

Pairwise Comparisons with Flexible Time-Dynamics

Lucas Maystre, Victor Kristof, Matthias Grossglauser

KDD Research Track 2 — August 6th, 2019



Pairwise comparison data

Association football example:

Team 1	Team 2	Score
France	Portugal	2-5
Luxembourg	Greece	0-3
...
Turkey	Slovakia	1-1
Bulgaria	Kosovo	0-1

Questions we might want to ask:

“How can we quantify the
skill of France?”

“How likely is South Korea
to **beat** Germany?”

→ we need **pairwise comparison** models.

Latent skill model

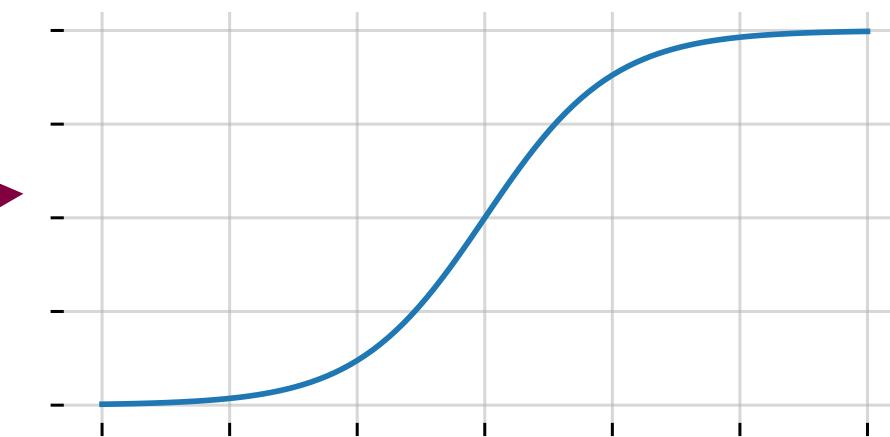
[Zermelo, 1928]
[Thurstone, 1927]

$$s_i \sim \mathcal{N}(0, \sigma^2)$$

latent skill of i

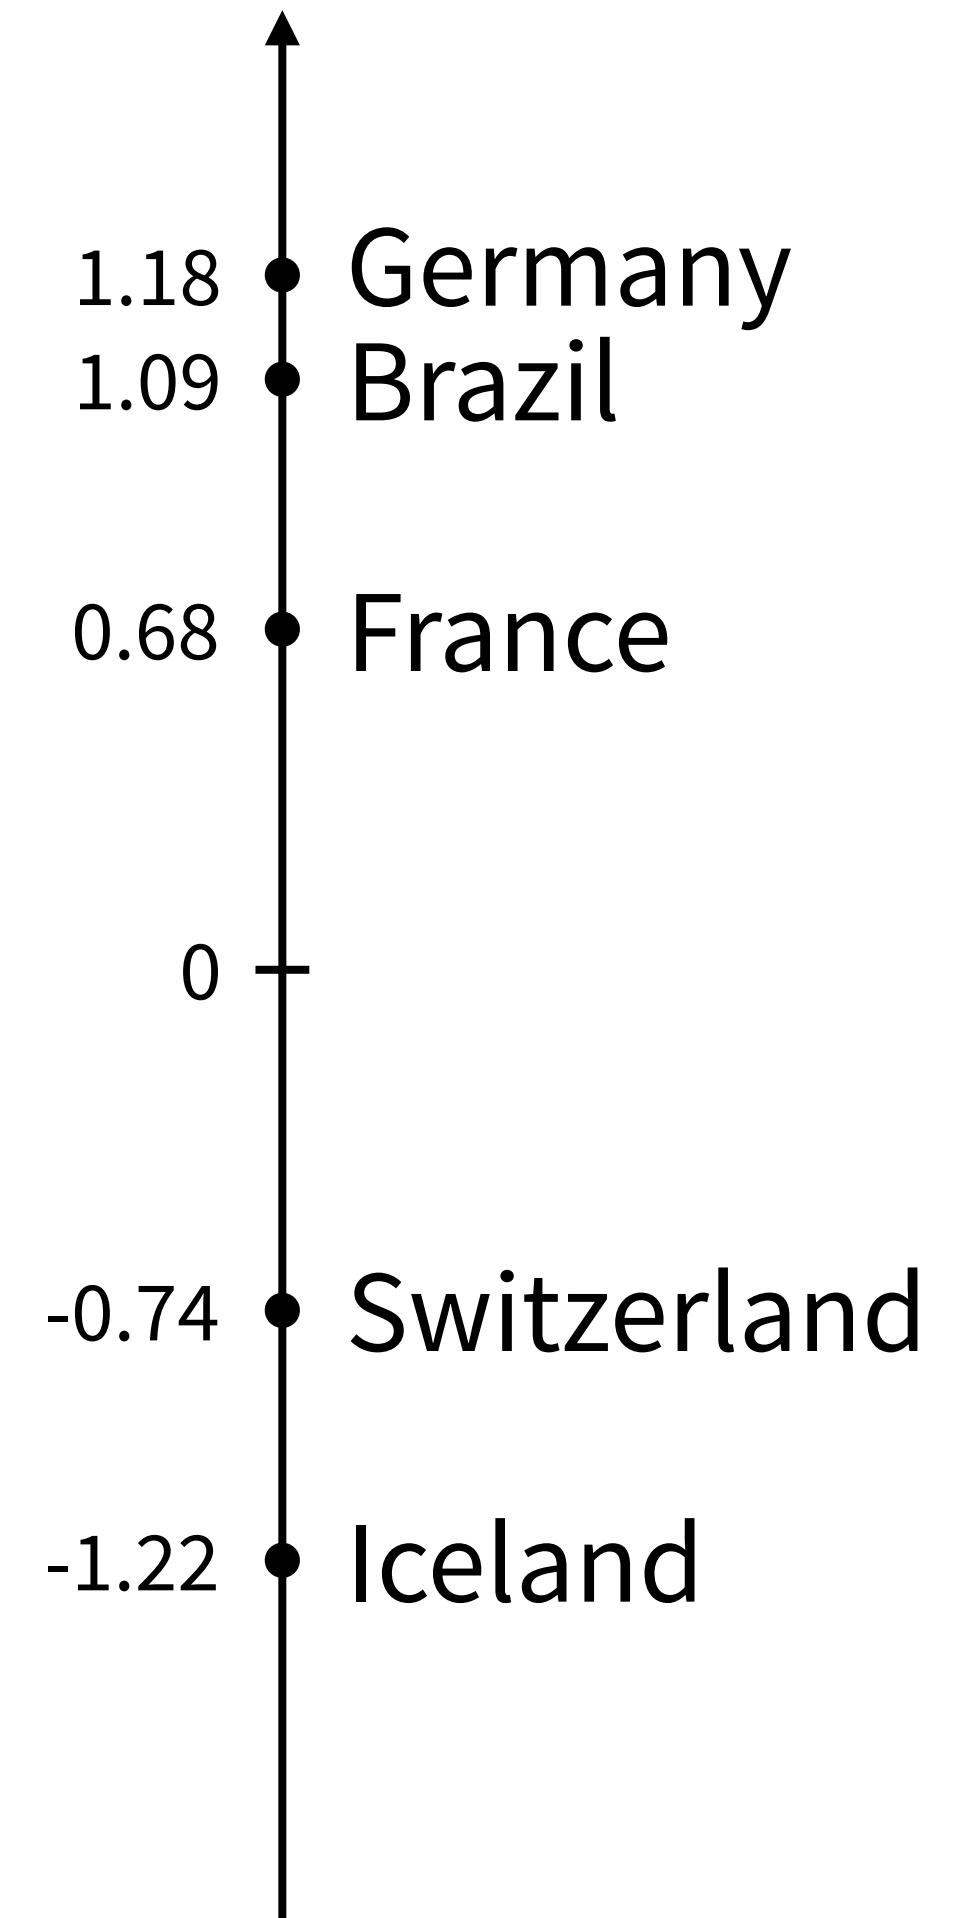
$$p(i \succ j) = \frac{1}{1 + \exp[-(s_i - s_j)]}$$

“ i wins over j ”



Given data \mathcal{D} , the posterior distribution is

$$p(\mathbf{s} | \mathcal{D}) \propto p(\mathbf{s}) \prod_{(i,j) \in \mathcal{D}} p(i \succ j)$$



Actual setting

Data come with a **timestamp**.

Date	Team 1	Team 2	Score
1923-09-15	France	Portugal	2-5
1923-09-16	Luxembourg	Greece	0-3
...
2018-06-21	Turkey	Slovakia	1-1
2018-06-21	Bulgaria	Kosovo	0-1

Questions we might want to ask:

“How strong was France
in **1972**? In **2018**? ”

“How likely is South Korea
to beat Germany **today**? ”

→ we need **dynamic** models.

