

NEURAL ORDINARY DIFFERENTIAL EQUATIONS

AI tea talk, April 28, 2023

Jakub Ševčík

ORDINARY DIFFERENTIAL EQUATIONS

- Vector valued \mathbf{y} changes in time respecting

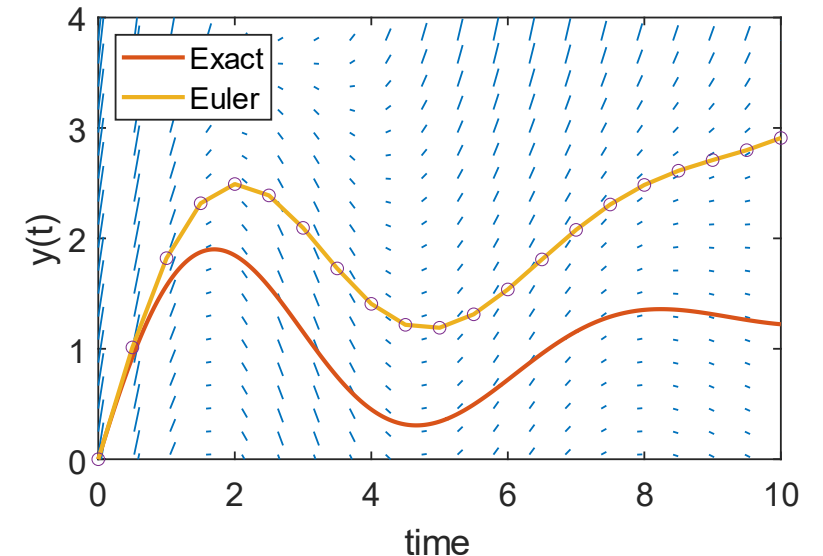
$$\frac{d\mathbf{y}}{dt}(t) = f_{\theta}(t, \mathbf{y}(t))$$

- Initial-value problem: given $\mathbf{y}(t_0)$, find:

$$\mathbf{y}(t_1) = \mathbf{y}(t_0) + \int_{t_0}^{t_1} f_{\theta}(t, \mathbf{y}(t)) dt$$

- Euler scheme

$$\mathbf{y}(t_0 + \Delta t) = \mathbf{y}(t_0) + \Delta t f(\mathbf{y}(t_0), t_0)$$



EULER AS AN ANALOGY TO RESNET

- Transformation of a hidden state in ResNets

$$\mathbf{y}_{k+1} = \mathbf{y}_k + f_{\theta_k}(k, \mathbf{y}_k)$$

- The difference $\mathbf{y}_{k+1} - \mathbf{y}_k$ can be interpreted as a discretization of the derivative and letting $\Delta t \rightarrow 0$


$$\lim_{\Delta t \rightarrow 0} \frac{\mathbf{y}_{k+1} - \mathbf{y}_k}{\Delta t} = \frac{d\mathbf{y}(t)}{dt} = f_{\theta}(t, \mathbf{y}(t))$$

- The Neural ODE can be seen as a continuous variant of ResNet

RESNET vs NEURAL ODE


```
def f(y, t,  $\theta$ ):  
    return nnet(y,  $\theta[t]$ )
```

```
def resnet(z):  
    for t in [1:T]:  
        y = y + f(y, t,  $\theta$ )  
    return y
```



```
def f(y, t,  $\theta$ ):  
    return nnet([y, t],  $\theta$ )
```

```
def neuralODE(y):  
    return ODEsolve(f, y0, t0, t1,  $\theta$ )
```



TRAINING OF NEURAL ODE

- Consider a scalar-valued loss function L , minimize

$$L(\mathbf{y}(t_1)) = L\left(\mathbf{y}(t_0) + \int_{t_0}^{t_1} f_{\theta}(\mathbf{y}(t), t) dt\right) = L(\text{ODESolve}(f_{\theta}, \mathbf{y}(t_0), t_0, t_1, \theta))$$

- How to get $\frac{\partial L(\mathbf{y}(t_1))}{\partial \theta}$?

