Beyond the ATE: Interpretable Modelling of Treatment Effects over Dose and Time

Julianna Piskorz, Krzysztof Kacprzyk, Harry Amad, Mihaela van der Schaar

UNIVERSITY OF CAMBRIDGE

van_der_Schaar \ LAB

Motivation

RQ: How does the treatment effect τ evolve over time, depending on the dose, a?

💡How long does it typically take for the treatment effect to exceed a clinically relevant *threshold* α?

💡 When, on average, does the treatment effect reach its *peak*, and how does this vary with the dose?

💡 What is the lowest dose ensuring that the effect has a *sustained benefit*?

To answer these question we want to model:

$$\tau_t(a) = \mathbb{E}[Y_t(a) - Y_t(0)]$$

Desiderata for a useful model: *interpretable, verifiable, editable*

SemanticODE for ATE Estimation

Treatment effects are not directly observable!

Dataset: $\mathscr{D} = \{\mathbf{X}_i, A_i, (T_{i,j}, Y_{i,j})_{j=1}^{m_i}\}_{i=1}^{n}$

To use SemanticODE, we construct *surrogate treatment effect:*

$$\tilde{\tau}_t(a) \approx Y_t(a) - Y_t(0)$$

- We train a **baseline trajectory model,** using untreated patients with $A = 0$, and utilising patient characteristics: $Y_t(0) \approx \hat{\varphi}_0(\mathbf{X}, T)$
- For all other patients in the dataset we then construct the surrogate treatment effect: $\tilde{\tau}_{i,j} = Y_{i,j} - \hat{\varphi}_0(\mathbf{X}_i, T_{i,j})$

Direct Semantic Modelling with SemanticODE

Composition map: describes the shape of the trajectory as a function of the dose, a.

Property maps: describe how the key properties of the identified composition (maxima, derivatives, asymptotes) depend on the dose a.

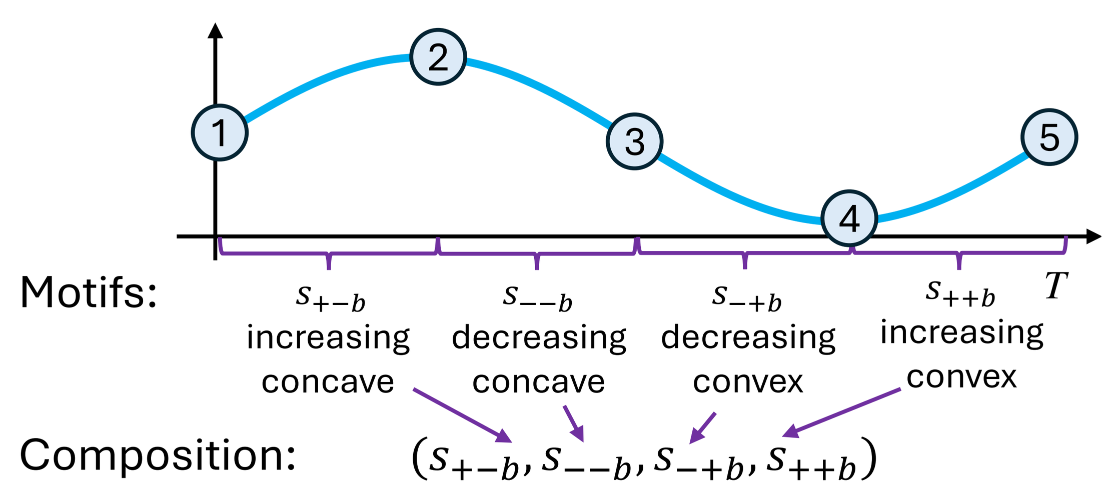

🚀 **This decomposition allows to easily impose inductive biases and edit the learned solutions to match the domain knowledge!**