This work: kickscore

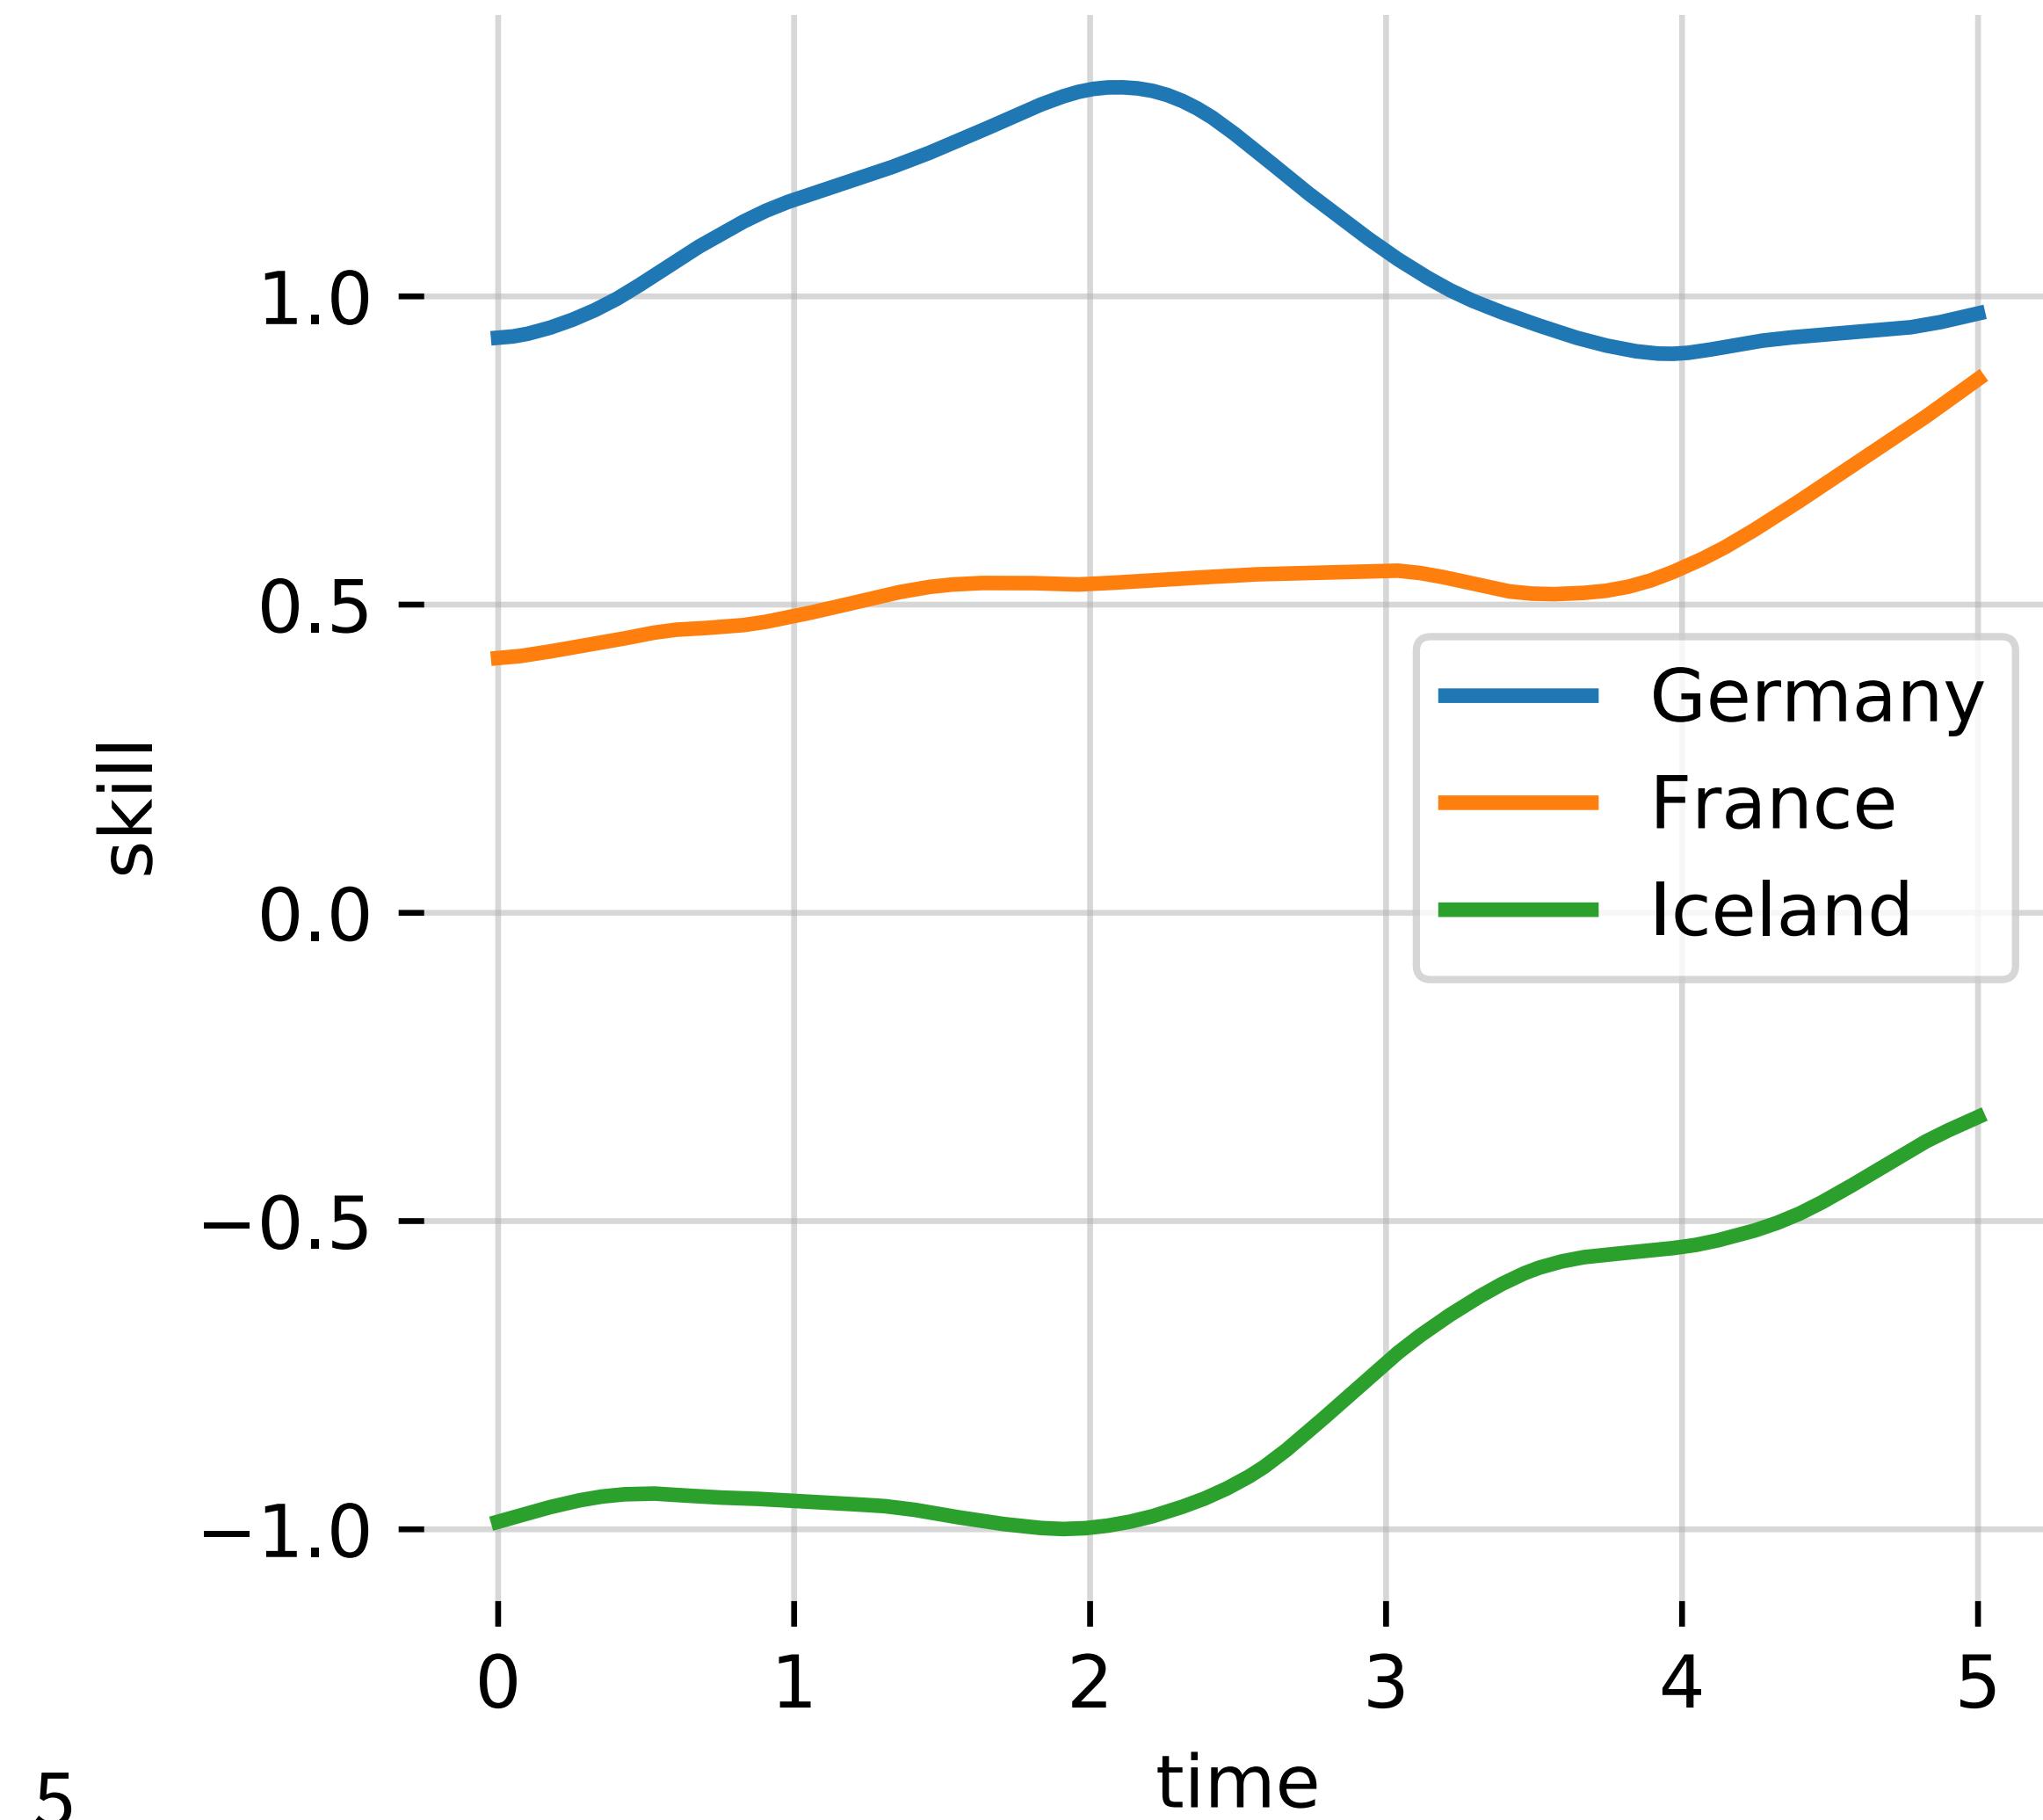
Skill becomes a (latent) **stochastic process**