- Discretise-then-optimise

- backpropagate through the internal operations of the differential equation solver

- Optimise-then-discretise (continuous adjoint method)

- produce a backwards-in-time differential equation and solve it

- Reversible ODE solver


[1] Chen, R. T., Rubanova, Y., Bettencourt, J., & Duvenaud, D. K. (2018). Neural ordinary differential equations. *Advances in neural information processing systems*, 31.

TRAINING OF NEURAL ODE

- Optimise-then-discretise (continuous adjoint method)
 - produce a backwards-in-time differential equation and solve it
 - Define function $\mathbf{a}(t) = \partial L / \partial \mathbf{y}(t)$, then

$$\frac{d\mathbf{a}(t)}{dt} = -\mathbf{a}(t)^T \frac{\partial f_{\theta}(t, \mathbf{y}(t))}{\partial \mathbf{y}}$$

$$\frac{\partial L}{\partial \theta} = -\int_{t_1}^{t_0} \mathbf{a}(t)^T \frac{\partial f_{\theta}(t, \mathbf{y}(t))}{\partial \theta} dt$$



- Solve ODE set with initial point $\left[\mathbf{y}(t_1), \frac{\partial L}{\partial \mathbf{y}(t_1)}, \mathbf{0}_{|\theta|} \right]^T$ and dynamics

$$\begin{bmatrix} f_{\theta}(t, \mathbf{y}(t)) \\ -\mathbf{a}(t)^T \frac{\partial f_{\theta}(t, \mathbf{y}(t))}{\partial \mathbf{y}} \\ -\mathbf{a}(t)^T \frac{\partial f_{\theta}(t, \mathbf{y}(t))}{\partial \theta} \end{bmatrix}$$

to get $\left[\mathbf{y}(t_0), \frac{\partial L}{\partial \mathbf{y}(t_0)}, \frac{\partial L}{\partial \theta} \right]^T$

[1] Chen, R. T., Rubanova, Y., Bettencourt, J., & Duvenaud, D. K. (2018). Neural ordinary differential equations. *Advances in neural information processing systems*, 31.

TRAINING OF NEURAL ODE

■ Discretise-then-optimize

- memory inefficient
- accurate and fast

■ ODE solver in an AD framework

■ Optimize-then-discretise


- memory efficient
- approximate and a little slow

■ ODE solver can be a black-box

[4] Kidger, P. (2022). On neural differential equations. *Doctoral thesis, University of Oxford. arXiv:2202.02435.*

LIMITATIONS OF NODE LEAD TO ANODE


- NeuralODE are not universal approximators
- Consider function $g: \mathbb{R} \rightarrow \mathbb{R}$ such that
 - $g(-1) = 1$
 - $g(1) = -1$



- Augmented NODE: augments the ODE space from \mathbb{R}^d to \mathbb{R}^{d+d_0}

$$\frac{d}{dt} \begin{bmatrix} \mathbf{y}(t) \\ \mathbf{h}(t) \end{bmatrix} = f_{\theta} \left(t, \begin{bmatrix} \mathbf{y}(t) \\ \mathbf{h}(t) \end{bmatrix} \right), \quad \begin{bmatrix} \mathbf{y}(0) \\ \mathbf{h}(0) \end{bmatrix} = \begin{bmatrix} \mathbf{x} \\ \mathbf{0} \end{bmatrix}$$

- Augmented NODEs are universal approximators



[2] Dupont, E., Doucet, A., & Teh, Y. W. (2019). Augmented neural odes. *Advances in neural information processing systems*, 32.

