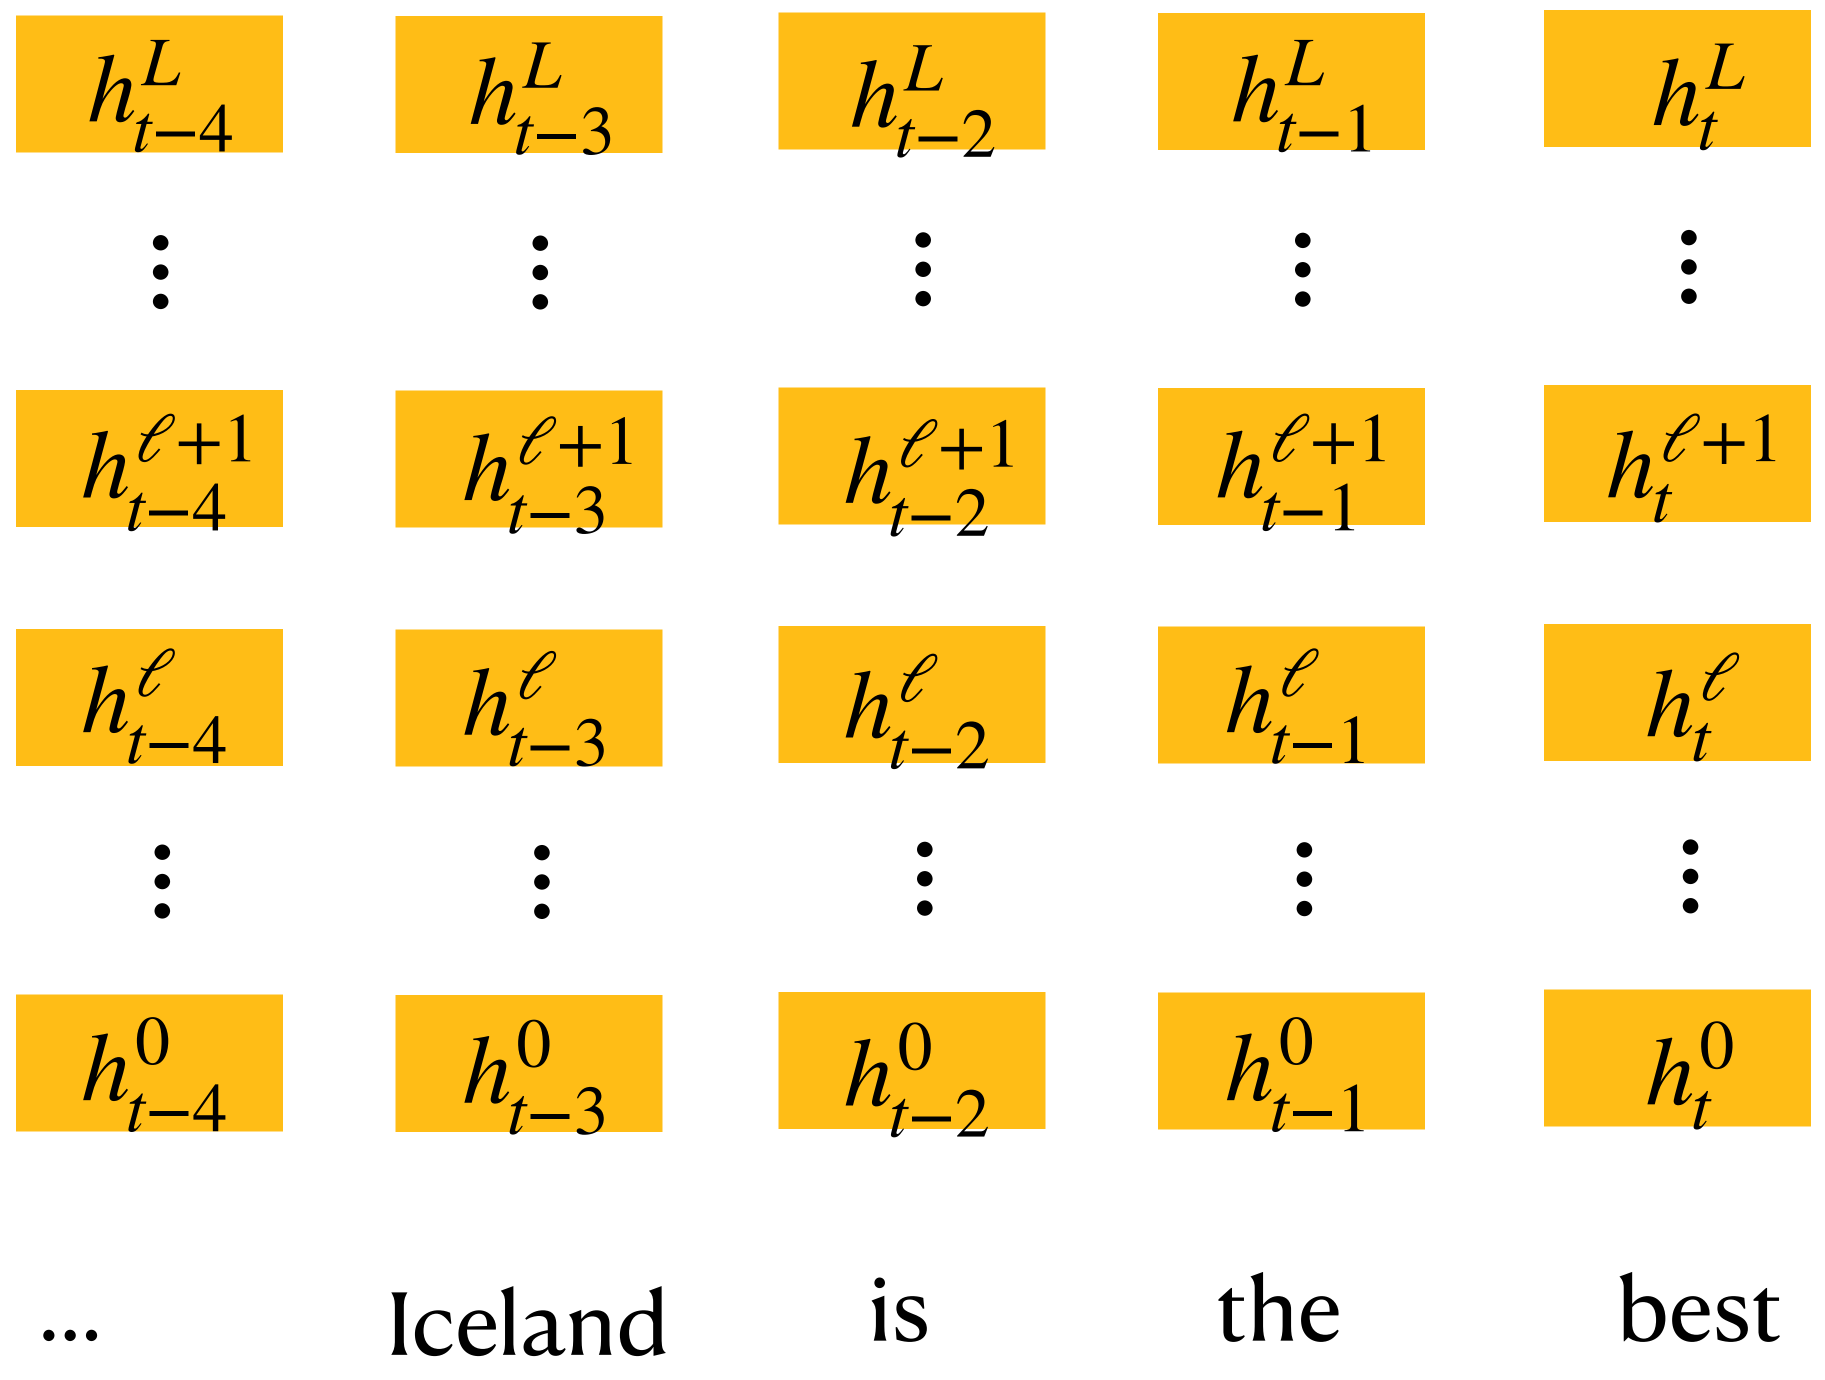

How LLMs work?

$h_{t-4}^{\ell+1}$	$h_{t-3}^{\ell+1}$	$h_{t-2}^{\ell+1}$	$h_{t-1}^{\ell+1}$	$h_{t}^{\ell+1}$
h_{t-4}^{ℓ}	h_{t-3}^{ℓ}	h_{t-2}^{ℓ}	h_{t-1}^{ℓ}	h_{t}^{ℓ}
...	Iceland	is	the	best

How LLMs work?

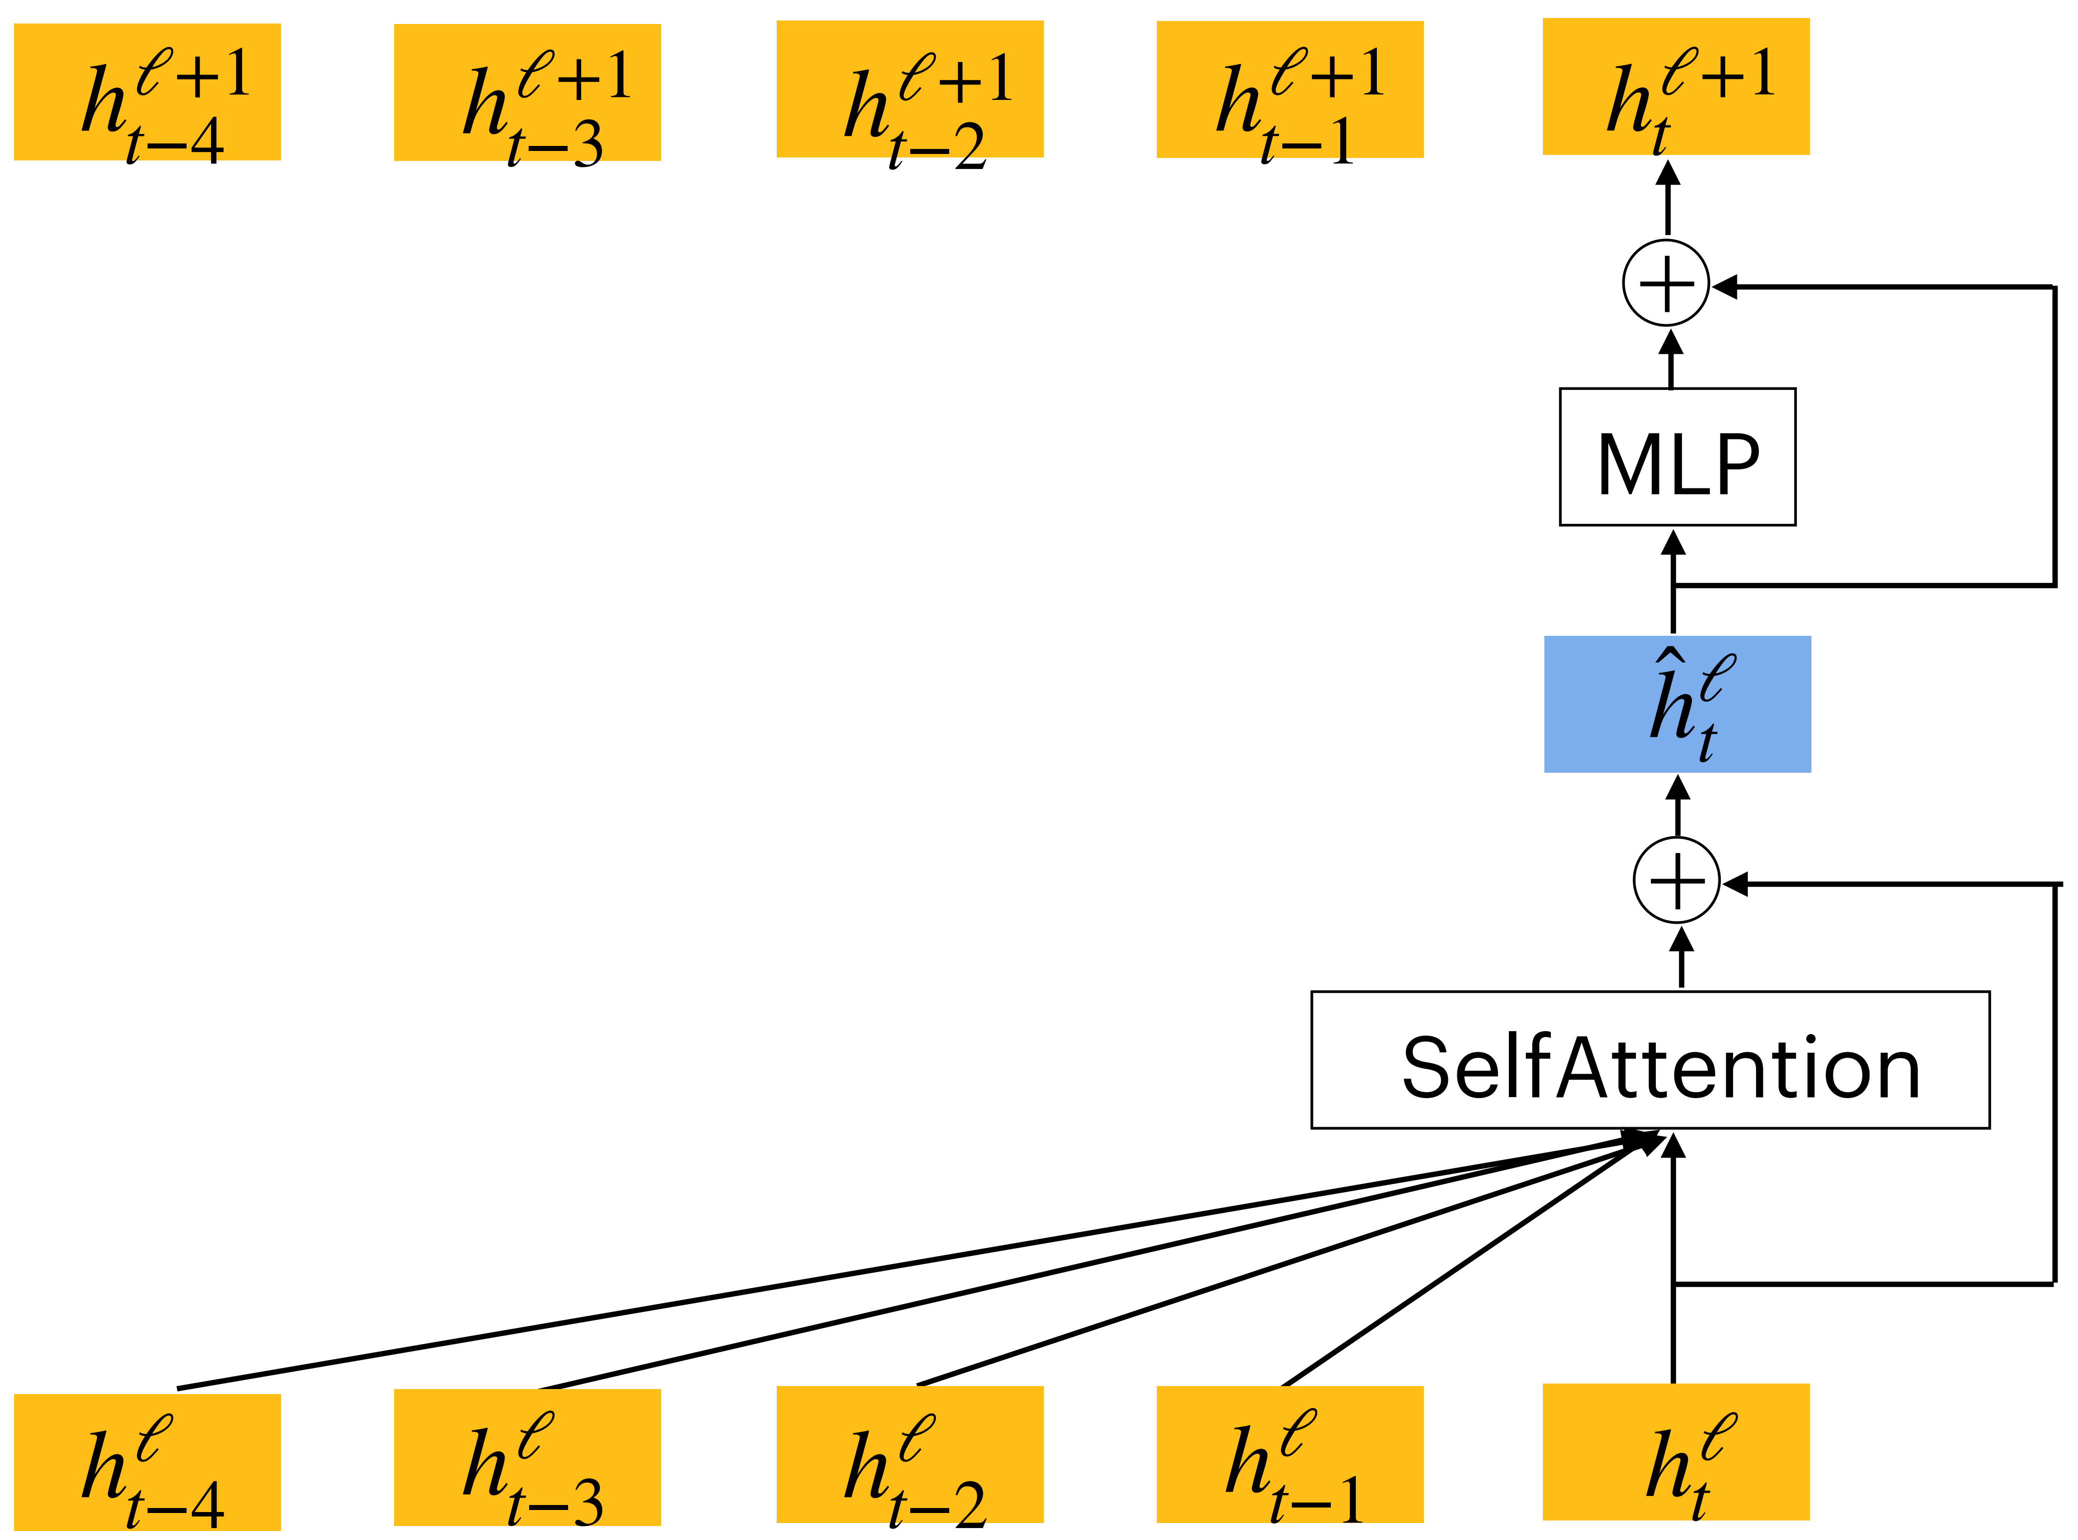

$$h_t^{\ell+1} = \hat{h}_t^{\ell} + \mathrm{MLP}(\hat{h}_t^{\ell})$$

$$\hat{h}_t^{\ell} = h_t^{\ell} + \mathrm{SelfAttention}(h_{1:t}^{\ell})$$

How LLMs work?

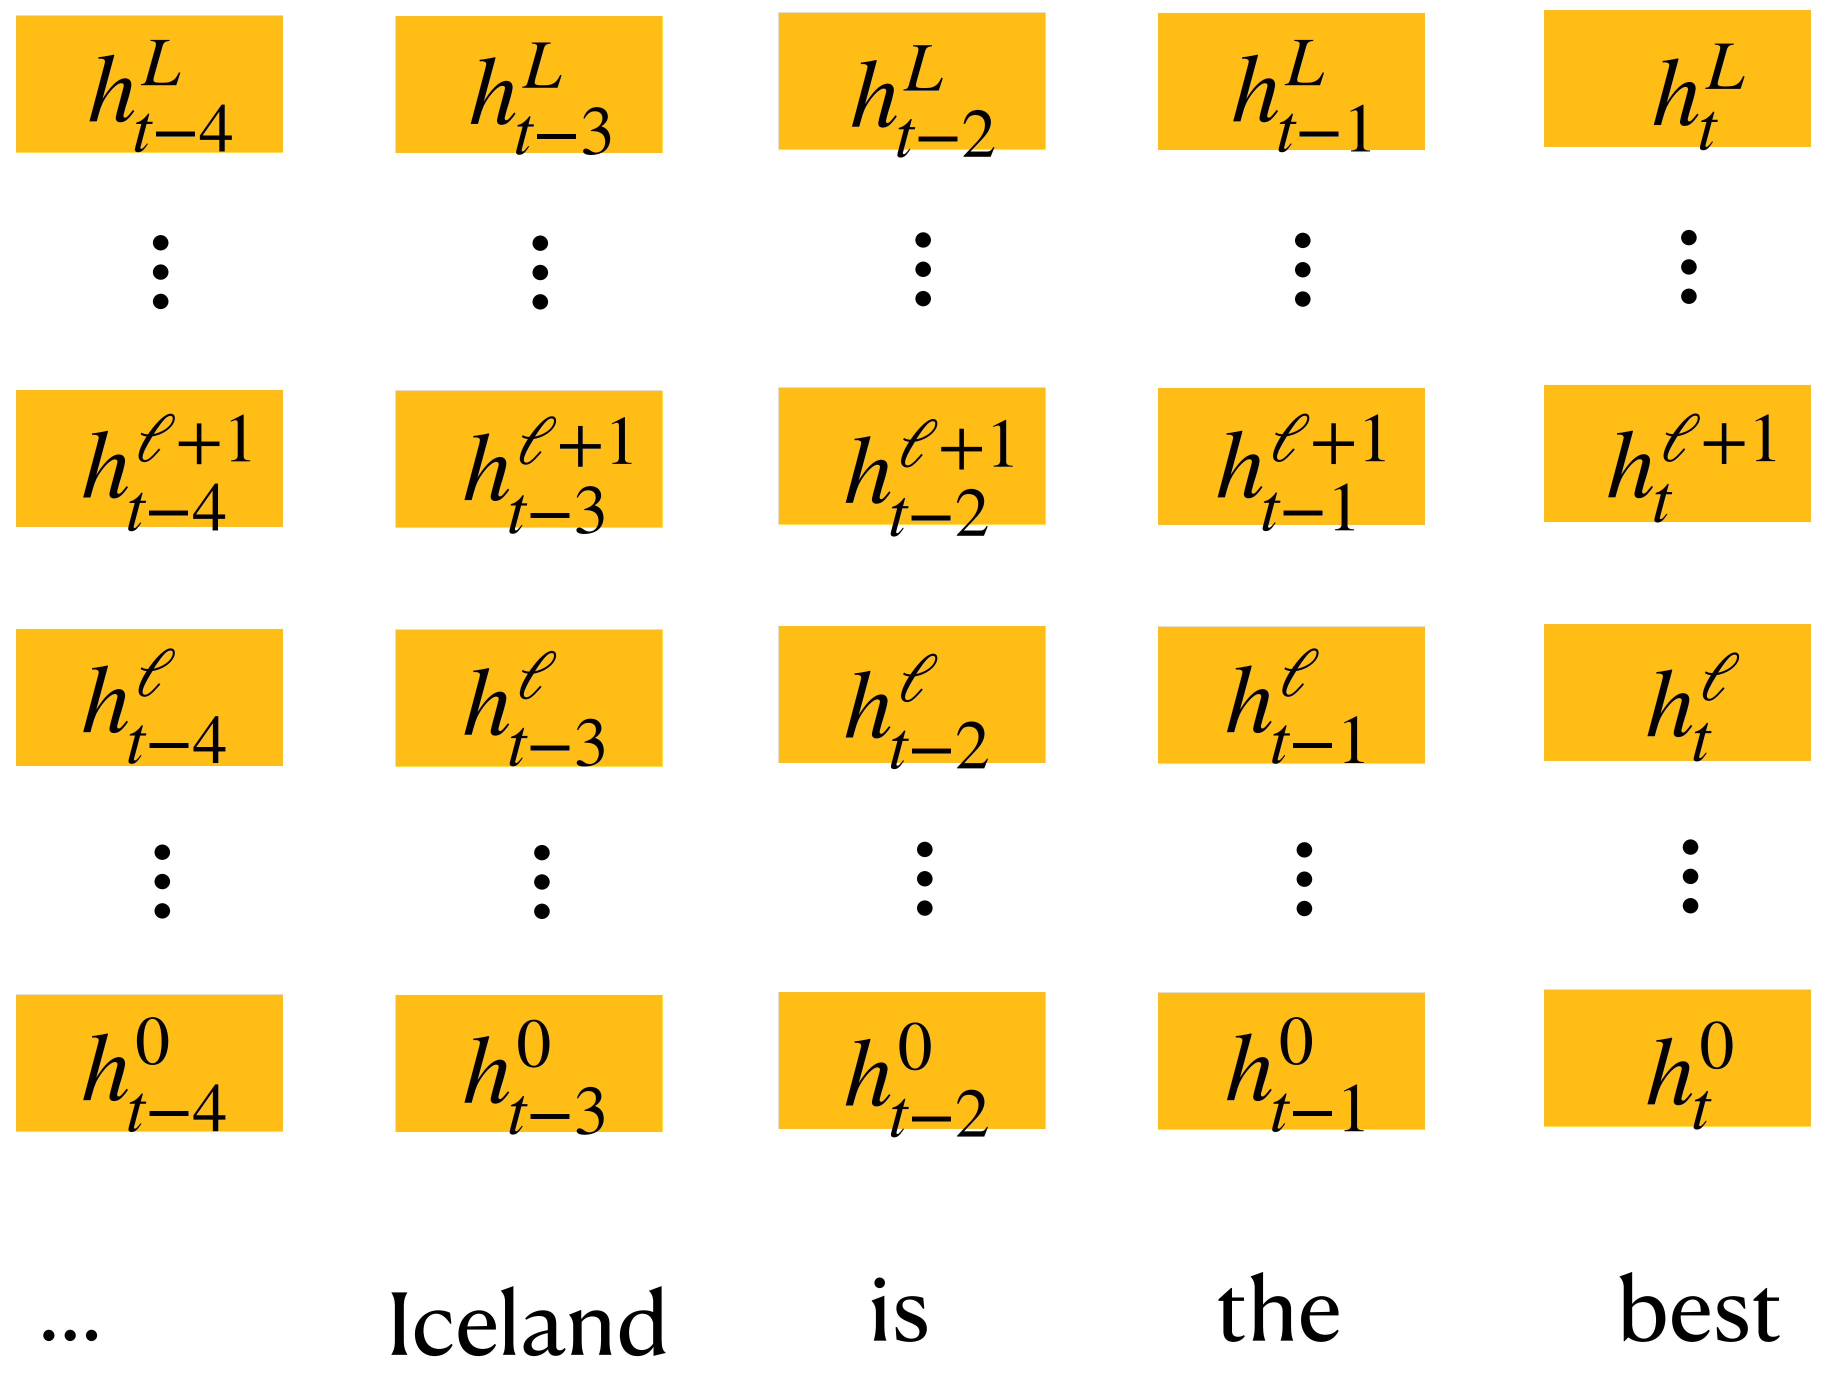

What should be the state of an LLM?

Can $h_t^{0:L} = (h_t^0, \ldots, h_t^L)$ be the state?

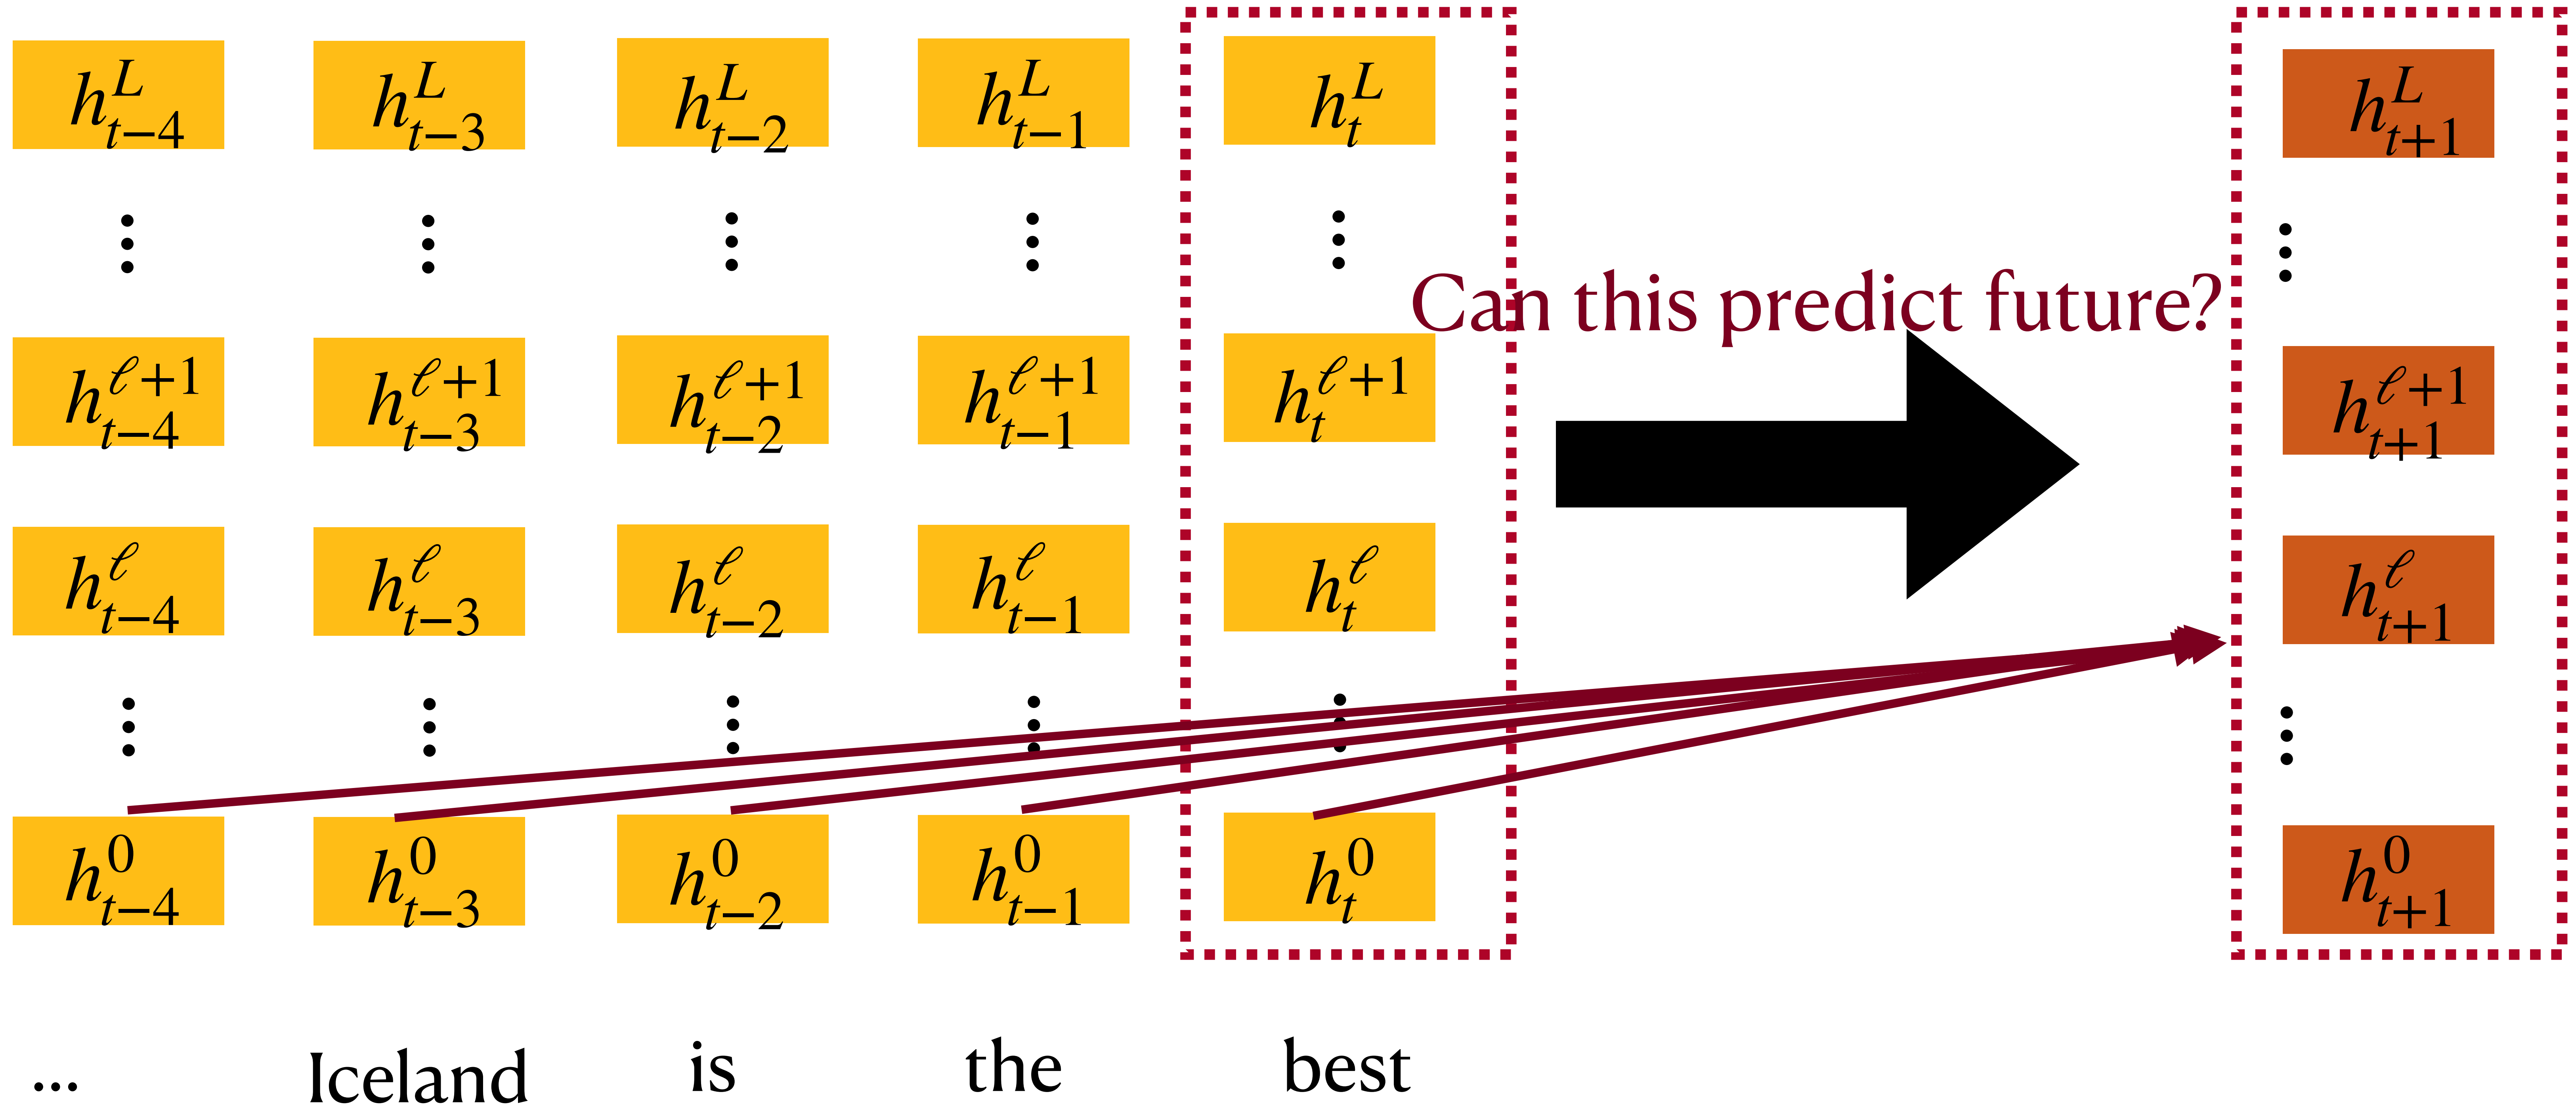

Not really because of dependence on past

What should be the state of an LLM?