$$s_i(t) \sim \text{GP}[0, k(t, t')]$$

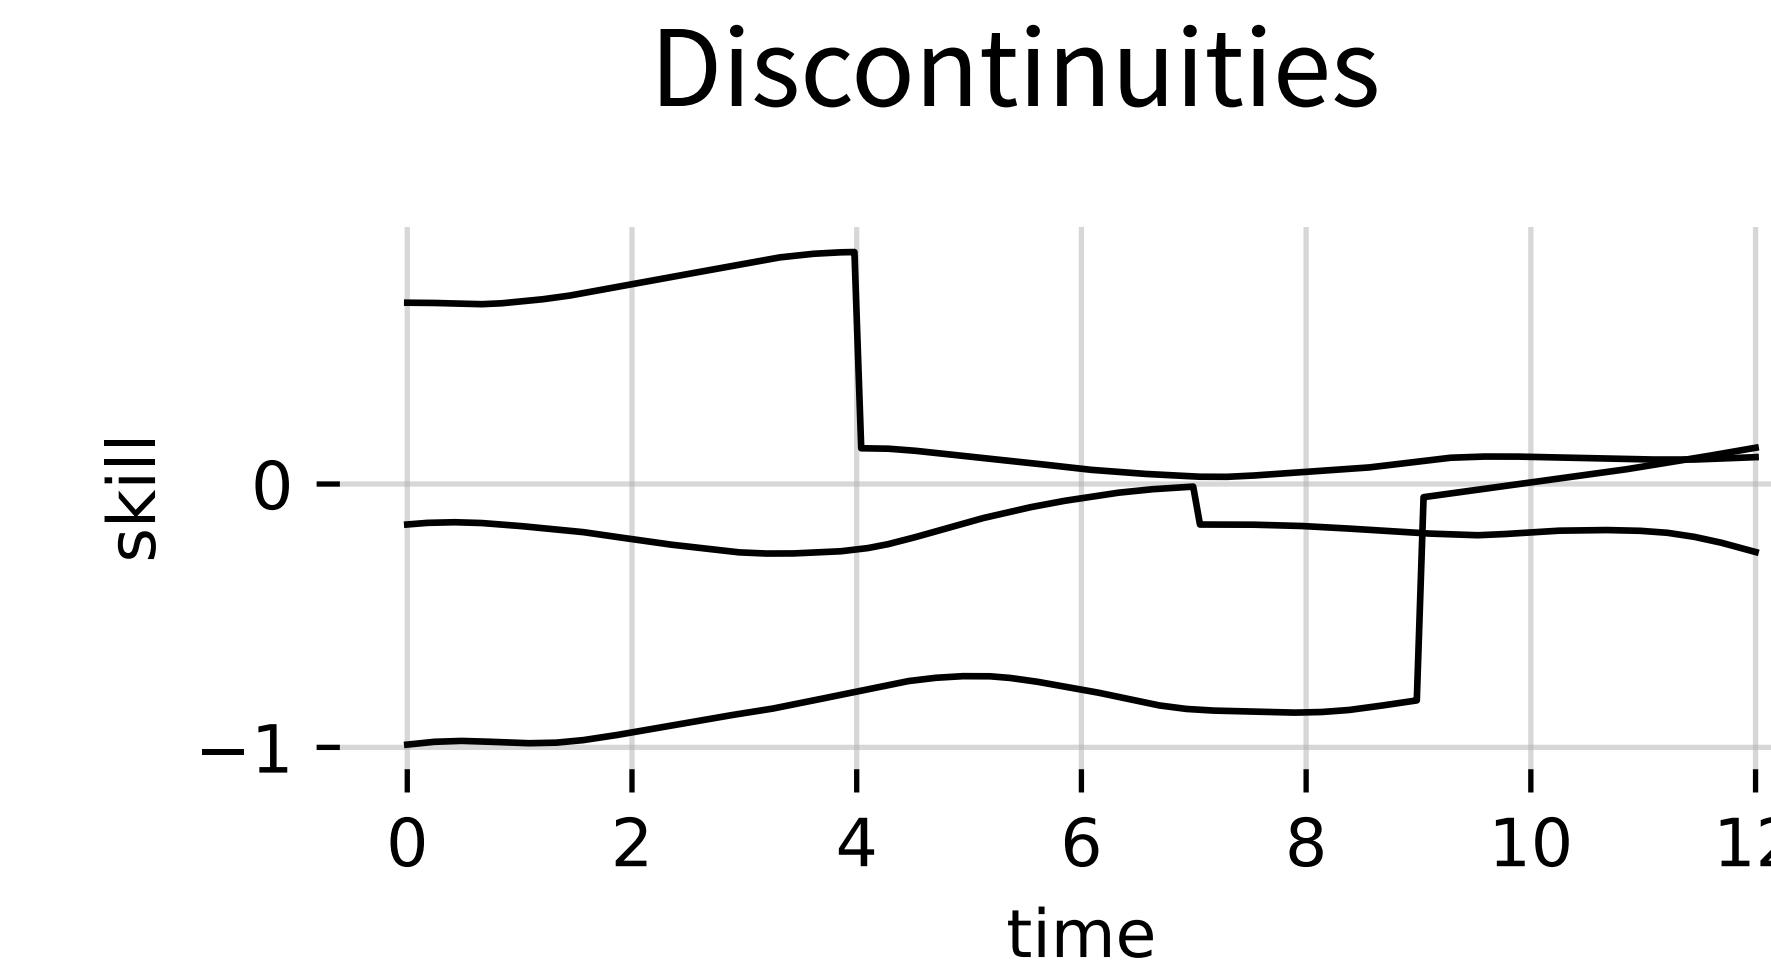
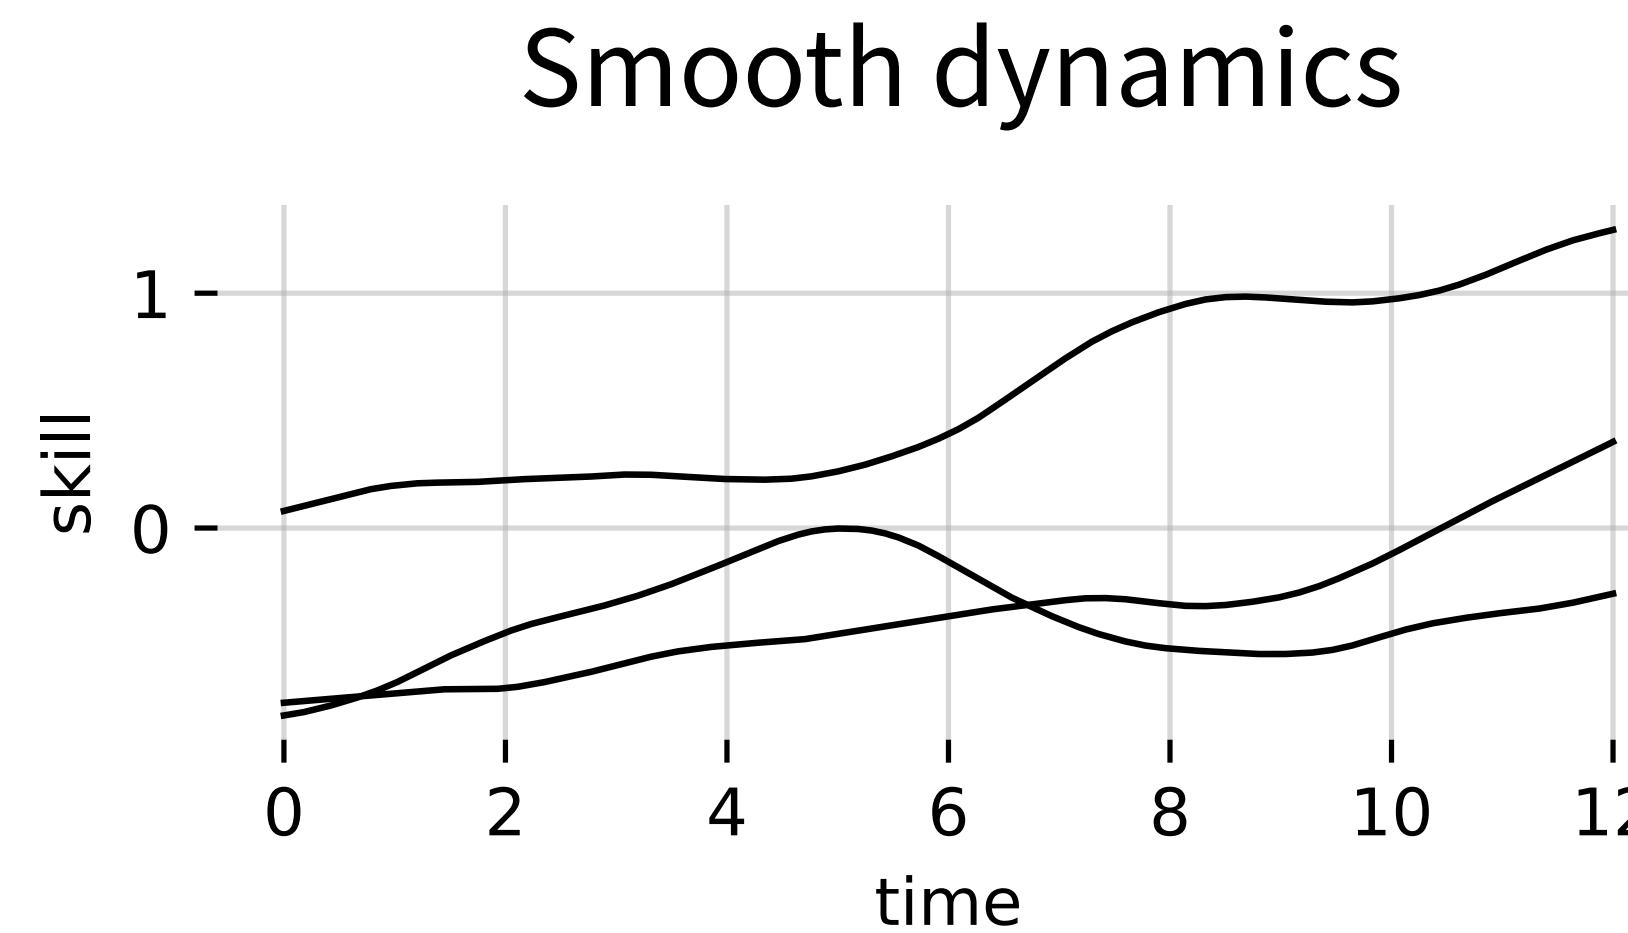