ANODES ARE UNIVERSAL APPROXIMATORS

$$x \xrightarrow{l_1} A_1 x + b_1 \xrightarrow{\text{ANODE}} A_2 y(T) + b_2 \xrightarrow{l_2} \text{output}$$

$\frac{dy}{dt} = \int_0^T f(t, y(t)) dt$

THEOREM

Fix $d, d_l, d_0 \in \mathbb{N}$ with $d_l \geq d + d_0$. For $f \in Lip(\mathbb{R} \times \mathbb{R}^{d_l}; \mathbb{R}^{d_l})$, $l_1 \in L_b(\mathbb{R}^d; \mathbb{R}^{d_l})$, $l_2 \in L_b(\mathbb{R}^{d_l}; \mathbb{R}^{d_0})$, let $\varphi_{f, l_1, l_2}: \mathbb{R}^d \rightarrow \mathbb{R}^{d_0}$ denotes the map $x \rightarrow z$ with

$$y(0) = l_1(x), \quad \frac{dy}{dt}(t) = f(t, y(t)), \quad \text{for } t \in [0, T], \quad z = l_2(y(T)).$$

Then

$$\left\{ \varphi_{f, l_1, l_2} \mid f \in Lip(\mathbb{R} \times \mathbb{R}^{d_l}; \mathbb{R}^{d_l}), l_1 \in L_b(\mathbb{R}^d; \mathbb{R}^{d_l}), l_2 \in L_b(\mathbb{R}^{d_l}; \mathbb{R}^{d_0}) \right\}$$

is a universal approximator for $C(\mathbb{R}^d; \mathbb{R}^{d_0})$.



[3] Zhang, H., Gao, X., Unterman, J., & Arodz, T. (2020). Approximation capabilities of neural ODEs and invertible residual networks. In *International Conference on Machine Learning* (pp. 11086-11095).

CASE STUDY: Neural ODE for Synchronous Drives

- Estimation of magnetic field of an IPMSM (Interior Permanent-Magnet Synchronous Motor)
- Synchronous machine electrical model in the rotating (d, q) - reference frame

$$\mathbf{u}_s = R\mathbf{i}_s + \omega_{el}J\boldsymbol{\psi}(\mathbf{i}_s, \mathbf{x}) + \frac{d}{dt}\boldsymbol{\psi}(\mathbf{i}_s, \mathbf{x})$$

- The aim is to get a good approximation of $\hat{\boldsymbol{\psi}} \approx \boldsymbol{\psi}(\mathbf{i}_s, \mathbf{x})$


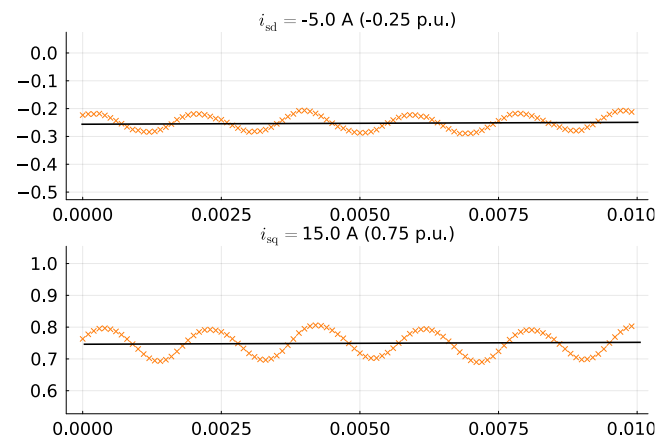



CASE STUDY: Neural ODE for Synchronous Drives

- Common approaches use look-up tables (LUT)
- It assumes steady regimes

$$\mathbf{u}_s = R\mathbf{i}_s + \omega_{el}J\boldsymbol{\psi}(\mathbf{i}_s, \mathbf{x}) + \cancel{\frac{d}{dt}\boldsymbol{\psi}(\mathbf{i}_s, \mathbf{x})}, \quad \text{since } \frac{d}{dt}\boldsymbol{\psi}(\mathbf{i}_s, \mathbf{x}) = 0$$