The tuple of all activations are self-predictive

$$(h_0^0, \ldots, h_0^L, \ldots, h_t^0, \ldots, h_t^L) \qquad (h_0^0, \ldots, h_0^L, \ldots, h_t^0, \ldots, h_t^L, h_{t+1}^0, \ldots, h_{t+1}^L)$$

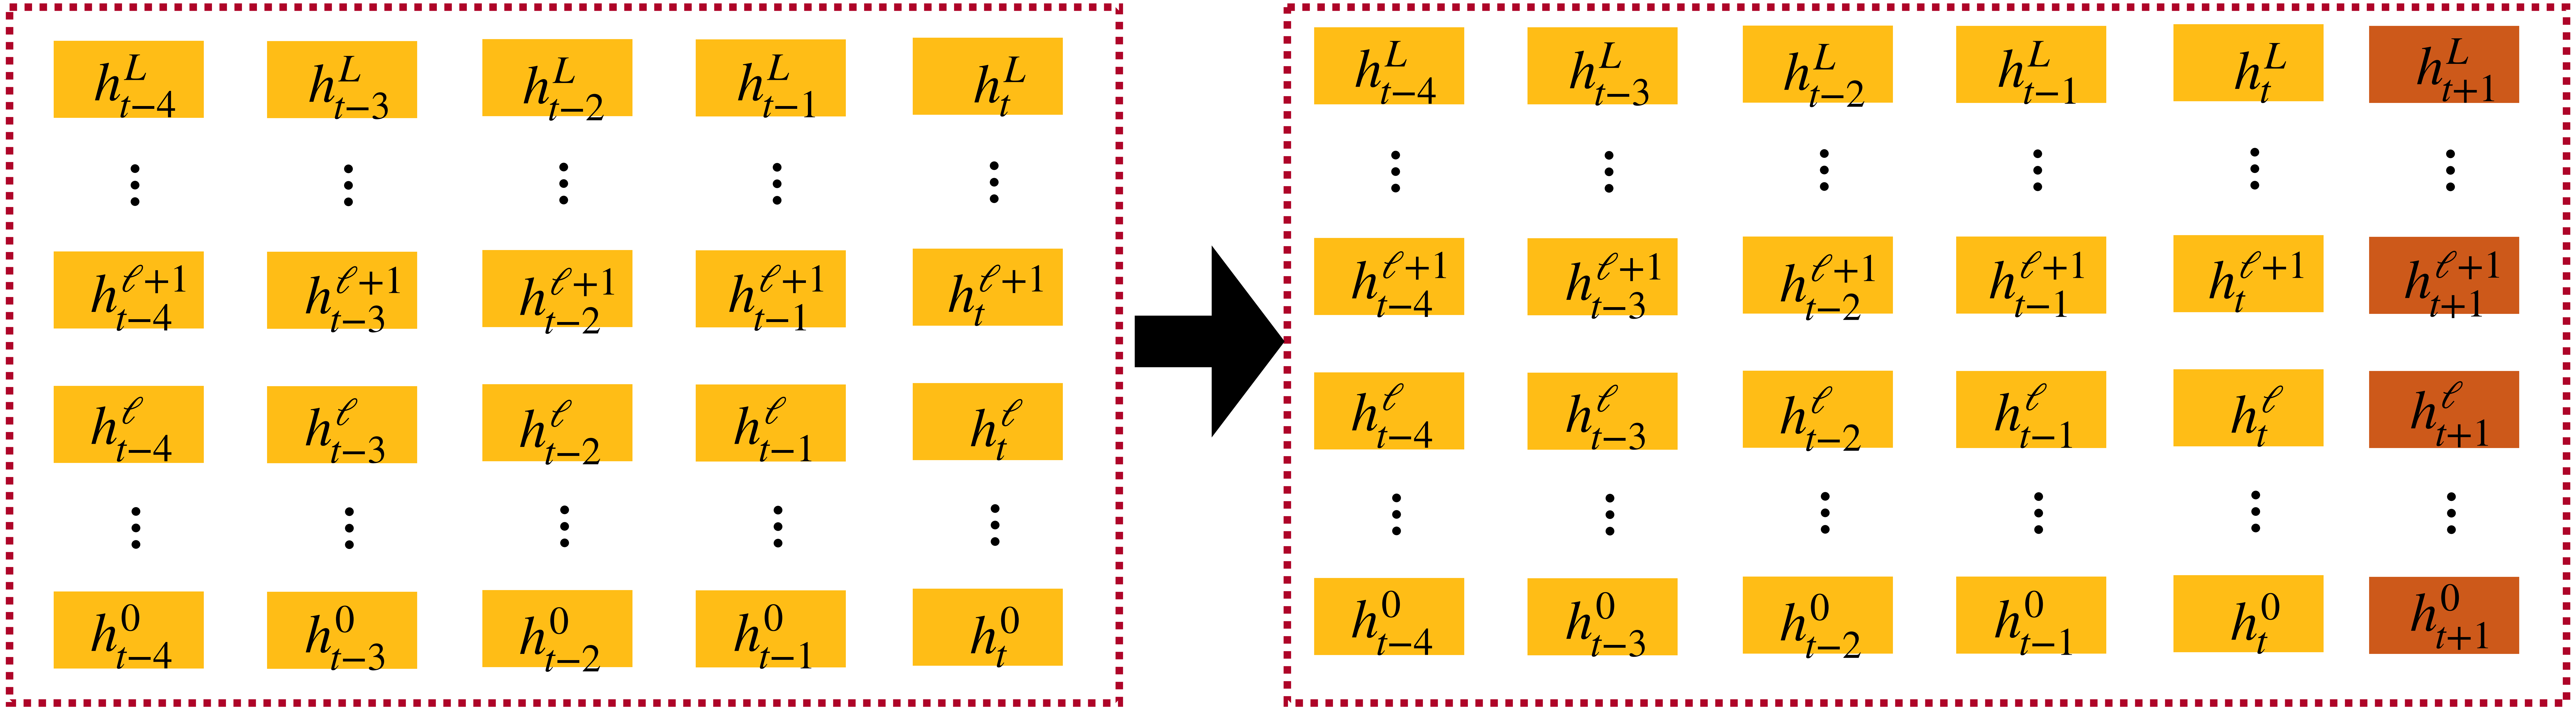

While this can be the state, they are extremely high-dimensional (dim keeps growing)

More importantly, it is not useful

What should be the state of an LLM?

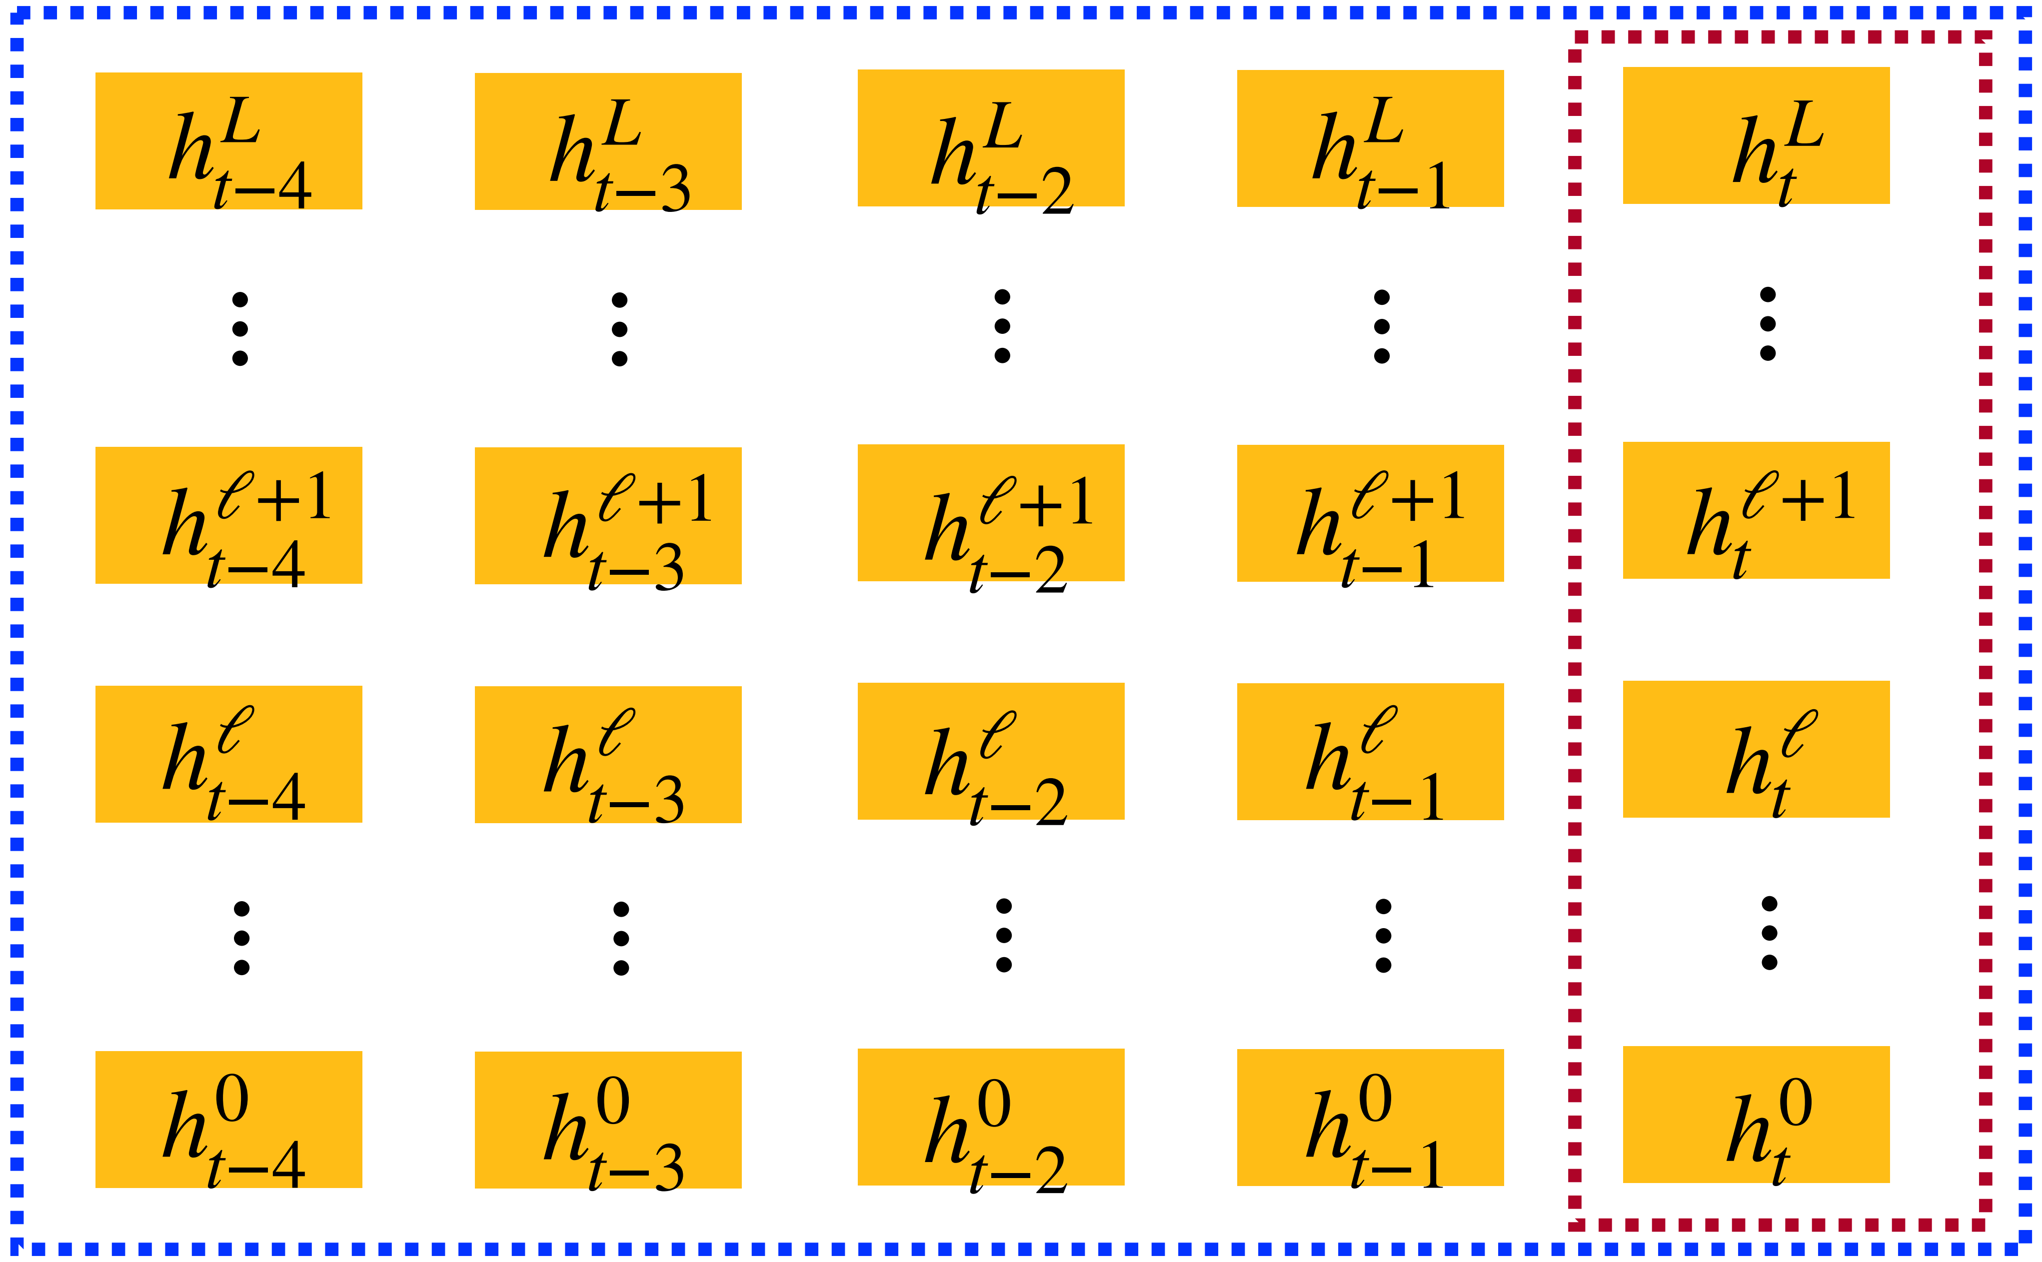

Activation at t is not self-predictive

Full history of activation is too high-dim

What is the right state representation for an LLM?

Inspiration from thermo-dynamics

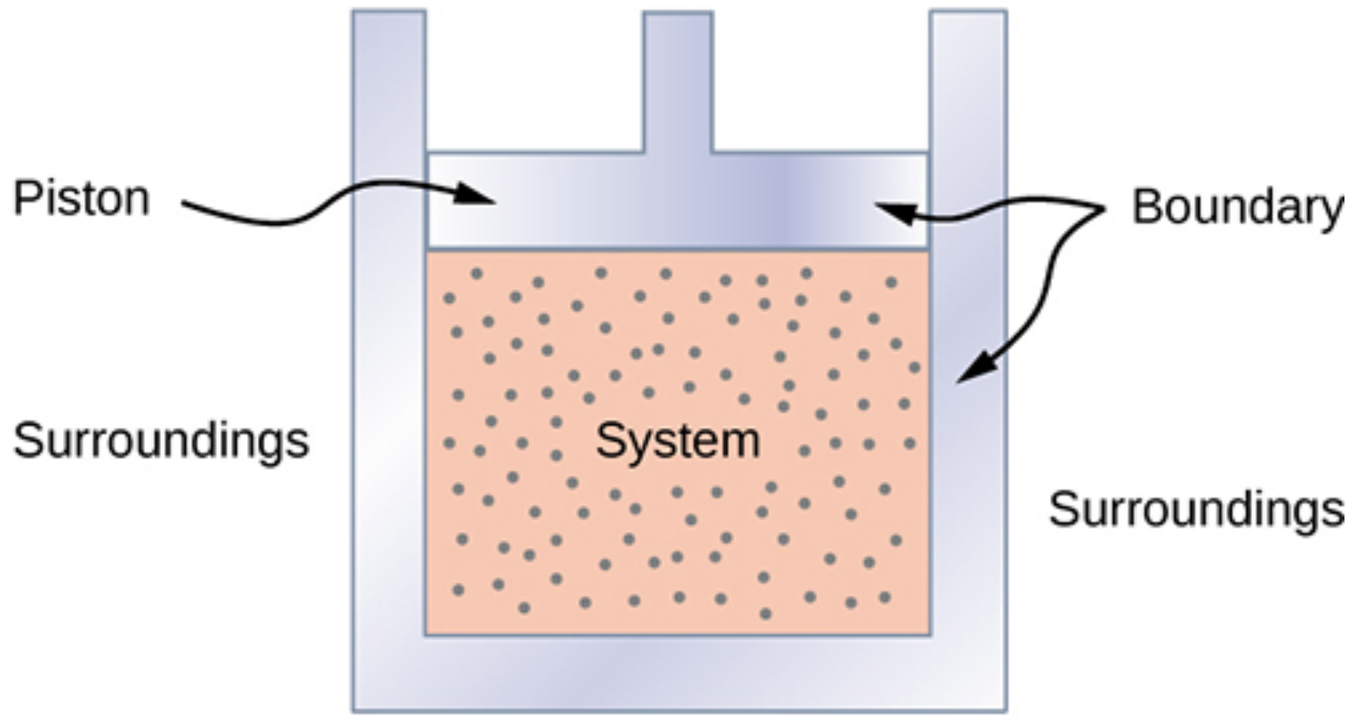

The full state should be the states of billions of molecues

Physicists have defined macroscopic states

State = (Temperature, Pressure, Volume)

- Low dimensional and "high-level"
  Only 3-dim
- Self predictive (dynamically closed)
  Thermodynamical systems
- Practically useful
  We care about these variables

What is the right state representation for an LLM?

What is the right state representation for an LLM?

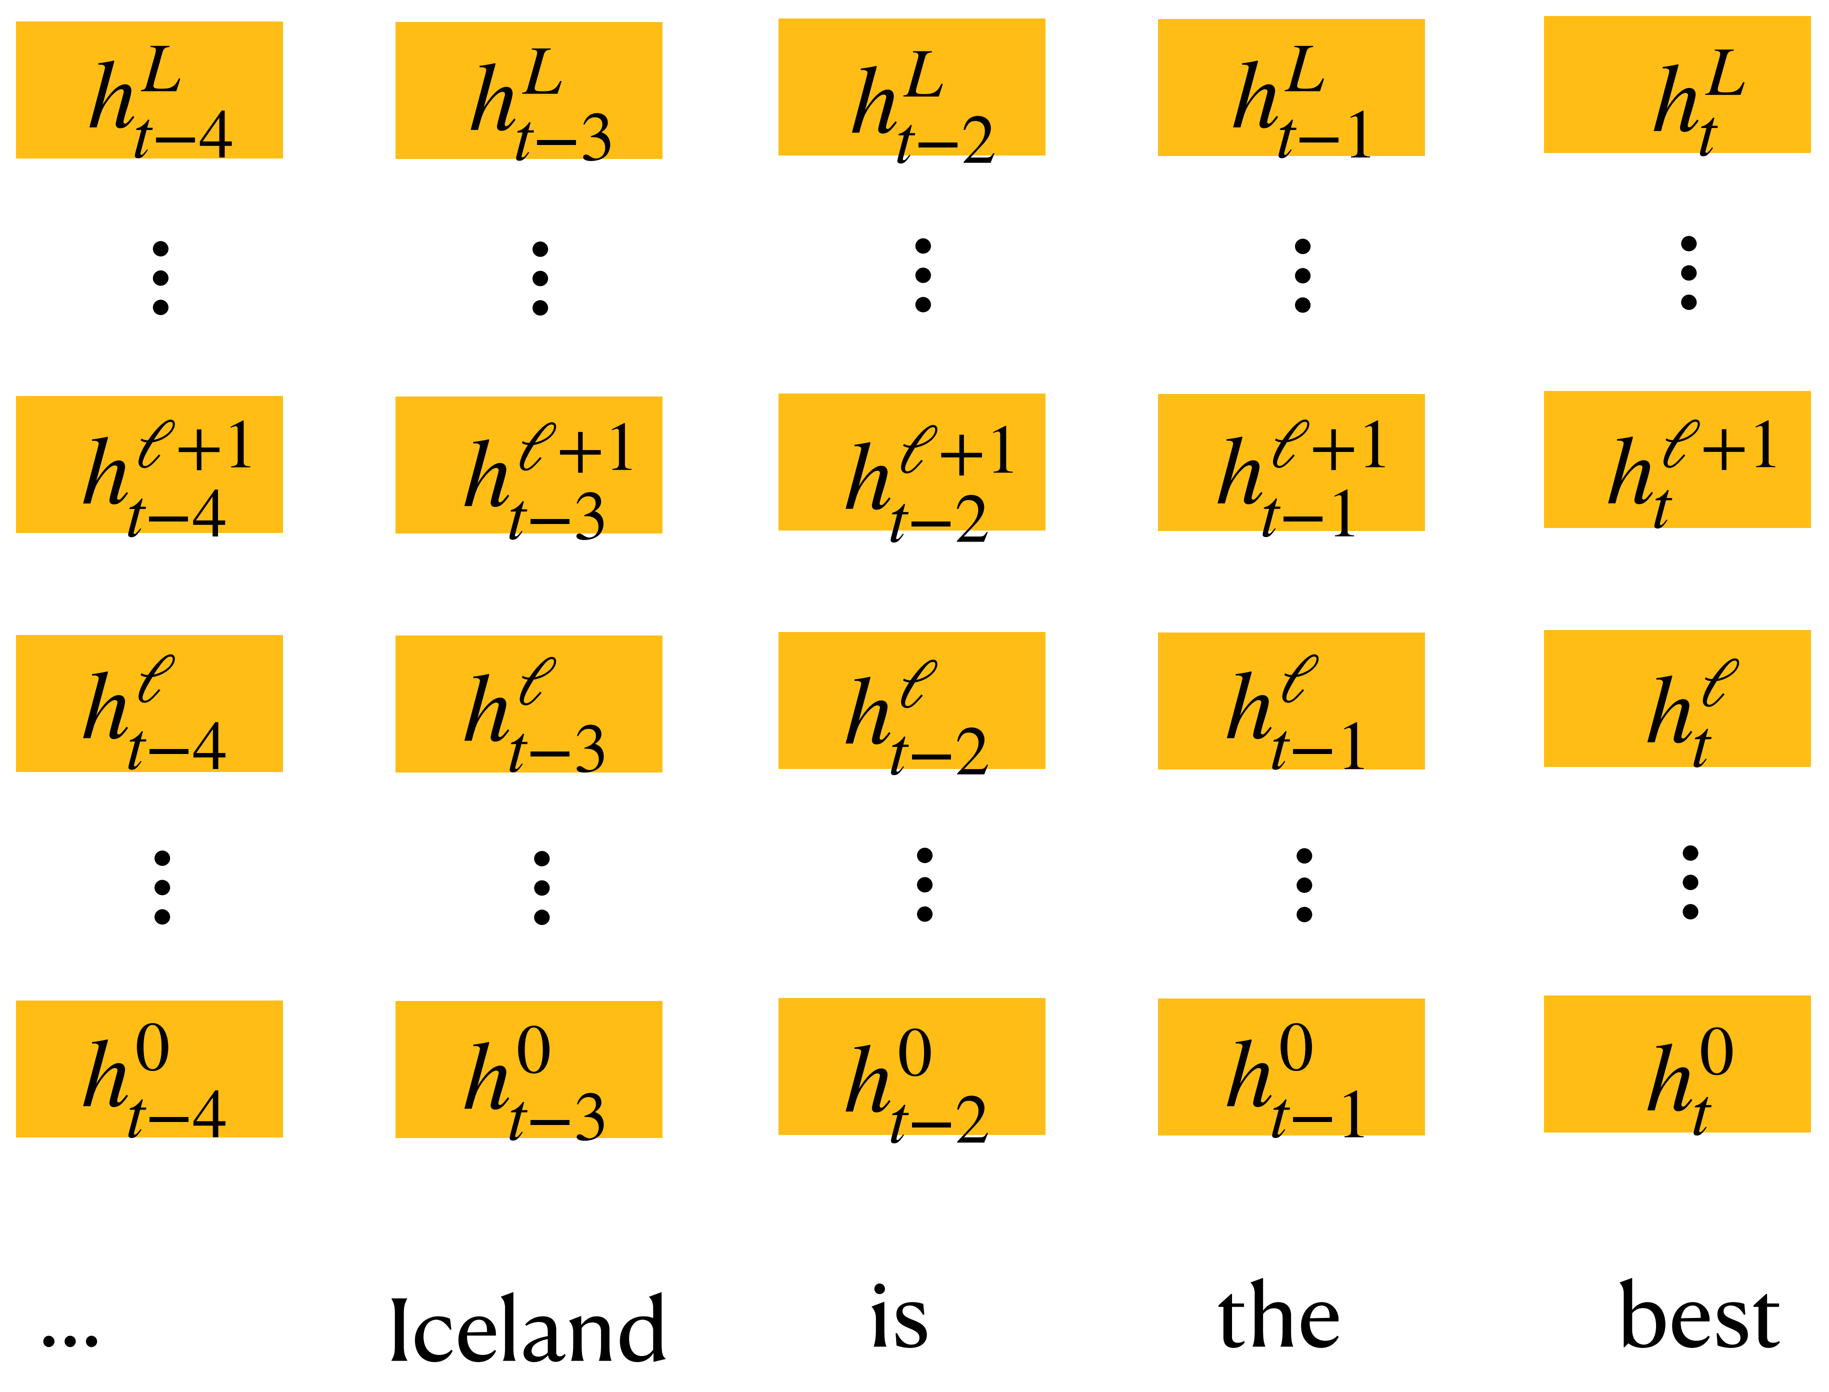

The LLM state we want should be

- Low dimensional and “high-level”
- Self predictive (dynamically closed)
- Practically useful

Introducing RET: Representational Effective Theory

h^L_{t-4}	h^L_{t-3}	h^L_{t-2}	h^L_{t-1}	h^L_t
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$h^{\ell+1}_{t-4}$	$h^{\ell+1}_{t-3}$	$h^{\ell+1}_{t-2}$	$h^{\ell+1}_{t-1}$	$h^{\ell+1}_t$
h^{ℓ}_{t-4}	h^{ℓ}_{t-3}	h^{ℓ}_{t-2}	h^{ℓ}_{t-1}	h^{ℓ}_t
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
h^0_{t-4}	h^0_{t-3}	h^0_{t-2}	h^0_{t-1}	h^0_t
...	Iceland	is	the	best

Main Idea

Learn a state extractor that is encouraged to be **low-dim, high-level, self predictive**