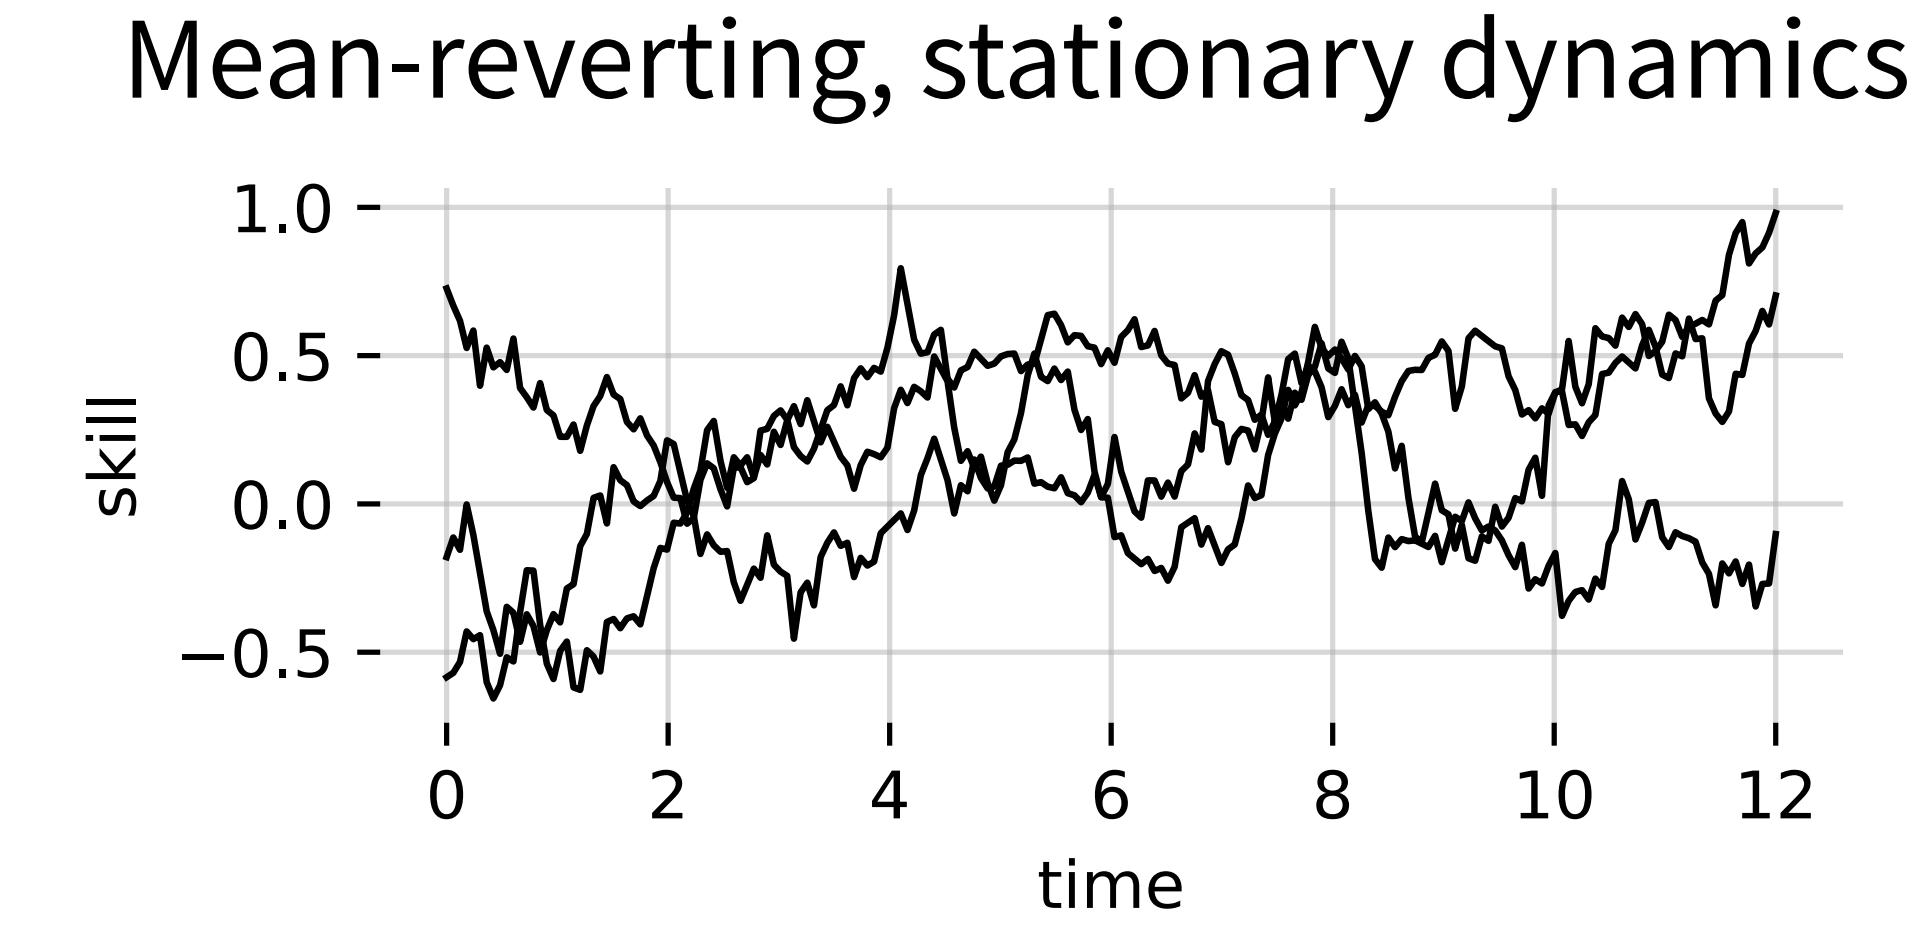
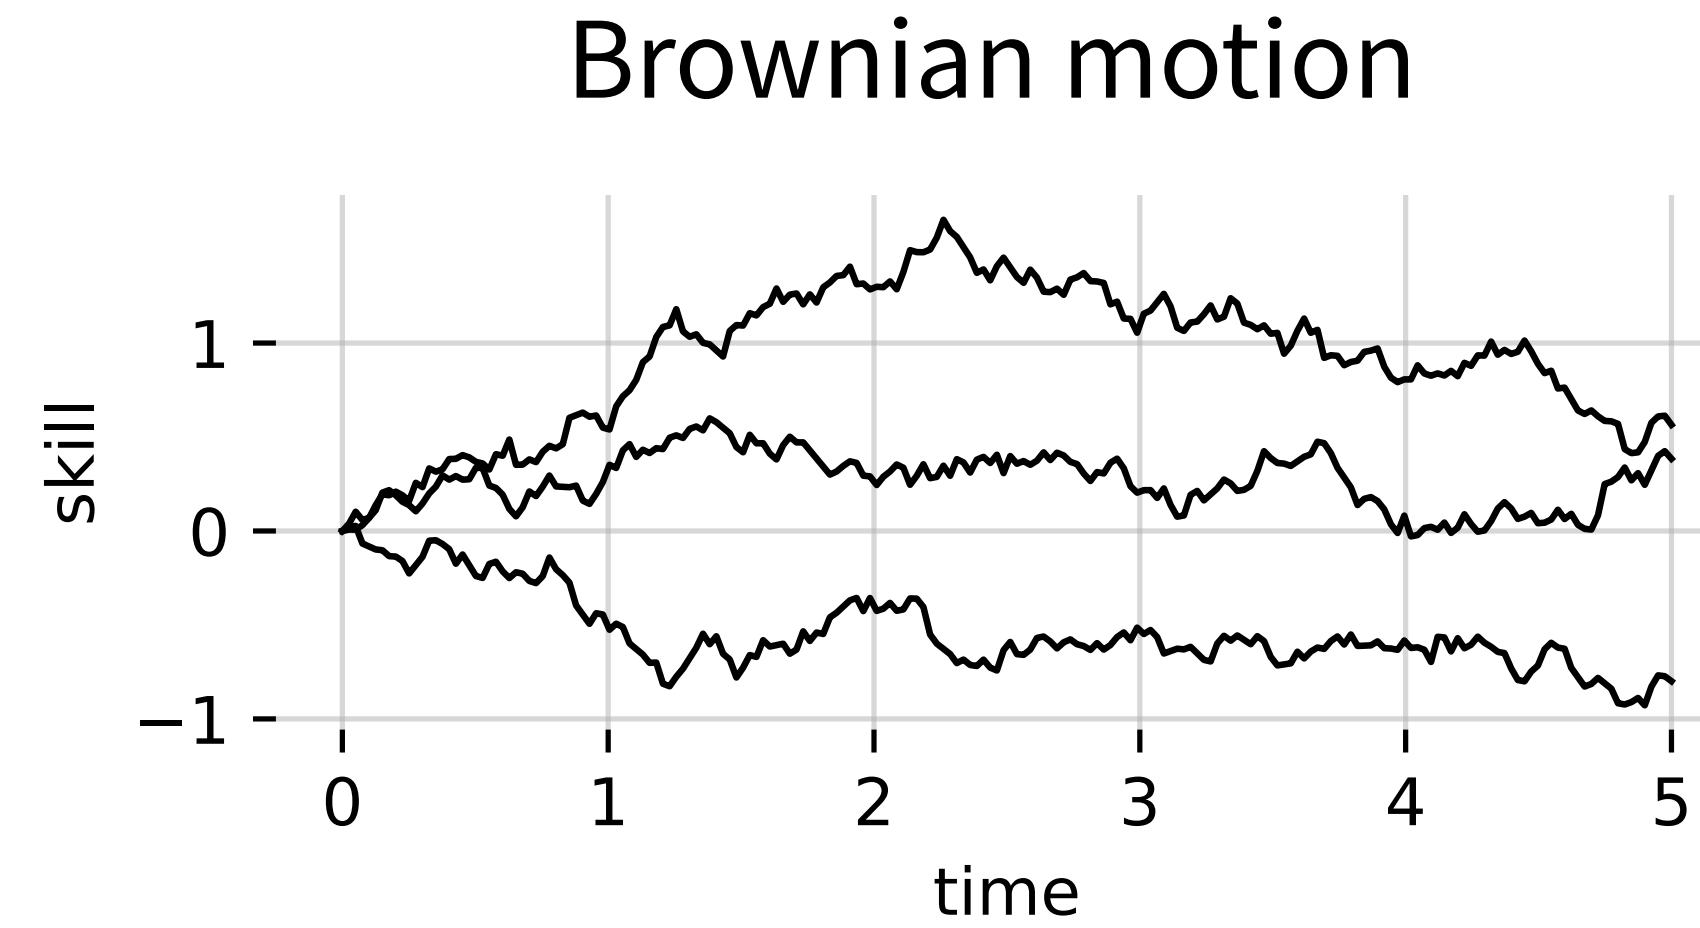
covariance function,
defines **time dynamics**

