Repeated decoding optimises neural representations and is a SIMPL approach to finding latent variables in spiking data

Tom George[1], Pierre Glaser[2], Kimberly Stachenfeld[3,4], Caswell Barry[5], Claudia Clopath[1,6]

[1]Sainsbury Wellcome Centre, UCL, London; *tom.george.20@ucl.ac.uk* [2]Gatsby, UCL, London [3]Columbia University [4]Google DeepMind [5]CDB, UCL [6]Imperial College London

1. ONE PROBLEM, TWO APPROACHES: A DICHOTOMY FOR NEURAL DATA ANALYSIS

Neuroscientists seek effective* methods to explain high-dimensional neural data:

1. **Tuning curves** fit spikes to behaviour** (e.g. place cells). This is easy and can work well, but misses information (e.g. if the animal isn't thinking of its current position***)
2. **Latent Variable Models (LVMs)** don't limit you to behaviour but are often hard to optimise, finnicky to tune, difficult to interpret or expensive to scale.

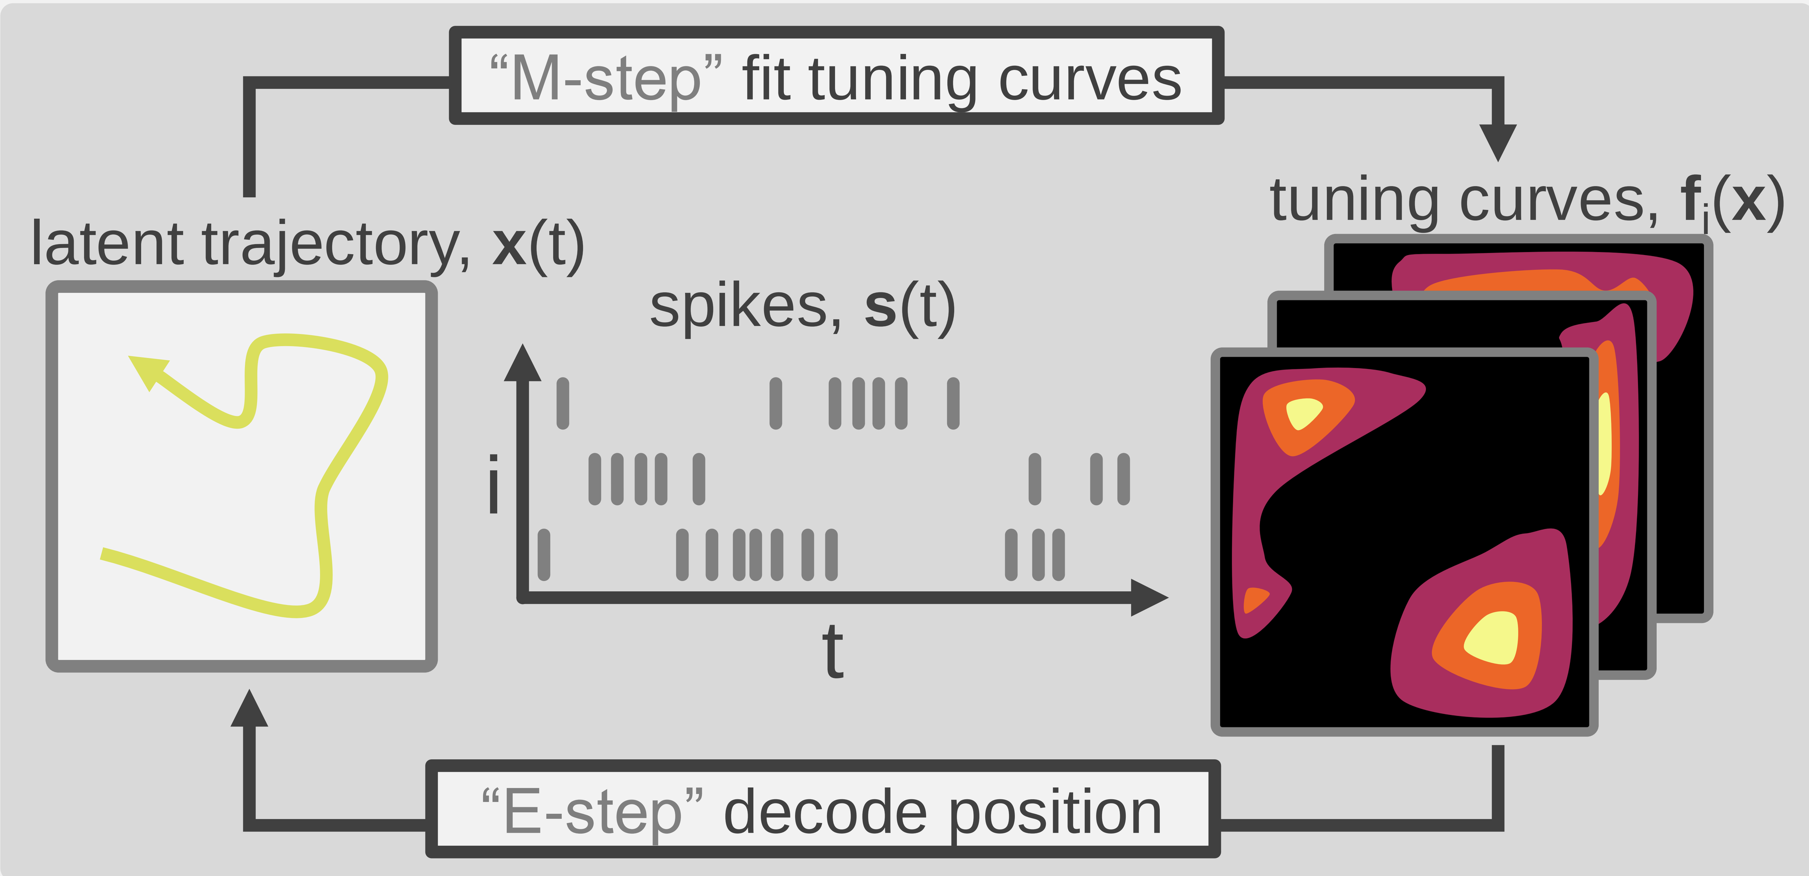

*fast, scalable, interpretable, identifiable and performant. **or stimulus, basically any available low-dimensional variable.
***replay, uncertainty, planning, offline states, sleep, consolidation etc. are all common examples

SIMPL is a hybrid technique: Inspired by expectation-maximisation[1] it fits tuning curves to behaviour (the "initial" latent estimate) and optimises them by iteratively decoding the latent and refitting the curves. This is:

- *Theory-backed:* decoding = an "E-step" and fitting = an "M-step"
- *Straightforward*: Uses common fitting and decoding techniques.
- *Cheap/Scalable*: Large datasets—200 cells, 1 hour, 10^6 spikes—in 1 CPU-min.
- *Interpretable*: Latent can be compared directly to behaviour, in the same units.