Introducing RET: Representational Effective Theory

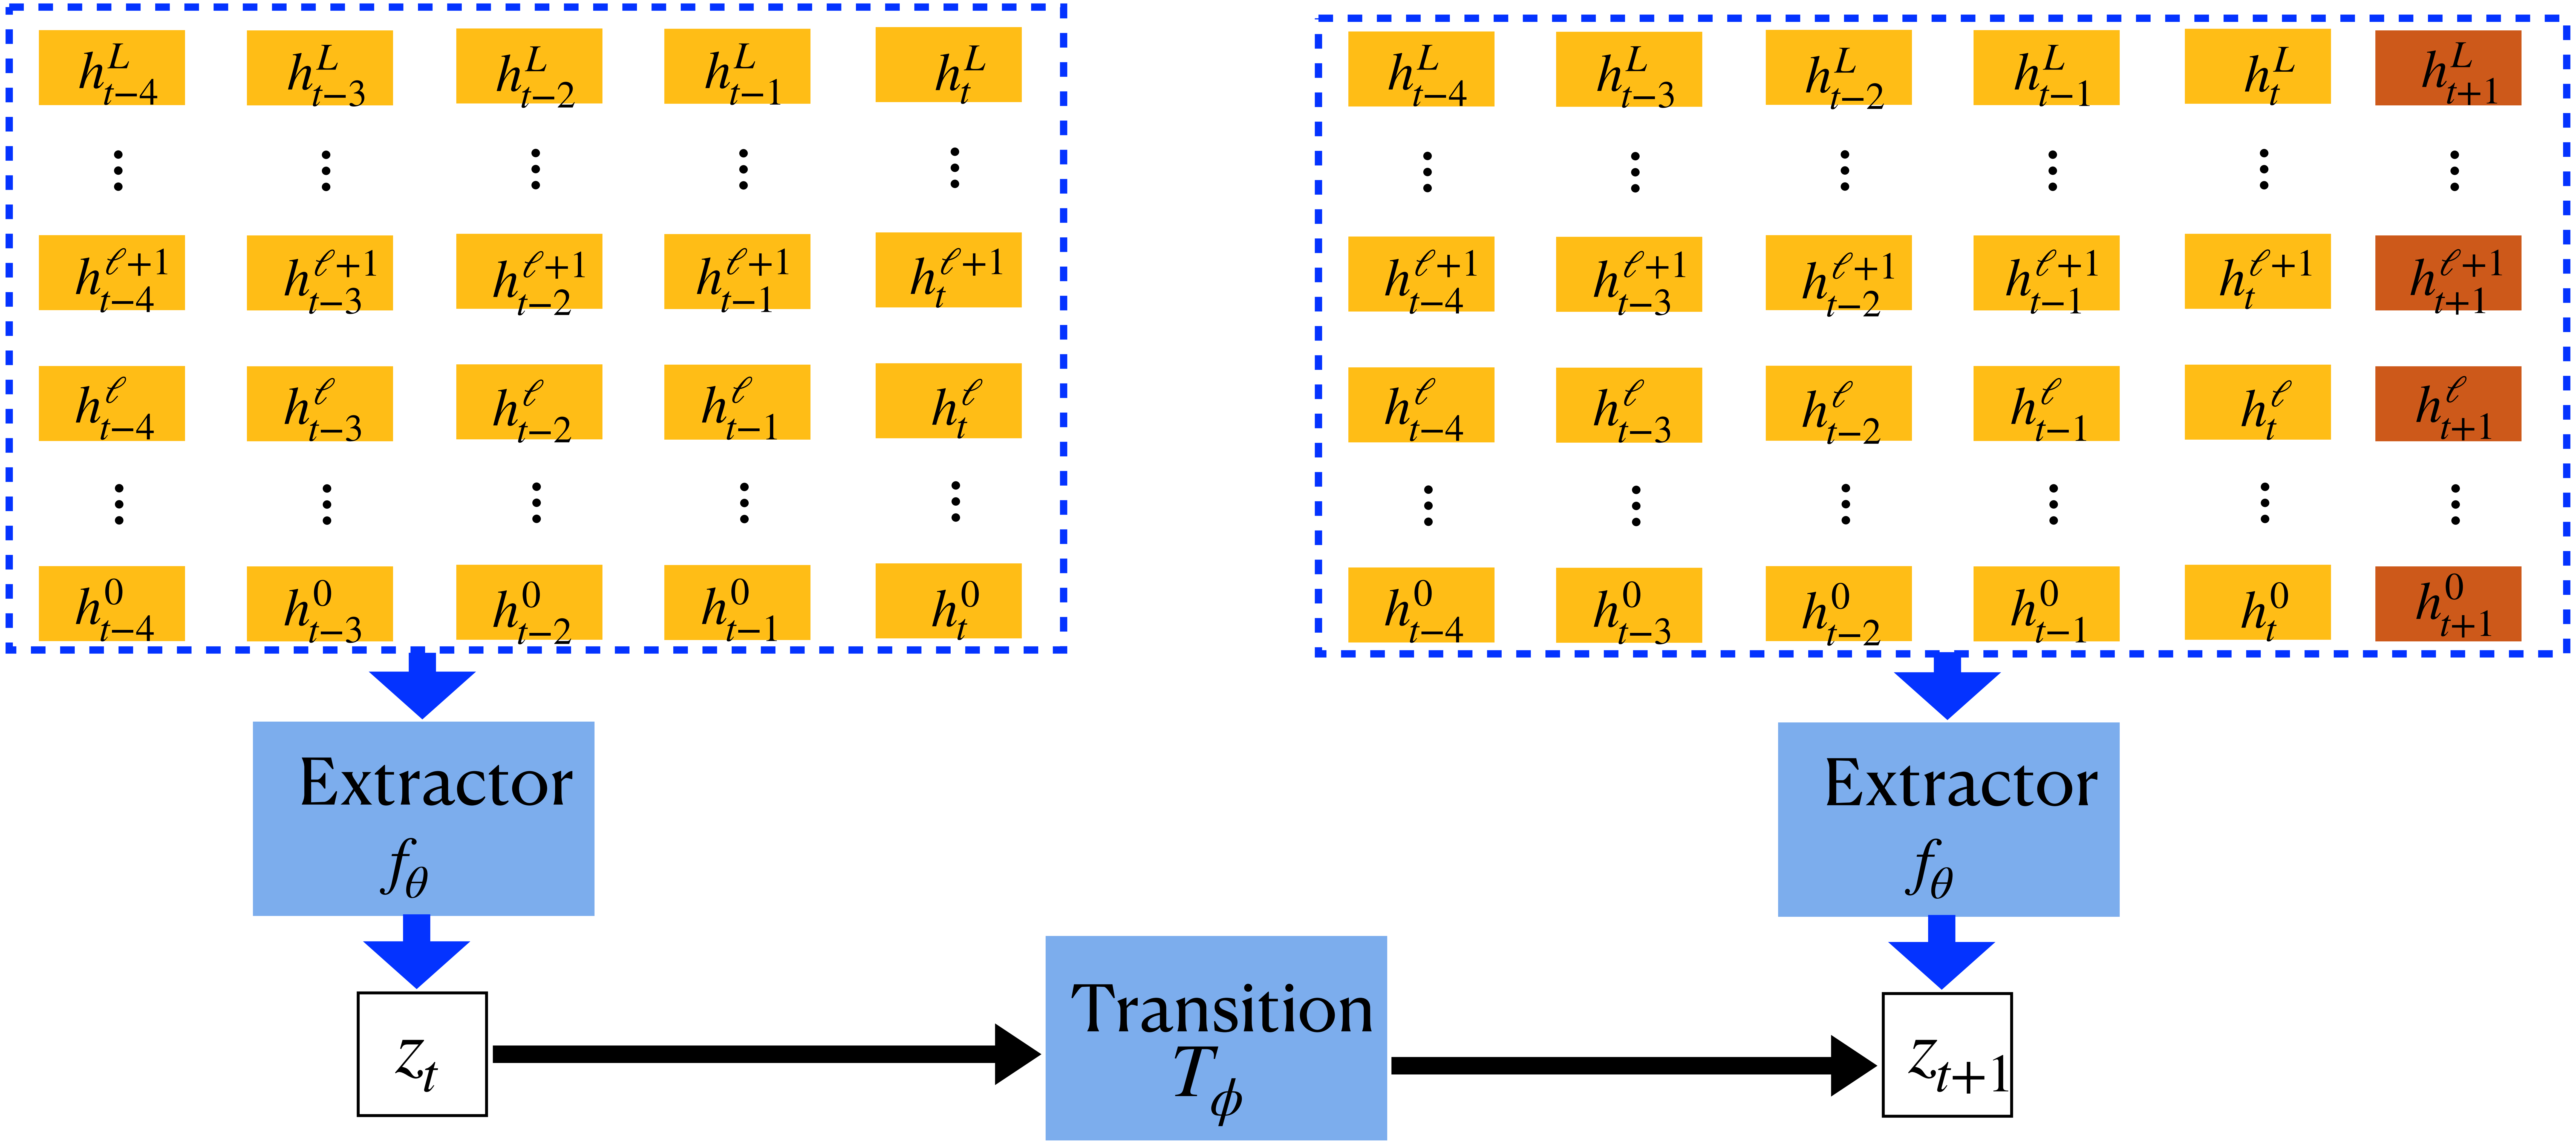

Introducing RET: Representational Effective Theory

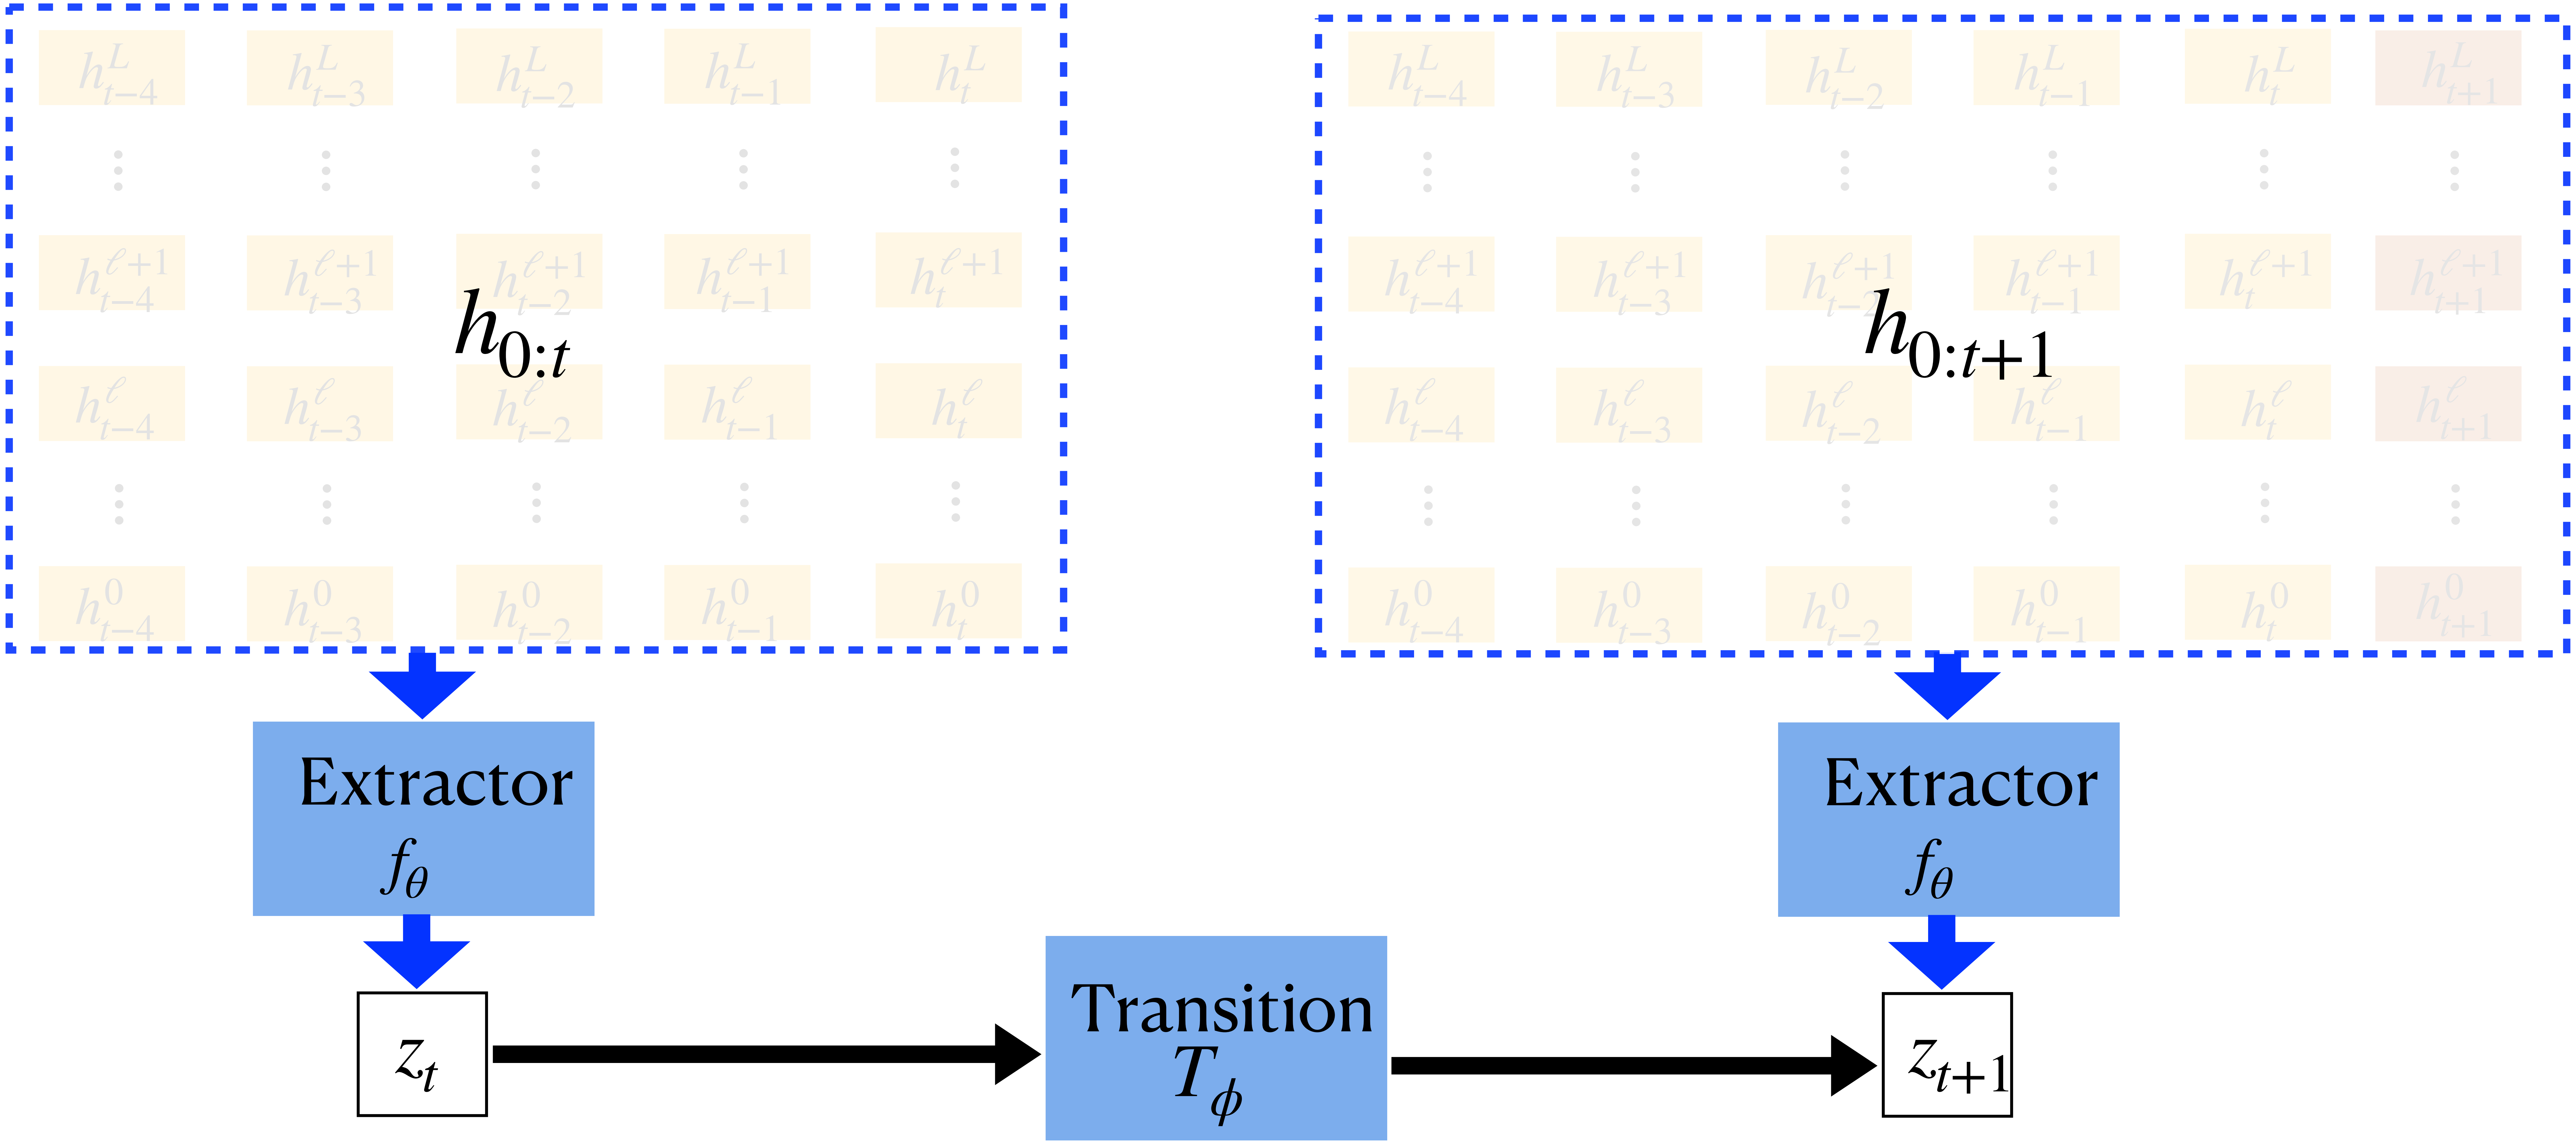

Introducing RET: Representational Effective Theory

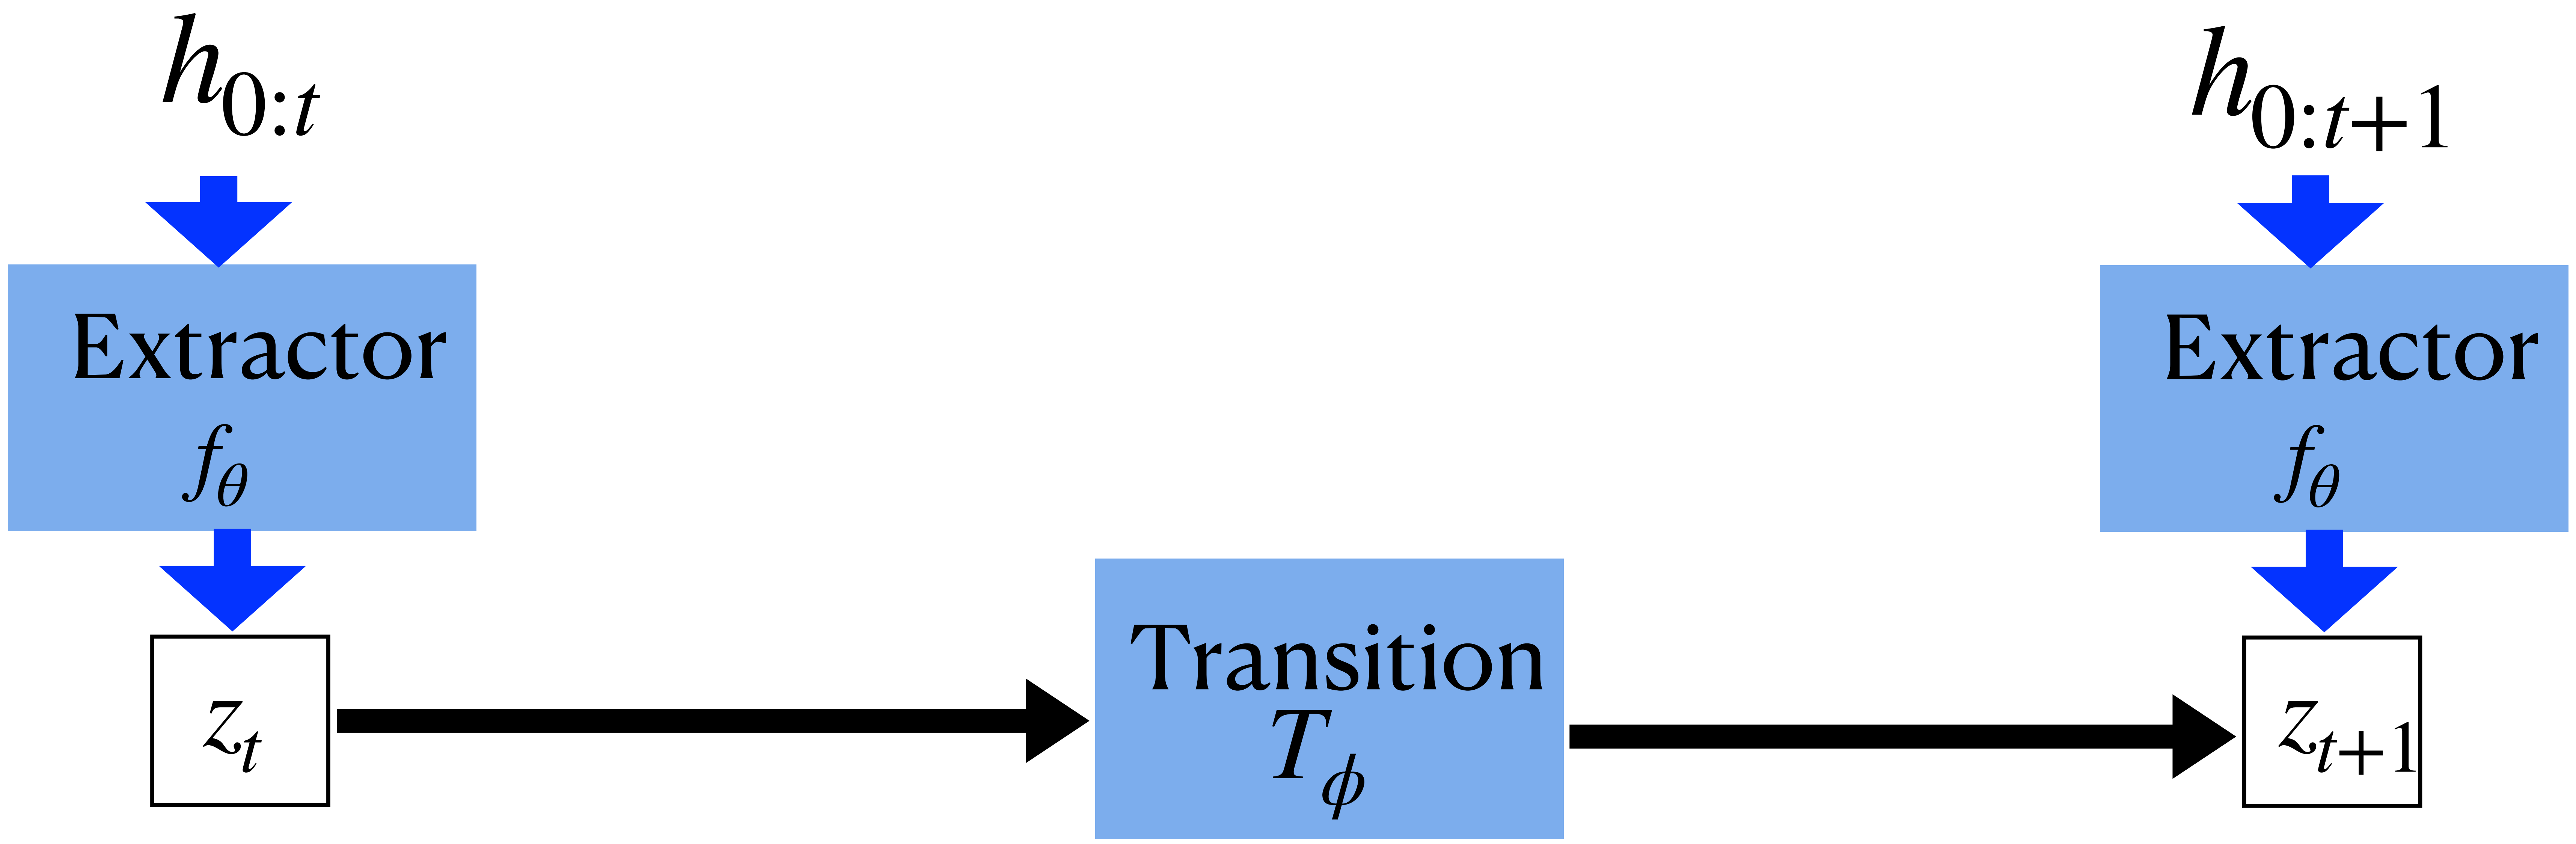

Introducing RET: Representational Effective Theory

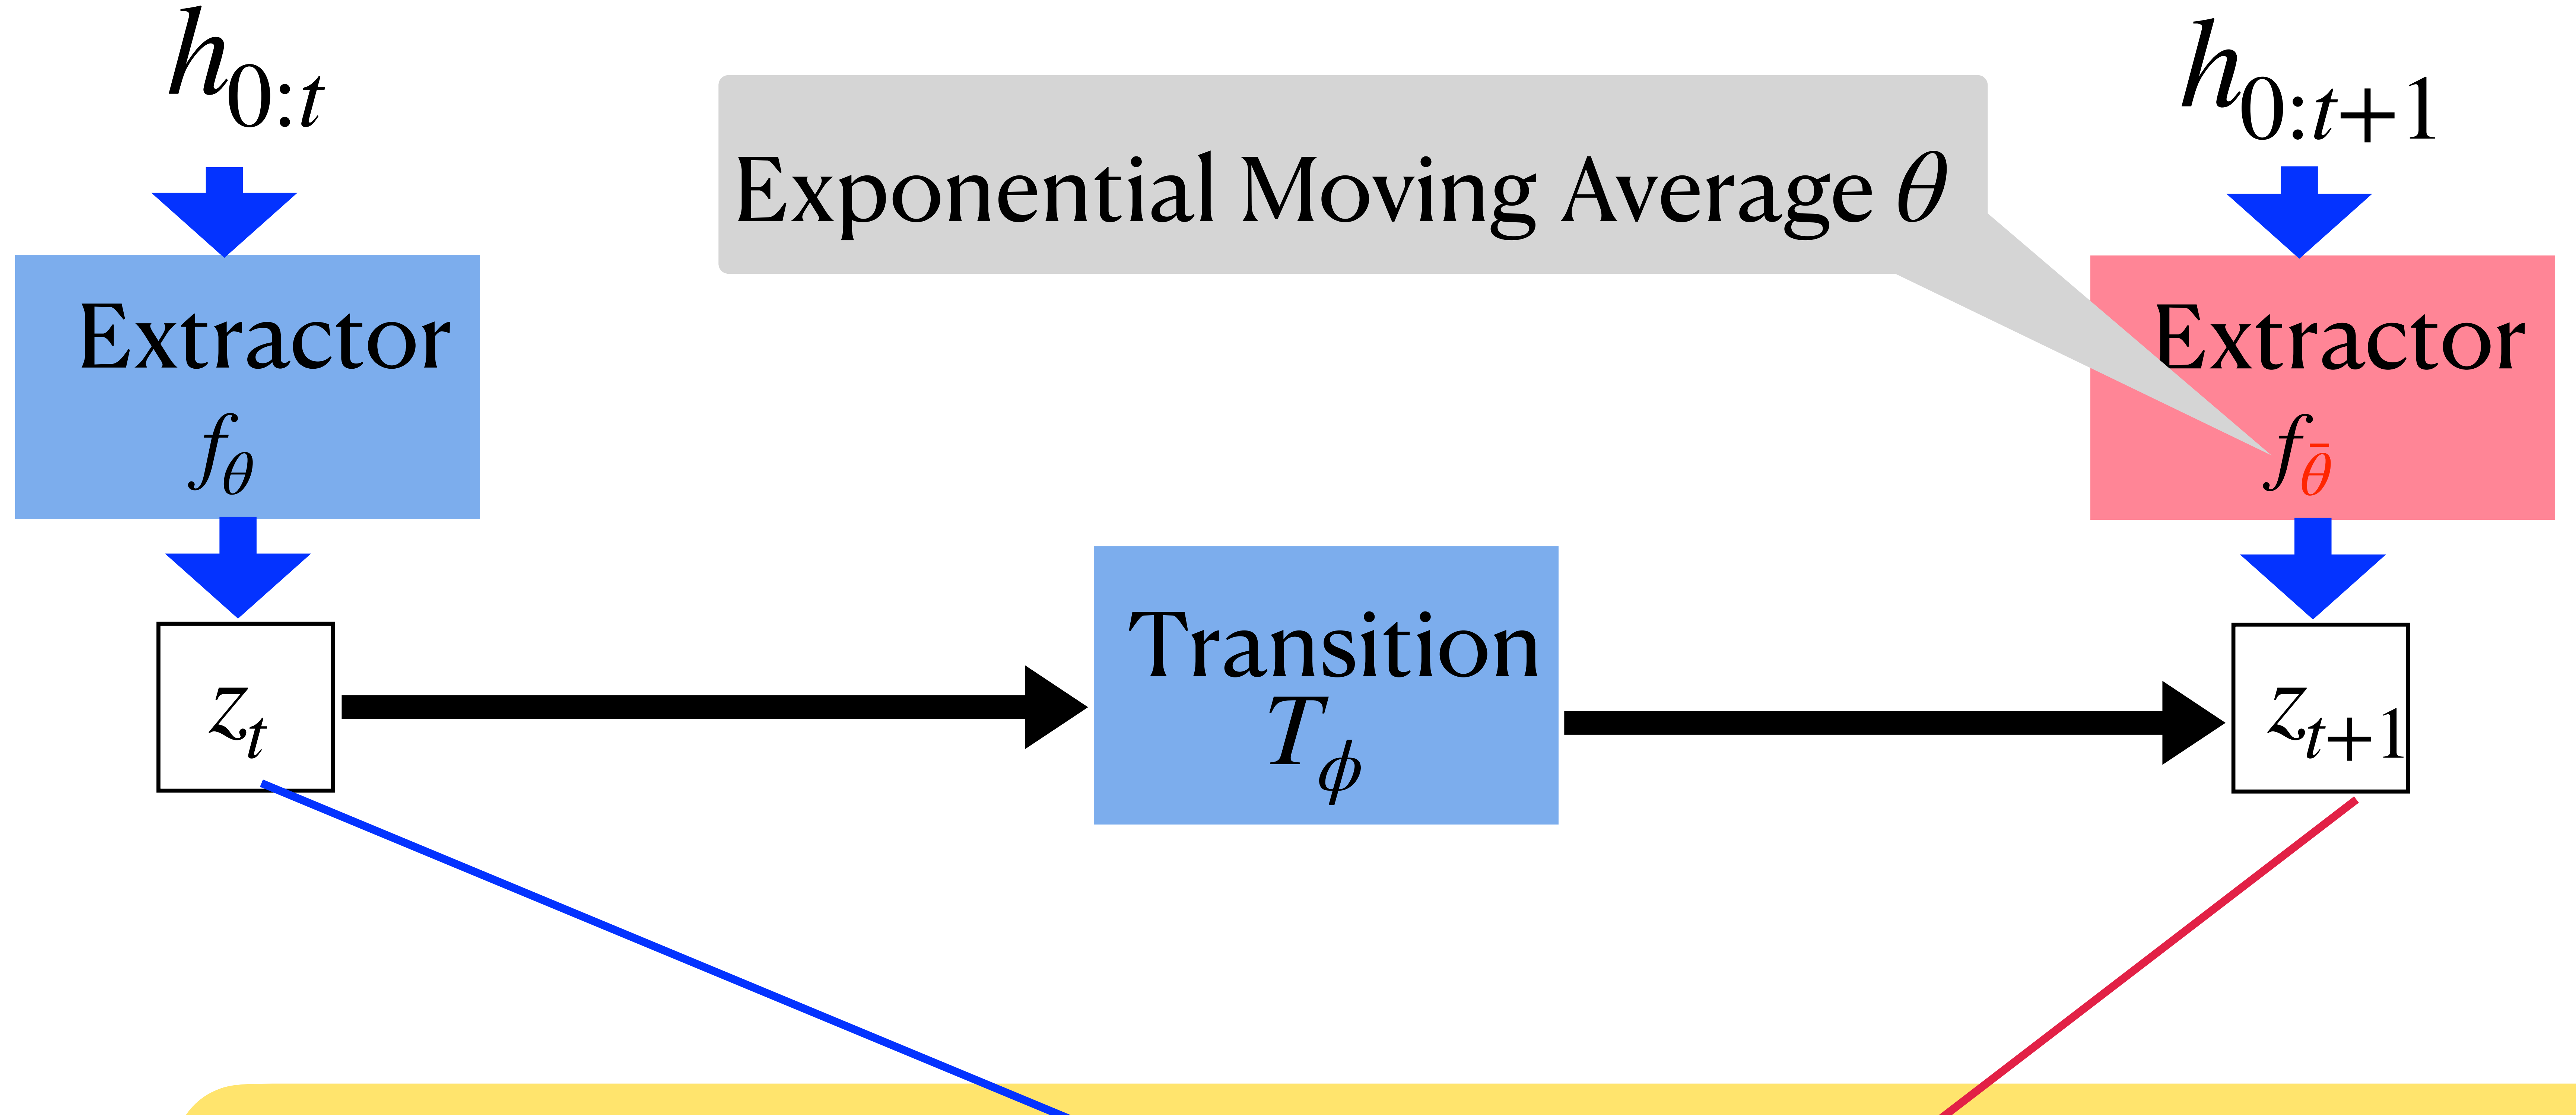

$$\mathscr{L}_{\text{pred}} = \|T_\phi(f_\theta(h_{0:t})) - f_{\bar{\theta}}(h_{0:t+1})\|_1.$$

Loss: Prediction Loss with EMA target

Introducing RET: Representational Effective Theory

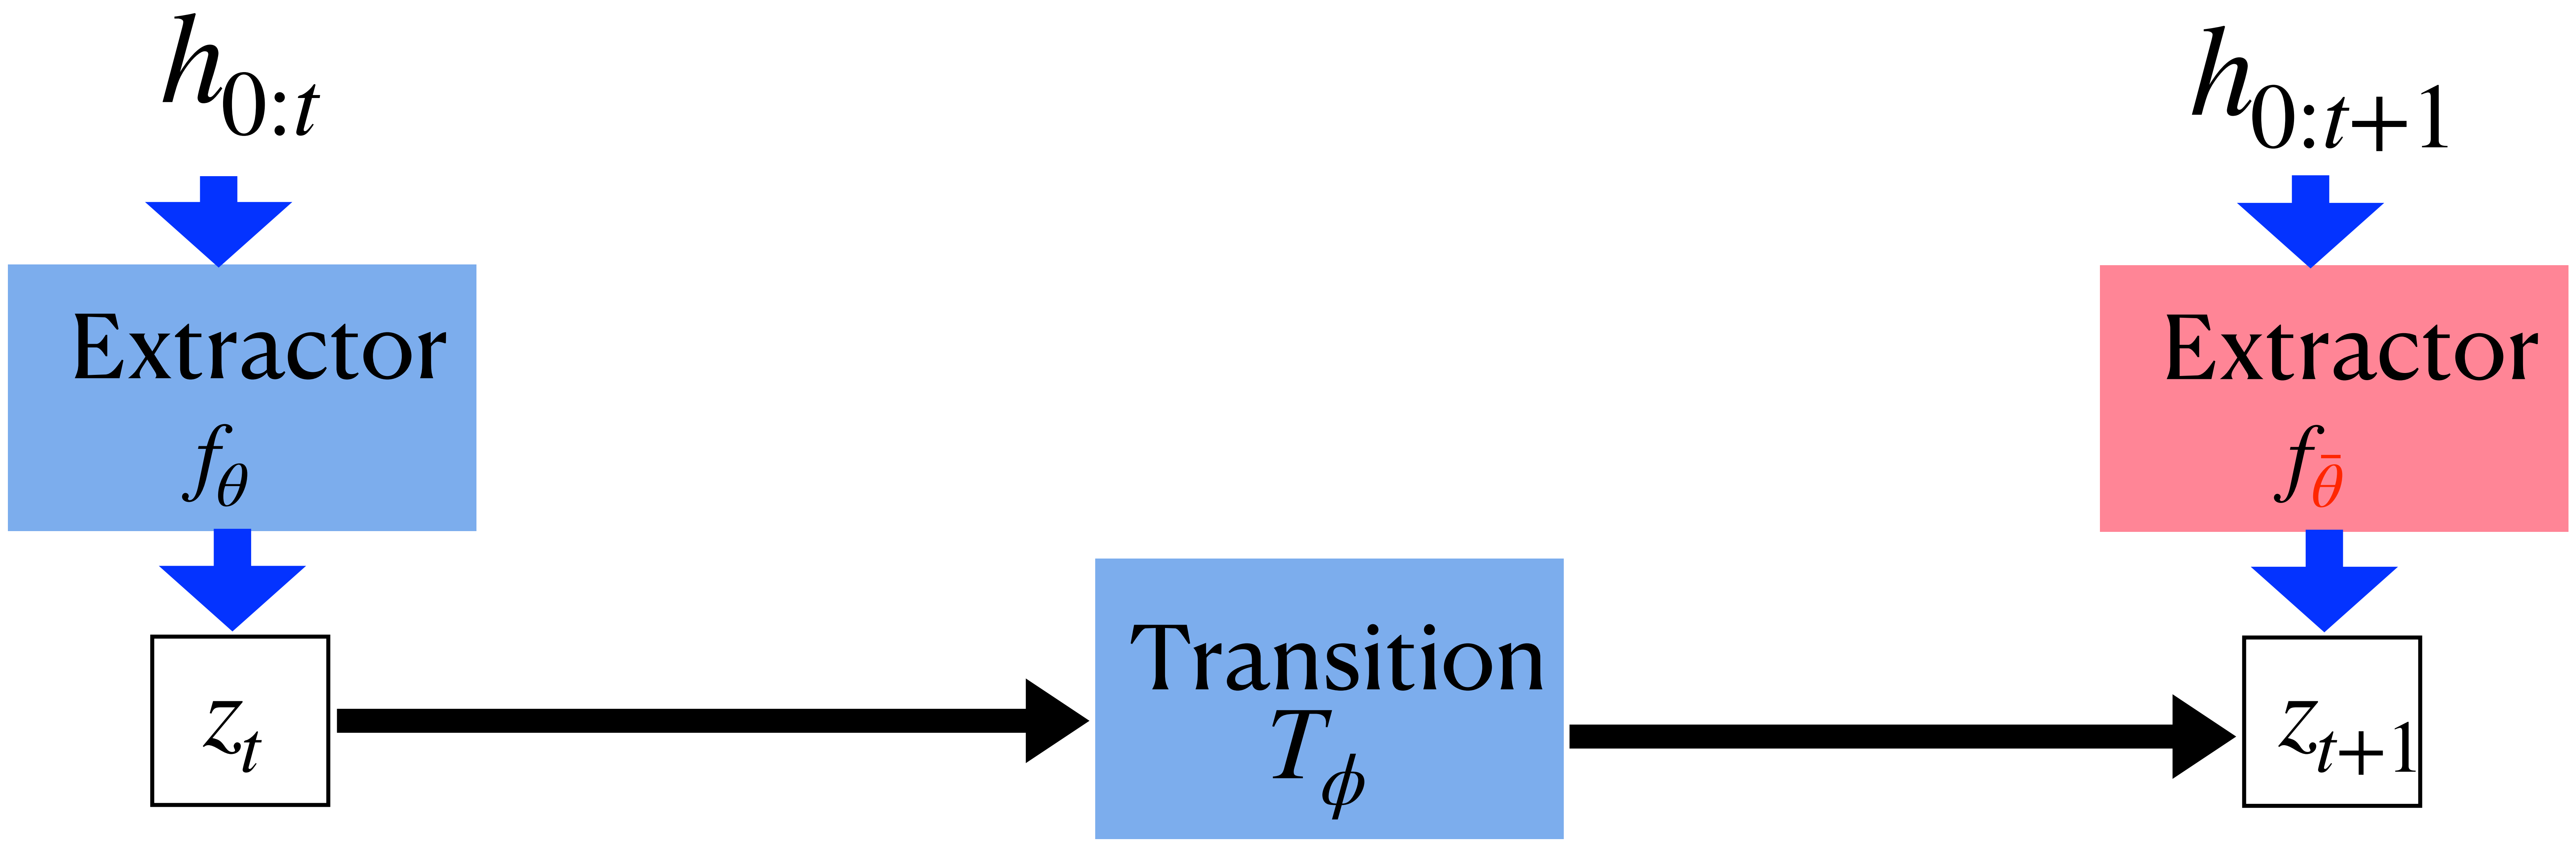

Rationale

- Low dim z: encourages low-dim summary
- Self-predictive: only keep what is important for the future behavior, ignoring microscopic details (like grammar, specific word use)
- EMA target: prevent representation collapse (similar to JEPA)