Obs. model is **conditionally parametric**:

$$p(i \succ j \mid t) = \frac{1}{1 + \exp\{-[s_i(t) - s_j(t)]\}}$$



Covariance functions



Outline

1

Inference
algorithm

2

Experimental
evaluation

Model inference

$$s_i(t) \sim \text{GP}[0, k(t, t')]$$

$$\mathbf{s}_i = [s_i(t_1) \quad \cdots \quad s_i(t_N)]$$

$$p(\mathbf{s}_1, \dots, \mathbf{s}_M \mid \mathcal{D}) \propto \prod_{i=1}^M p(\mathbf{s}_i) \prod_{(i,j,t) \in \mathcal{D}} p(i \succ j \mid t)$$

$$\min_{\{\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i\}} \text{KL}(q \| p)$$

$$\approx q(\mathbf{s}_1, \dots, \mathbf{s}_M) \doteq \prod_i \mathcal{N}(\mathbf{s}_i \mid \boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i)$$

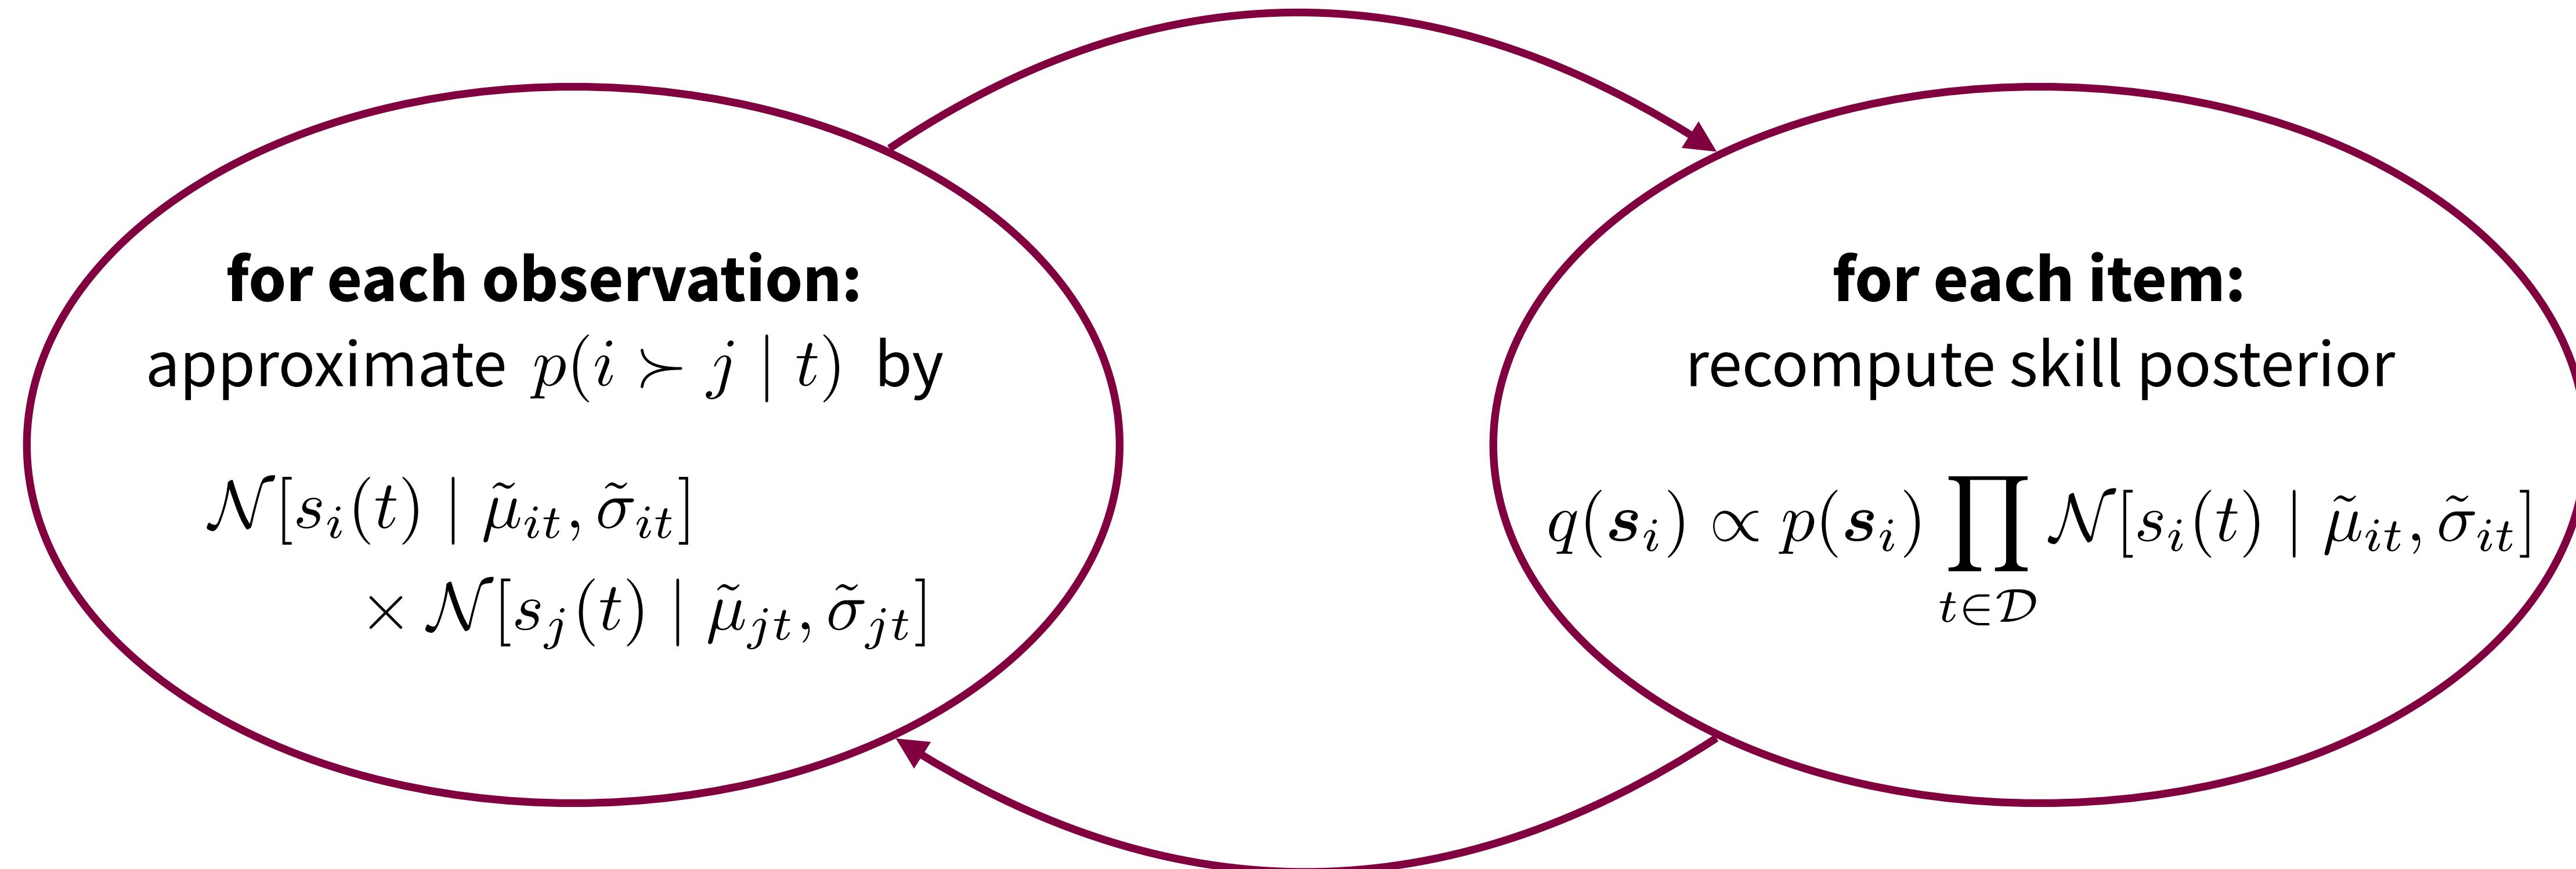
$$q(\mathbf{s}_i)/p(\mathbf{s}_i)$$

Alternative viewpoint:

$$q(\mathbf{s}_i) \propto p(\mathbf{s}_i) \overline{\prod_{t \in \mathcal{D}} \mathcal{N}[s_i(t) \mid \tilde{\mu}_{it}, \tilde{\sigma}_{it}]}$$

Inference algorithm

Converges in a few
linear time iterations

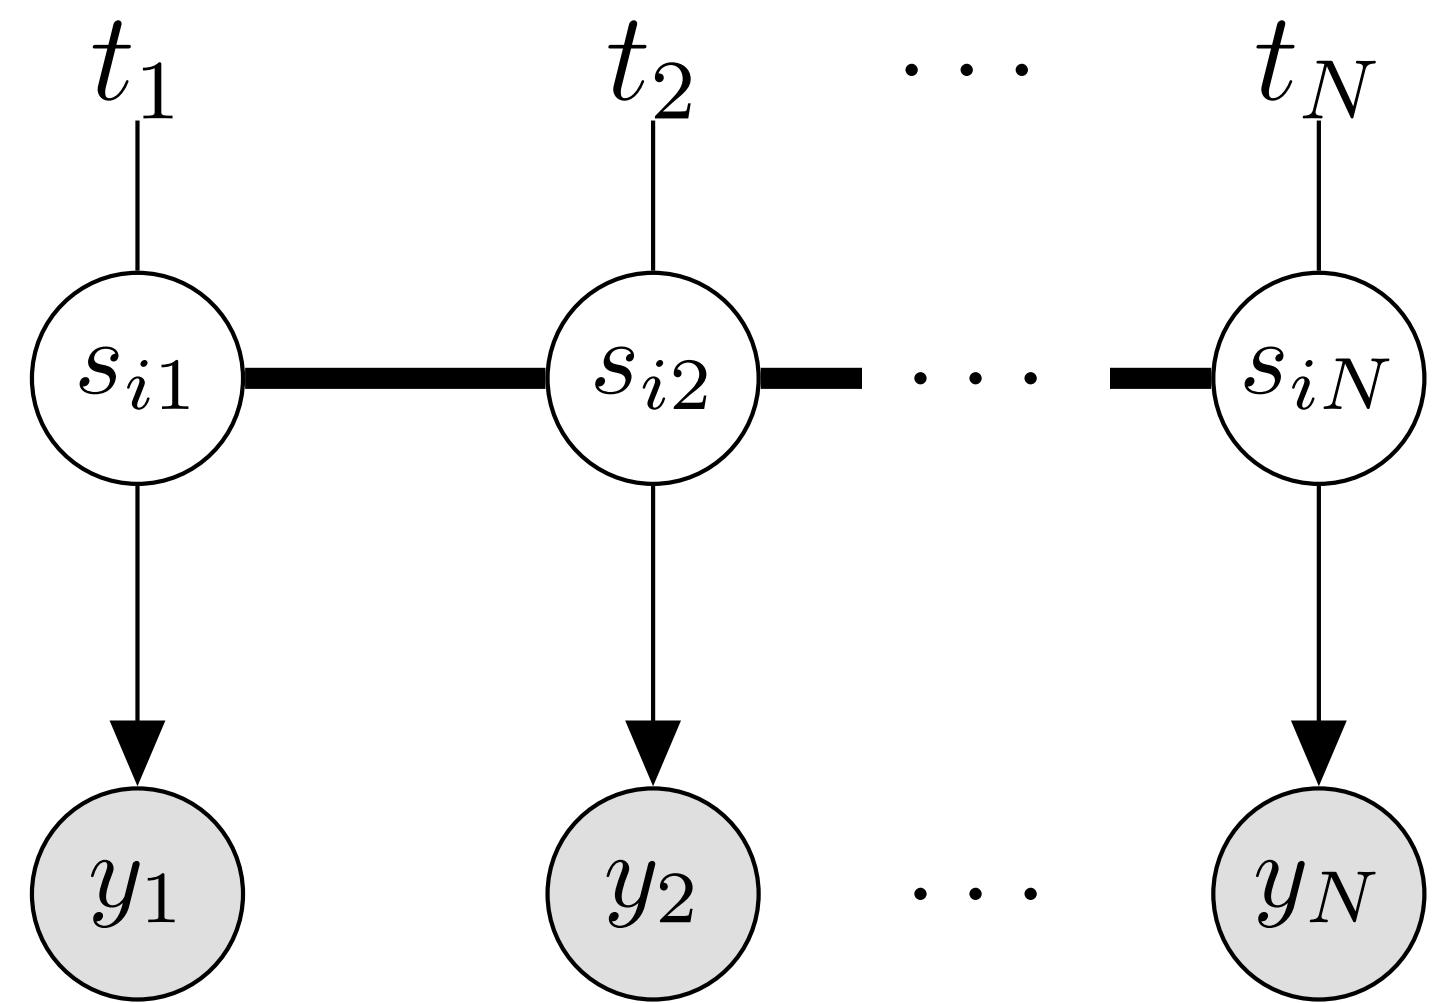


Using **EP** [Minka, 2001] or
CVI [Khan et al., 2017]