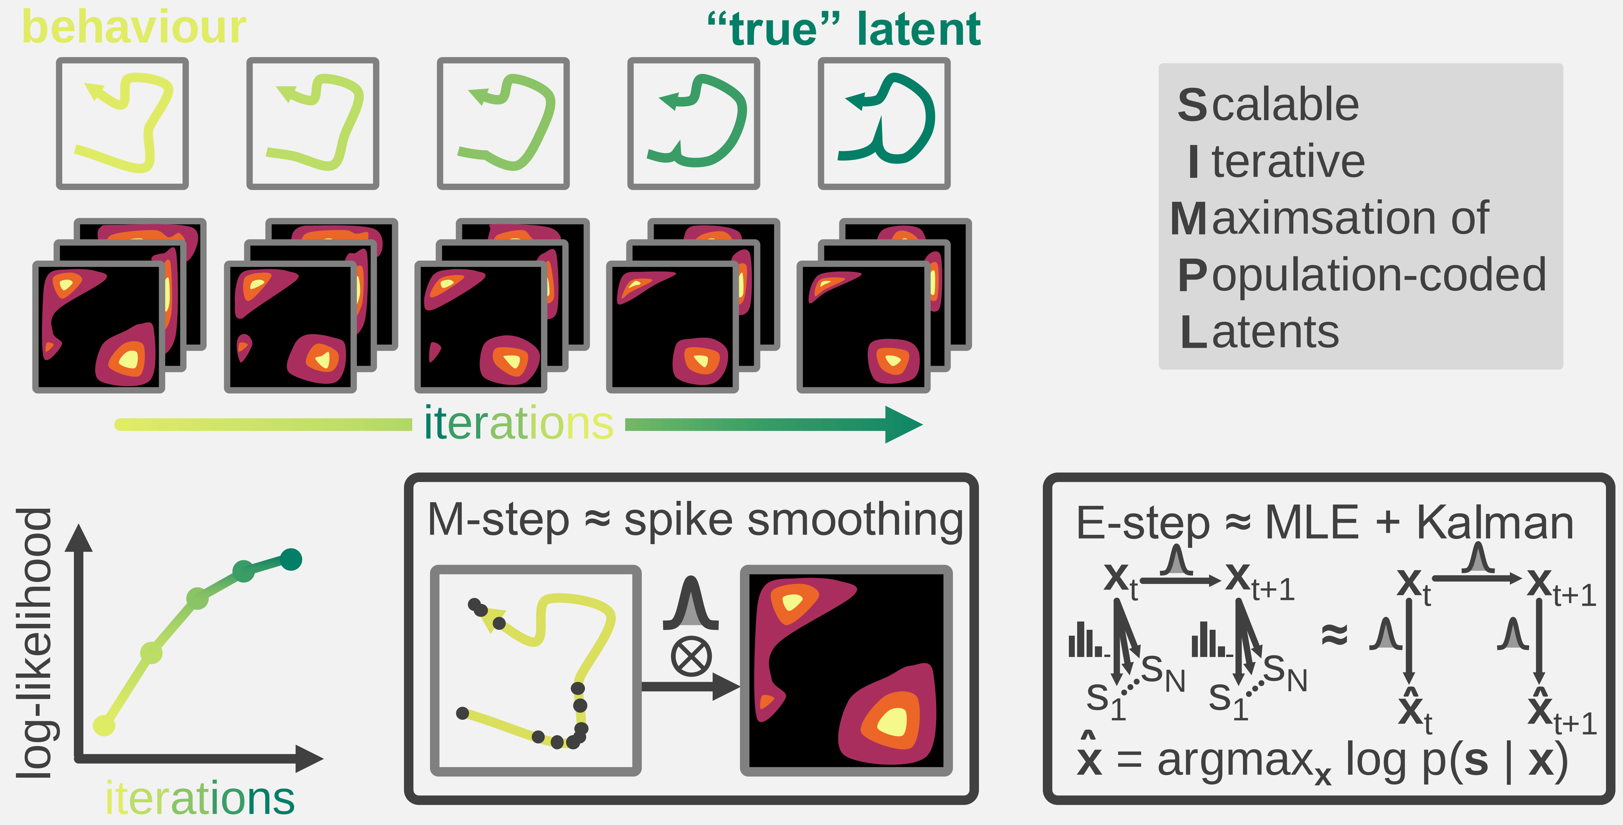

Scalable
I terative
Maximisation of
Population-coded
Latents

2.1 SYNTHETIC GRIDS (DATASET 1)

We simulate[2] grid cells for an agent *uncertain* of its own position

where the animal is
≉
where the animal thinks it is

This discrepancy, δ, blurs grid structure. Can SIMPL recover it?