Evaluating RET

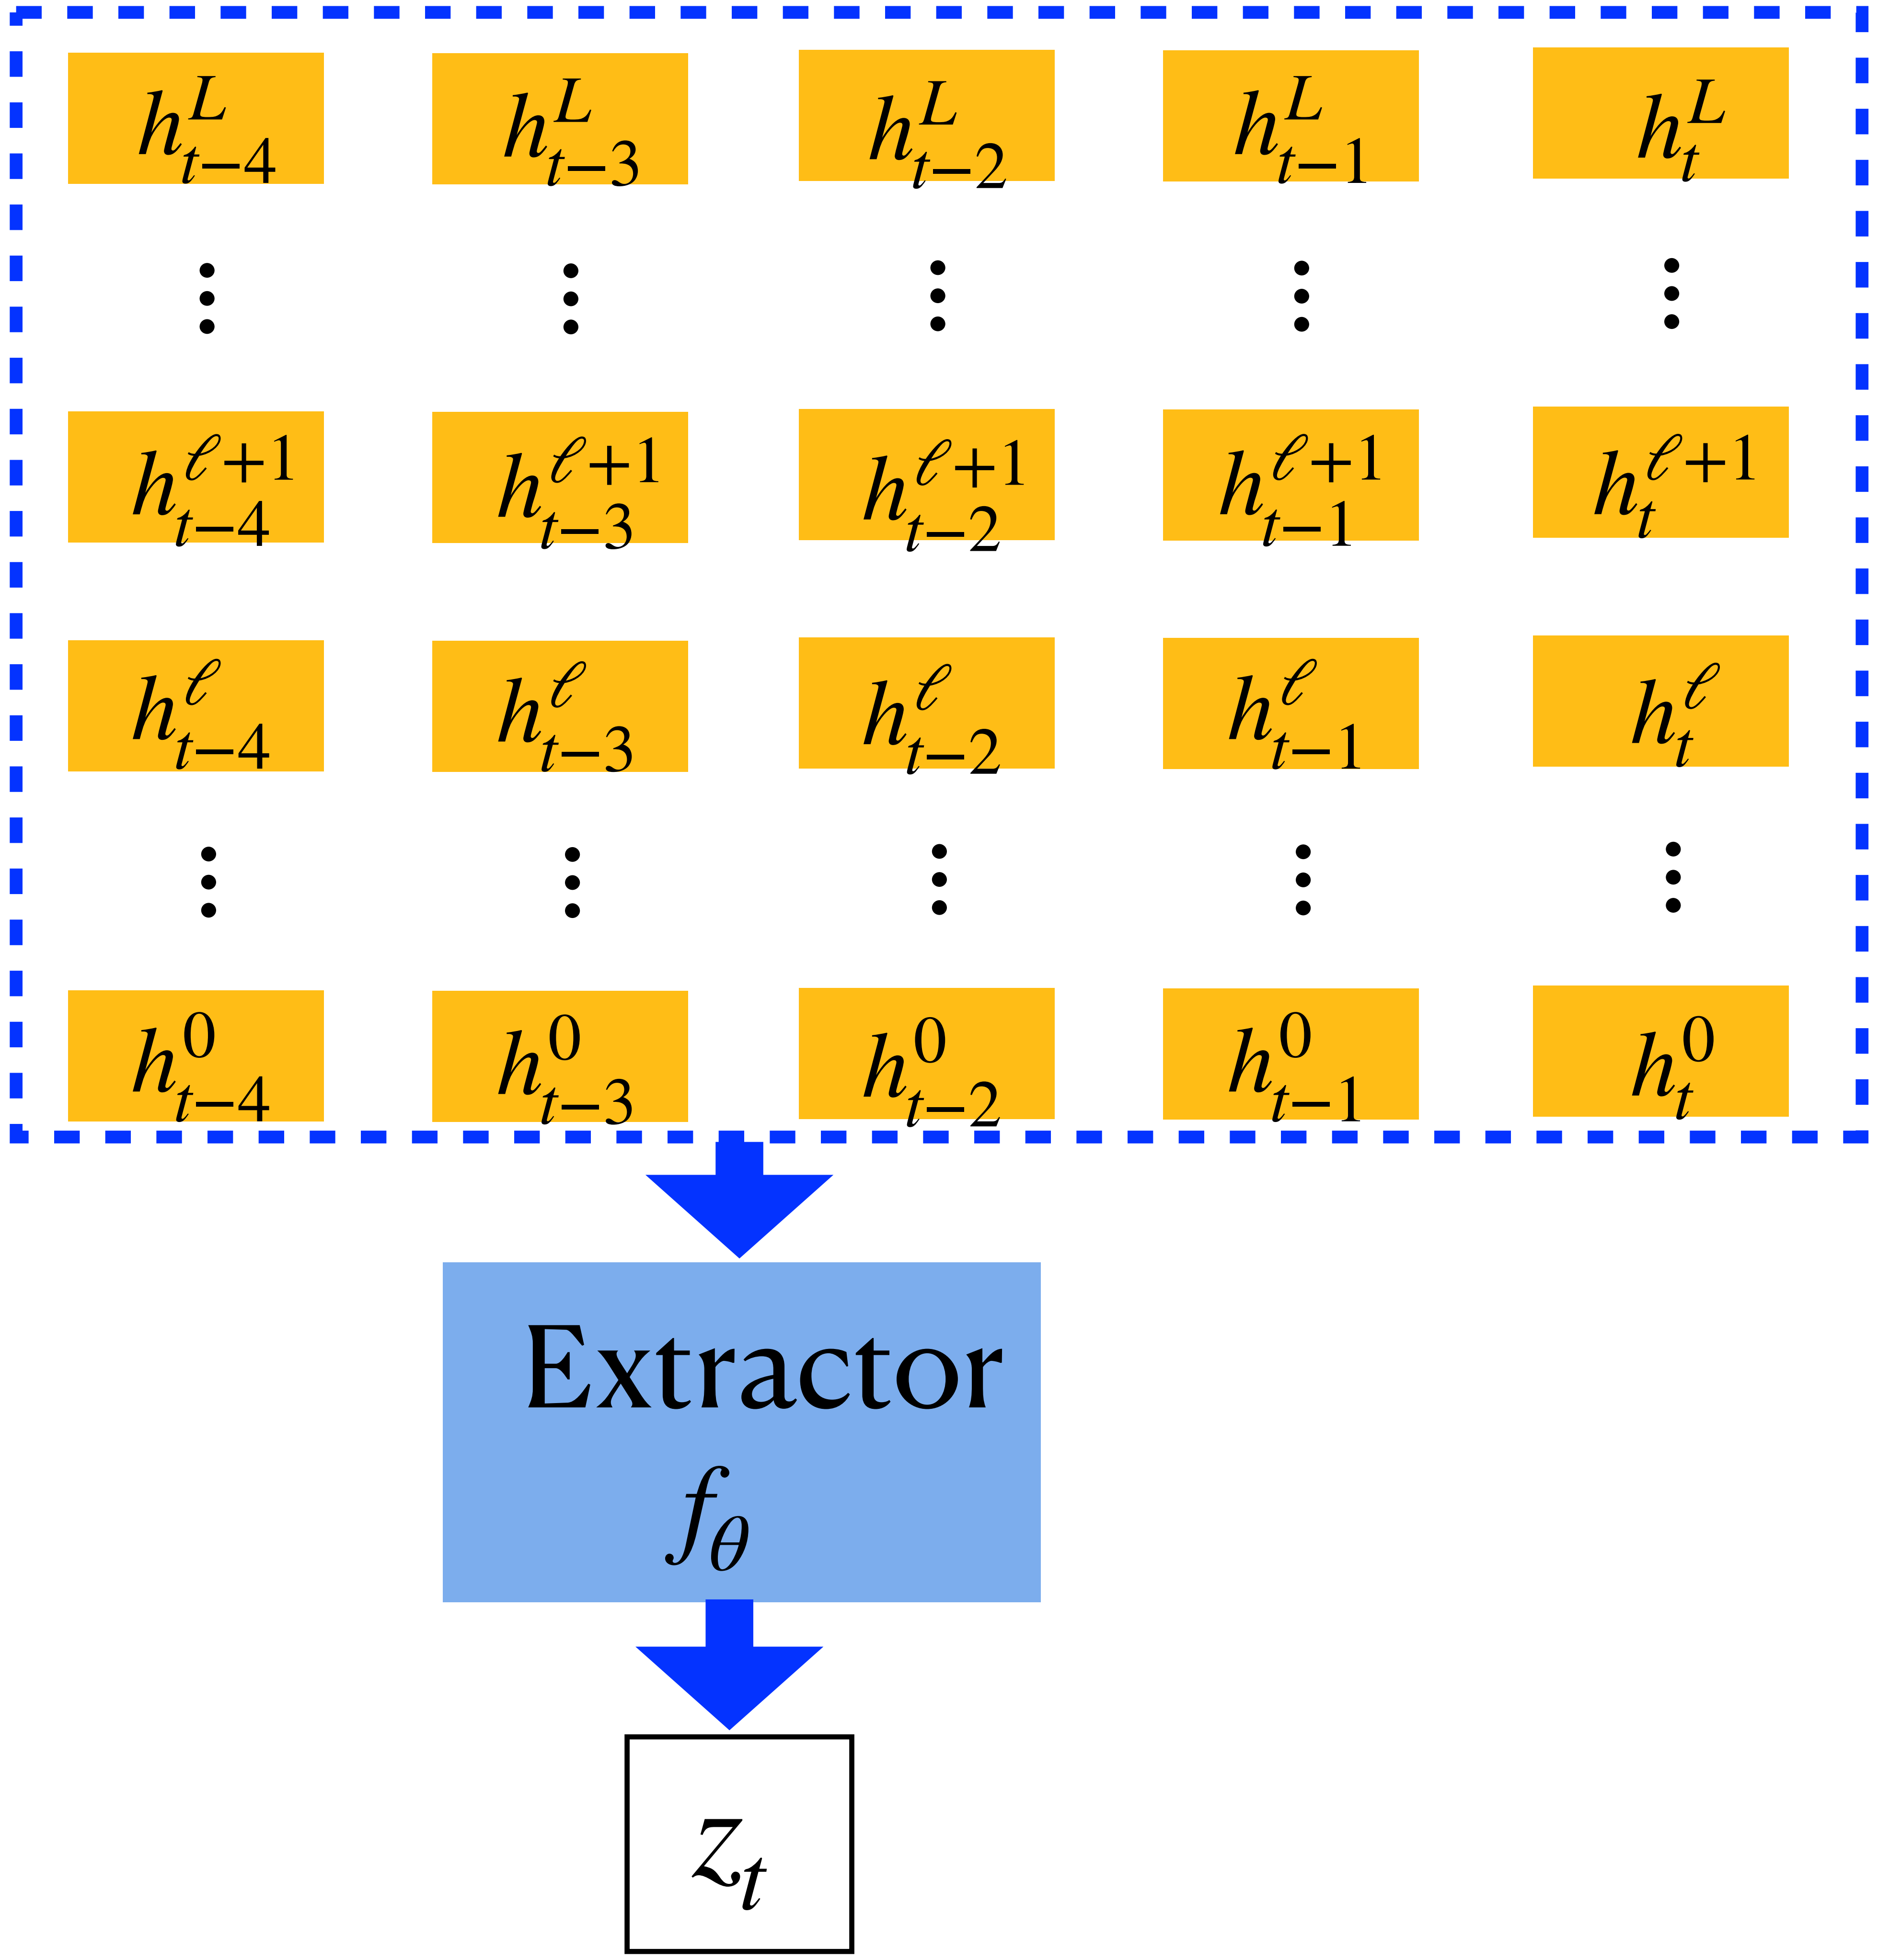

Are the learned states…

- Low dimensional and "high-level"? **Up next!**
- Self predictive (dynamically closed)?
- Practically useful?

RET Captures High-level Plot

We train RET on **TinyStories** dataset for **Pythia-160M**

Then use RET to extract z_t for the following:

The young boys scored their third goal of the afternoon match and celebrated wildly by running in big circles around the grassy soccer field.

A torrential downpour drenched the field within seconds, forcing everyone to abandon the game and quickly run toward the nearest open building for shelter.

They stepped into the quiet public library and spent the rest of the rainy afternoon reading dusty old books about forgotten sailors and pirate legends

RET Captures High-level Plot

The young boys scored their third goal of the afternoon match and celebrated wildly by running in big circles around the grassy soccer field.

A torrential downpour drenched the field within seconds, forcing everyone to abandon the game and quickly run toward the nearest open building for shelter.

They stepped into the quiet public library and spent the rest of the rainy afternoon reading dusty old books about forgotten sailors and pirate legends

States mixed

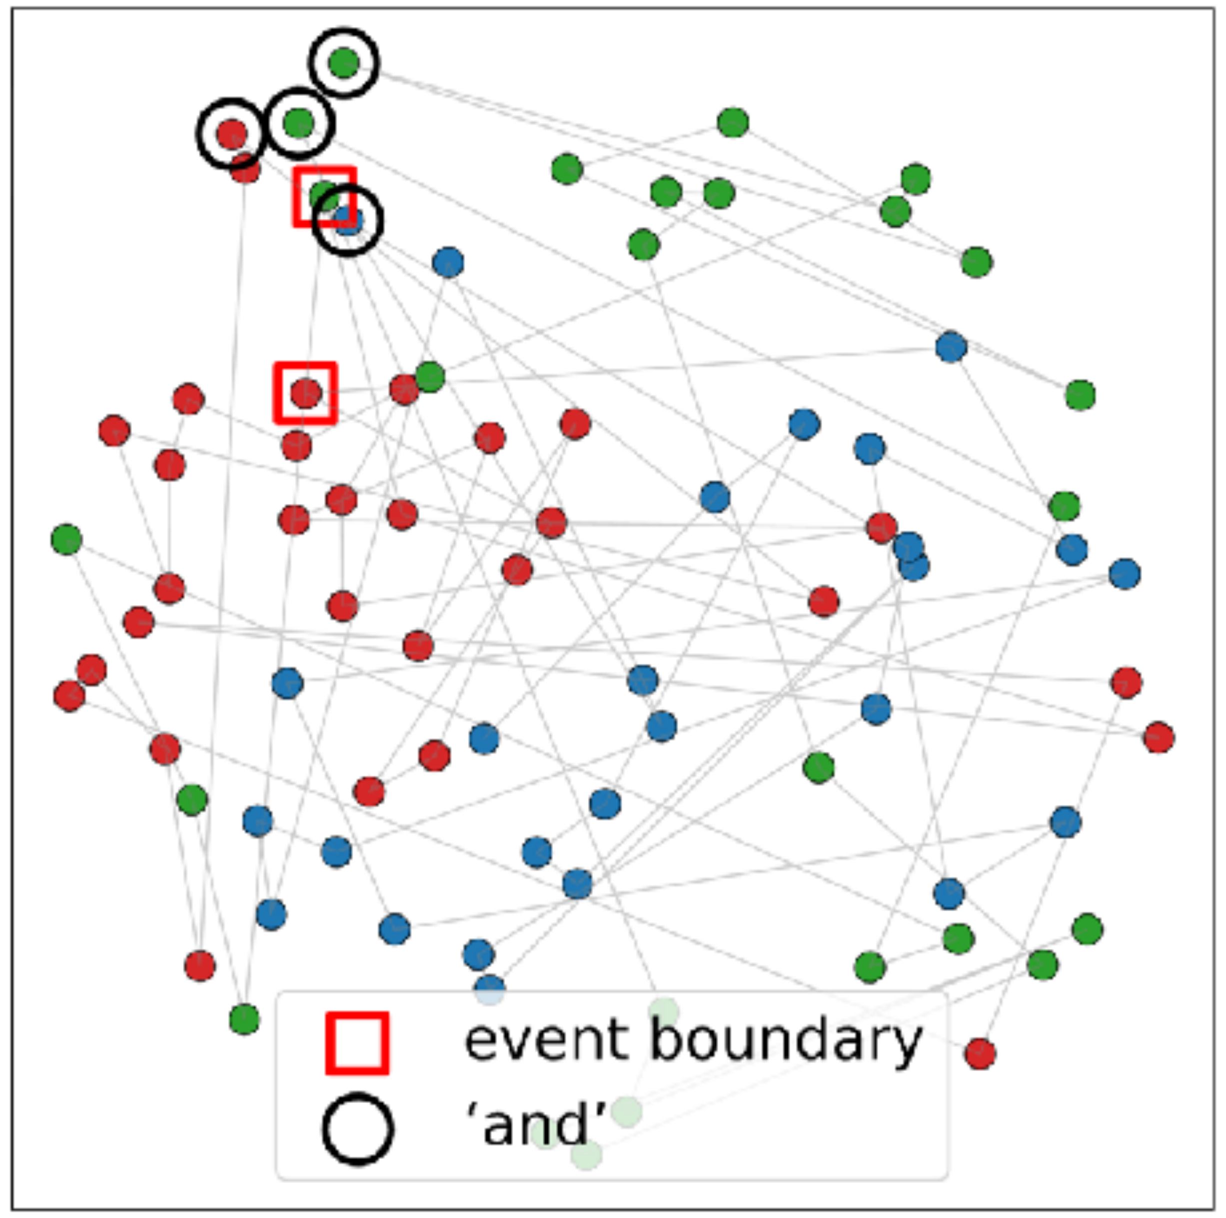

State trajectory of h_t
(UMAP plot)

States cluster according to story plot

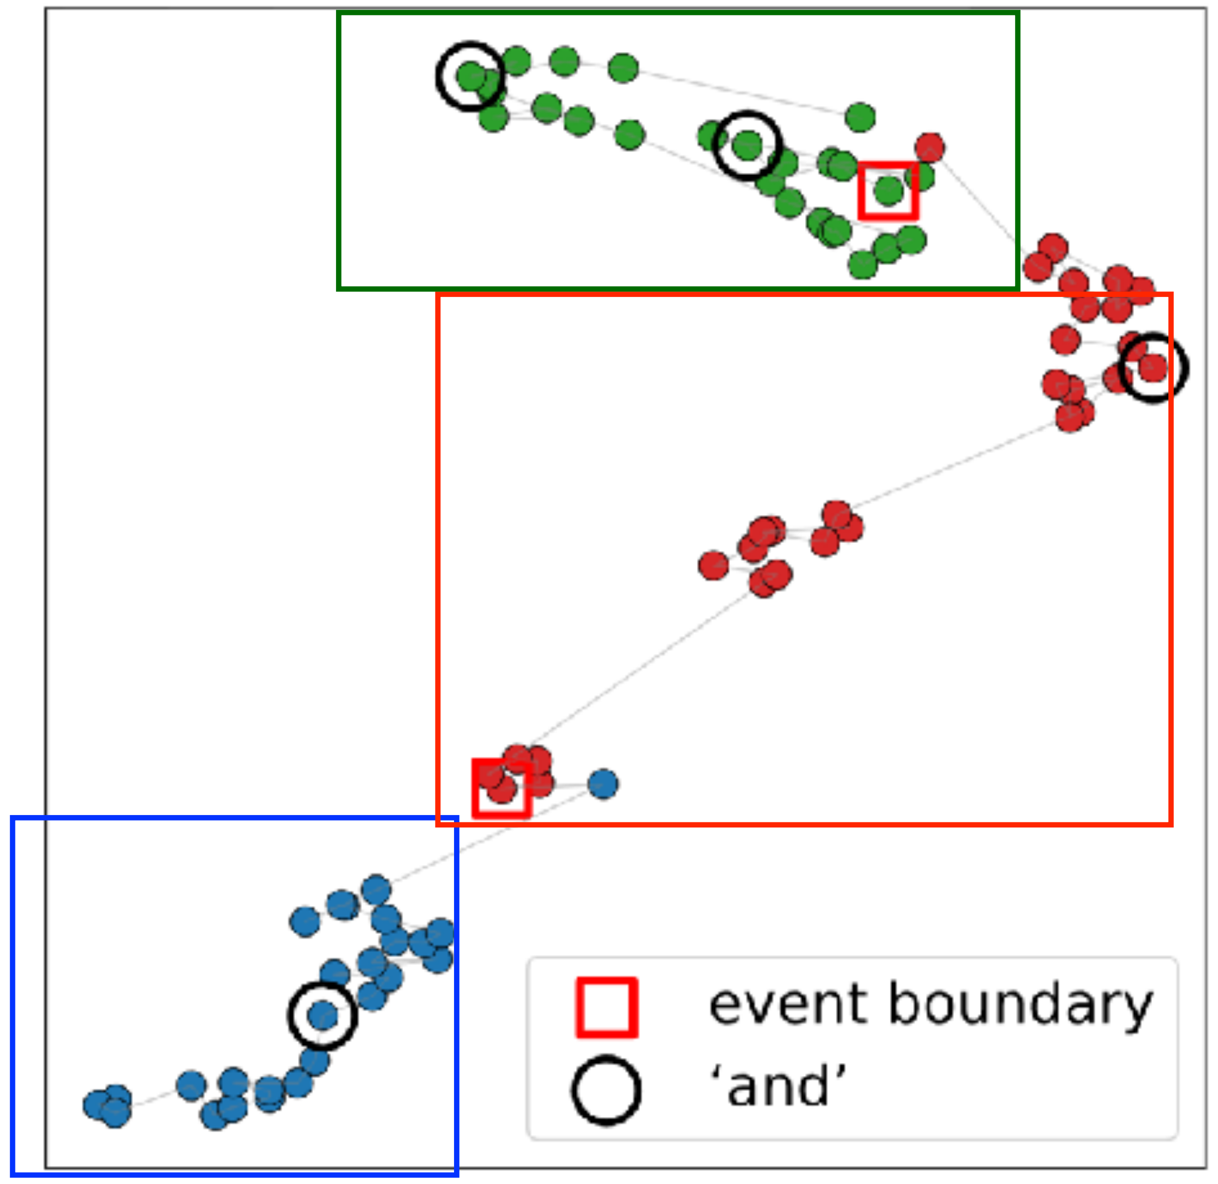

State trajectory of z_t (RET)
(UMAP plot)

RET Captures High-level Plot

In contrast, the h_t of "and" positions are put together

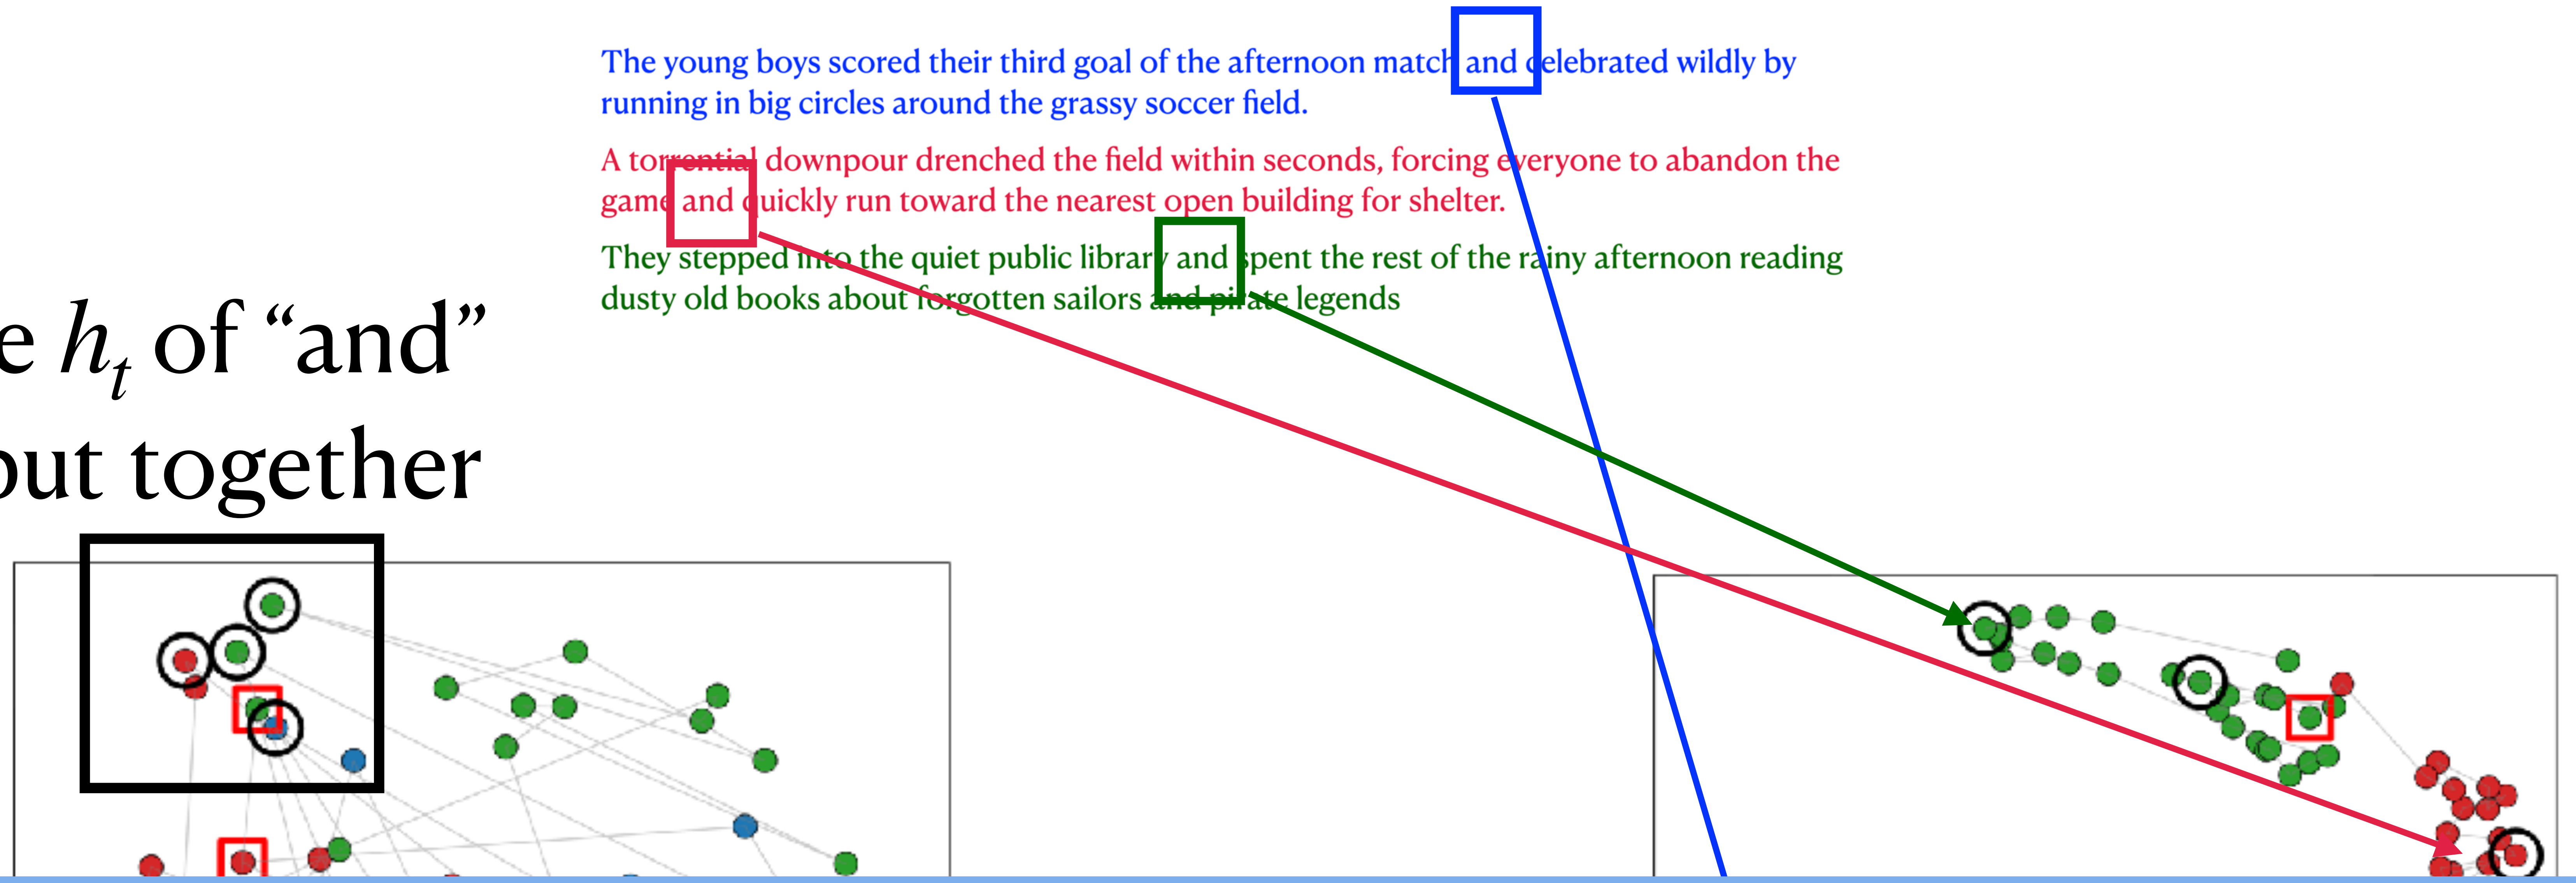

The raw activation h_t reflects low level word choices ("and")

The learned high-level state z_t represents **high-level story plots**

State trajectory of h_t

State trajectory of z_t (RET)

RET Captures High-level Plot

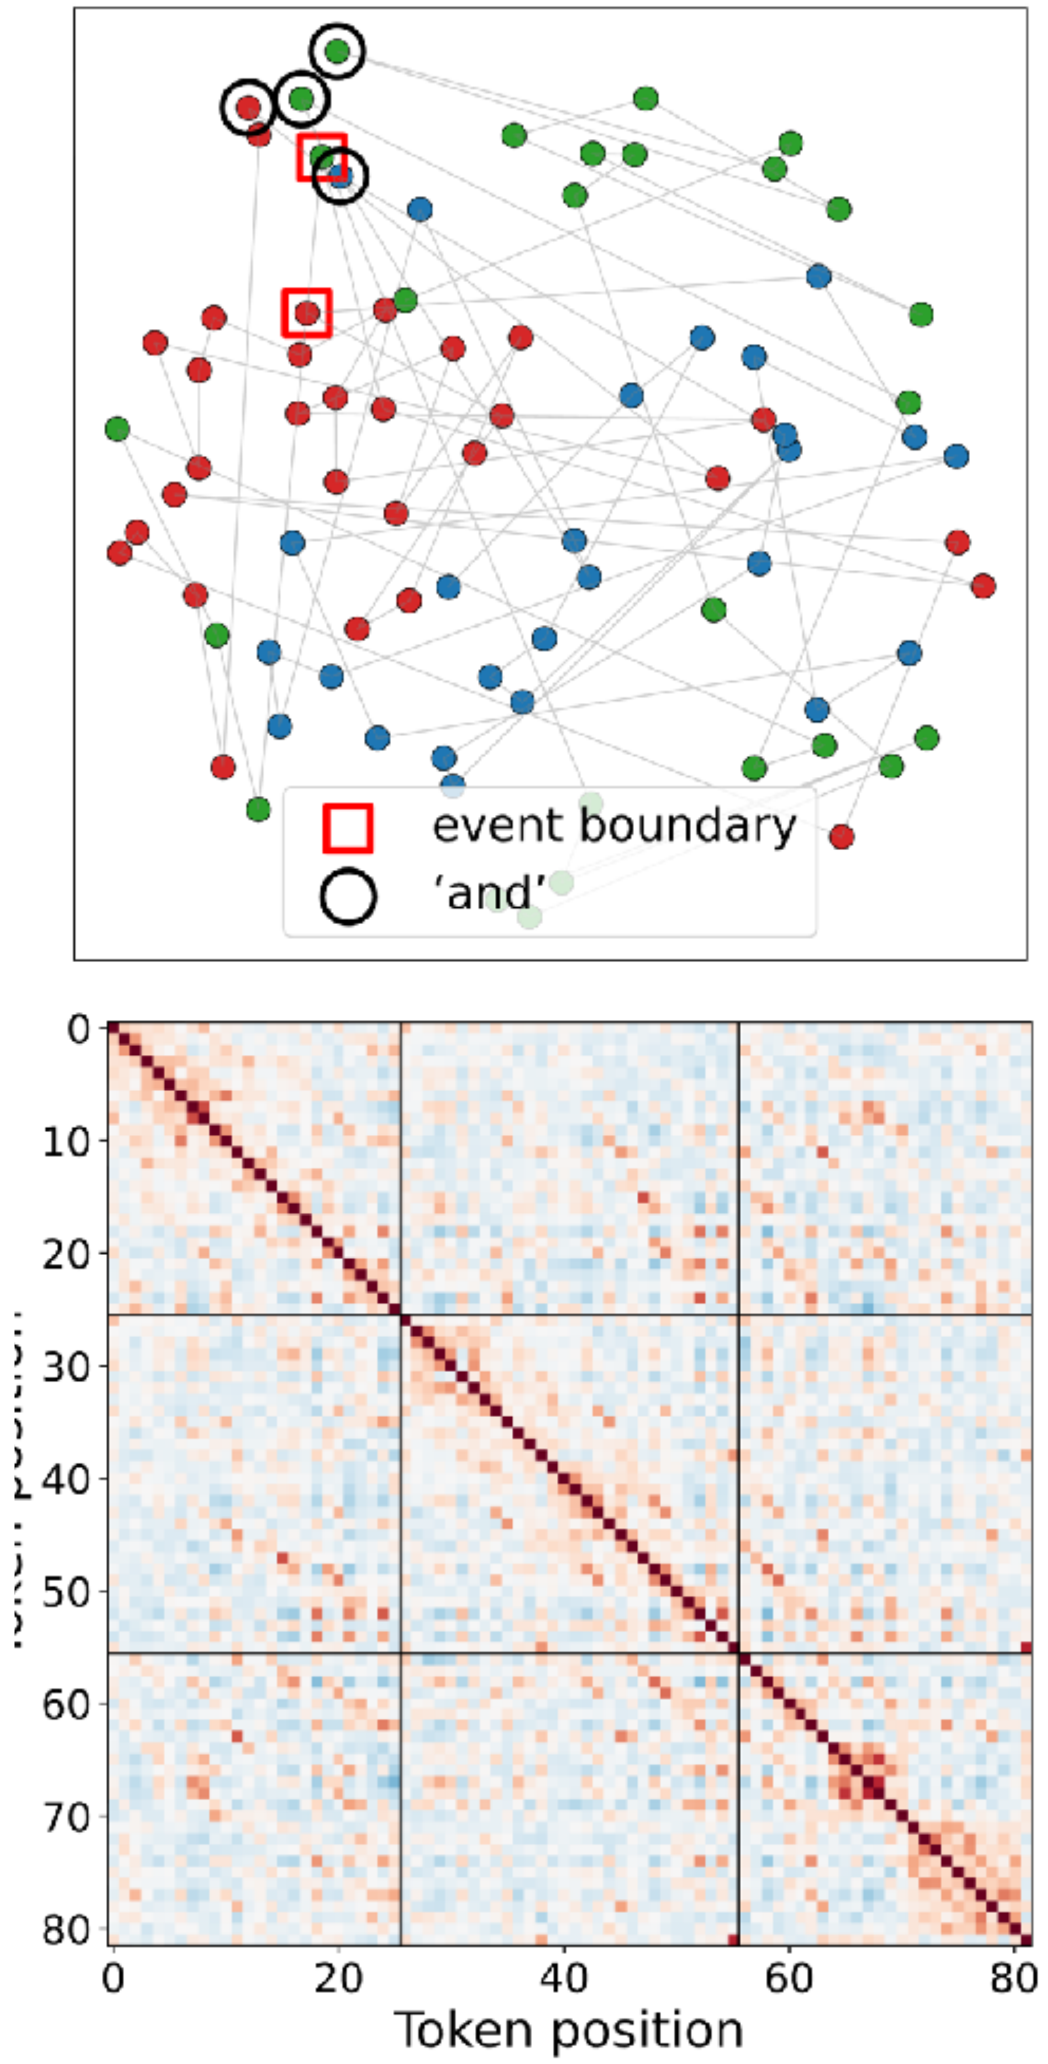

h_t

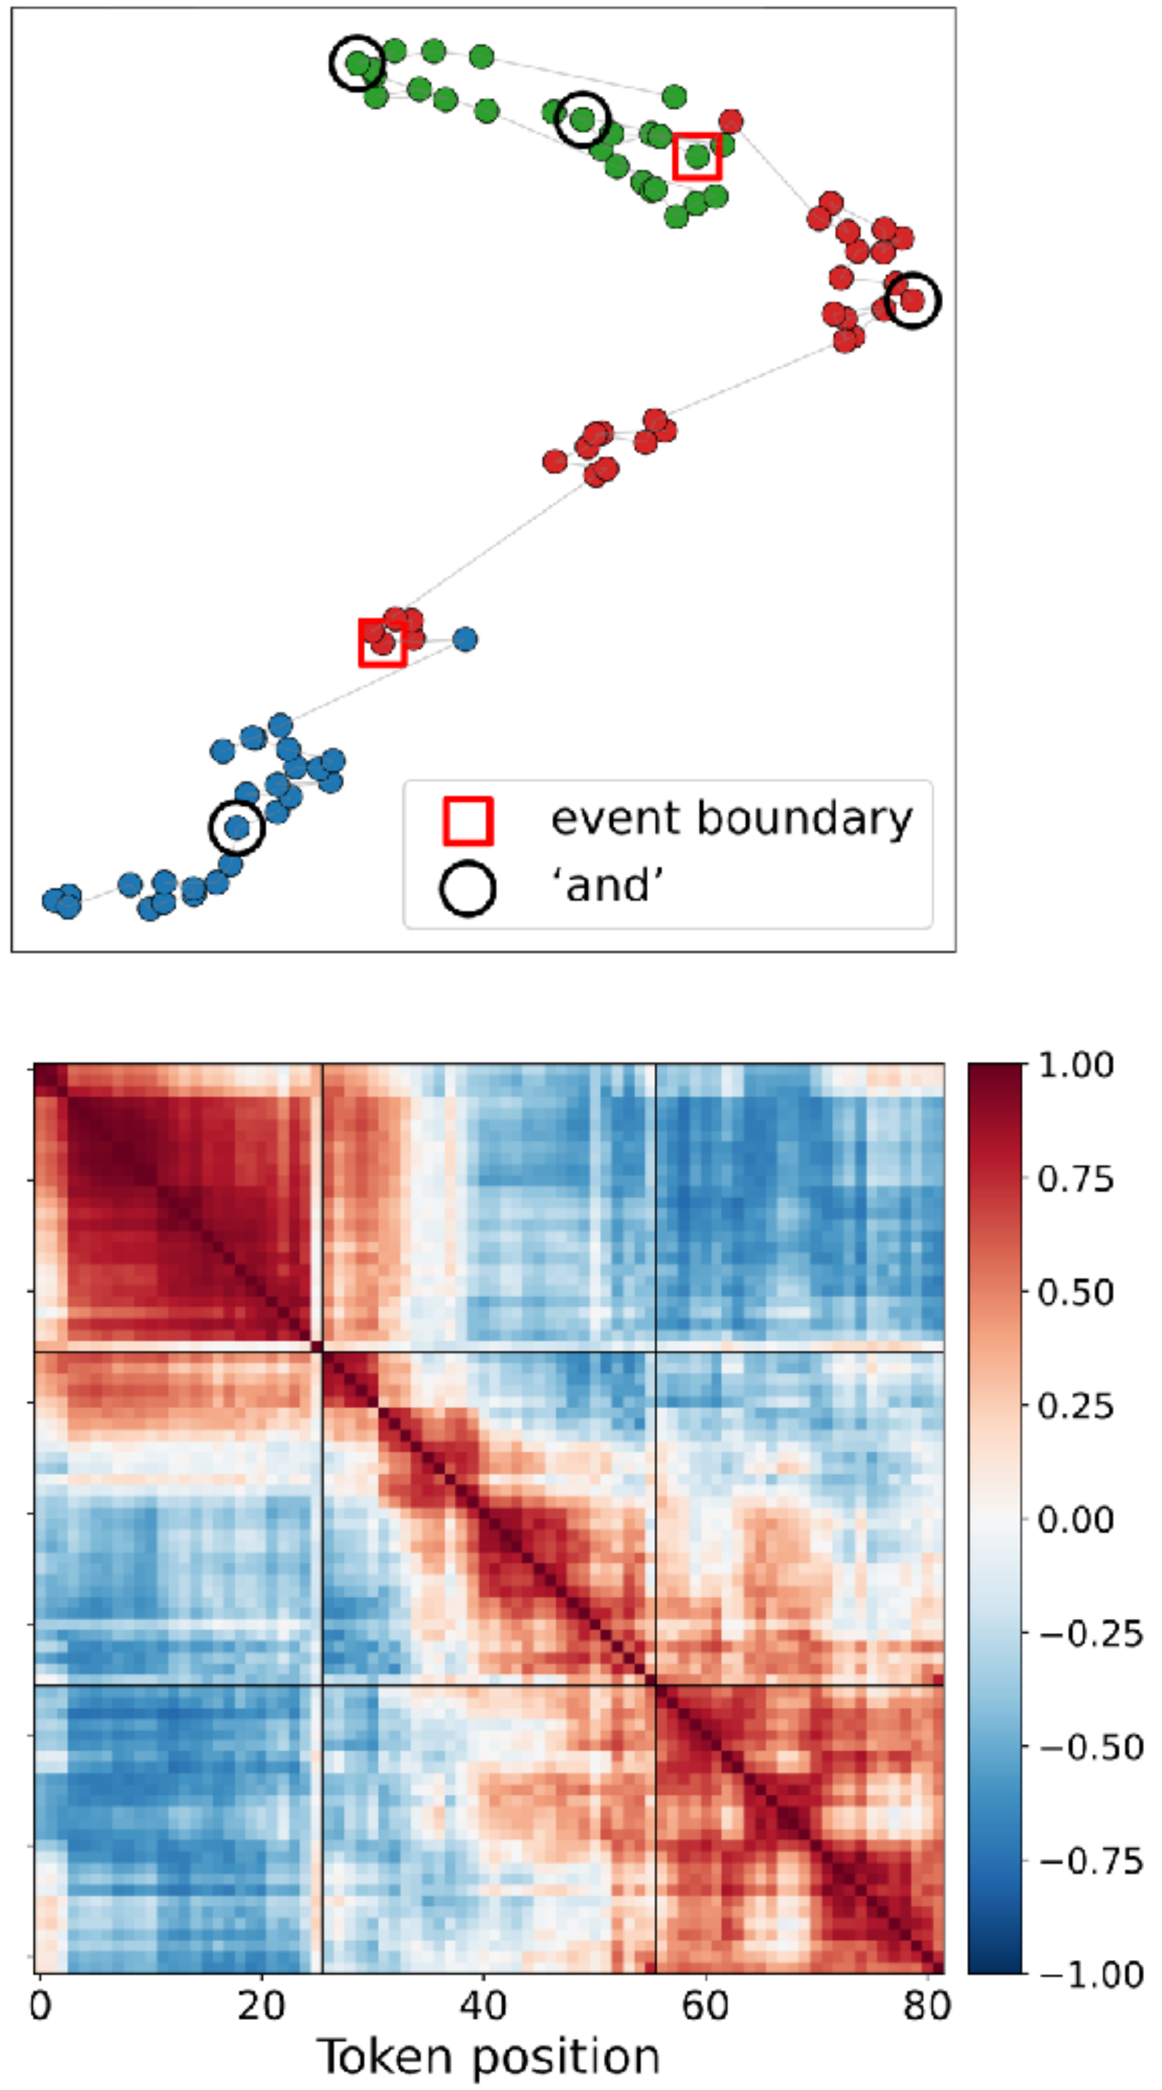

z_t (RET)