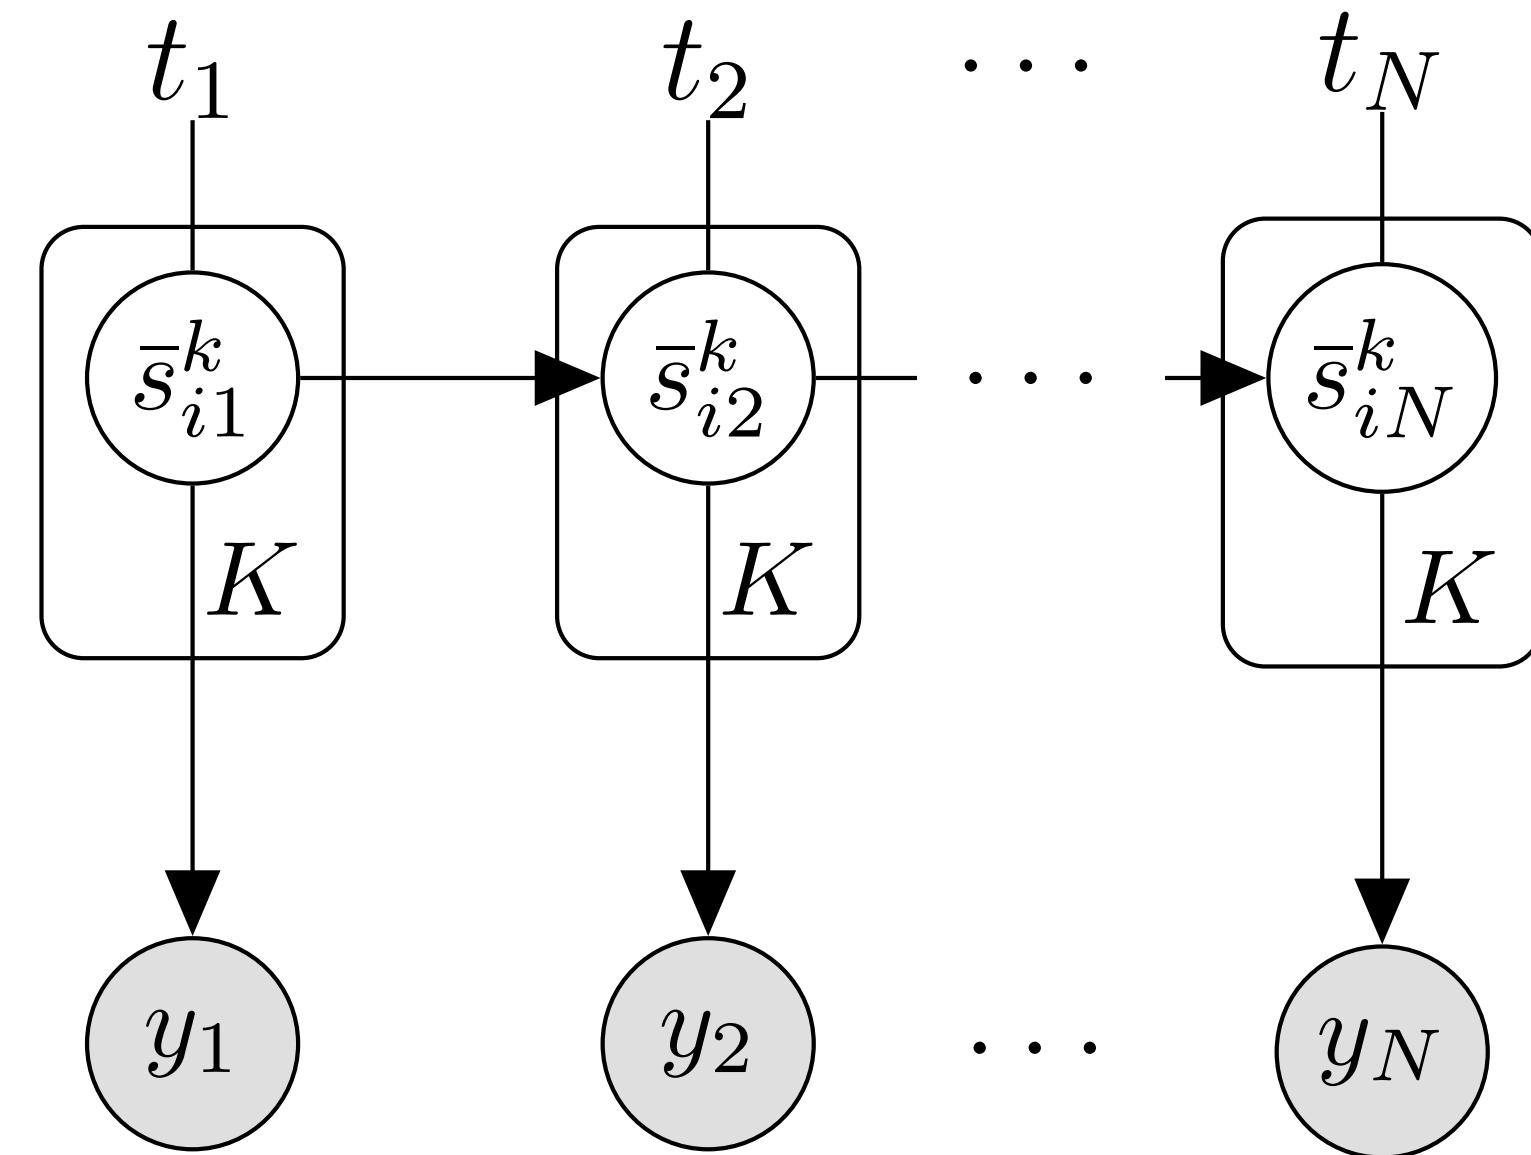
Using **SSM** reformulation
[Hartikainen & Särkka., 2010]

SSM reformulation

[O'Hagan, 1978]
[Hartikainen & Särkkä, 2010]



$s_i(t)$ is a Gaussian process



$\bar{s}_i(t)$ is a **Gauss-Markov** process

Experimental evaluation

Dataset	N	Timespan
ATP tennis	618,934	1991-2017
NBA basketball	67,642	1946-2018
FIFA football	19,158	1908-2018
Chess	7,169,202	1475-2017
StarCraft WoL	61,657	—
StarCraft HotS	28,582	—

Measure **predictive performance**,
compare against:

- Static

$$s_i(t) \equiv s_i$$

- Elo rating system

$$s_i \leftarrow s_i + \lambda \frac{\partial}{\partial s_i} \log p(i \succ j)$$

- Trueskill

$$s_i(t + \Delta t) = s_i(t) + \mathcal{N}[0, \sigma^2 \Delta t]$$

Predictive log-loss

Dataset	Constant	Elo	TrueSkill	kickscore	Covariance function
ATP tennis	0.581	0.563	0.563	0.552	Affine + Wiener
NBA basketball	0.692	0.634	0.634	0.630	Constant + Matérn 1/2
World football	0.929	0.950	0.937	0.926	Constant + Matérn 1/2
ChessBase small	1.030	1.035	1.030	1.026	Constant + Wiener

kickscore **outperforms
baselines** on all datasets

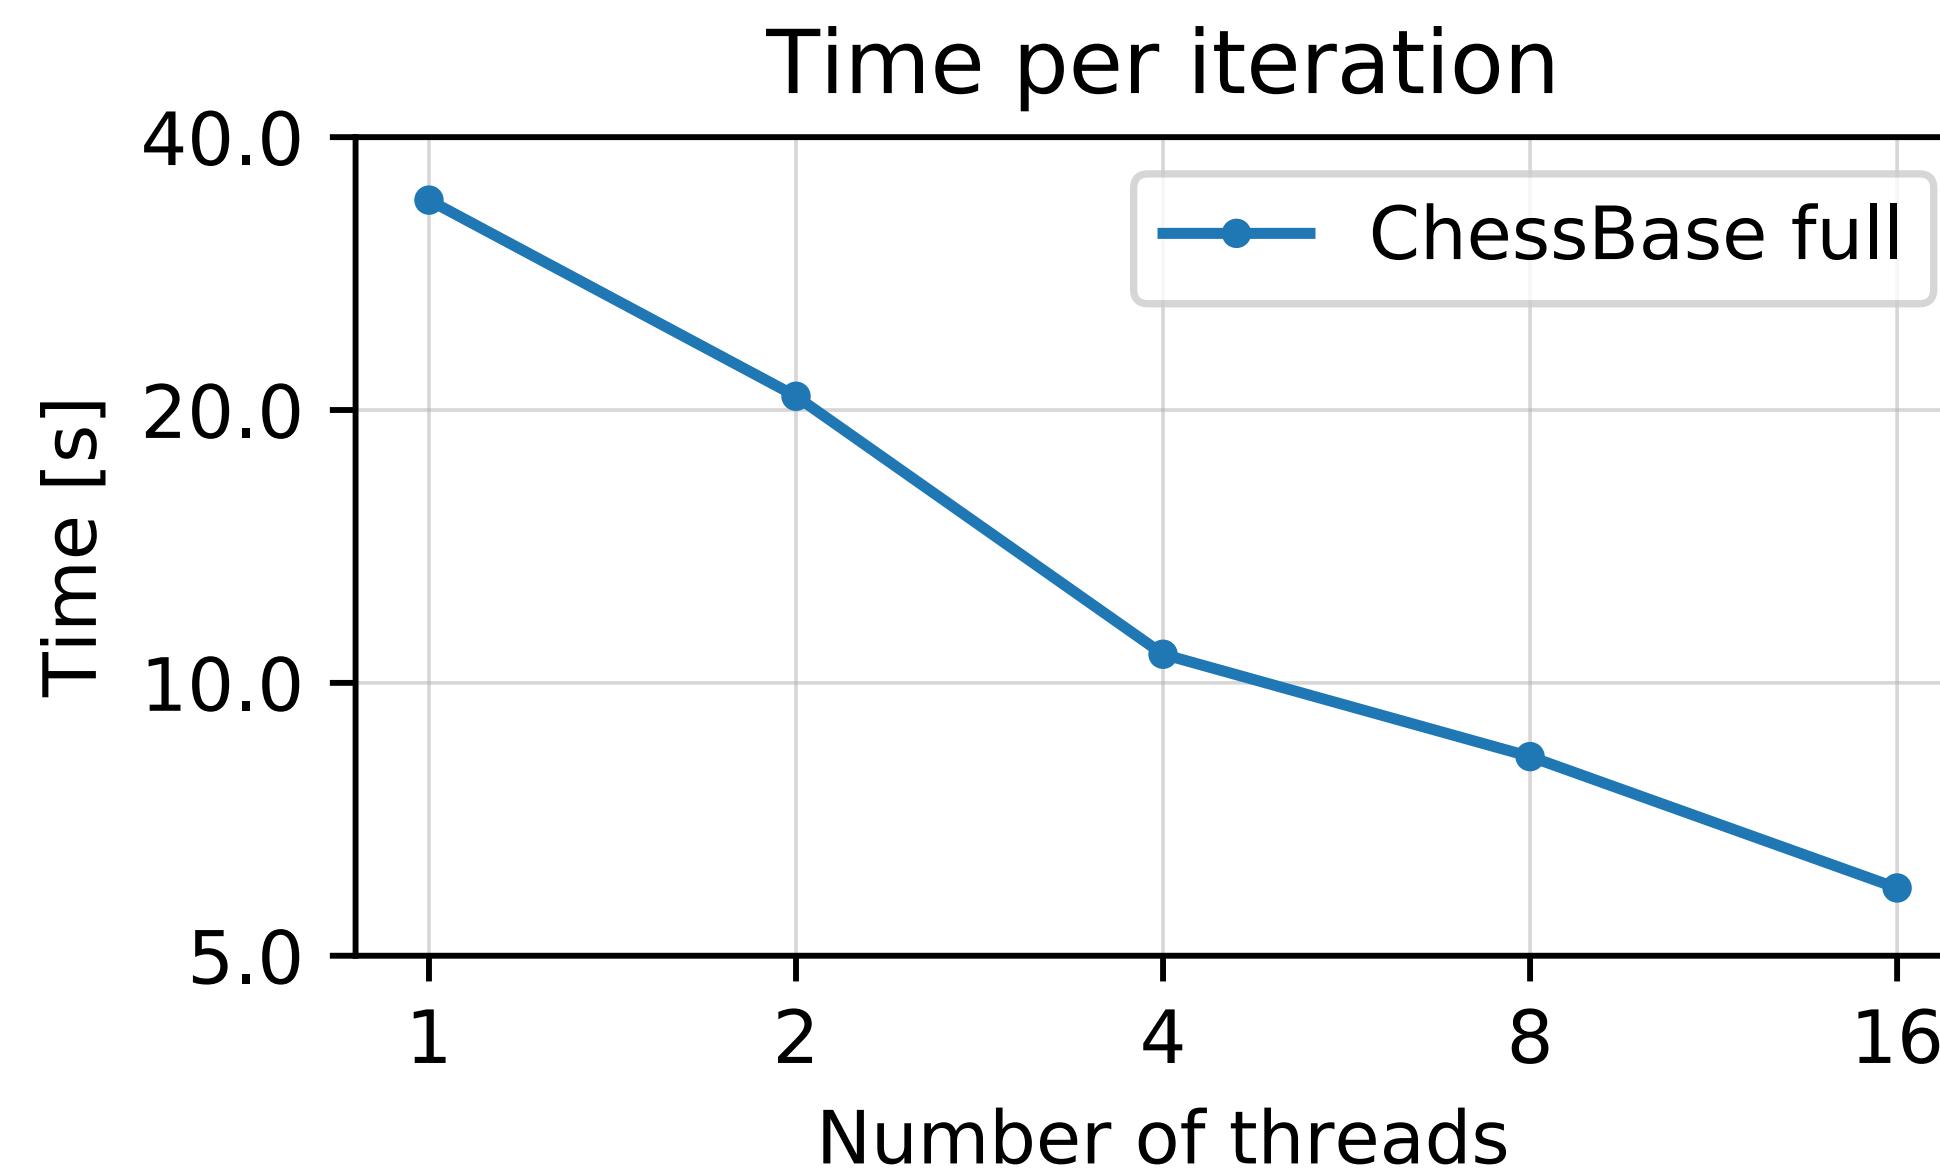
Best dynamics **vary**
across datasets

Experimental evaluation

Multi-threaded implementation in Go

→ **linear speed-up** with # threads

→ **scales to 7M** observations seamlessly



Impact of variational approximation

$$p(s_1, \dots, s_M \mid \mathcal{D}) \propto \prod_{i=1}^M p(s_i) \prod_{(i,j,t) \in \mathcal{D}} p(i \succ j \mid t)$$