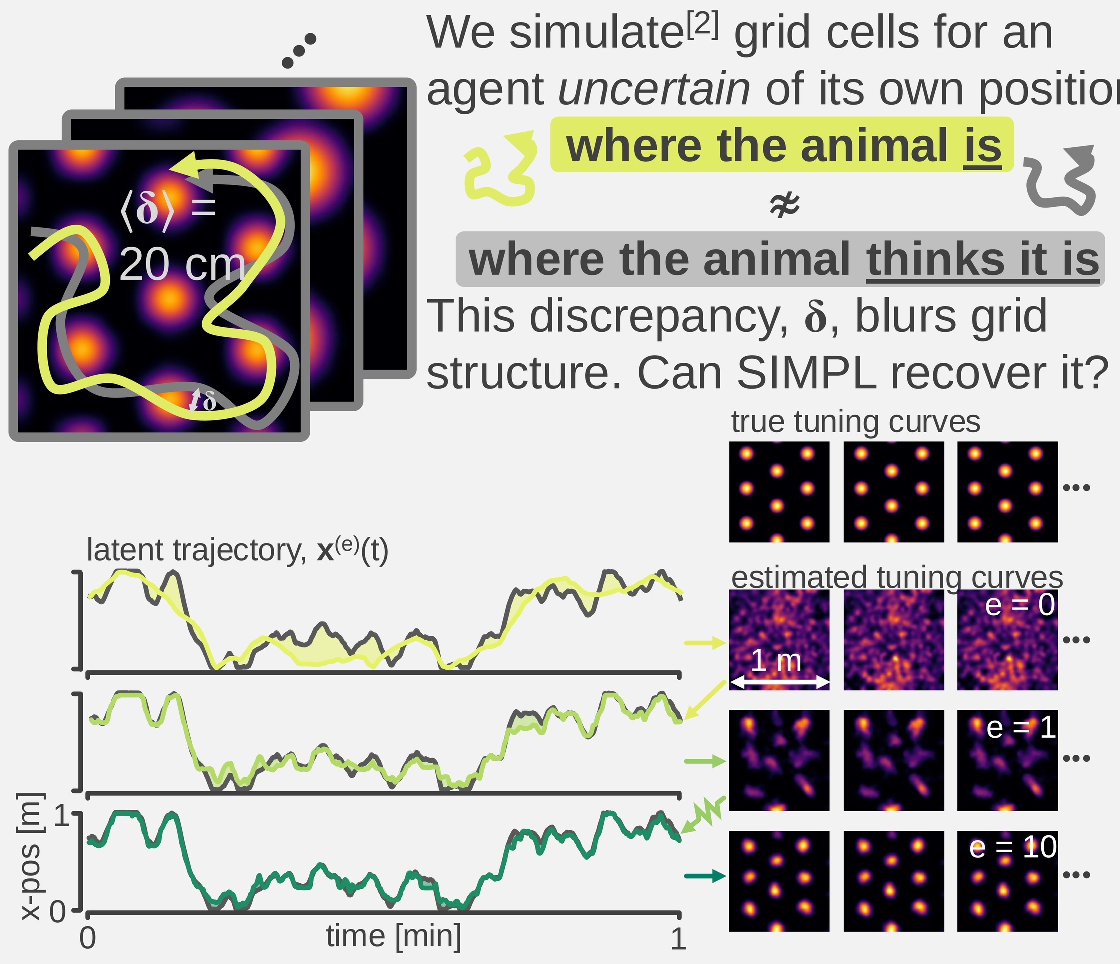

2.2 HIPPOCAMPUS (DATASET 2)

We ran SIMPL on a real place cell dataset[3].

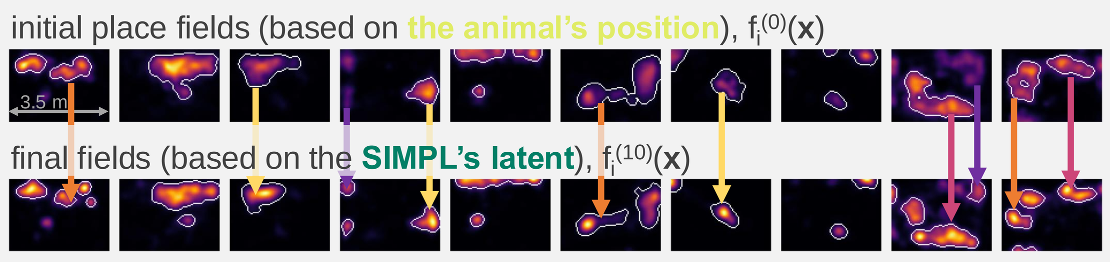

Optimised fields were similar to, but accounted for *more* variance than, those based on the animal's position. They had 19% more fields that were 25% smaller, fired 45% faster and were less noisy.