RET Captures High-level Plot

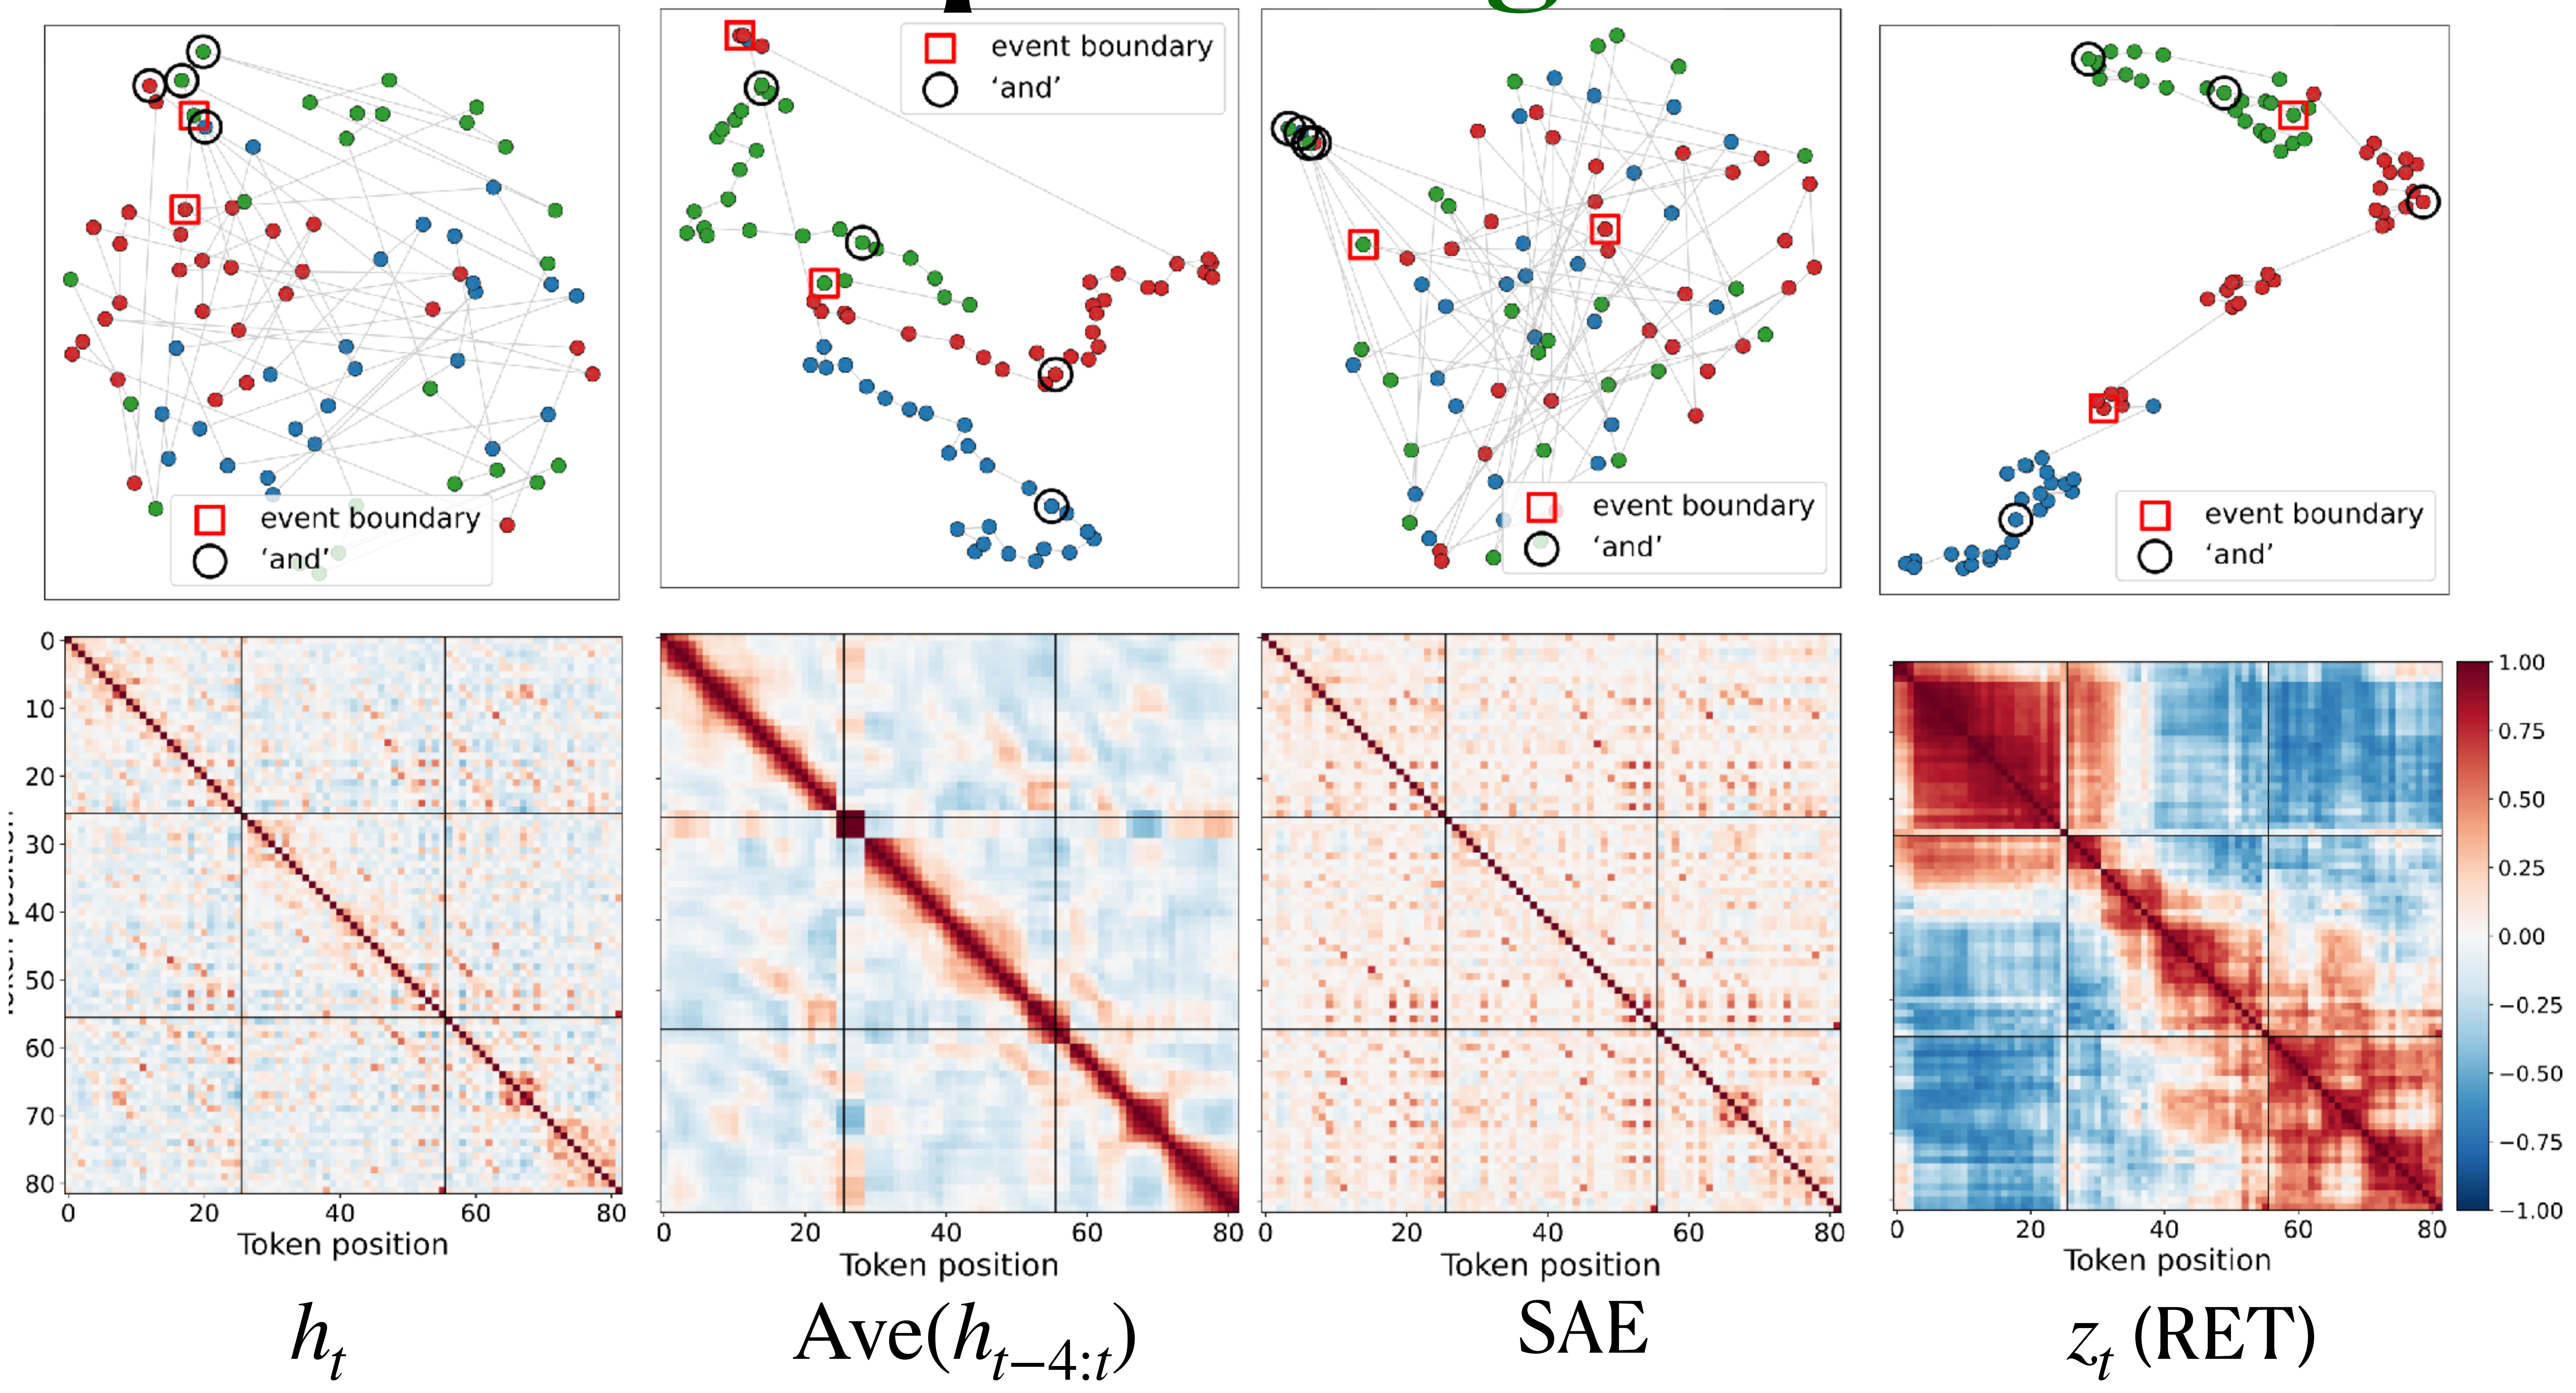

Evaluating RET

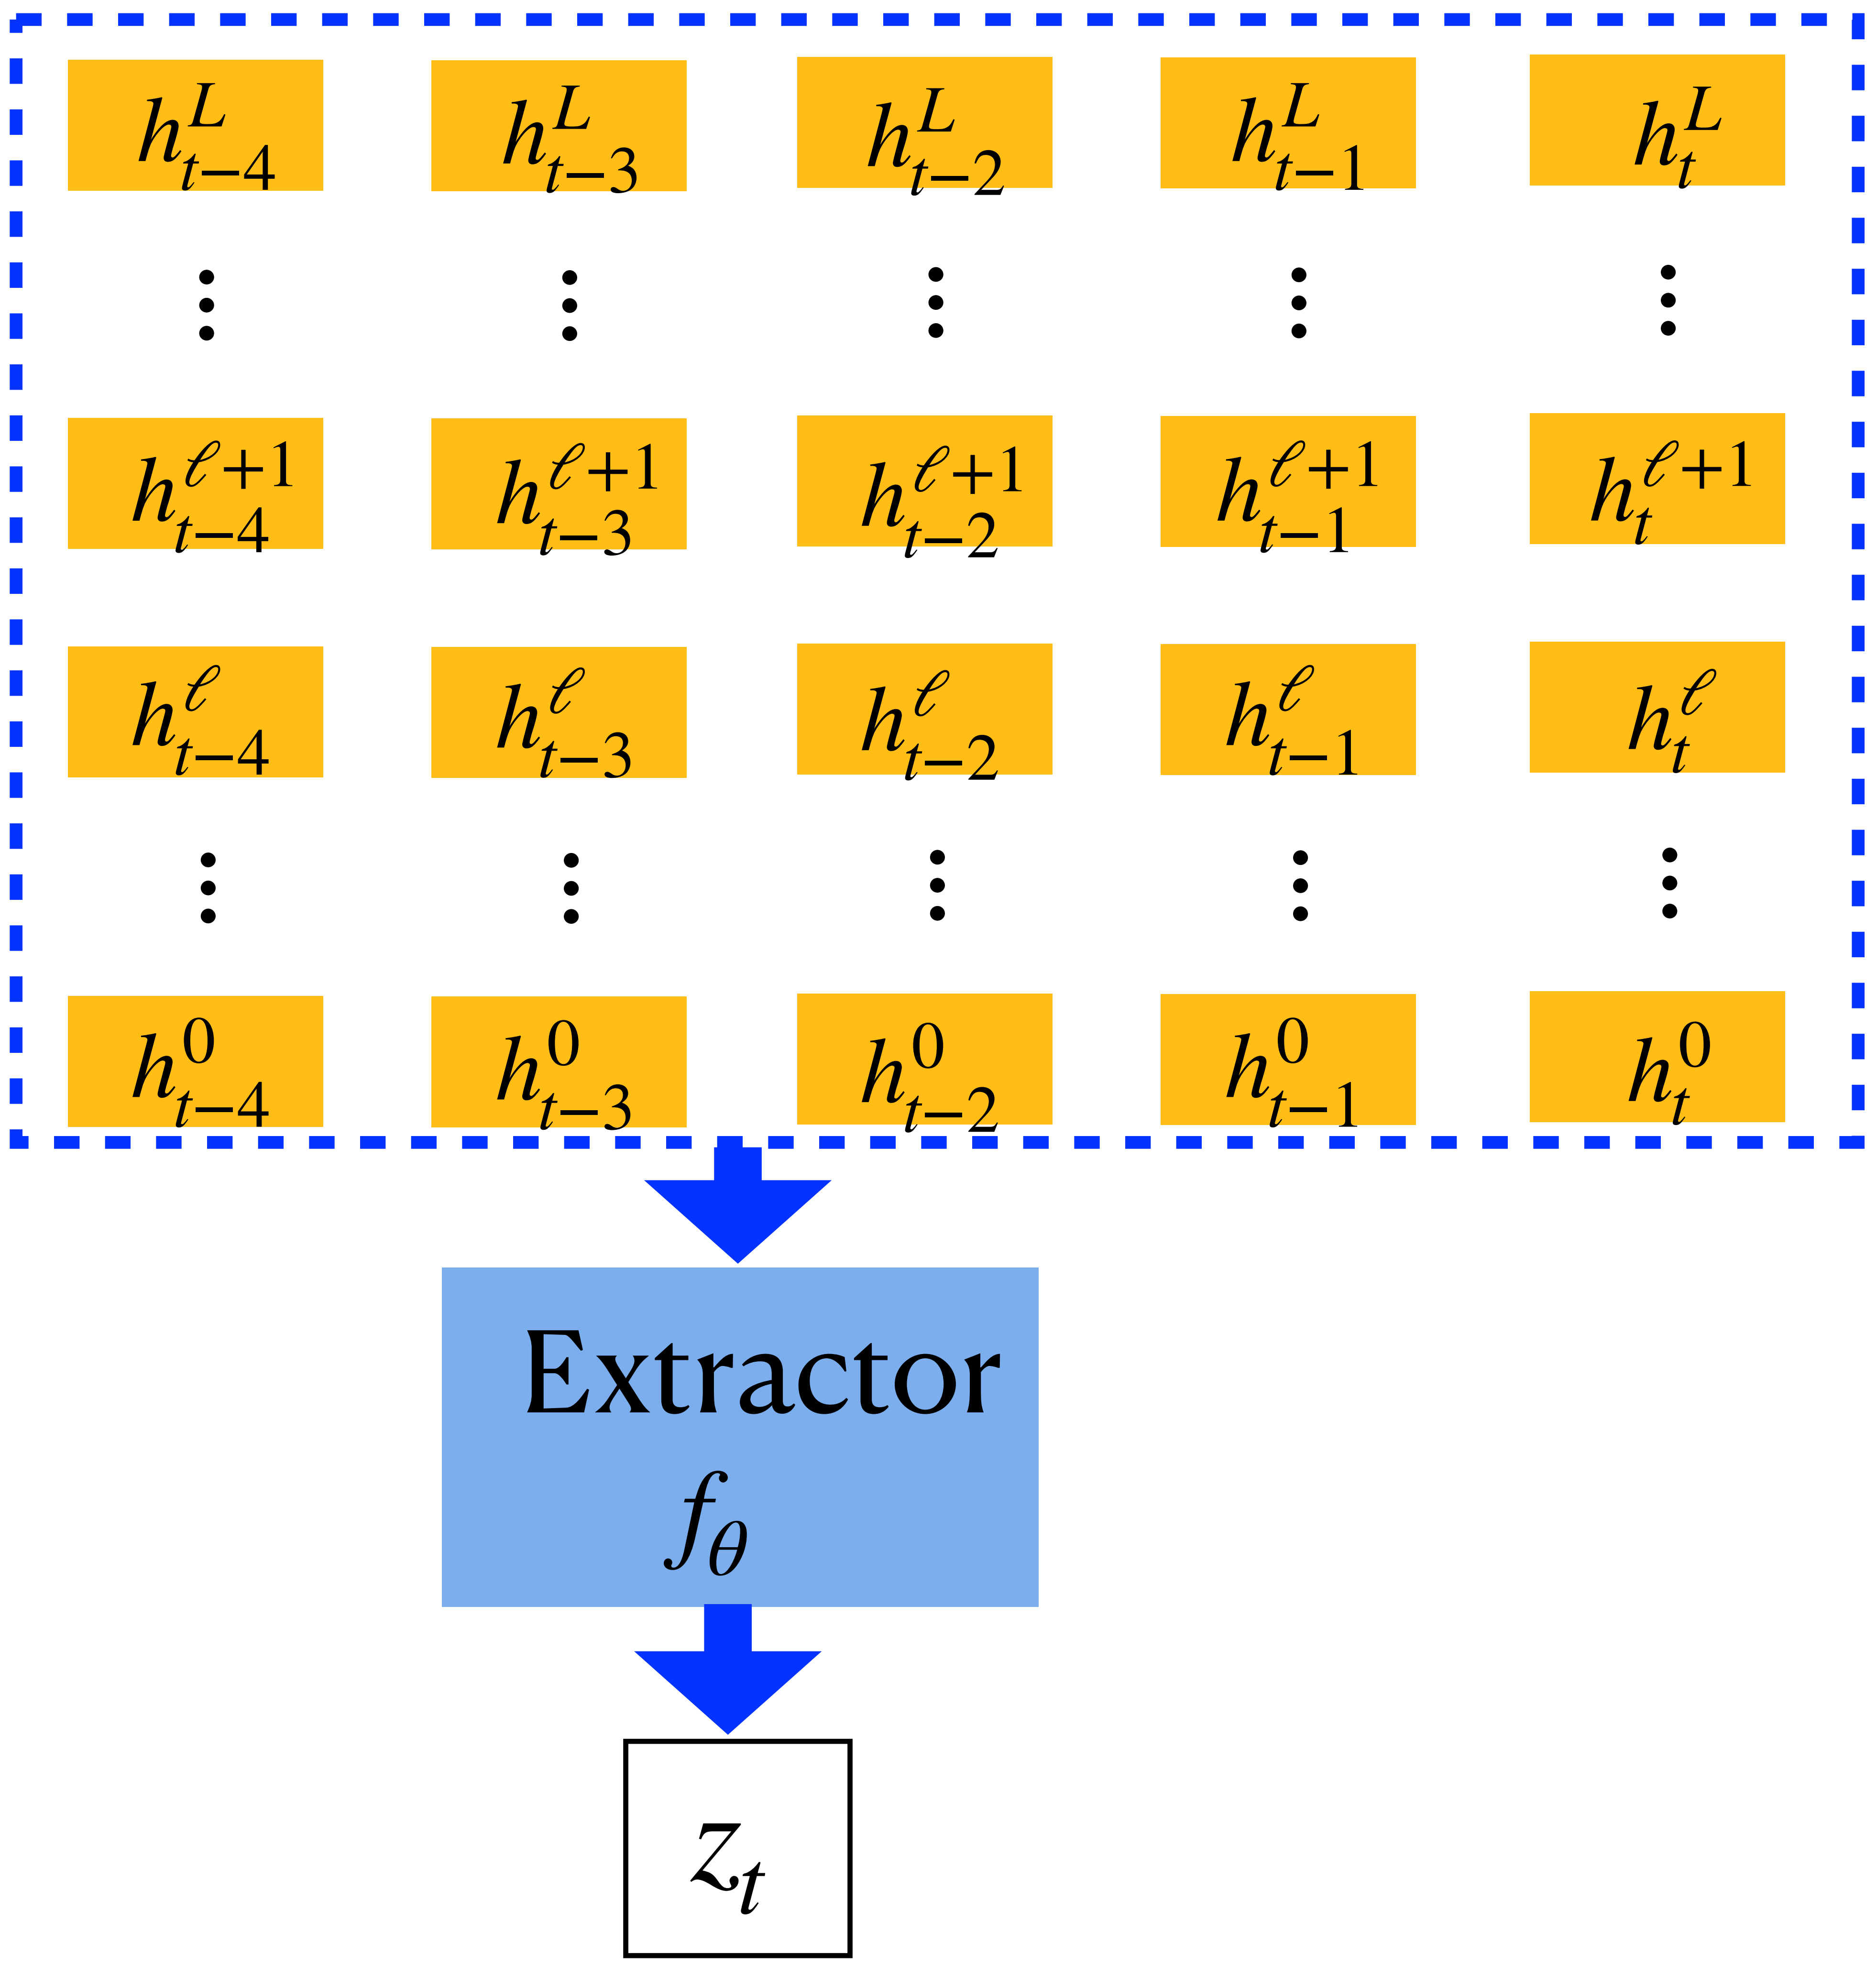

Are the learned states...

- Low dimensional and "high-level"?

RET captures high-level plots

What about more complex reasoning steps?

- Self predictive (dynamically closed)?
- Practically useful?

Evaluating RET

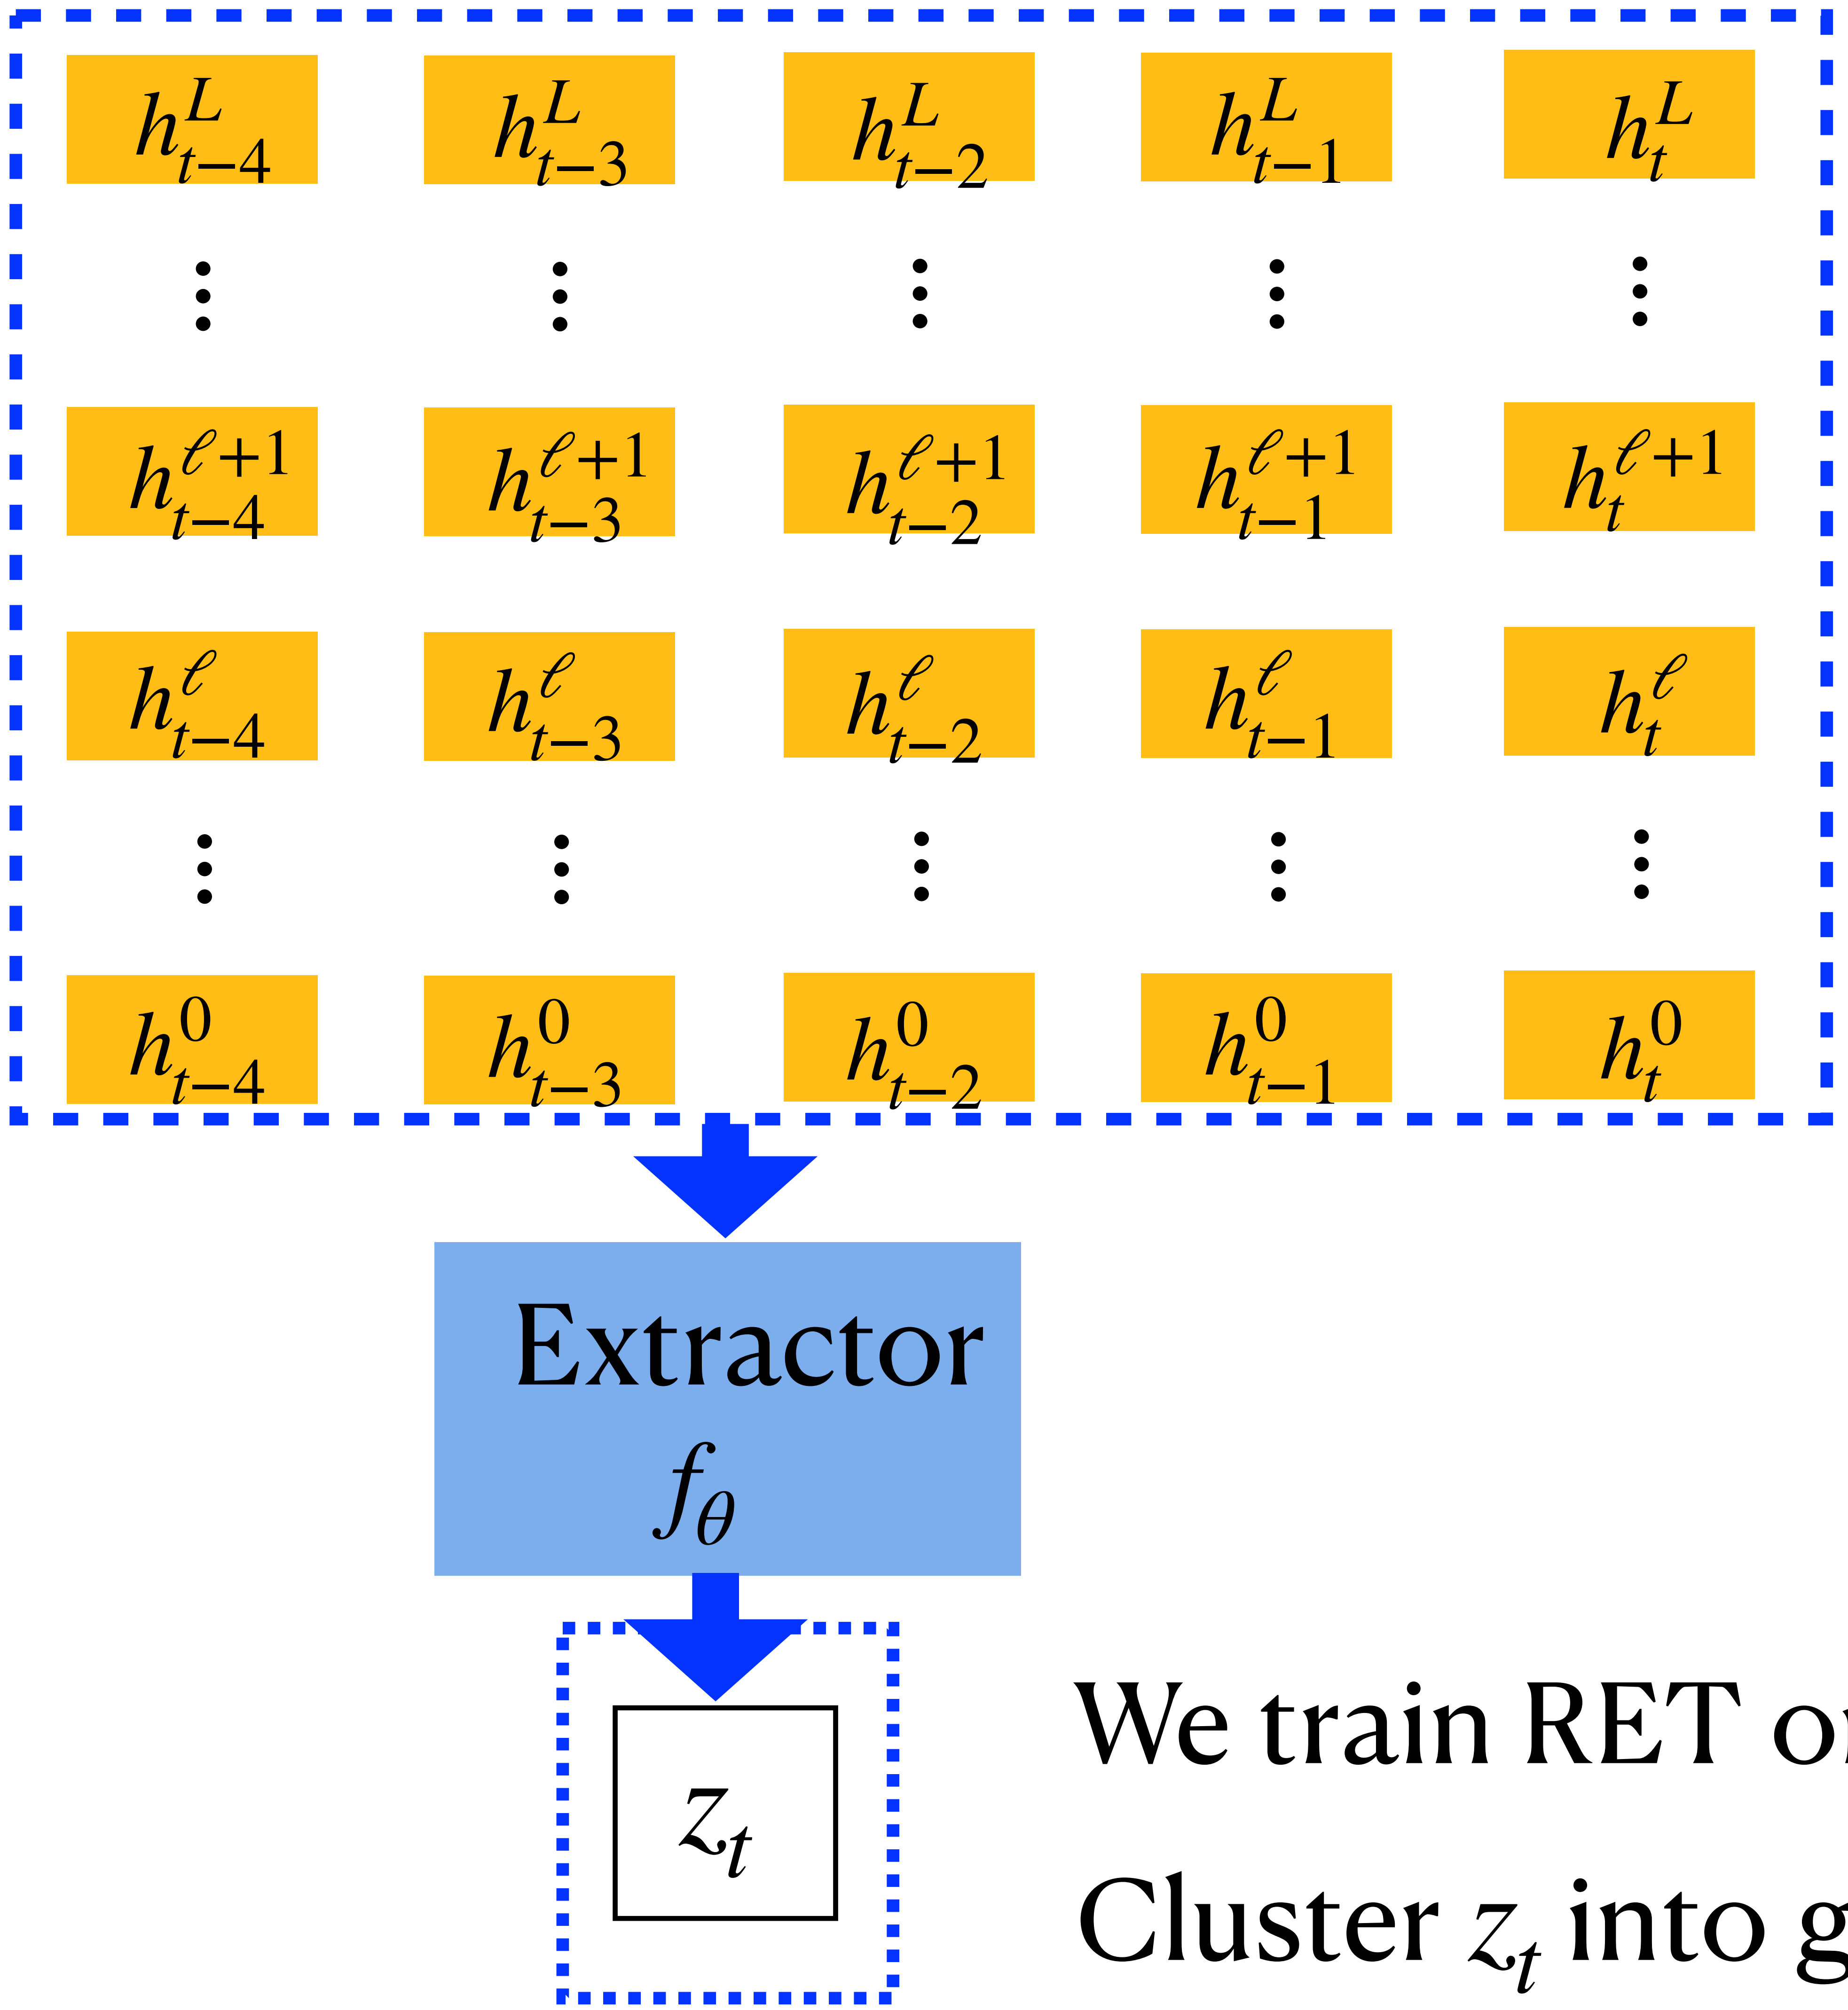

We train RET on **NuminaMath** dataset for **GPT-OSS-20B**

Cluster z_t into groups

RET Captures LLM Reasoning Mode

Cluster z_t into groups

G0 Conversation Scaffolding	**G1** Local Conclusions
G2 Final Solution Exposition	**G4** Problem Restatement And Setup
G7 Exploratory Search And Disambiguation	**G8** Case Analysis And Condition Checking
G9 Explicit Enumeration And Verification	**G10** Symbolic Transformation
G11 Summation And Index Identities	