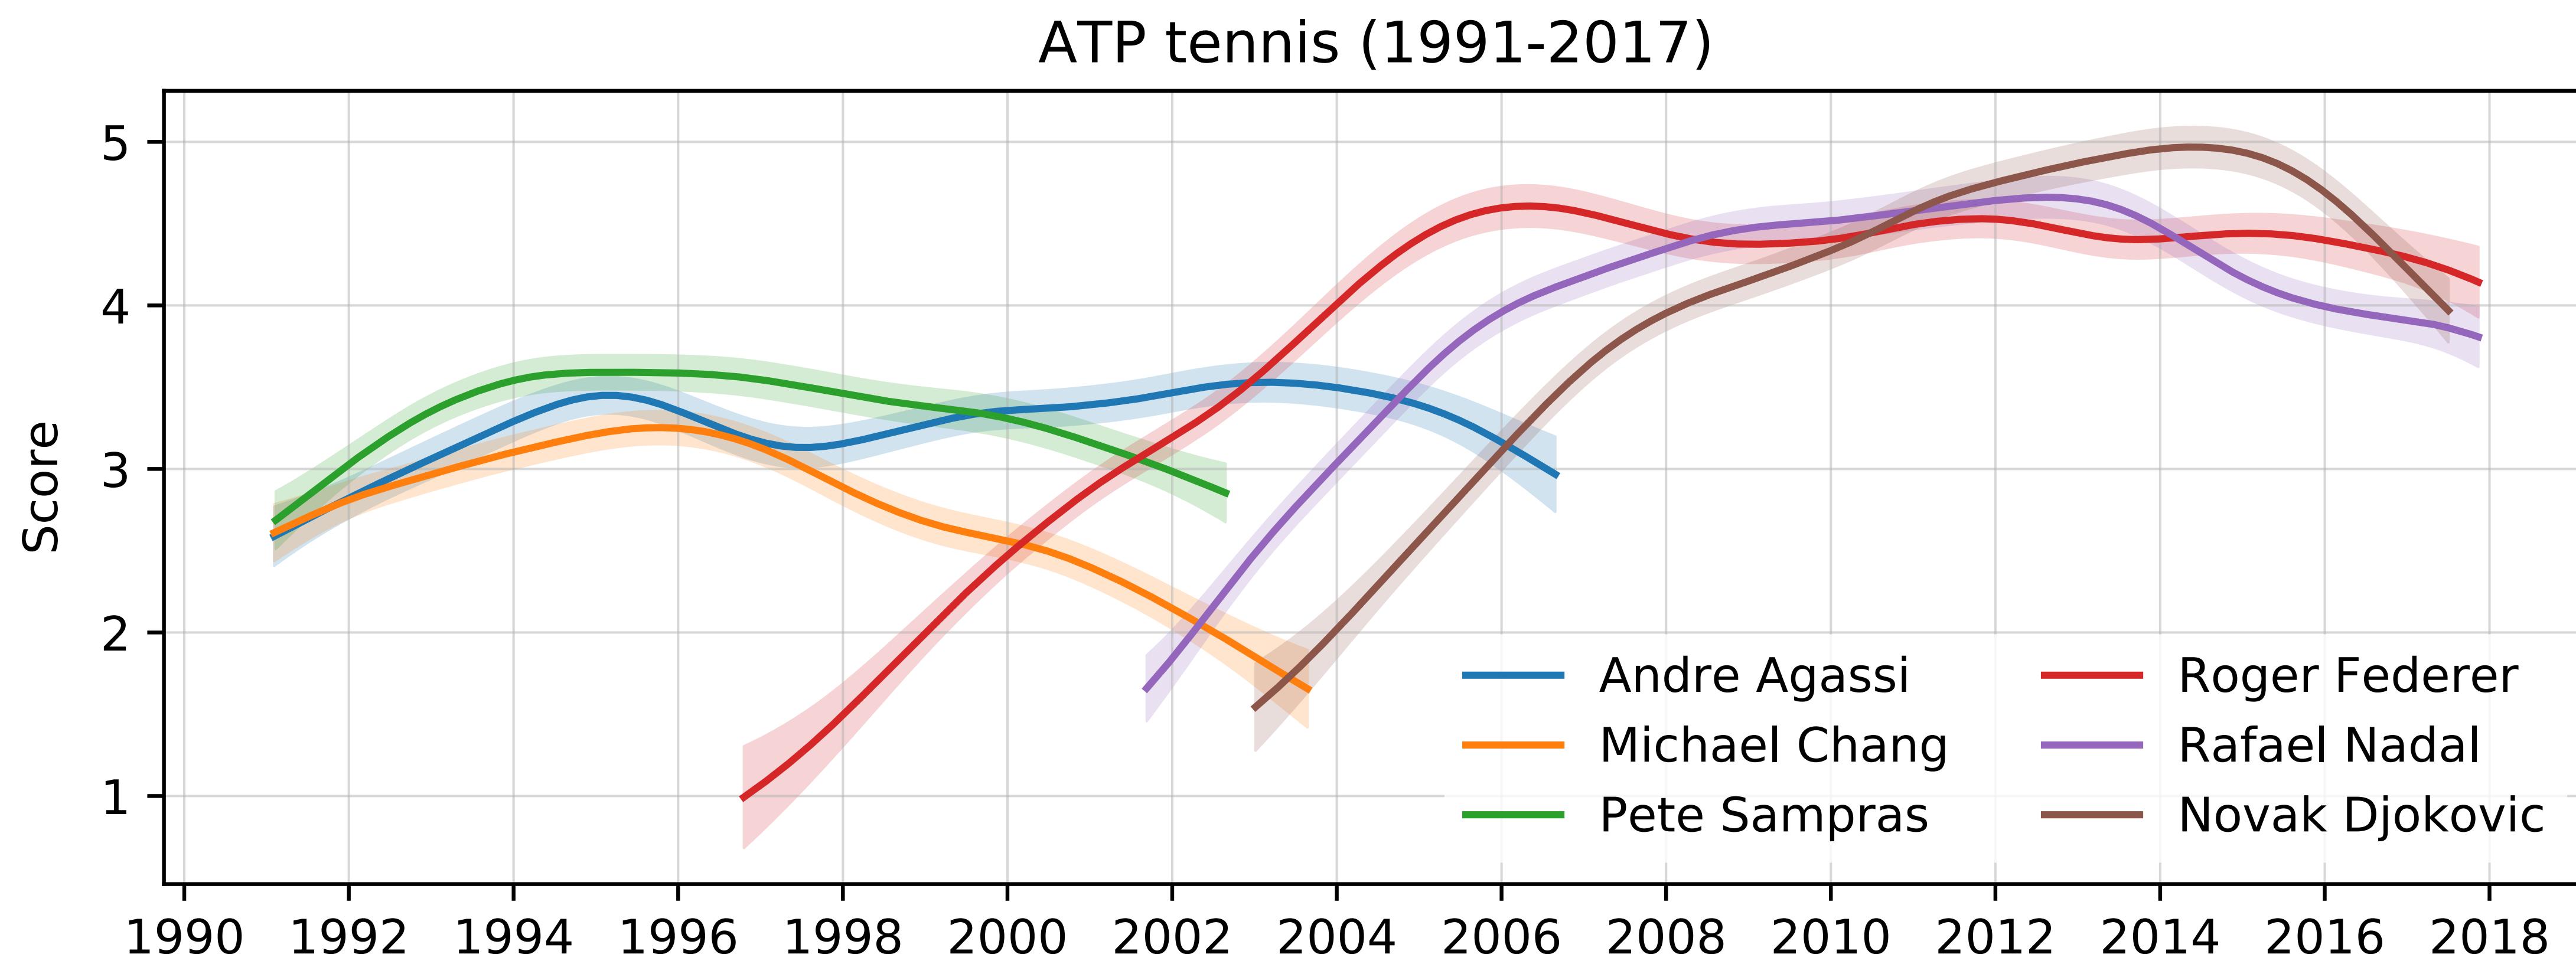
$$\approx q(s_1, \dots, s_M) \doteq \prod_i \mathcal{N}(s_i \mid \mu_i, \Sigma_i)$$

predictive accuracy is **preserved**

Extensions of the model

- handle **intransitive data**
- accommodate other **observations** **likelihoods** (Poisson, Gaussian, ...)

ATP tennis



Online resources

The screenshot shows a GitHub repository page for 'kickscore'. The title bar indicates the repository is 'lucasmaystre/kickscore: Pairwise skill rating system'. The page content is the 'README.rst' file. It features a large heading 'kickscore', a build status badge (build passing), and a codecov coverage badge (86%). A text block states: 'kickscore is the dynamic skill rating system powering [Kickoff.ai](#)'. Below this, another text block explains: 'In short, kickscore can be used to understand & visualize the skill of players (or teams) competing in pairwise matches, and to predict outcomes of future matches. It extends the [Elo rating system](#) and [TrueSkill](#)'. At the bottom, there is a line chart titled 'Evolution of the skill of NBA teams (1946-2019)' showing the skill levels of three teams: LAL (blue), CHI (orange), and BOS (green) over time.