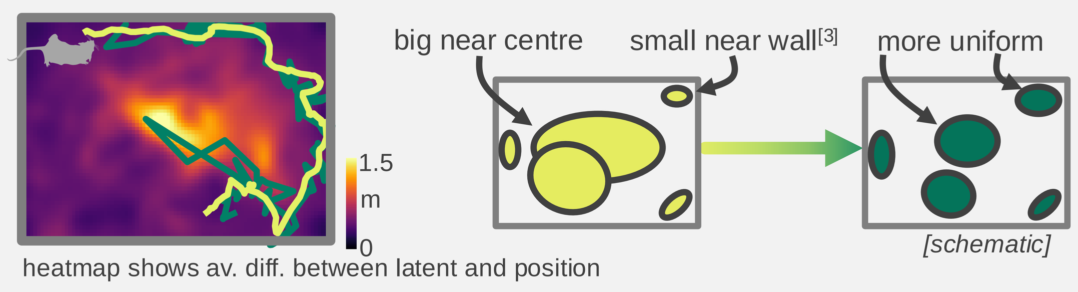

SIMPL's latent tracked position closely but exposed biased discrepancies between latent and position.

2.3 MOTOR CONTROL (DATASET 3)

Somatosensory spikes for a motor task[4].
Run SIMPL with 4D latent: $\mathbf{x}^{(0)}=[x,y,v_x,v_y]$
Latent reveals behaviour-correlated internal dynamics. "Hand fields" tighten substantially.