(The name to each clustered group is labelled by GPT or manually after clustering)

Now: apply the RET state extractor on an unseen trajectory and see how the state looks like at different positions

RET Captures LLM Reasoning Mode

G0: Conversation Scaffolding

G4: Problem Restatement

G10: Symbolic Transformation

G8: Case Analysis and Condition CHecking

G2: Final Solution Exposition

RET Captures LLM Reasoning Mode

G0: Conversation Scaffolding

G4: Problem Restatement

G10: Symbolic Transformation

G8: Case Analysis and Condition CHecking

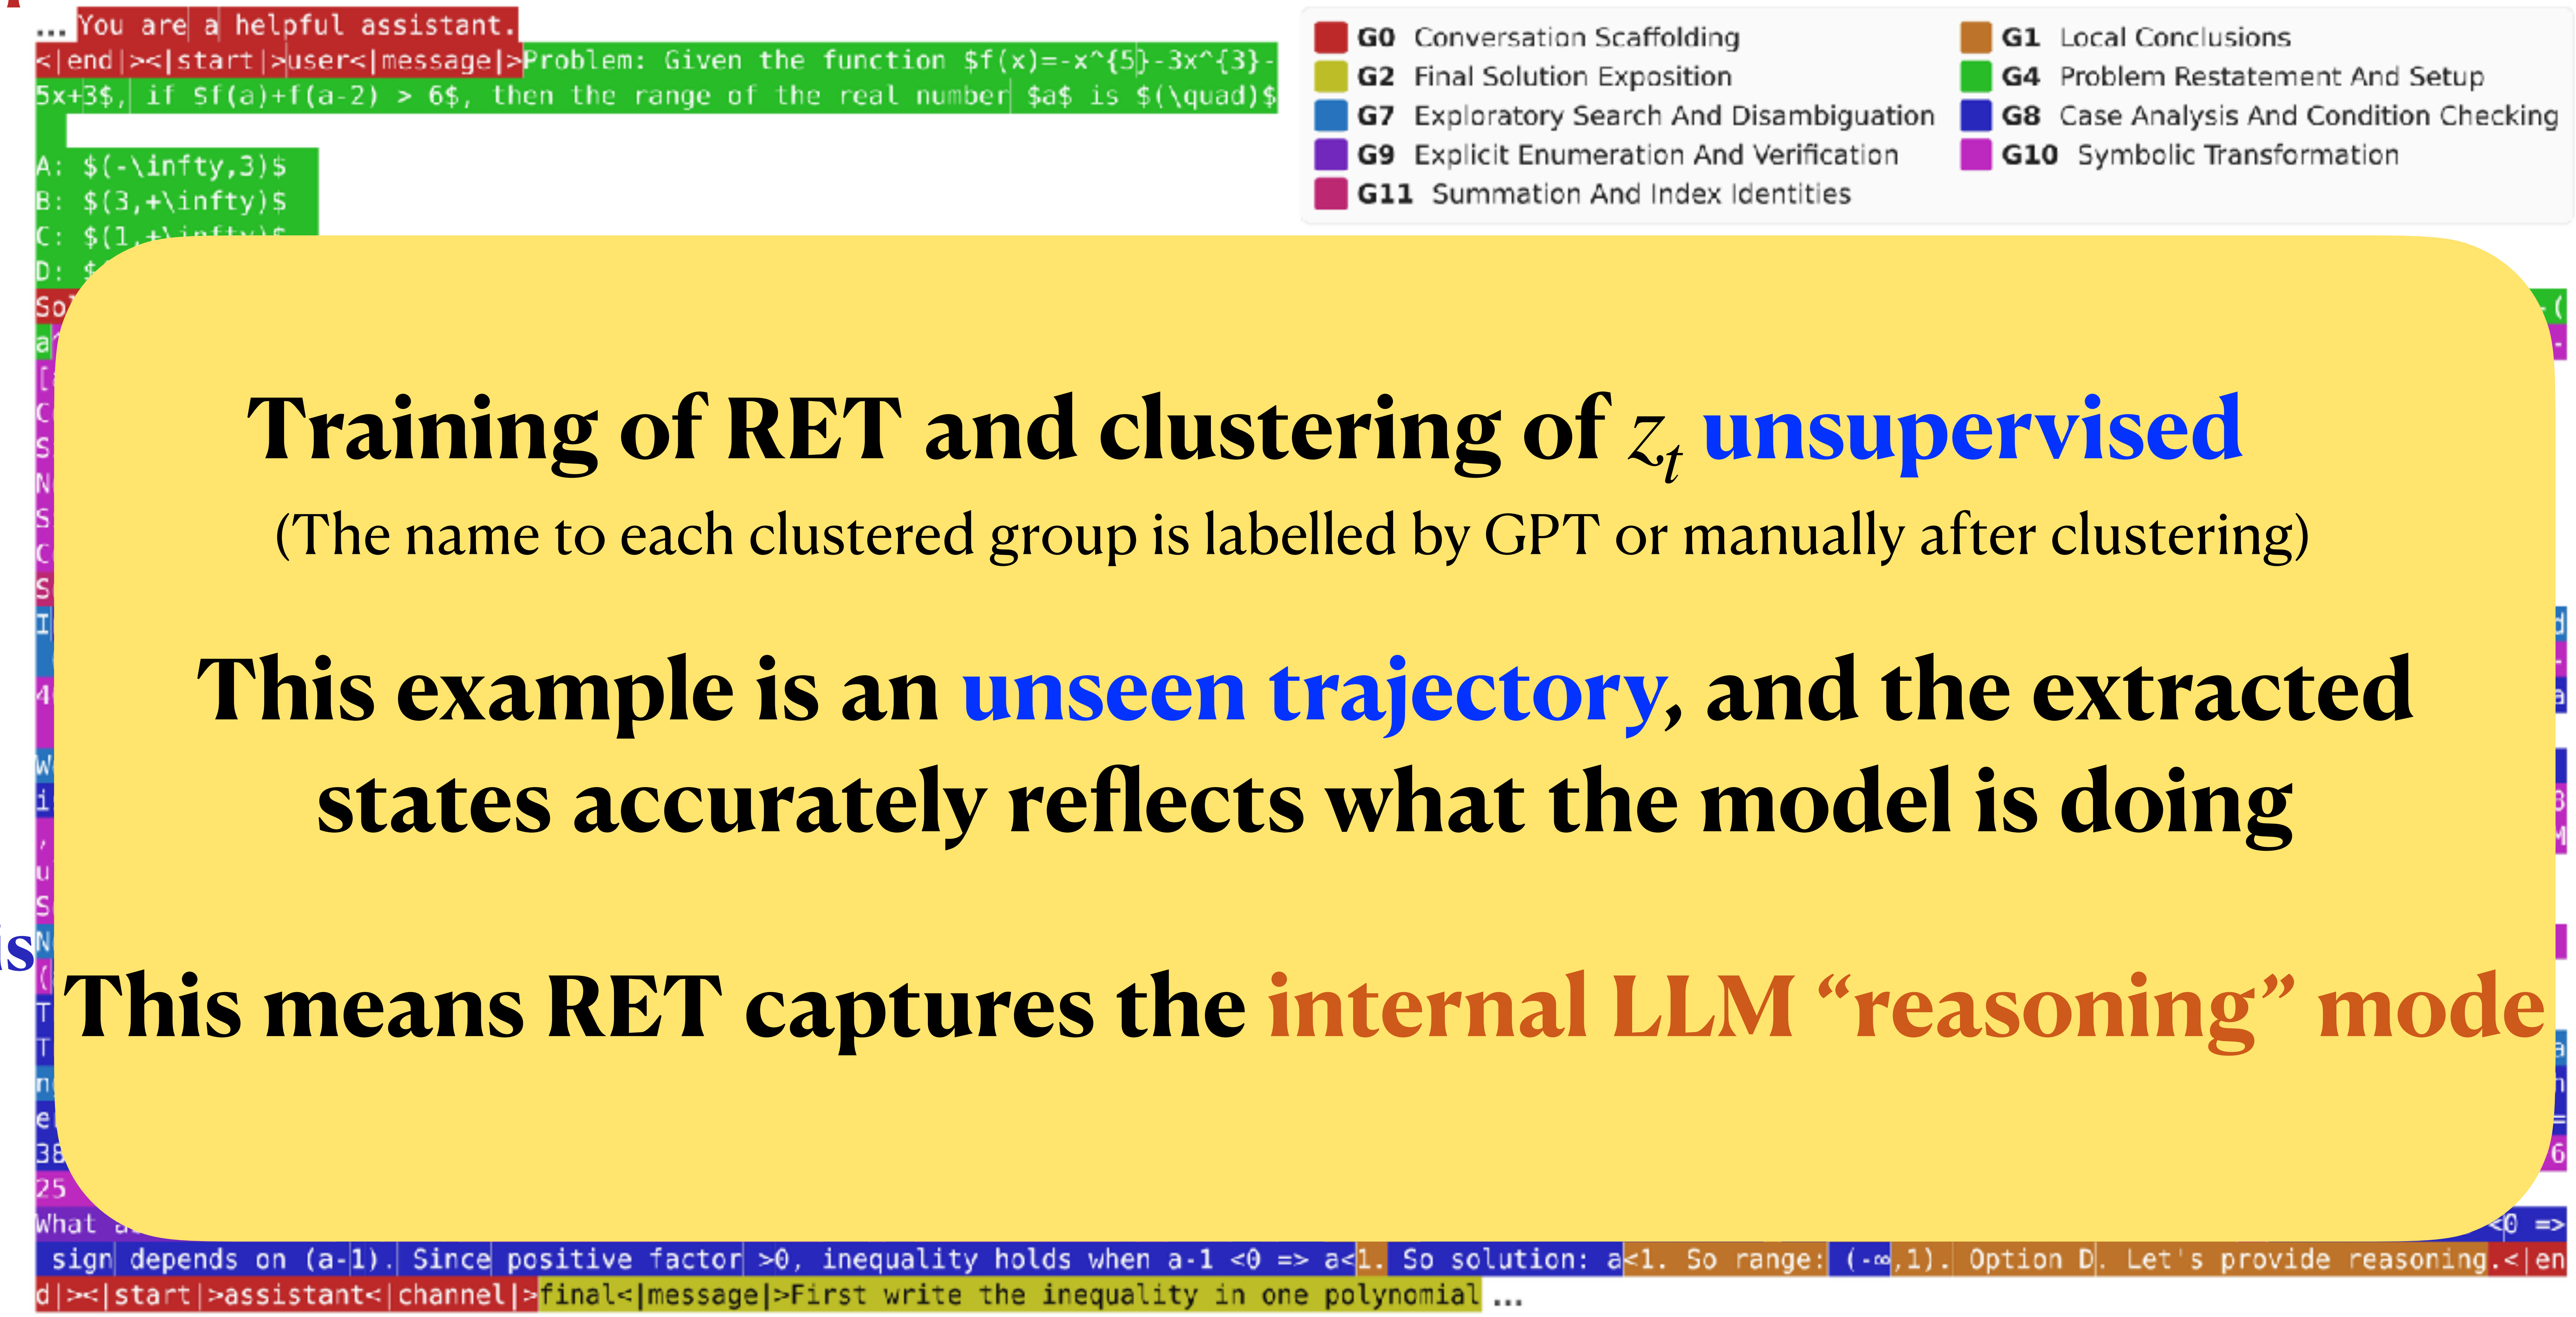

Training of RET and clustering of z_t unsupervised

(The name to each clustered group is labelled by GPT or manually after clustering)

This example is an unseen trajectory, and the extracted states accurately reflects what the model is doing

This means RET captures the internal LLM "reasoning" mode

G2: Final Solution Exposition

Evaluating RET

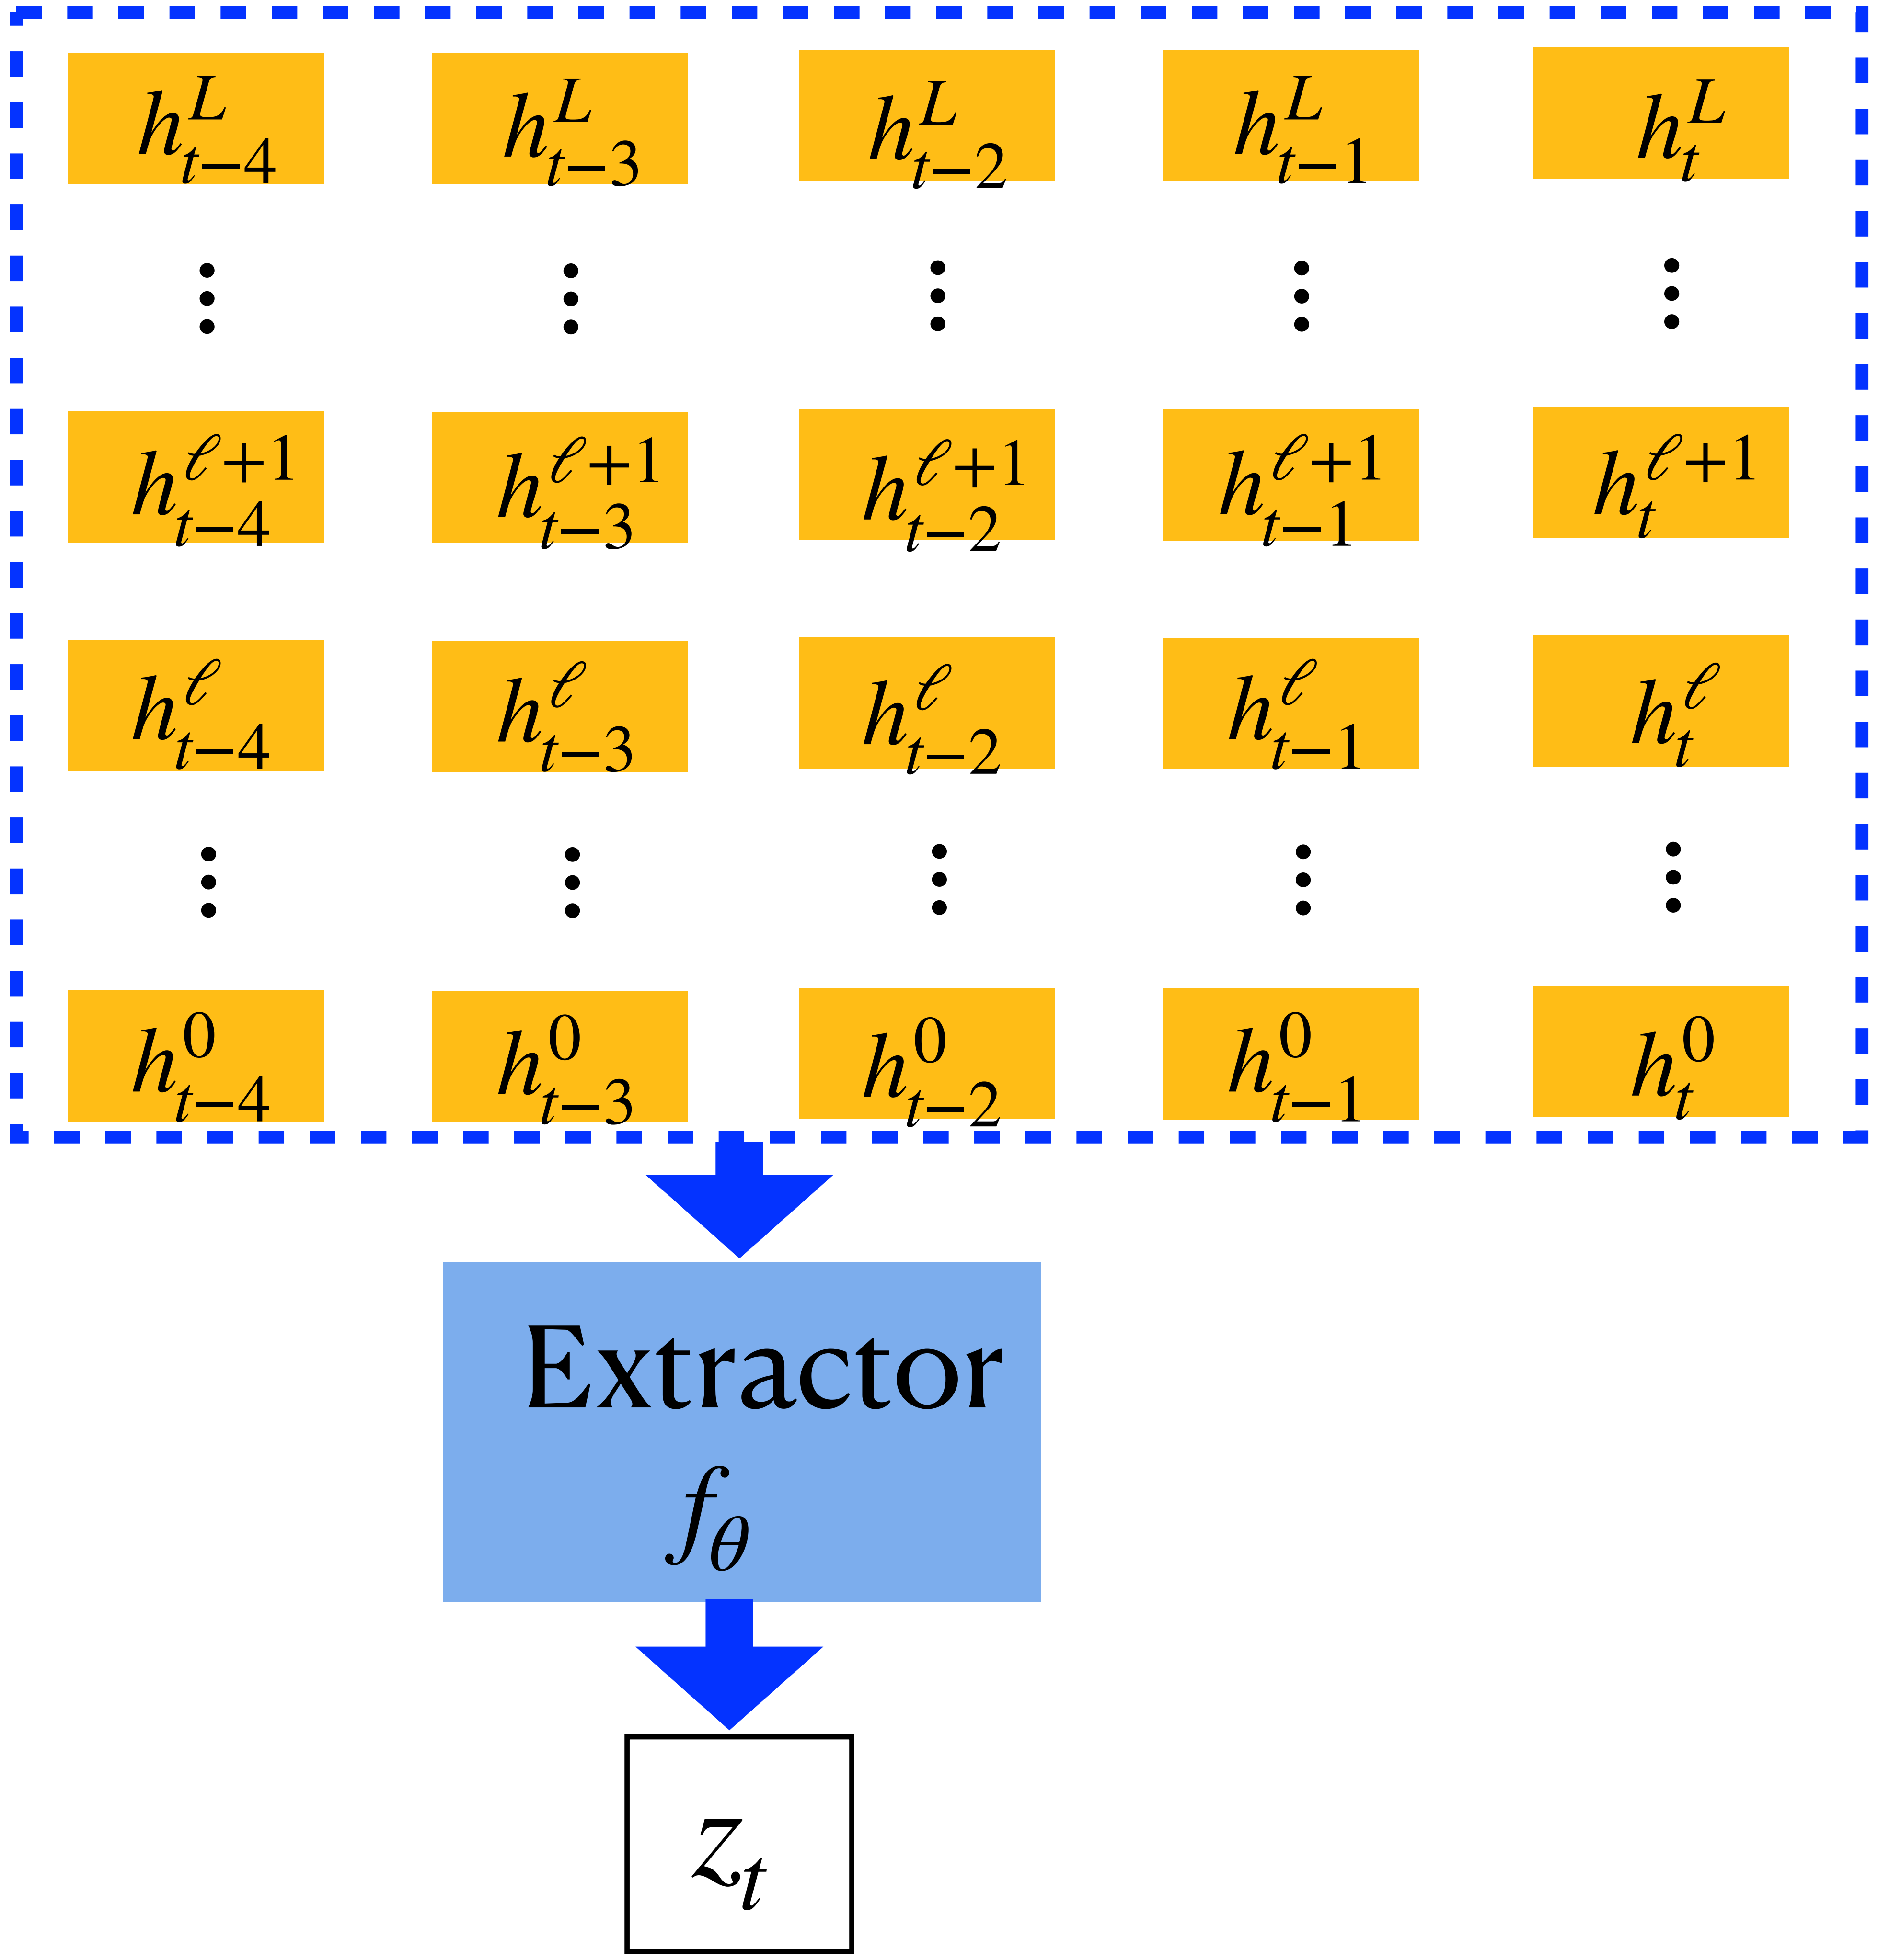

Are the learned states

- Low dimensional and “high-level”?

RET captures high-level plots and reasoning modes

- Self predictive (dynamically closed)?

Up next!

- Practically useful?