- The data for a constant set point $[i_{sdw}, i_{sqw}]$ are clustered and a common flux-linkage map (LUT) is obtained using the OLS (simple averaging)
- $2N^2$ parameters with inputs \mathbf{i}_s , $2N^{2+|\mathbf{x}|}$ parameters when additional inputs are added

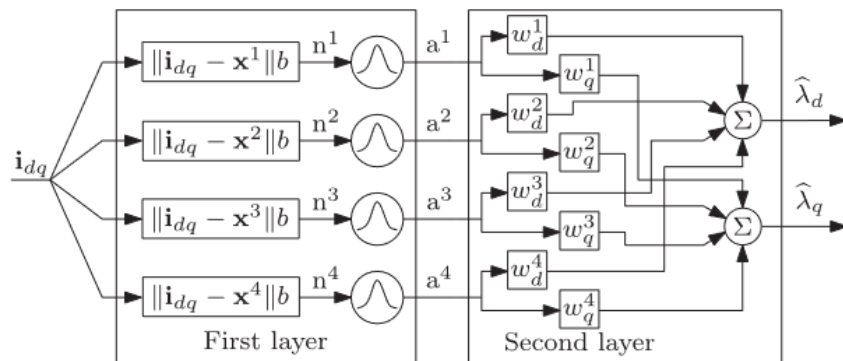


CASE STUDY: Neural ODE for Synchronous Drives

- Instead of a common LUT, we can use Radial Basis Functions (RBF)

$$\hat{\psi}_{RBF,sd}(\mathbf{i}_s) = \sum_{k=1}^K w_{d,k} \exp(-(b\|\mathbf{i}_s - \mathbf{x}_k\|)^2), \quad \hat{\psi}_{RBF,sq}(\mathbf{i}_s) = \sum_{k=1}^K w_{q,k} \exp(-(b\|\mathbf{i}_s - \mathbf{x}_k\|)^2),$$

where \mathbf{x}_k is a fixed center of the k -th RBF, and $w_{d,k}$ and $w_{q,k}$ are unknown weights of the model.



- Advantages of RBF model:

- Weights $w_{d,k}$ and $w_{q,k}$ can be estimated using OLS from steady model
- Weights $w_{d,k}$ and $w_{q,k}$ can be estimated using OLS even from dynamic model (considering $\frac{d}{dt} \hat{\psi}_{RBF}(\mathbf{i}_s)$)

[5] Ortombina, L., Pasqualotto, D., Tinazzi, F., & Zigliotto, M. (2020). Magnetic model identification of synchronous motors considering speed and load transients. IEEE Transactions on Industry Applications, 56(5), 4945-4954.

CASE STUDY: Neural ODE for Synchronous Drives

■ From voltage equation

$$\mathbf{u}_s = R\mathbf{i}_s + \omega_{el}J\boldsymbol{\psi}(\mathbf{i}_s, \mathbf{x}) + \frac{d}{dt}\boldsymbol{\psi}(\mathbf{i}_s, \mathbf{x}), \quad \text{using } \frac{d\boldsymbol{\psi}(\mathbf{i}_s, \mathbf{x})}{dt} = \frac{\partial\boldsymbol{\psi}(\mathbf{i}_s, \mathbf{x})}{\partial\mathbf{i}_s} \frac{d\mathbf{i}_s}{dt} + \frac{\partial\boldsymbol{\psi}(\mathbf{i}_s, \mathbf{x})}{\partial\mathbf{x}} \frac{d\mathbf{x}}{dt}$$

we can express the dynamic equation for current (if $\mathbf{x} = \emptyset$)