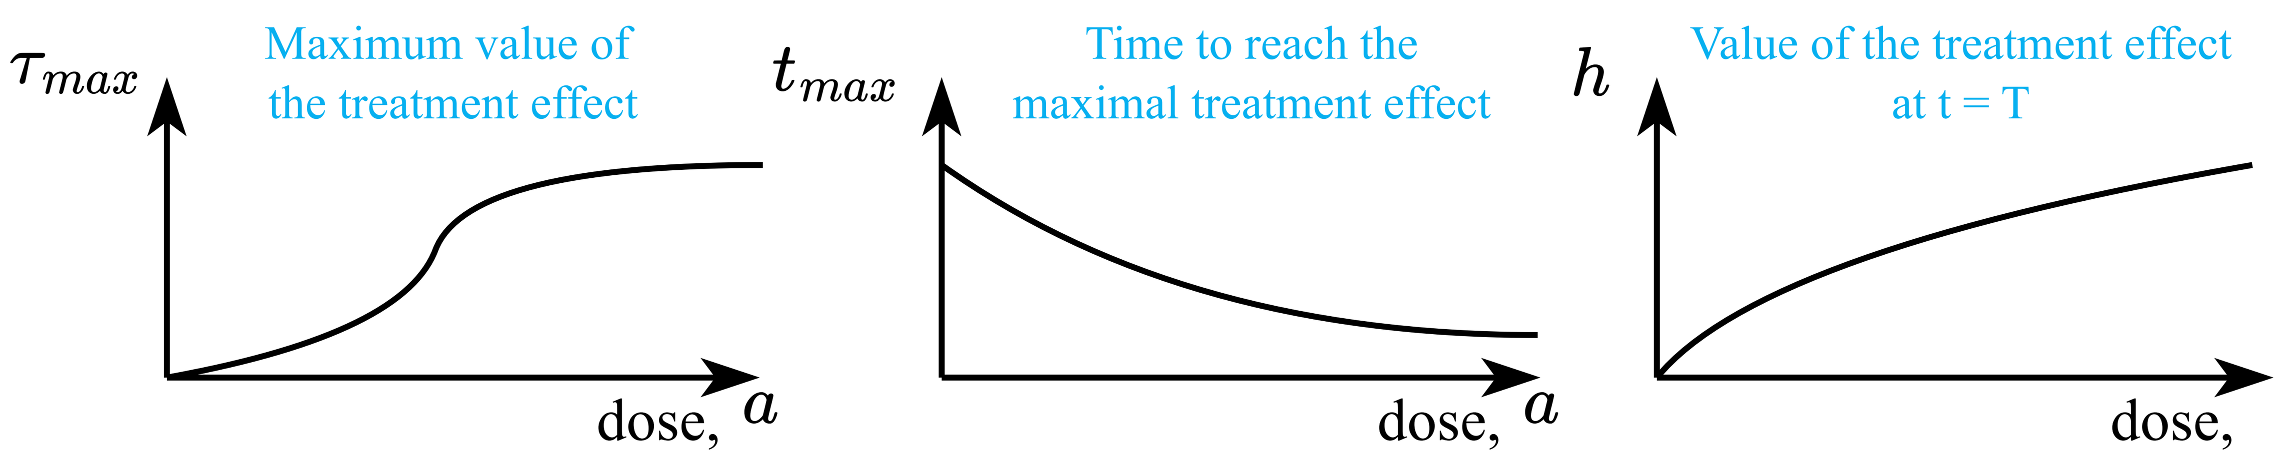

Numerical Experiments on Synthetic and Semi-Synthetic Datasets

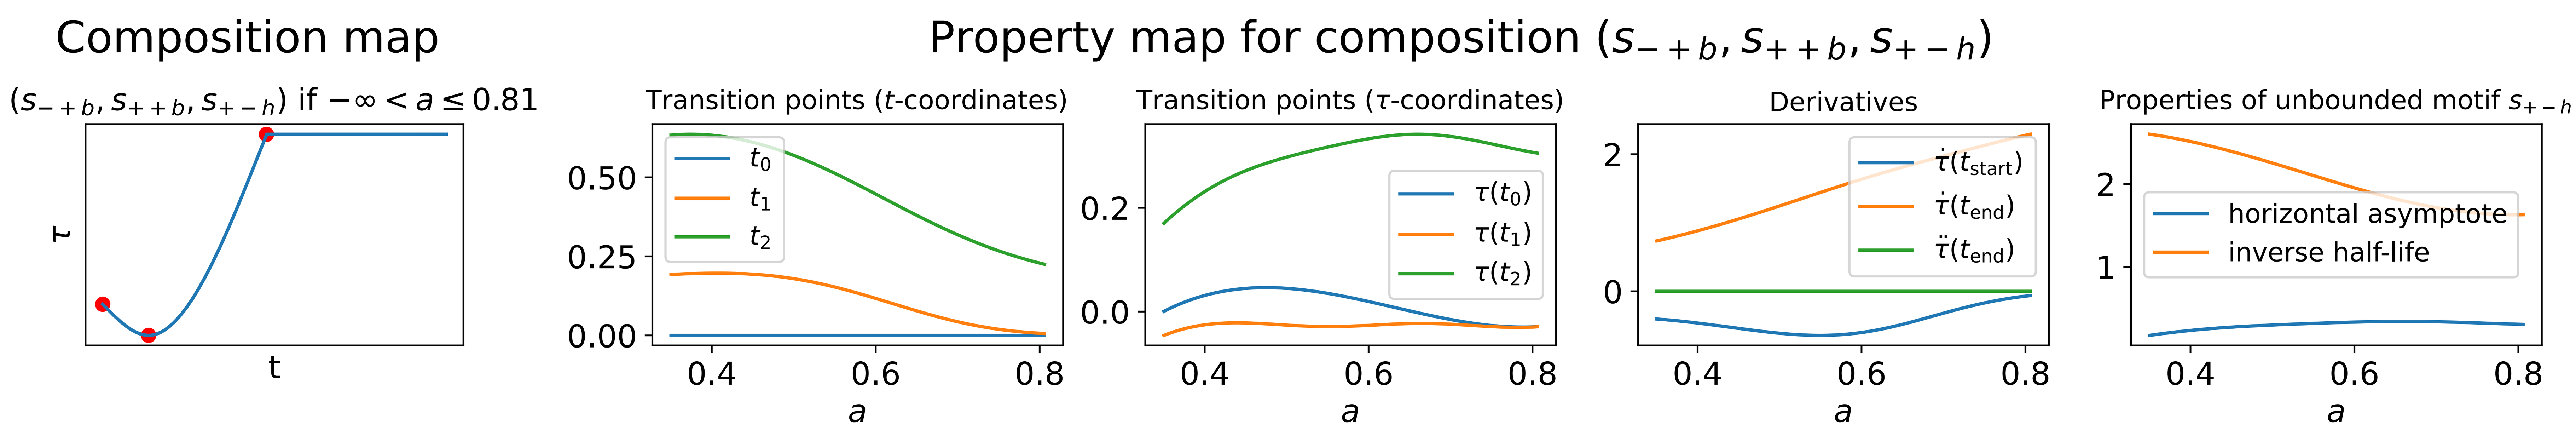

📊 **SemanticATE outperforms baselines!** By allowing to impose the necessary inductive biases, SemanticATE achieves incomparable performance both in-distribution ($t \in [0, T]$) and out-of-distribution ($t > T$).

Method	PK-random		PK-real		IHDP-based	
	In-domain	Out-domain	In-domain	Out-domain	In-domain	Out-domain
SINDy	2.81 ± 1.04	3.88 ± 5.62	$> 10^6$	$> 10^3$	0.46 ± 0.12	2.32 ± 0.54
WSINDy	0.53 ± 0.21	7.96 ± 3.64	1.31 ± 0.50	5.21 ± 1.27	0.54 ± 0.04	25.16 ± 31.20
NeuralODE	0.50 ± 0.03	0.17 ± 0.08	0.54 ± 0.12	0.21 ± 0.13	0.39 ± 0.19	1.37 ± 1.15
XGBoost	0.40 ± 0.15	0.11 ± 0.04	0.43 ± 0.35	0.08 ± 0.02	0.36 ± 0.14	0.20 ± 0.09
PolyReg	**0.27 ± 0.06**	211.05 ± 82.49	**0.34 ± 0.10**	197.25 ± 69.20	0.05 ± 0.01	4.31 ± 1.83
SemanticATE	0.41 ± 0.28	**0.04 ± 0.01**	0.37 ± 0.29	**0.02 ± 0.01**	**0.04 ± 0.02**	**0.02 ± 0.02**

📊 **SemanticATE allows to easily impose inductive biases and modify obtained solutions, without sacrificing accuracy!** The practitioner can impose known inductive biases (IB) by modifying the set of possible compositions, or directly changing the property maps if they don't agree with domain knowledge

Method	PK-random	PK-real	IHDP-based
Base	0.25 ± 0.09	0.22 ± 0.10	0.04 ± 0.02
Base + IB	0.25 ± 0.10	0.23 ± 0.09	0.06 ± 0.01
Base + IB + edits	0.23 ± 0.08	0.23 ± 0.09	0.08 ± 0.03

Direct semantic modelling

This paper

Enjoyed our work? Let's connect!

jp2048@cam.ac.uk, kk751@cam.ac.uk

in/julianna-piskorz, in/krzysztof-kacprzyk