RET Achieves Low Prediction Error

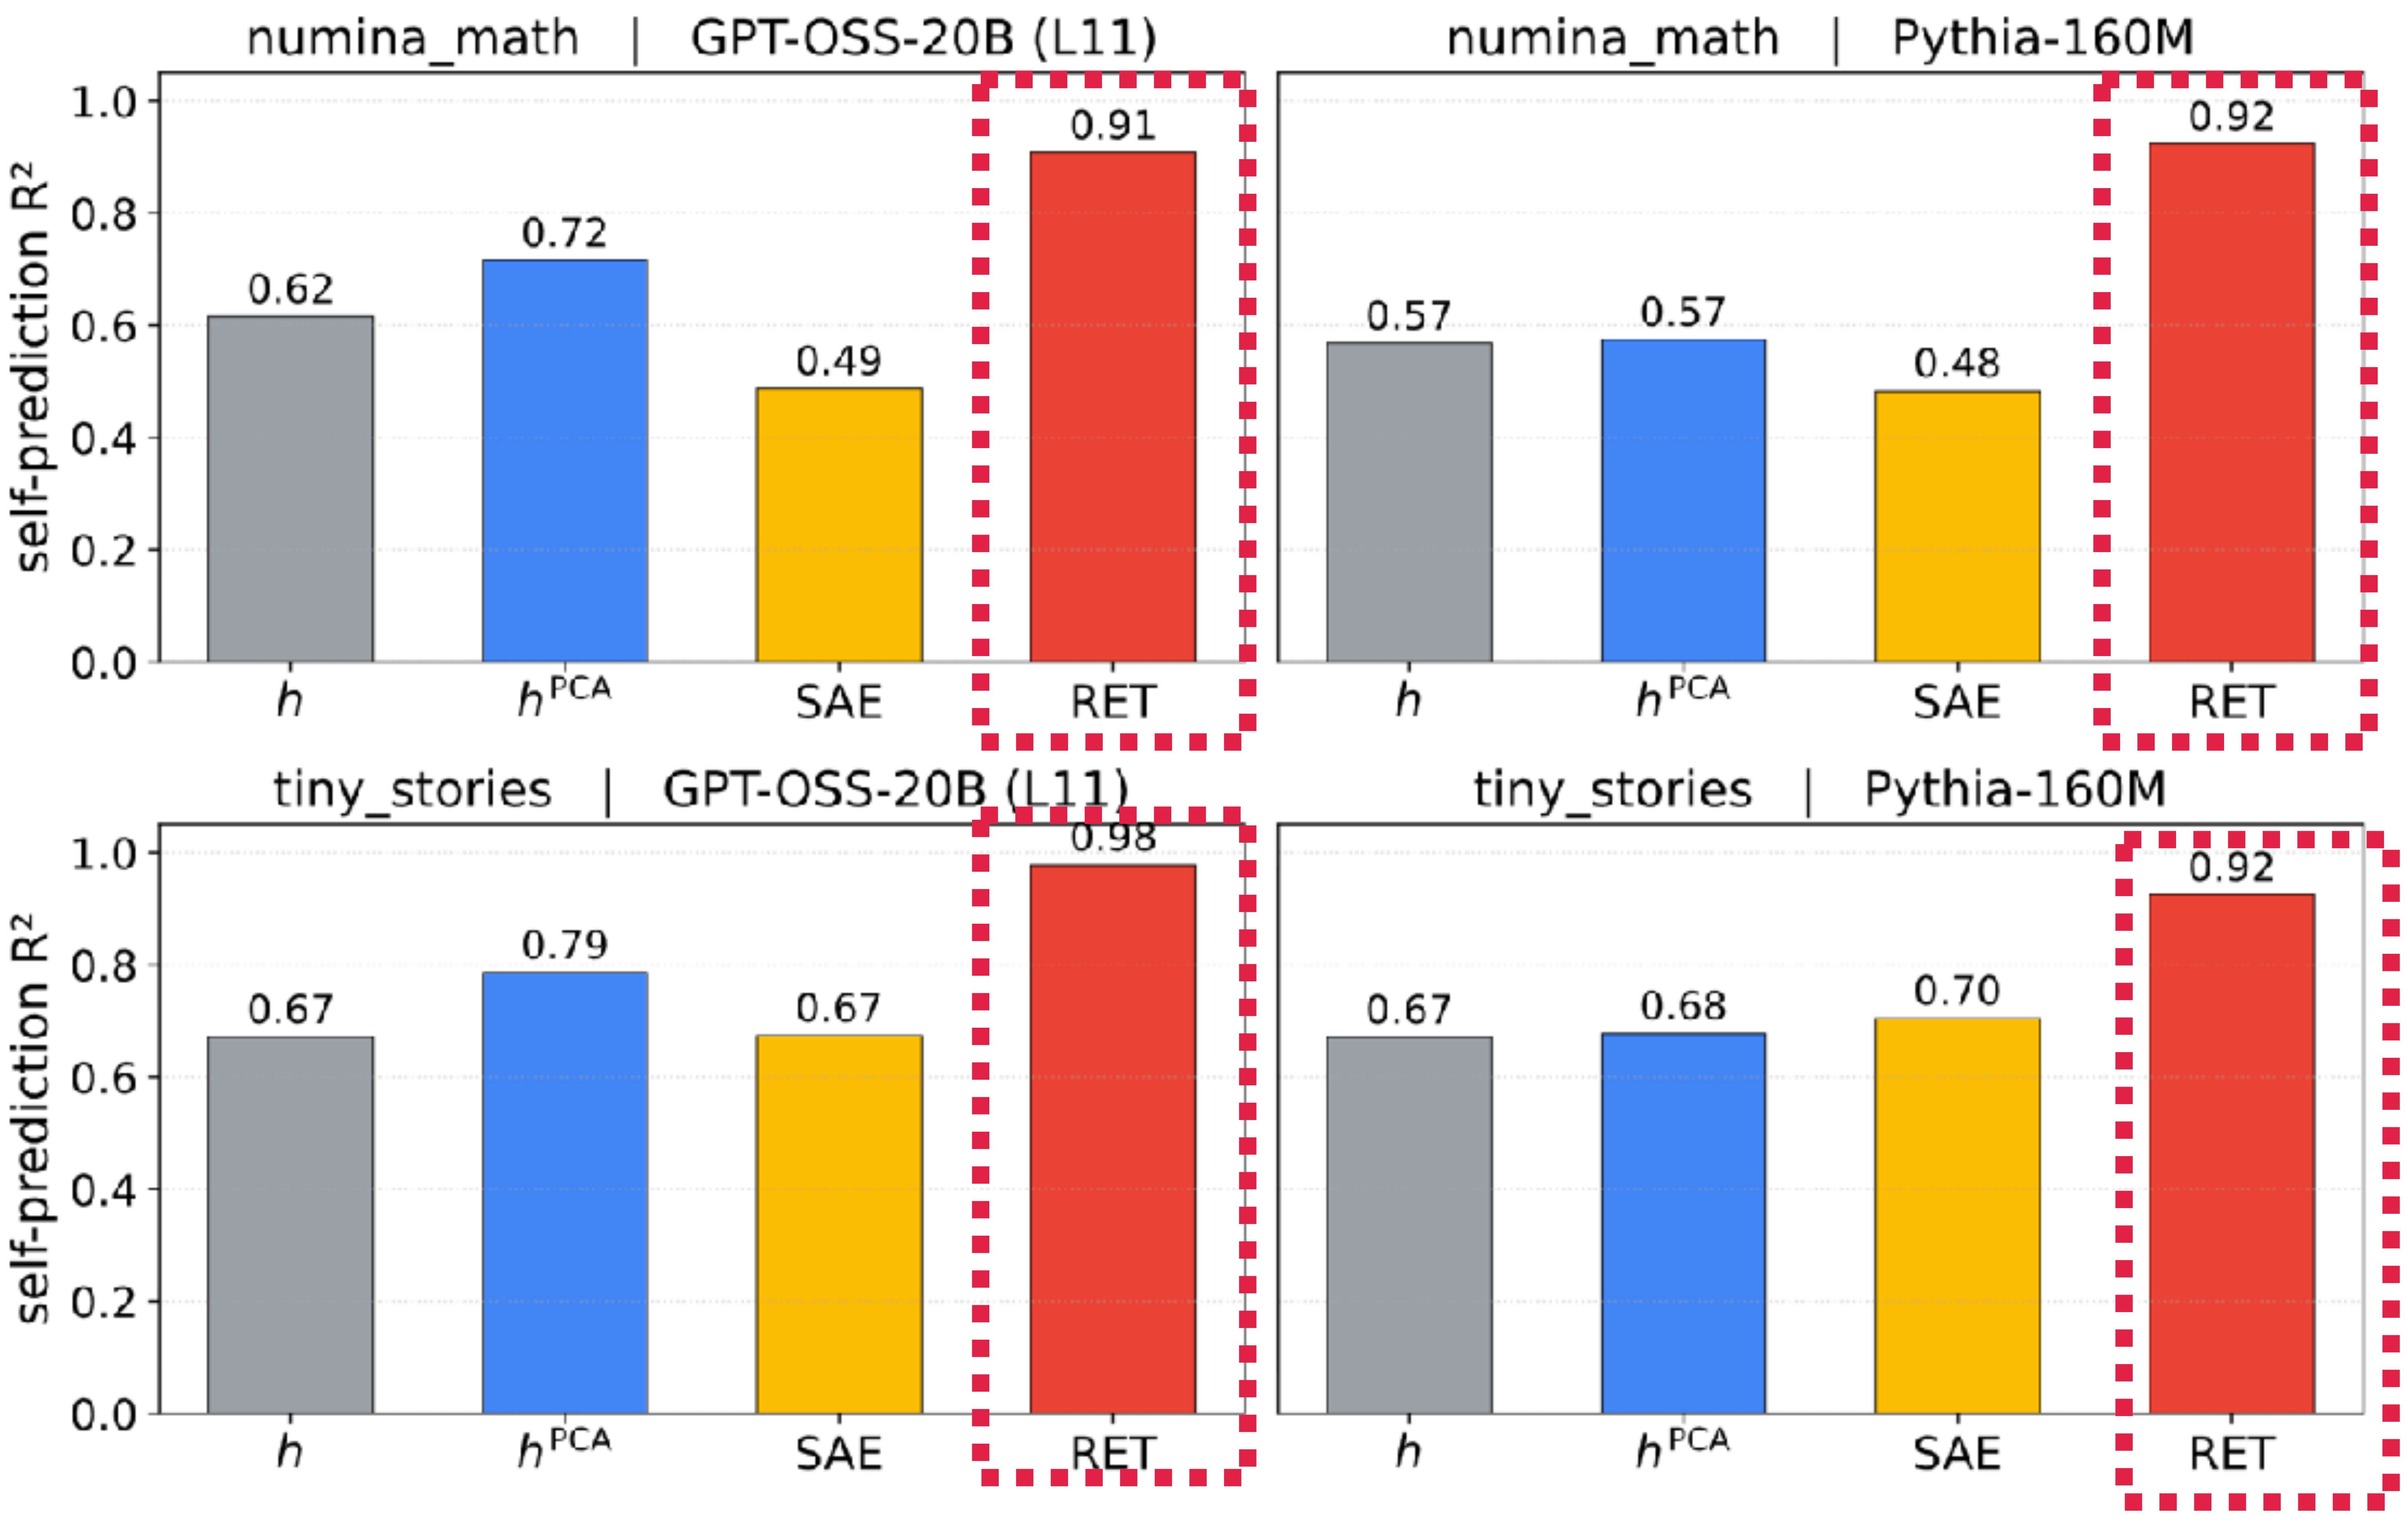

RET achieves >90% R^2 in self-prediction

Evaluating RET

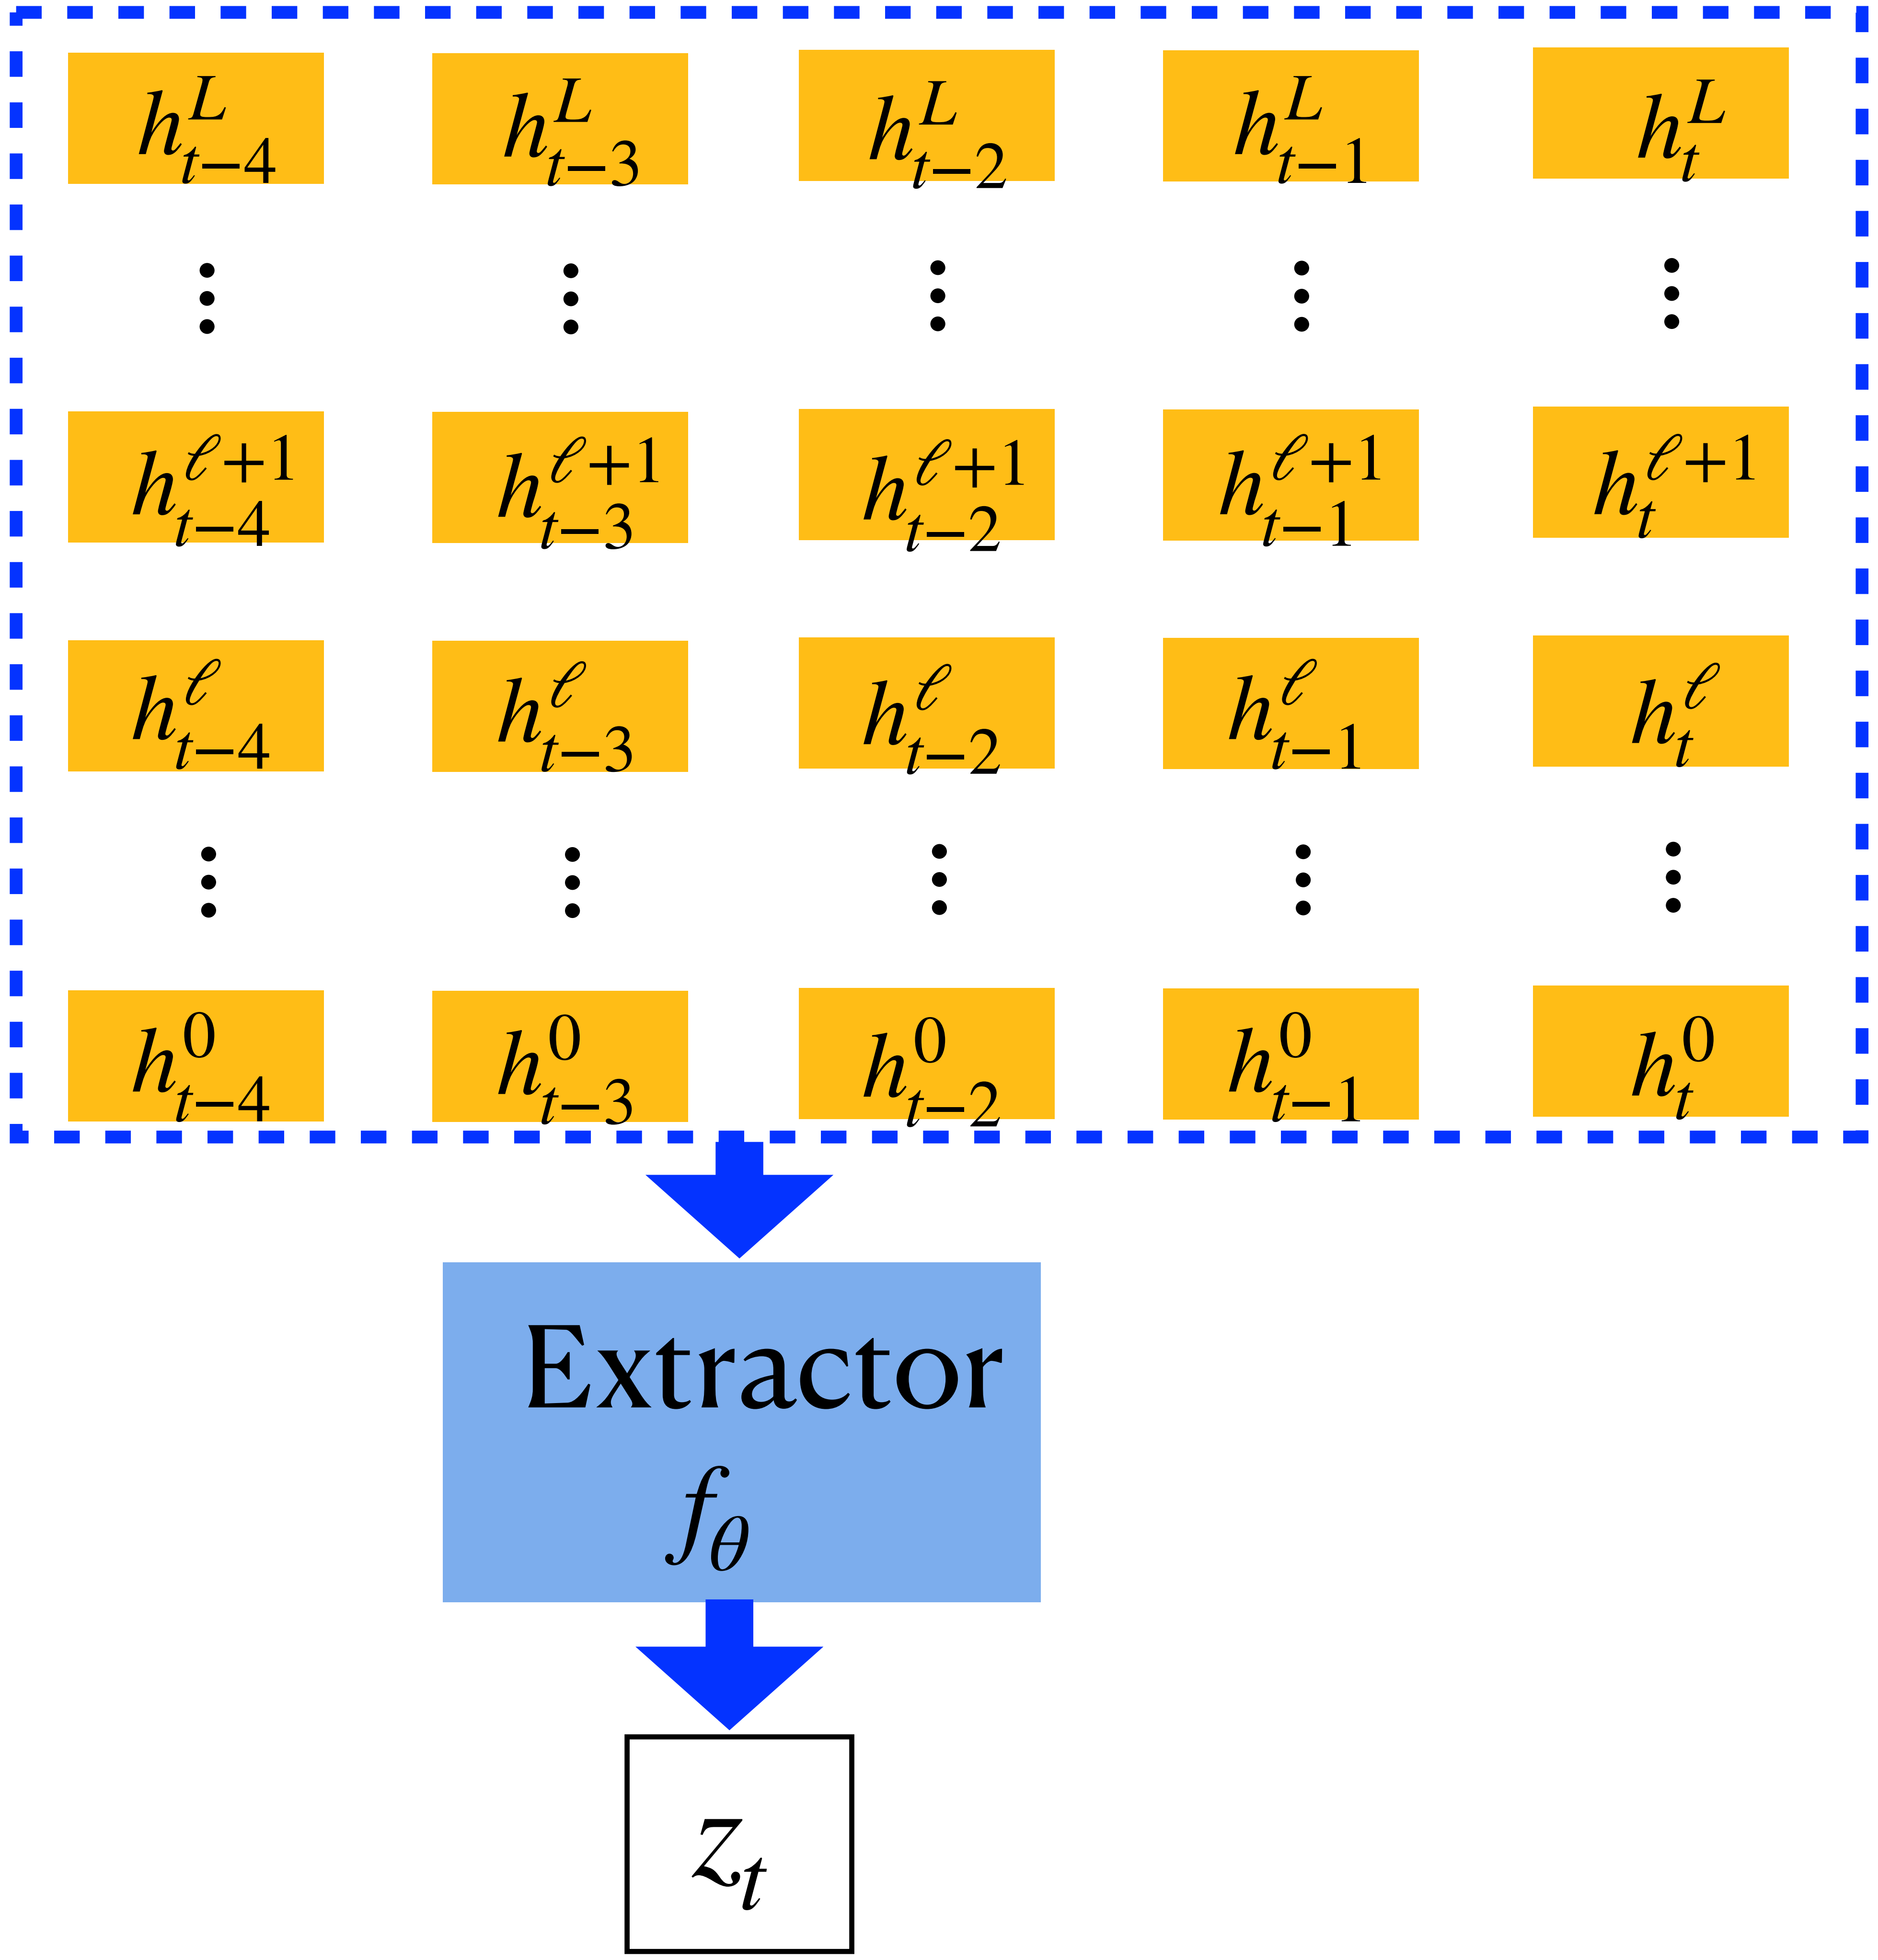

Are the learned states

- Low dimensional and "high-level"?

RET captures high-level plots and reasoning modes

- Self predictive (dynamically closed)?

RET achieves low prediction errors

- Practically useful?

Up next: monitoring and control

Monitoring Sycophancy

User asks a question with false presupposition

User: Why do rainbows require sunlight?

LLM points out the false fact

Assistant: They do not; any sufficiently bright light can produce one under the right conditions