$$\frac{d}{dt}\hat{\mathbf{i}}_s = \text{inv} \left(\underbrace{\left. \frac{\partial\hat{\boldsymbol{\psi}}(\mathbf{i}_s)}{\partial\mathbf{i}_s} \right|_{\hat{\mathbf{i}}_s}}_{L(\hat{\mathbf{i}}_s)} \right) \cdot \left(\mathbf{u}_s - R\hat{\mathbf{i}}_s - \omega_{el}J\hat{\boldsymbol{\psi}}(\hat{\mathbf{i}}_s) \right)$$

or more generally

$$\frac{d}{dt}\hat{\mathbf{i}}_s = \text{inv} \left(\left. \frac{\partial\hat{\boldsymbol{\psi}}(\mathbf{i}_s, \mathbf{x})}{\partial\mathbf{i}_s} \right|_{\hat{\mathbf{i}}_s, \hat{\mathbf{x}}} \right) \cdot \left(\mathbf{u}_s - R\hat{\mathbf{i}}_s - \omega_{el}J\hat{\boldsymbol{\psi}}(\hat{\mathbf{i}}_s, \hat{\mathbf{x}}) - \left. \frac{\partial\hat{\boldsymbol{\psi}}(\mathbf{i}_s, \mathbf{x})}{\partial\mathbf{x}} \right|_{\hat{\mathbf{i}}_s, \hat{\mathbf{x}}} \frac{d}{dt}\hat{\mathbf{x}} \right)$$

CASE STUDY: Neural ODE for Synchronous Drives

- Dynamic equation for stator current with an unknown term $\hat{\boldsymbol{\psi}}(\mathbf{i}_s, \mathbf{x})$ that is modeled as NN
 - Sometimes called universal differential equation [6]
- Training:
 - Statically, i.e. minimizing $\left\| \hat{\boldsymbol{\psi}}(\mathbf{i}_s, \mathbf{x}) - \frac{J^{-1}(\mathbf{u}_s - R\mathbf{i}_s)}{\omega_{el}} \right\|_2$
 - Dynamically, i.e. ODE solution for current vs. measurement ($\Delta t = 10^{-4}$)
- $\hat{\boldsymbol{\psi}}(\mathbf{i}_s, \omega_{el})$ vs. $\hat{\boldsymbol{\psi}}(\mathbf{i}_s, \omega_{el}, \sin(K_s\vartheta), \cos(K_s\vartheta))$

[6] Rackauckas, C., Ma, Y., Martensen, J., Warner, C., Zubov, K., Supekar, R., ... & Edelman, A. (2020). Universal differential equations for scientific machine learning. *arXiv preprint arXiv:2001.04385*.