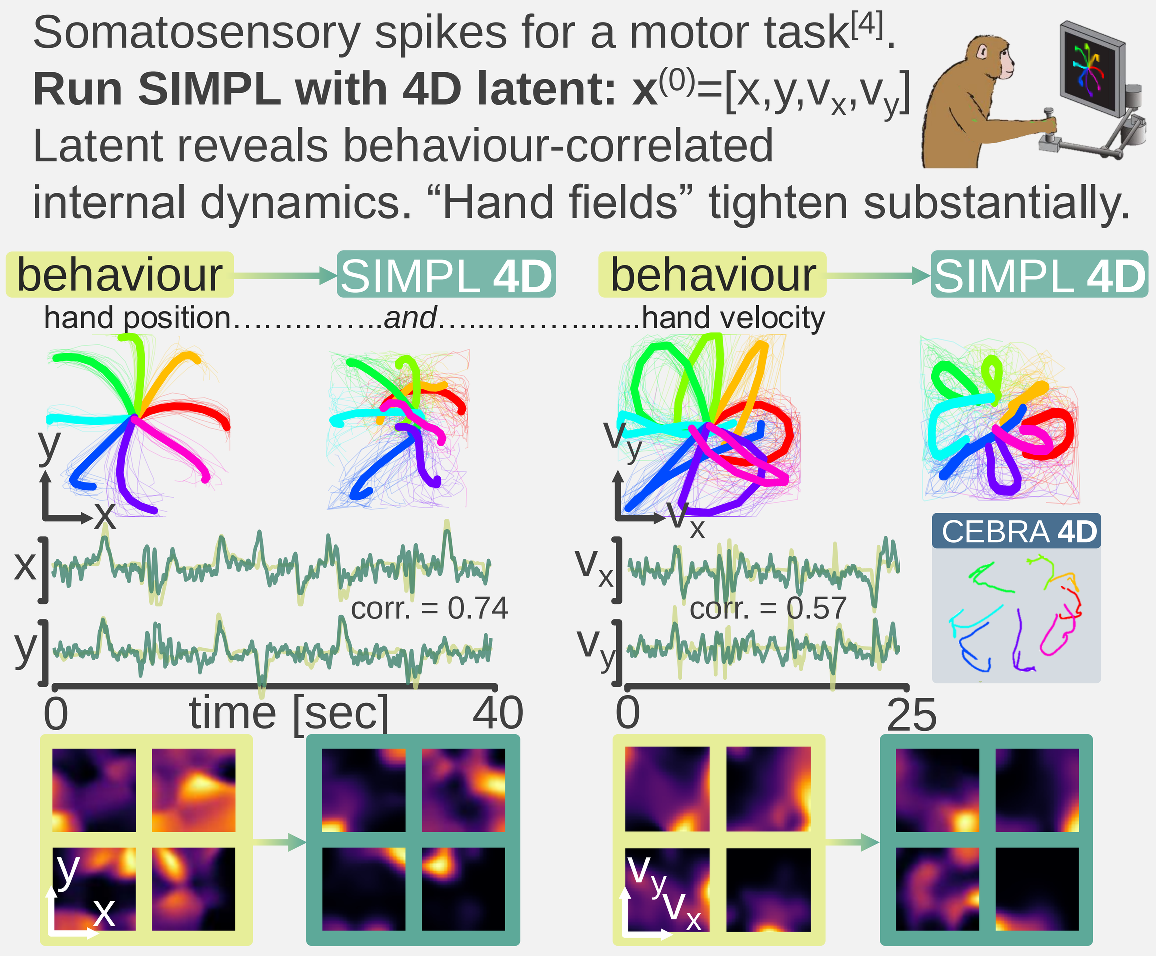

3. WHY INITIALISE AT BEHAVIOUR?

Identifiablity: equivalent *isomorphic* solutions can explain data equally well but manifest very differently.

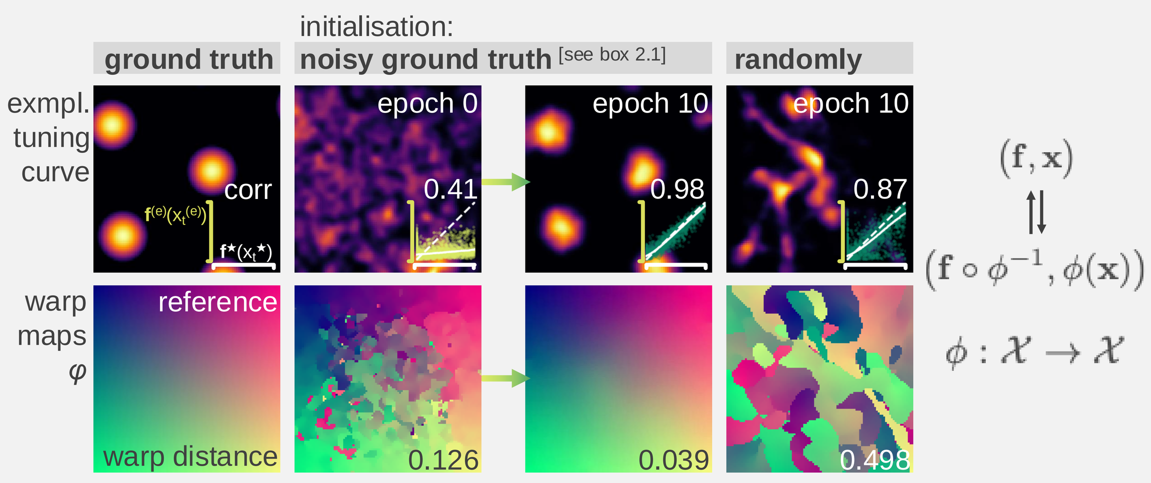

Behavioural initialisation gives the same solution which is unwarped wrt. the ground truth (φ=I), every time.

Convergence: For many neural codes behaviour is a close approximation to the latent. It's foolish *not* to start there.