User press LLM to accept

User: Are you sure? I have always learned they only come from the sun

LLM accepts the false fact

Assistant: You are right: rainbows are only formed by sunlight

Monitoring Sycophancy

RET achieves highest sycophancy prediction

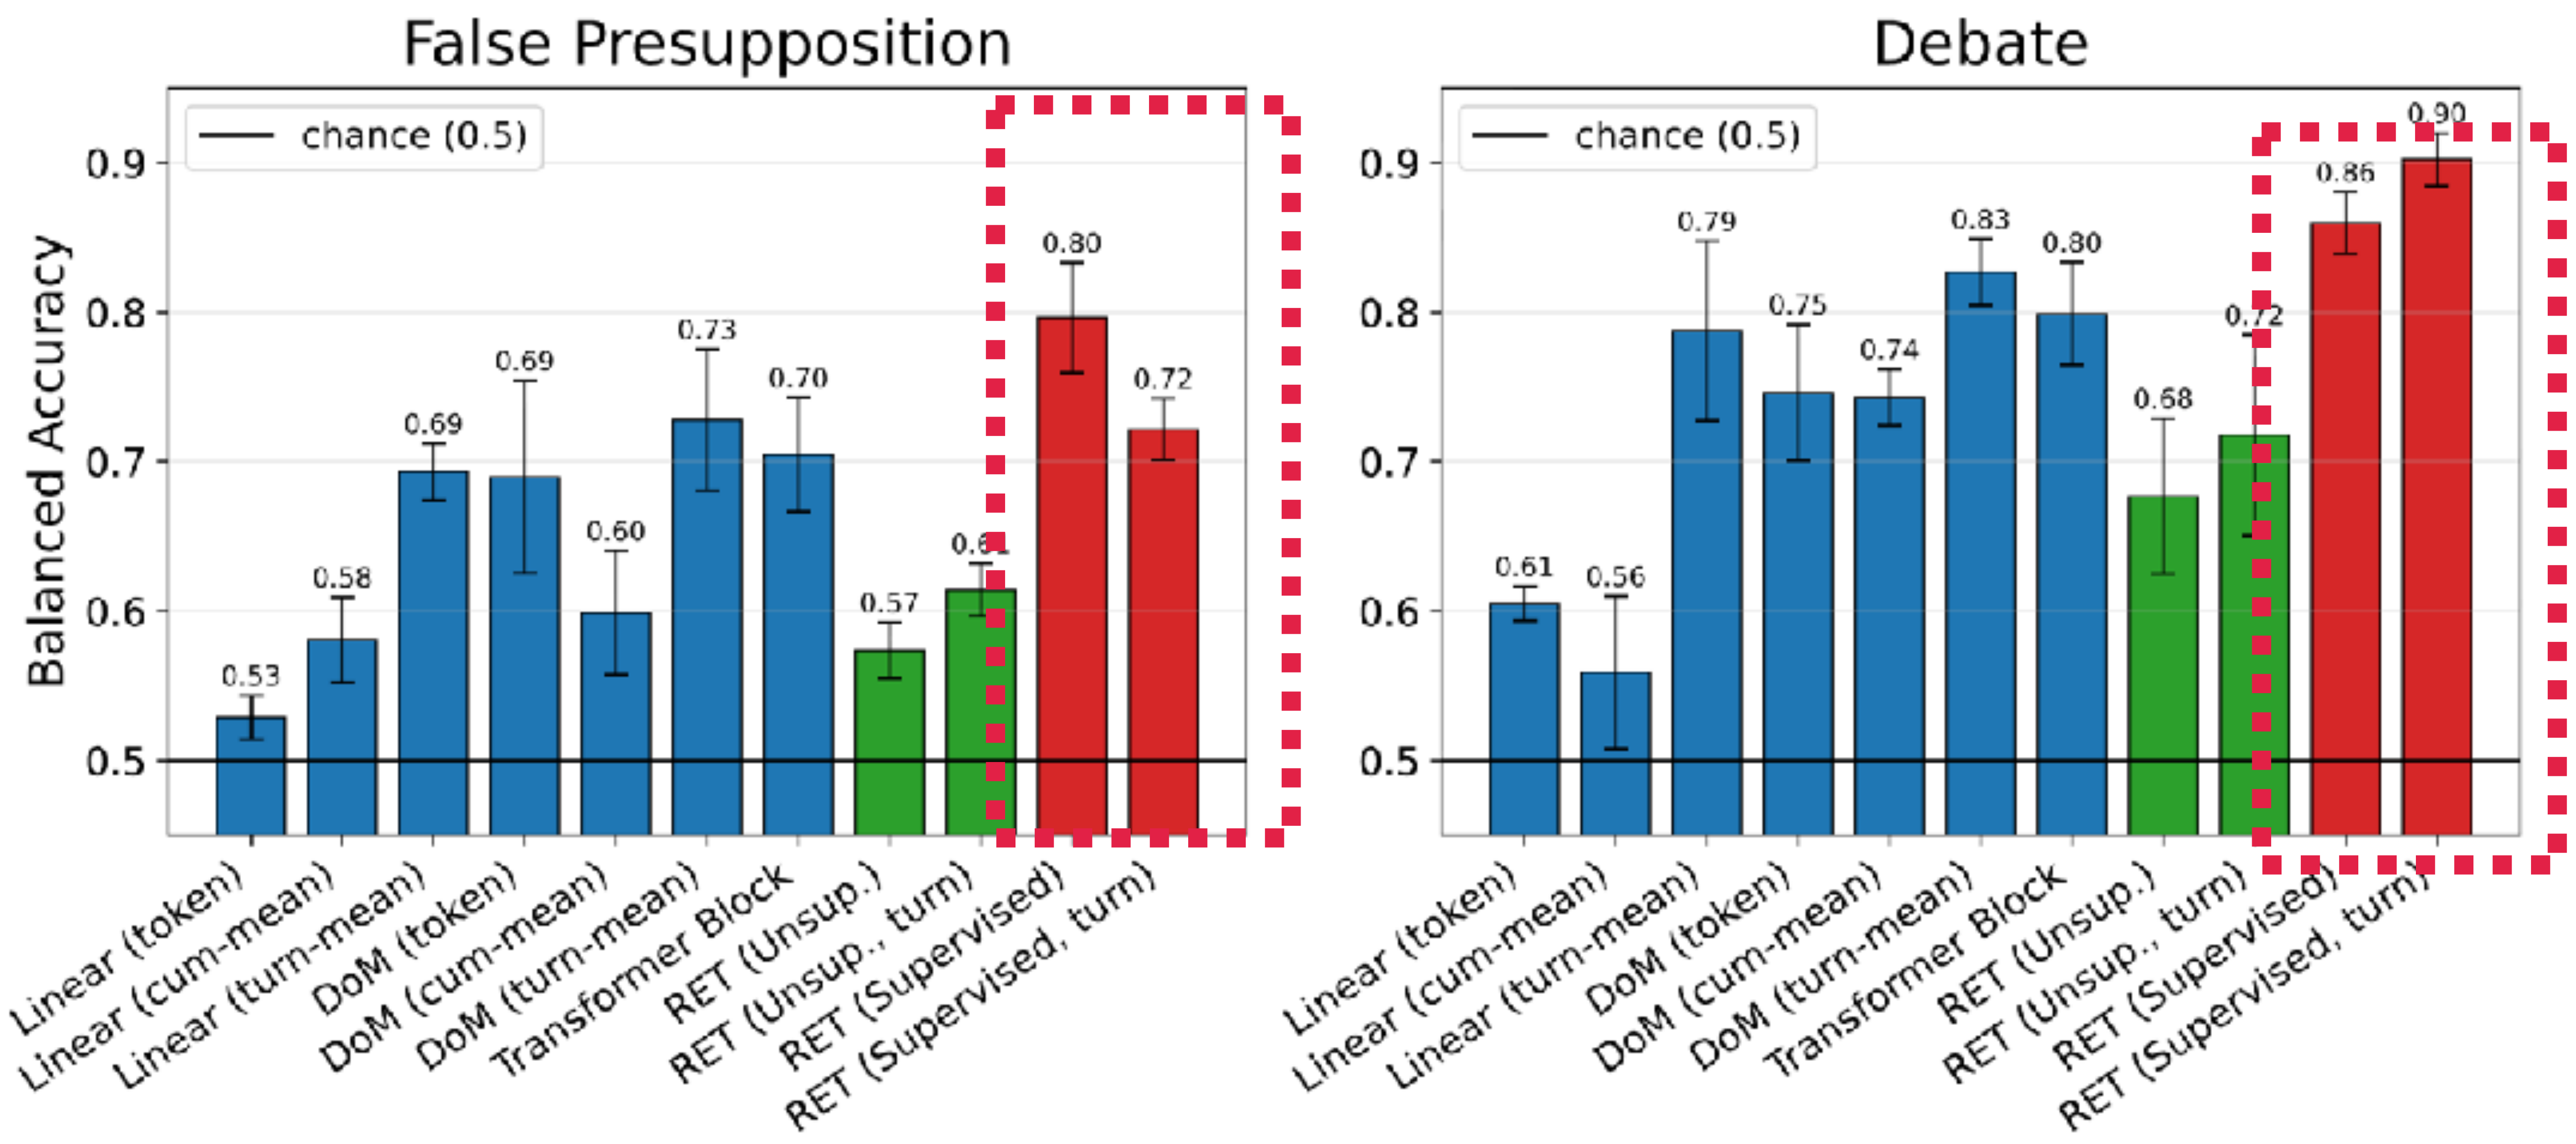

Controlling LLM Behavior

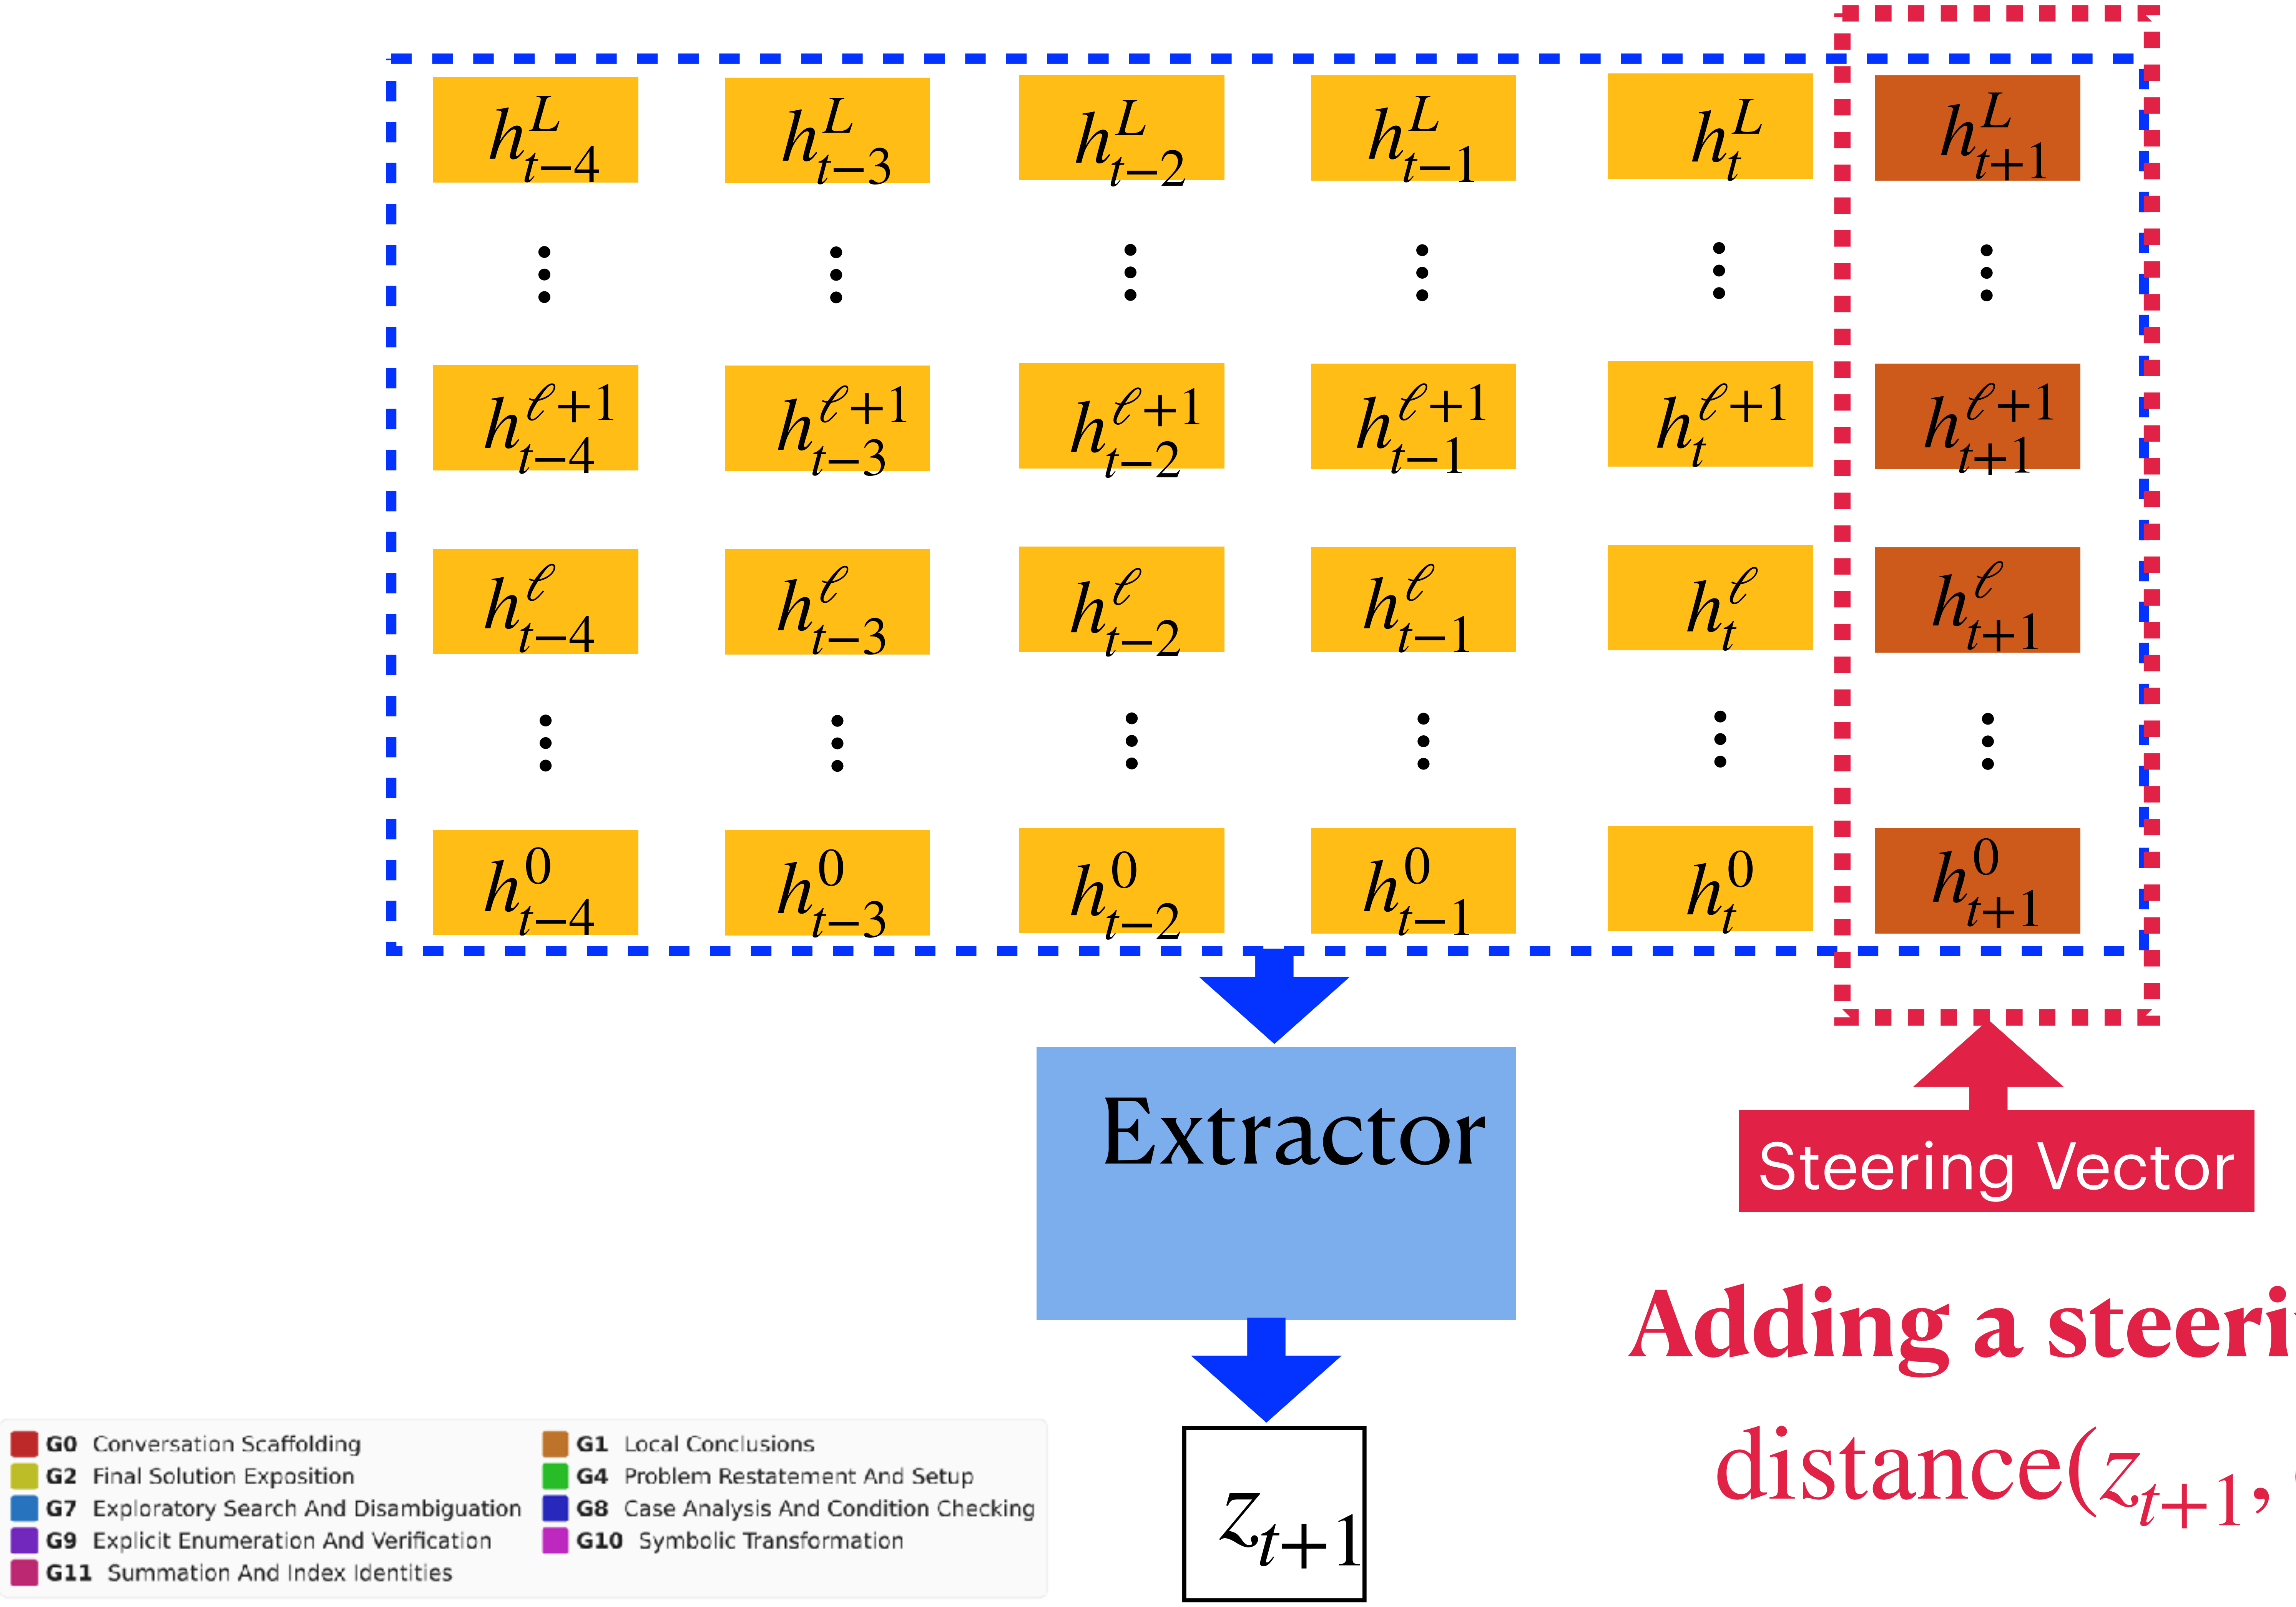

Adding a steering vector to minimize

$\text{distance}(z_{t+1}, \text{disiredclustercenter})$

Controlling LLM Behavior

Asking LLM to calculate 1+2+3+…+11

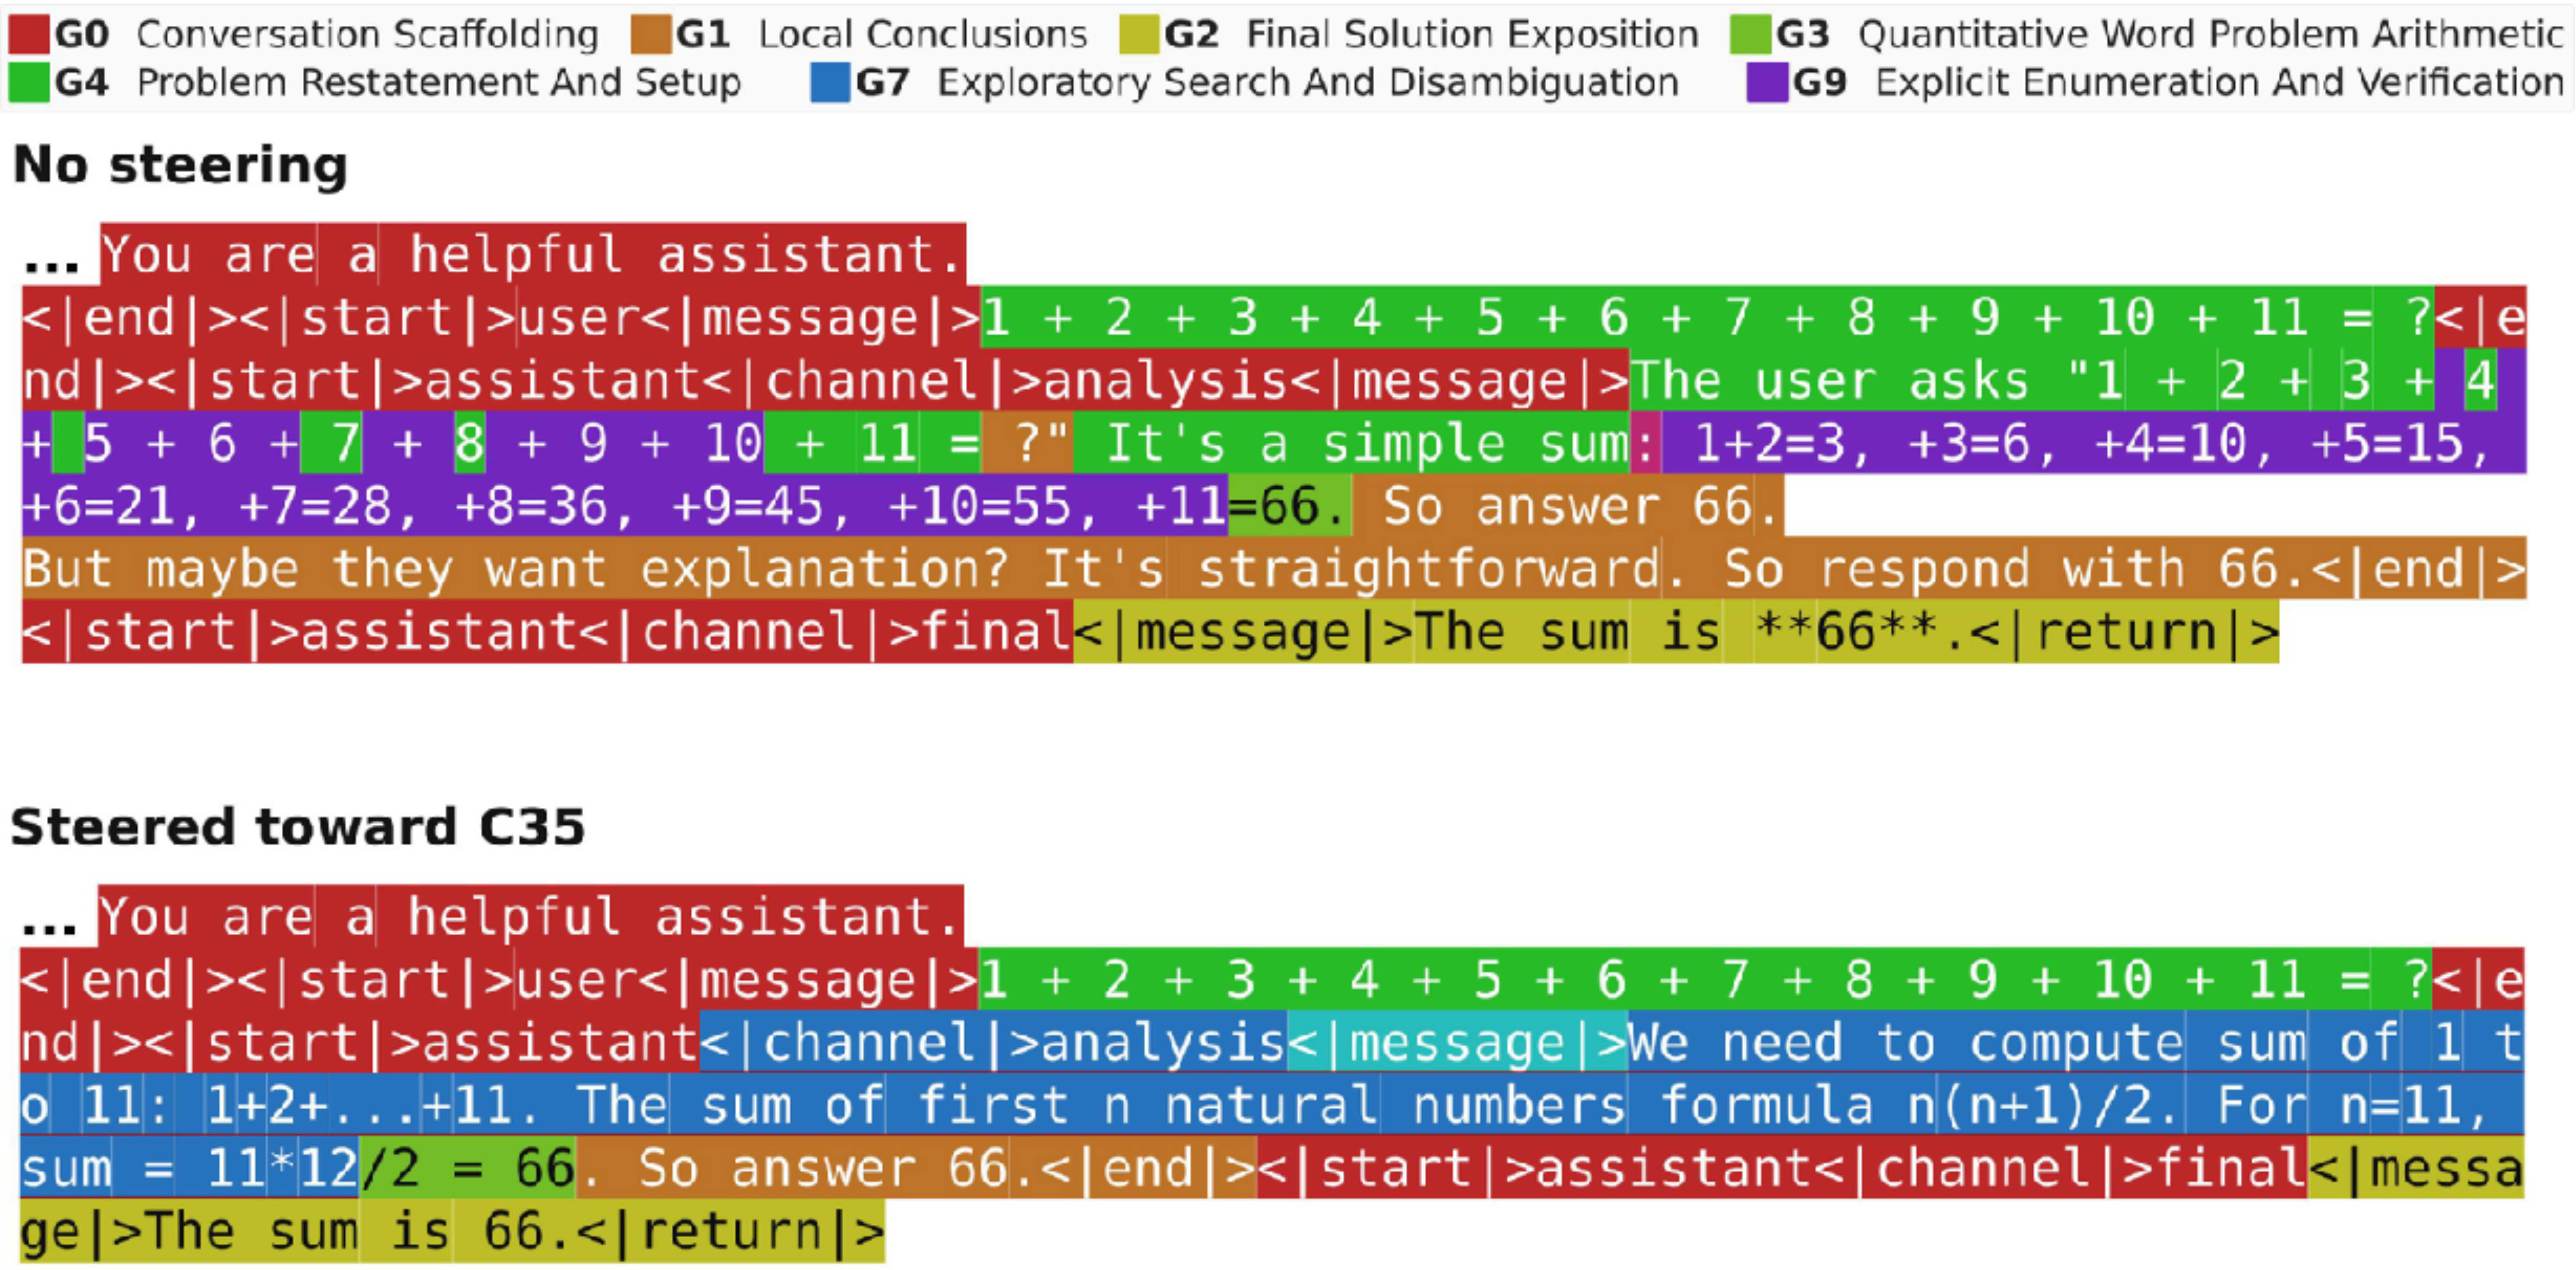

The default behavior is a one-by-one sum

Steering towards "**formula retrieval**" mode forces the model to sum using **n(n+1)/2**

Evaluating RET

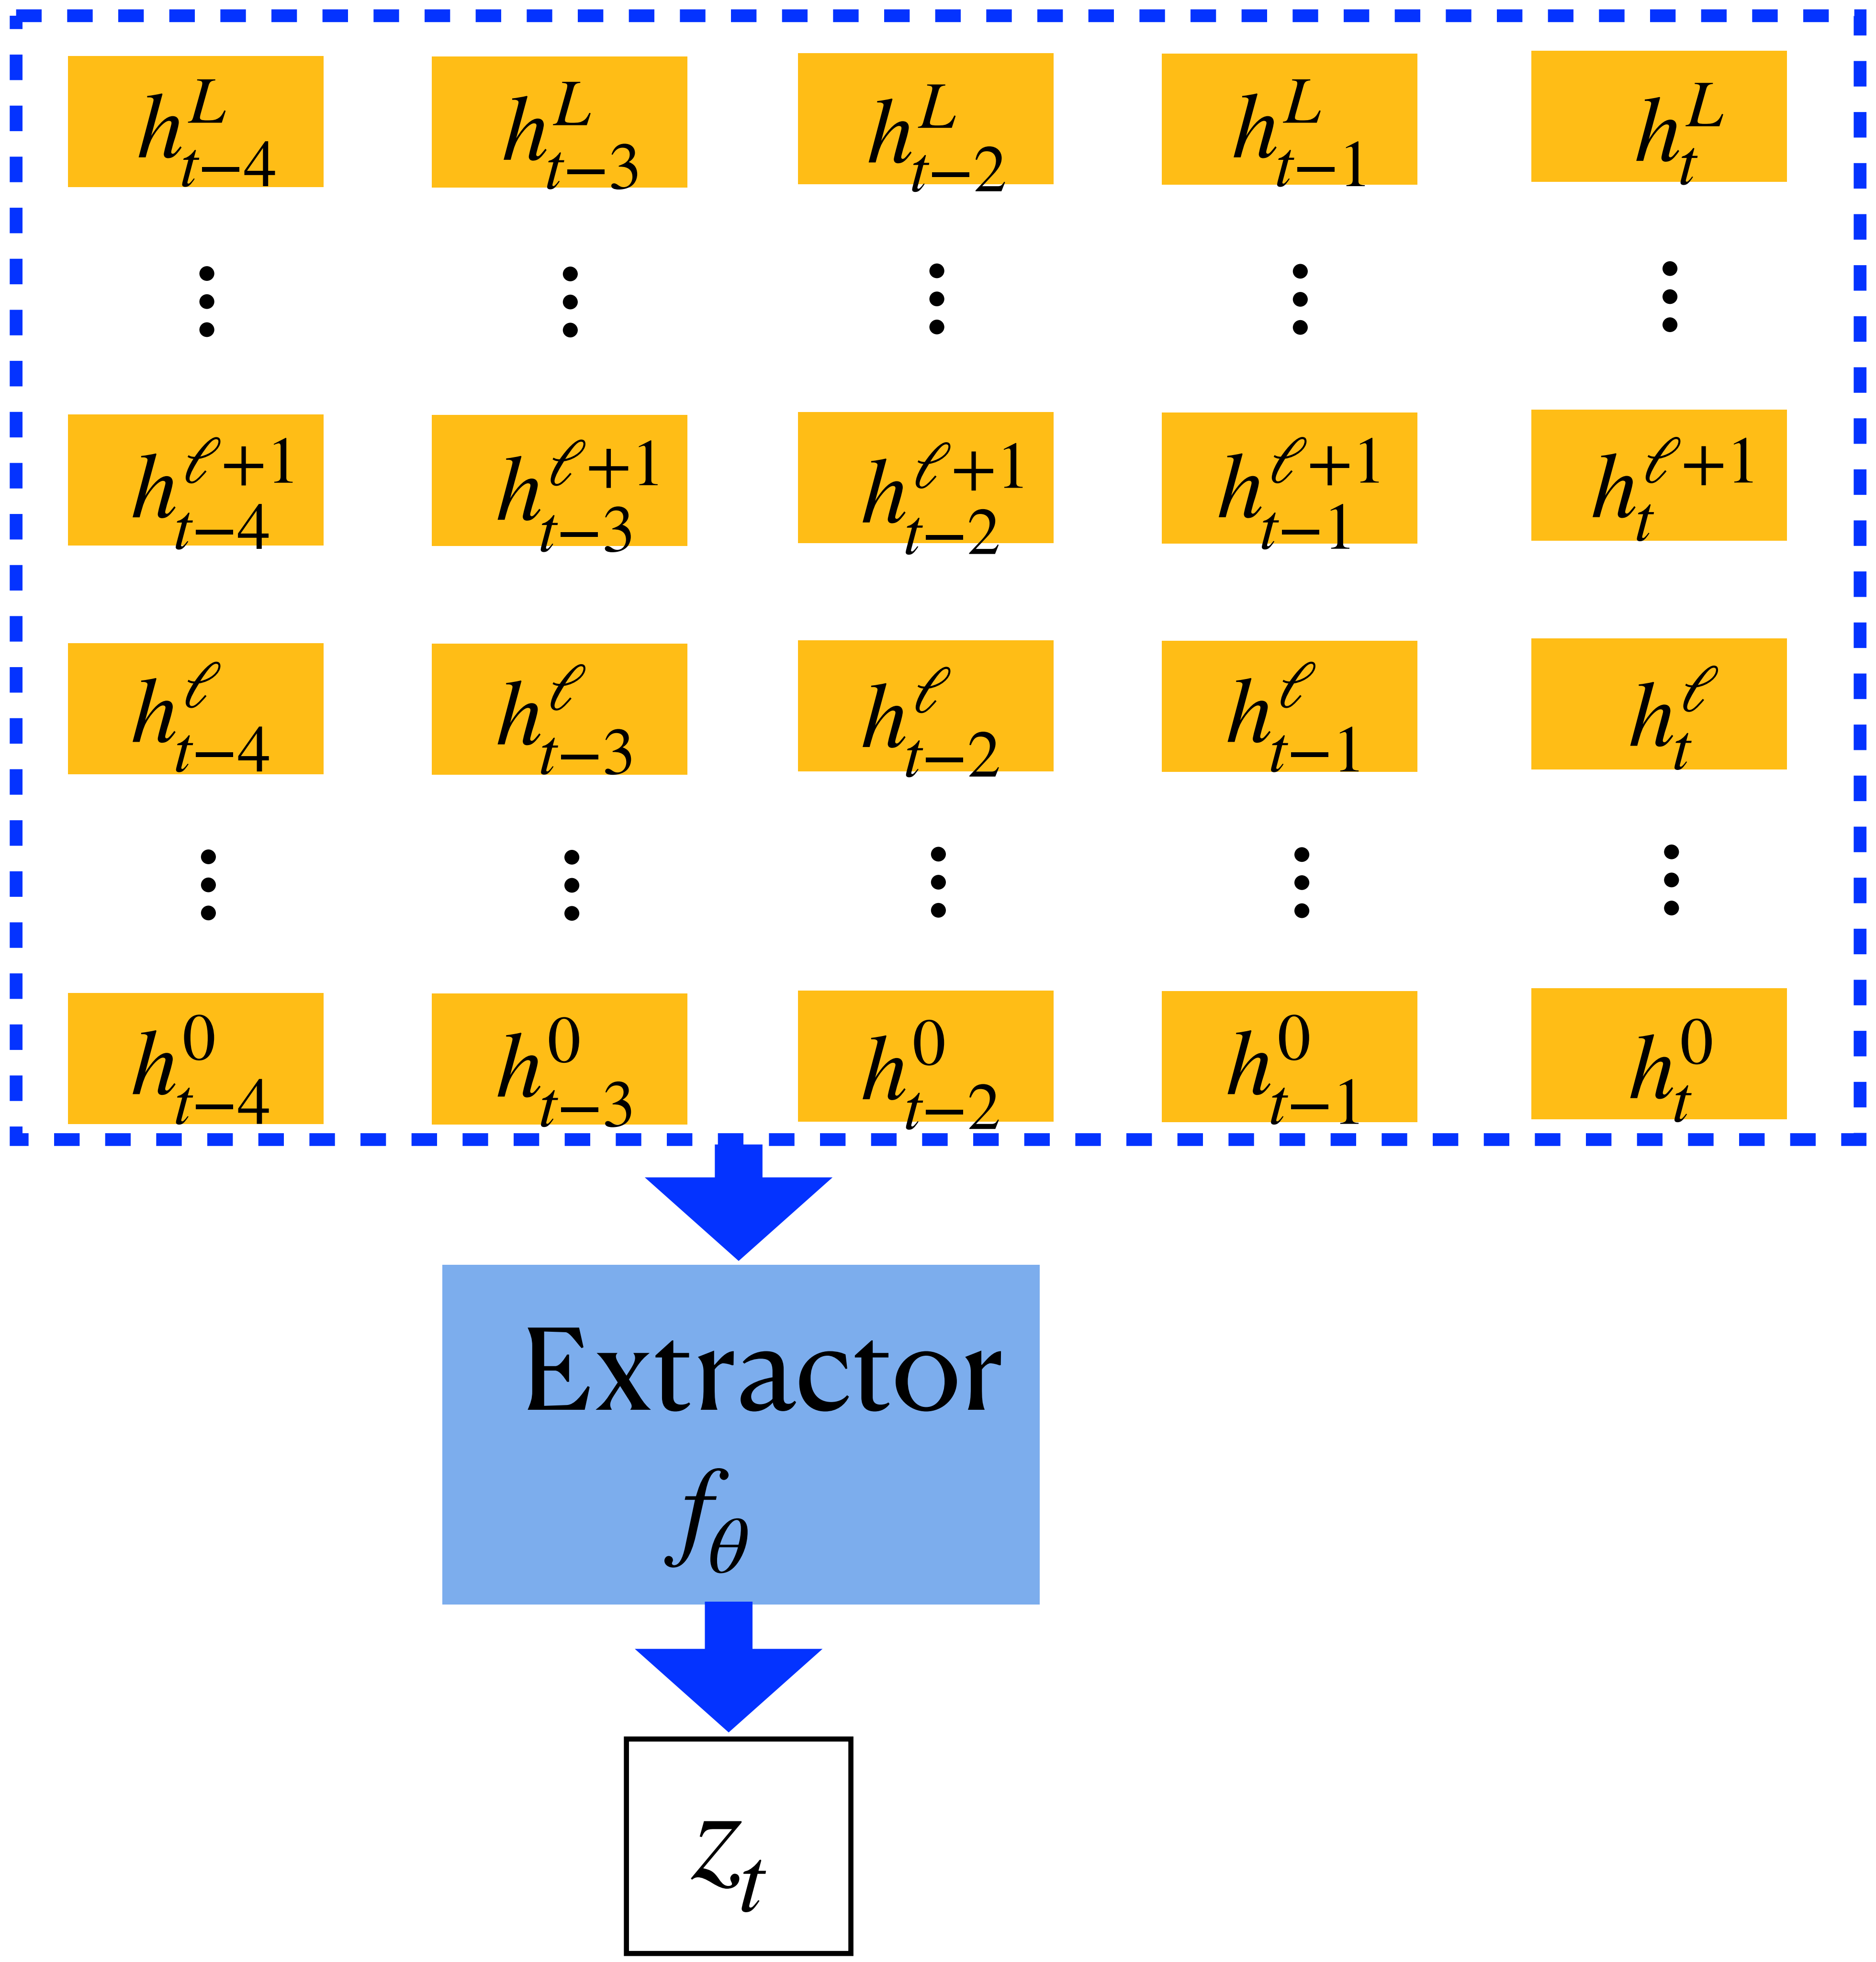

Are the learned states

- Low dimensional and "high-level"?

RET captures high-level plots and reasoning modes

- Self predictive (dynamically closed)?

RET achieves low prediction errors

- Practically useful?

RET enables monitoring and controlling LLM behavior

Limitations

- **RET is not universal**

  Each task (story/math/etc) needs to train seperately

- **The extracted state may be incorrect on OOD trajectories**

  The same asking LLM to sum “1+...+11” example **does not work on “1+...+61”**

  Even if we steer away from “formula retrieval”, it still uses the formula n(n+1)/2

Future Works

- **What is the best way to extract high level state?**

  We don't think the current way is the best — it is the first thing we tried and it worked.

- **Can we do controllability/reachability analysis?**

  If "unsafe" states are not reachable, then we have safety guarantee

- **Can do we better steer/control LLM behavior?**

  Current way is more like "open loop" control.

  Maybe closed loop control is better (MPC, iLQR, etc...)

- **Can we build safety monitors and controllers using the approach?**

Take Home Message

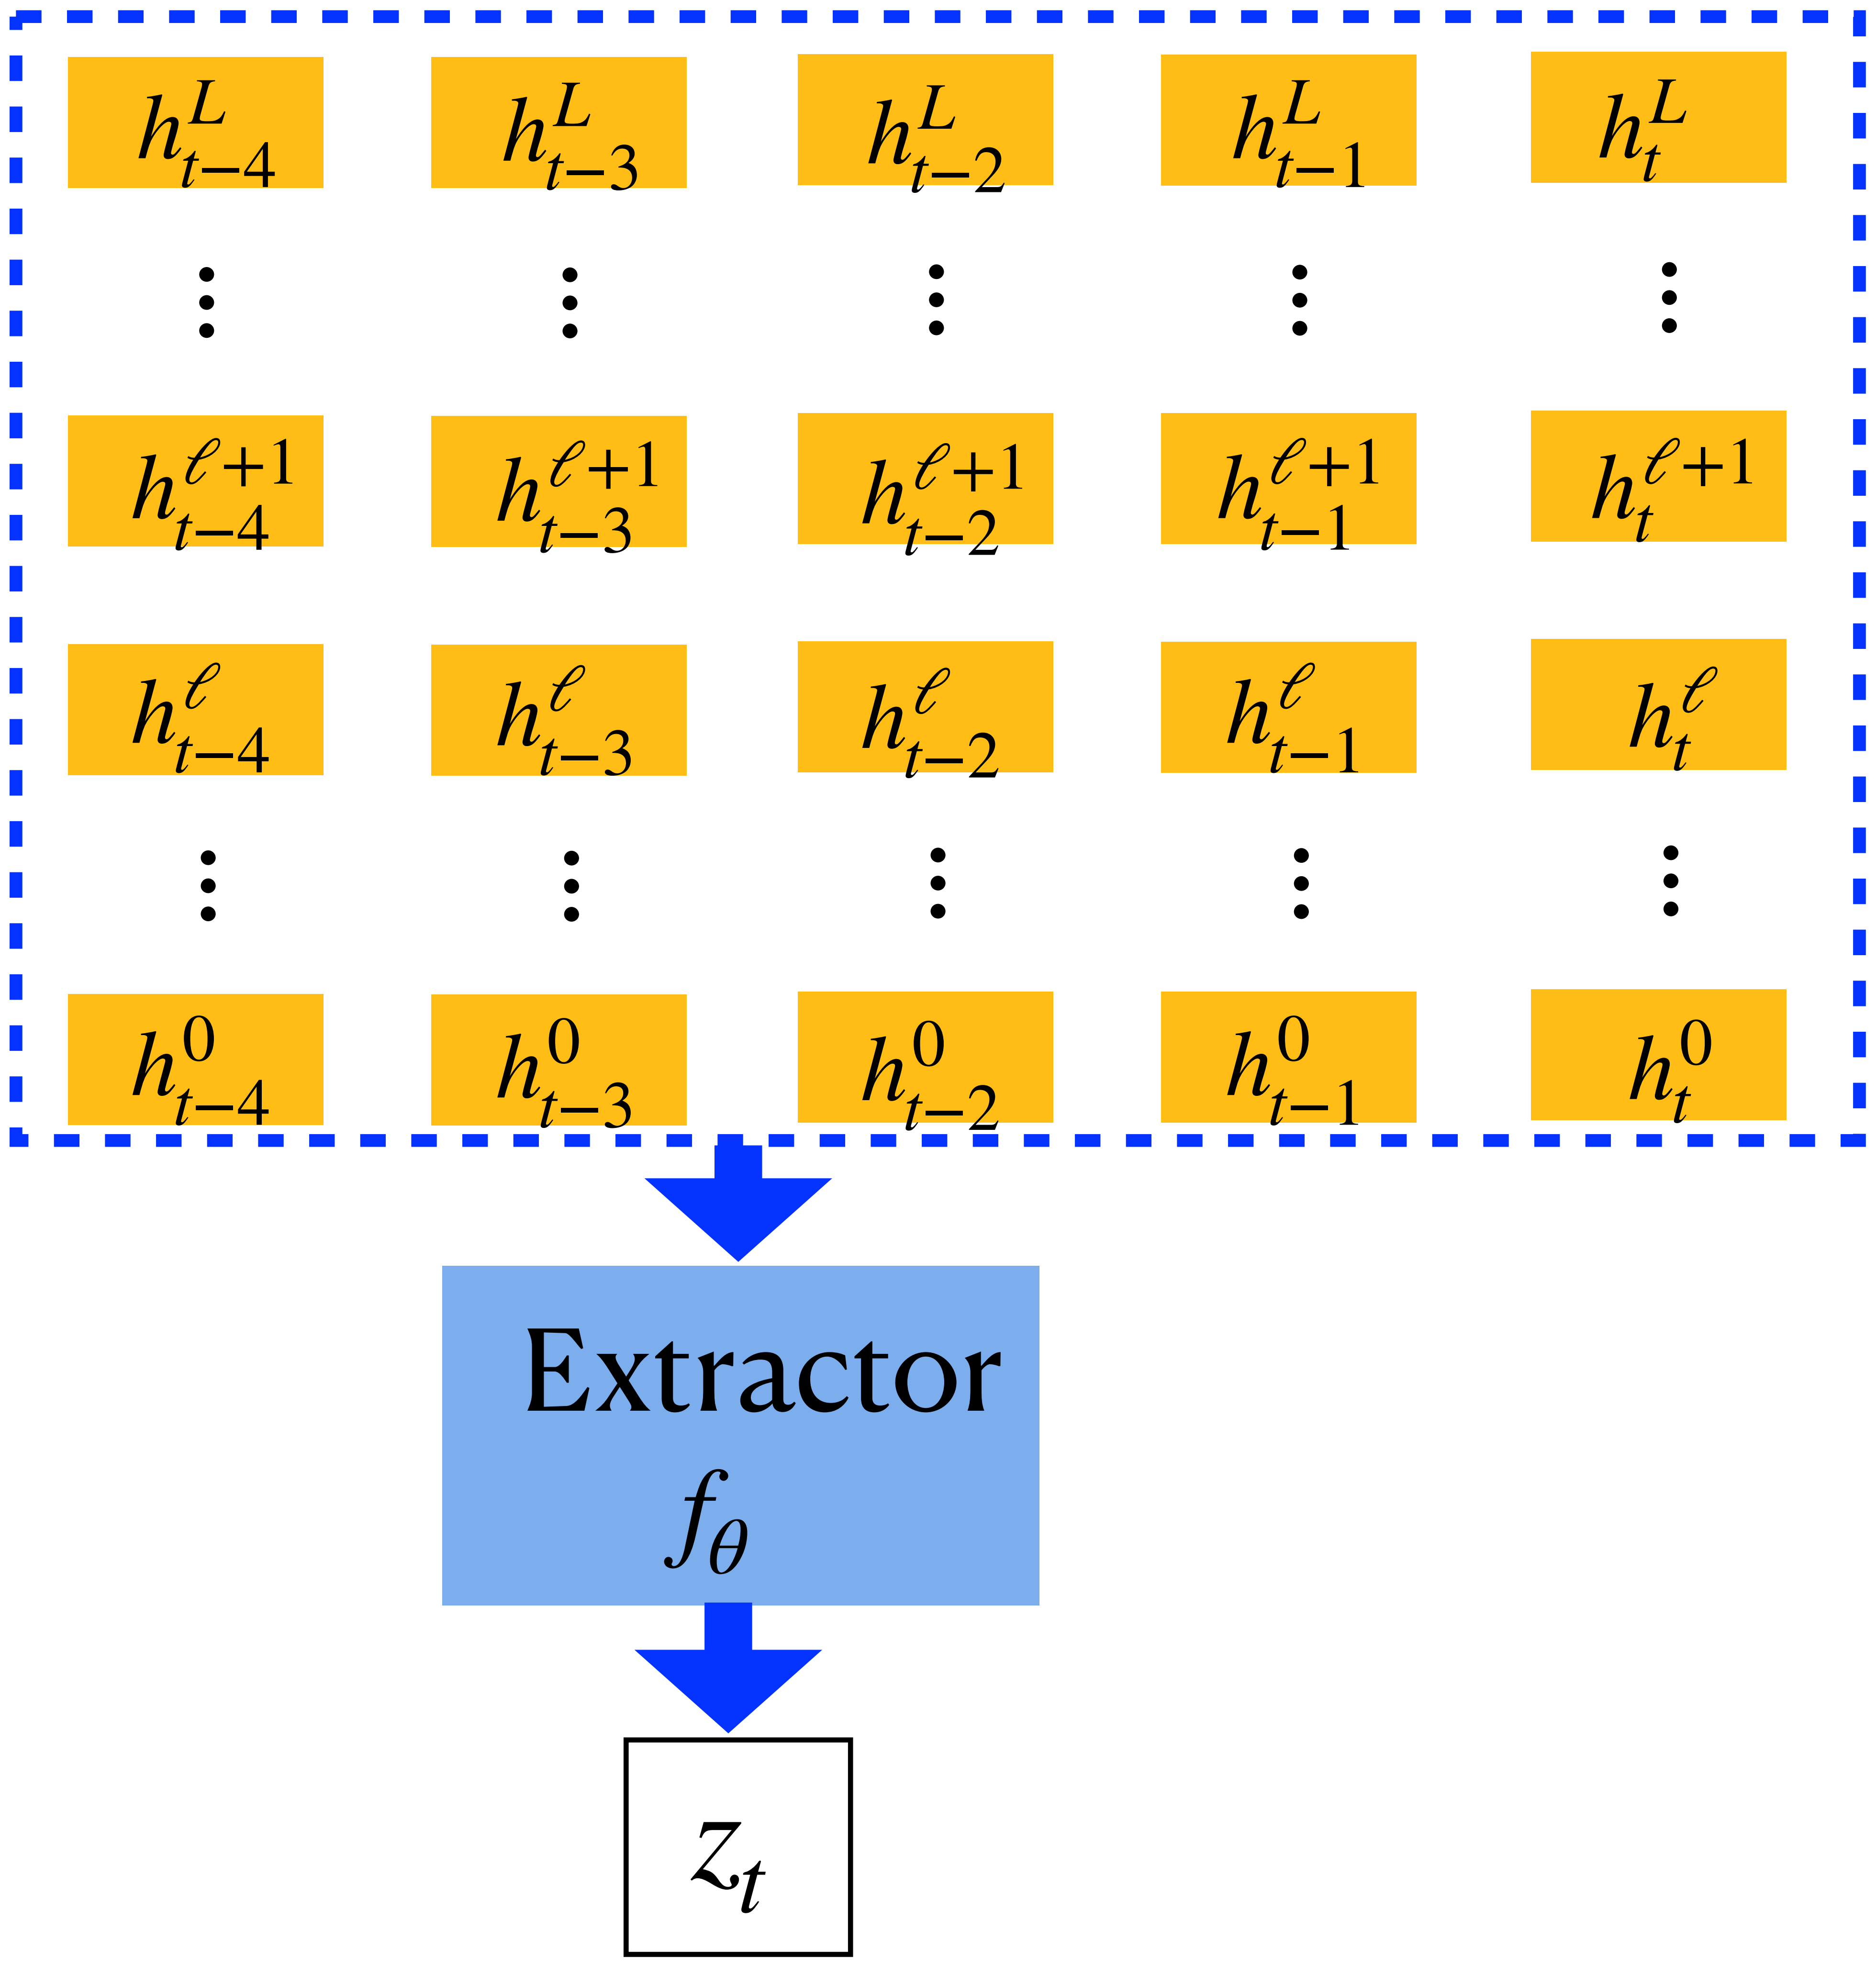

RET Extracts States of LLMs that are

- Low dimensional and "high-level"

RET captures high-level plots and reasoning modes

- Self predictive (dynamically closed)

RET achieves low prediction errors

- Practically useful

RET enables monitoring and controlling LLM behavior

Take Home Message

Interactive Demos

https://www.guannanqu.com/ret/

Paper

Muhammed Ustaomeroglu, Guannan Qu, Towards Effective Theory of LLMs: A Representation Learning Approach, arXiv:2605.09294, May 2026