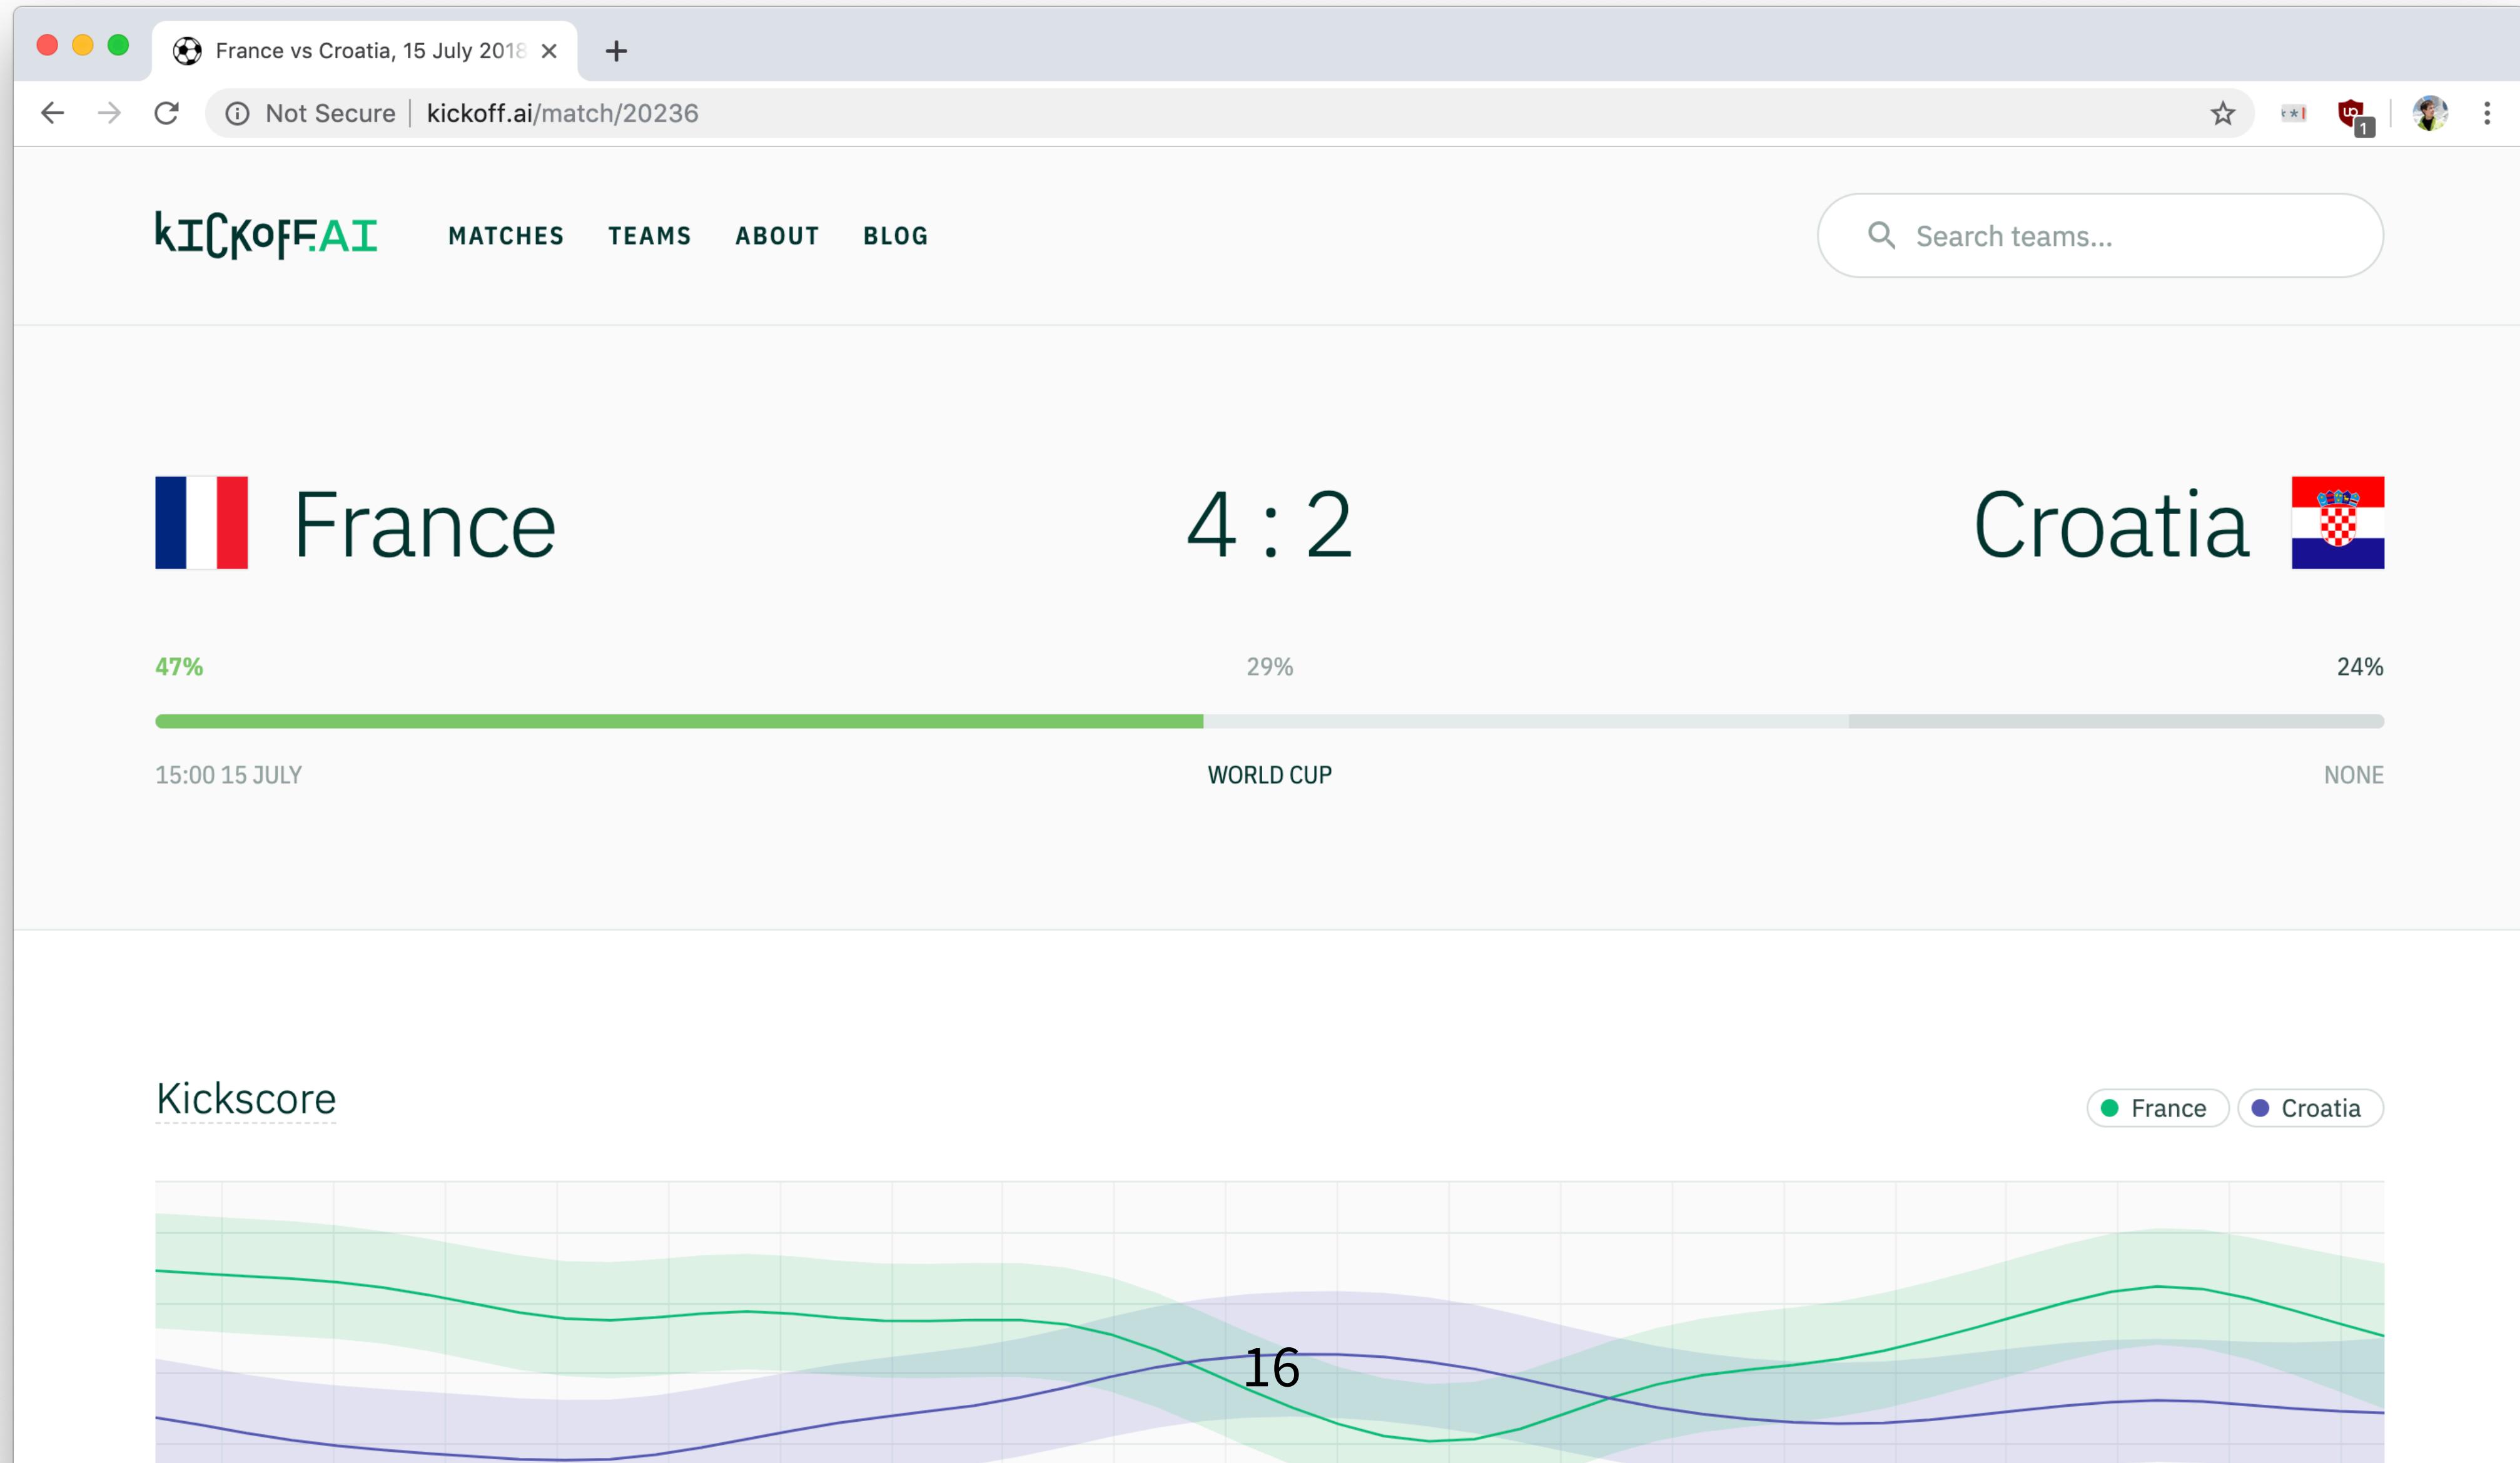
Evolution of the skill of NBA teams (1946-2019)

LAL
CHI
BOS

1953 1963 1973 1983 1993 2003 2013

Online resources

<https://kickoff.ai>



Conclusion

Key contributions:

- Pairwise comparison model with **flexible time-dynamics**
- **Linear-time** inference algorithm

Try it out:

```
pip install kickscore
```