4. COMPARISONS AND BENCHMARKS

On the synthetic dataset SIMPL finds latents *faster* and *better* than other NN or GP techniques.

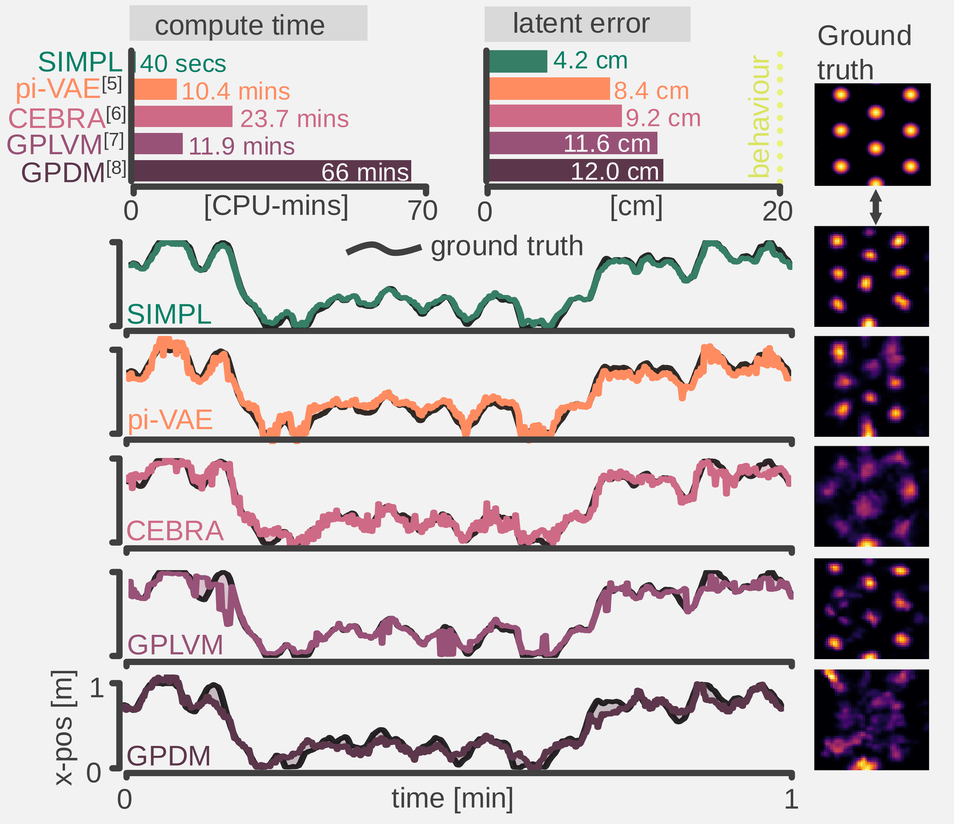

SIMPL satisfies a unique set of desiderata:

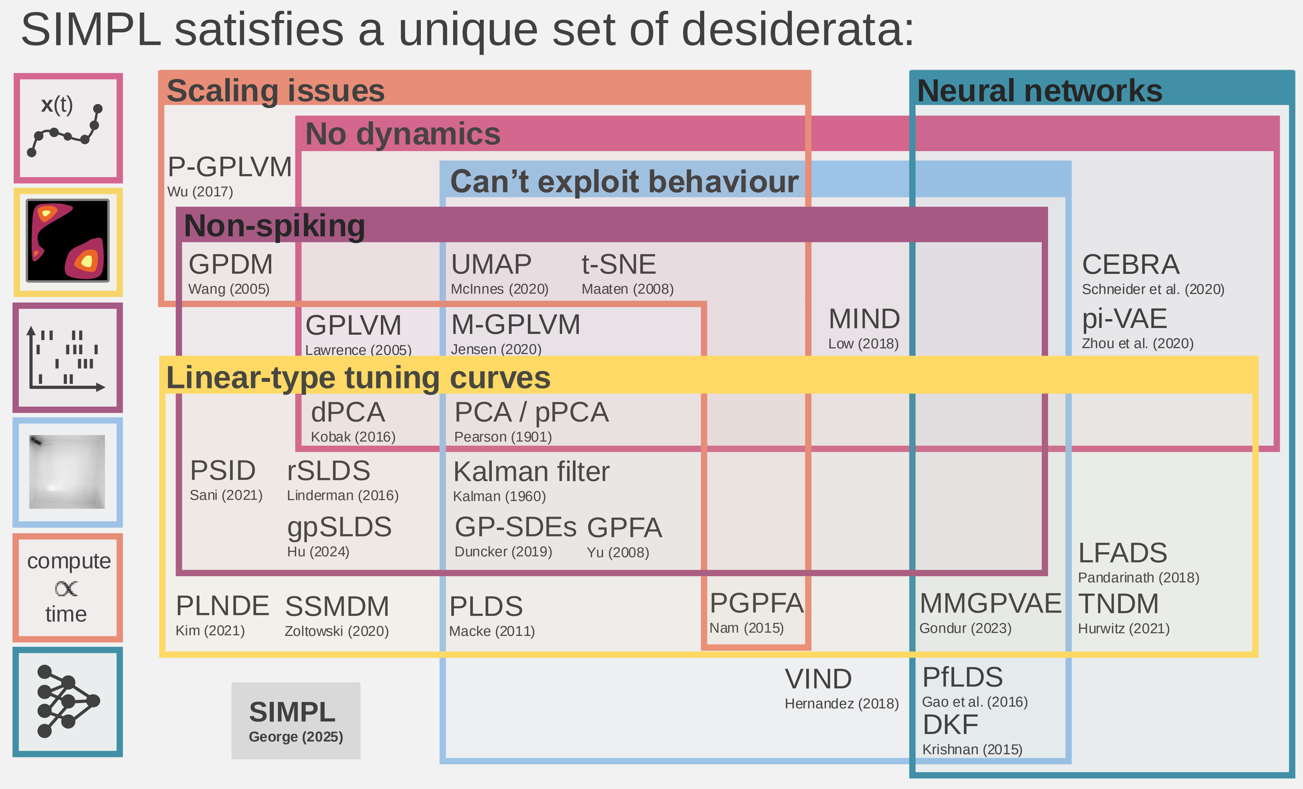

Furthermore, existing methods are also often **complicated**, **lack useable code** or **require GPUs** creating a barrier to their adoption. Pragmatism matters.

5. CONCLUSIONS

1. **Tuning curves are fast and interpretable** but misleading whenever behaviour ≉ latent.
2. **LVMs correct for this** at the cost of interpretability, identifiability and complexity.
3. **SIMPL blends these approaches.** It is simple, theoretically justified, performant.
4. **SIMPL works across domains** including spatial and motor datasets.
5. **SIMPL outperforms other techniques** at a fraction of the cost.
6. **Behavioural initialisation is a neat trick** for identifiability and improved optimization.

Github:

ICLR paper:

References: [1] Dempster et al. (1977), Maximum likelihood from incomplete data via the EM algorithm. **[2] George et al. (2024),** RatInABox, a toolkit for modelling locomotion and neuronal activity in continuous environments. **[3] Tanni et al. (2022),** State transitions in the statistically stable place cell population correspond to rate of perceptual change. **[4] Chowdhury et al. (2020),** Area 2 of primary somatosensory cortex encodes kinematics of the whole arm. **[5] Zhou et al. (2020),** Learning identifiable and interpretable latent models of high-dimensional neural activity using pi-vae. **[6] Schneider and Lee et al. (2023),** Learnable latent embeddings for joint behavioural and neural analysis. **[7] Lawrence (2003),** Gaussian process latent variable models for visualisation of high dimensional data. **[8] Wang et al. (2005),** Gaussian process dynamical models.

CLOPATH LAB Computational Neuroscience · Sainsbury Wellcome Centre · barry lab · Google DeepMind · GATSBY · Columbia University