CASE STUDY: Neural ODE for Synchronous Drives

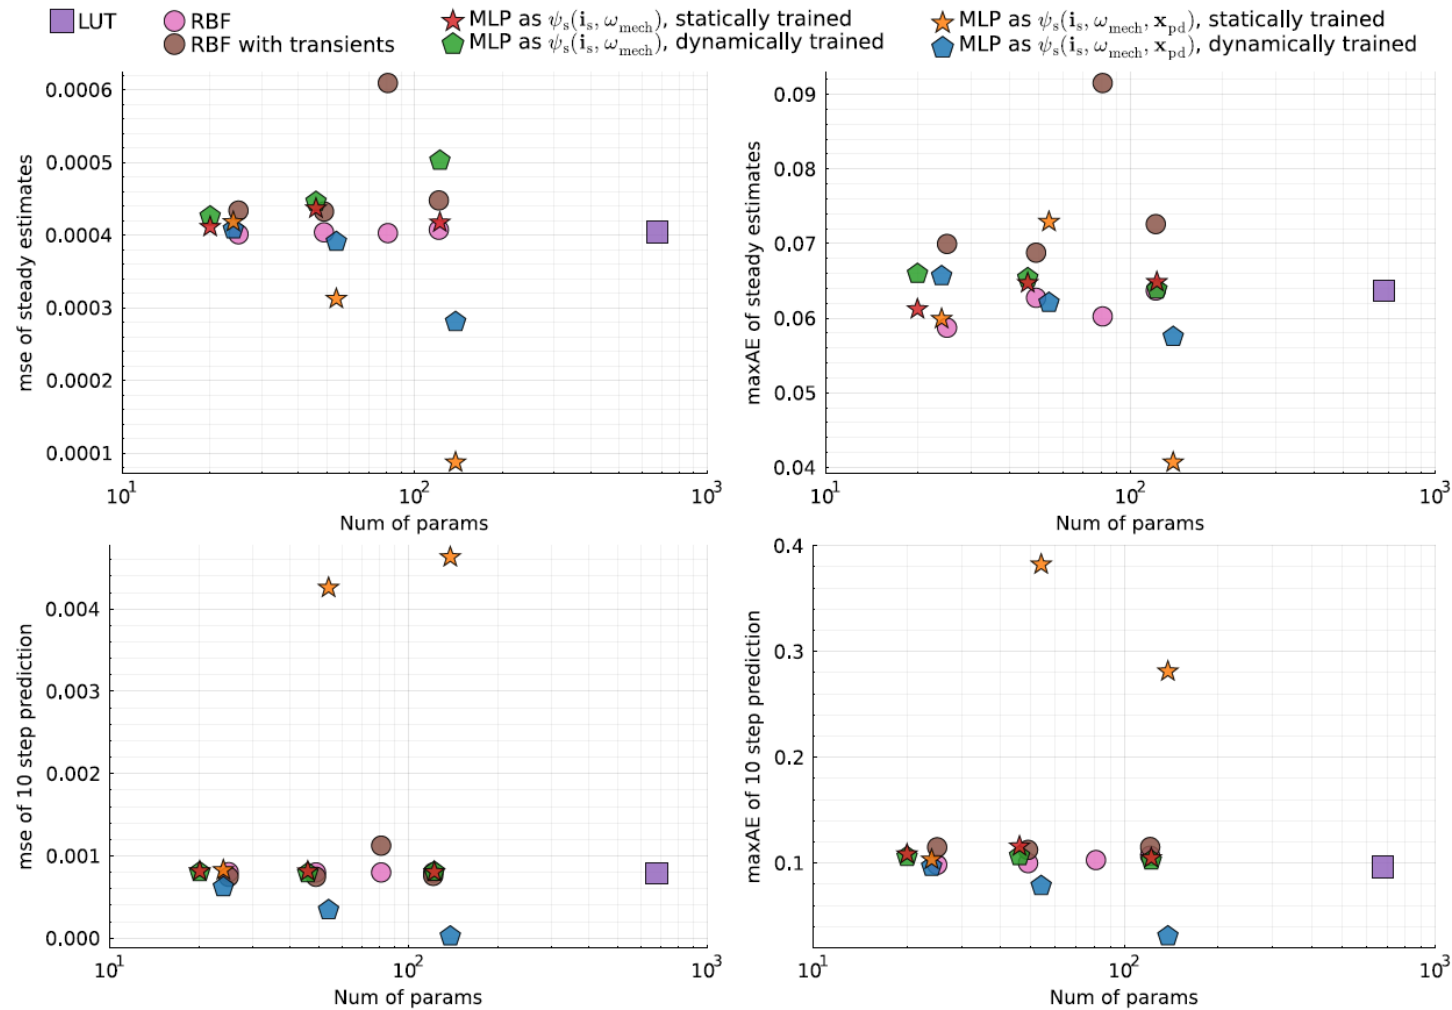


Fig. 4. Estimation errors of IPMSM's flux linkages models using validation data set with the constant speed $\omega_{mech}^* = -1300$ rpm. MSE in the left column, maxAE in the right column, errors of steady evaluation in the top row, errors of 10-step prediction in the bottom row. (10-step predictions of RBF-transient model with 9x9 grid has a great maxAE error (6.48) and therefore it is not displayed in the graph.)

CASE STUDY: Neural ODE for Synchronous Drives

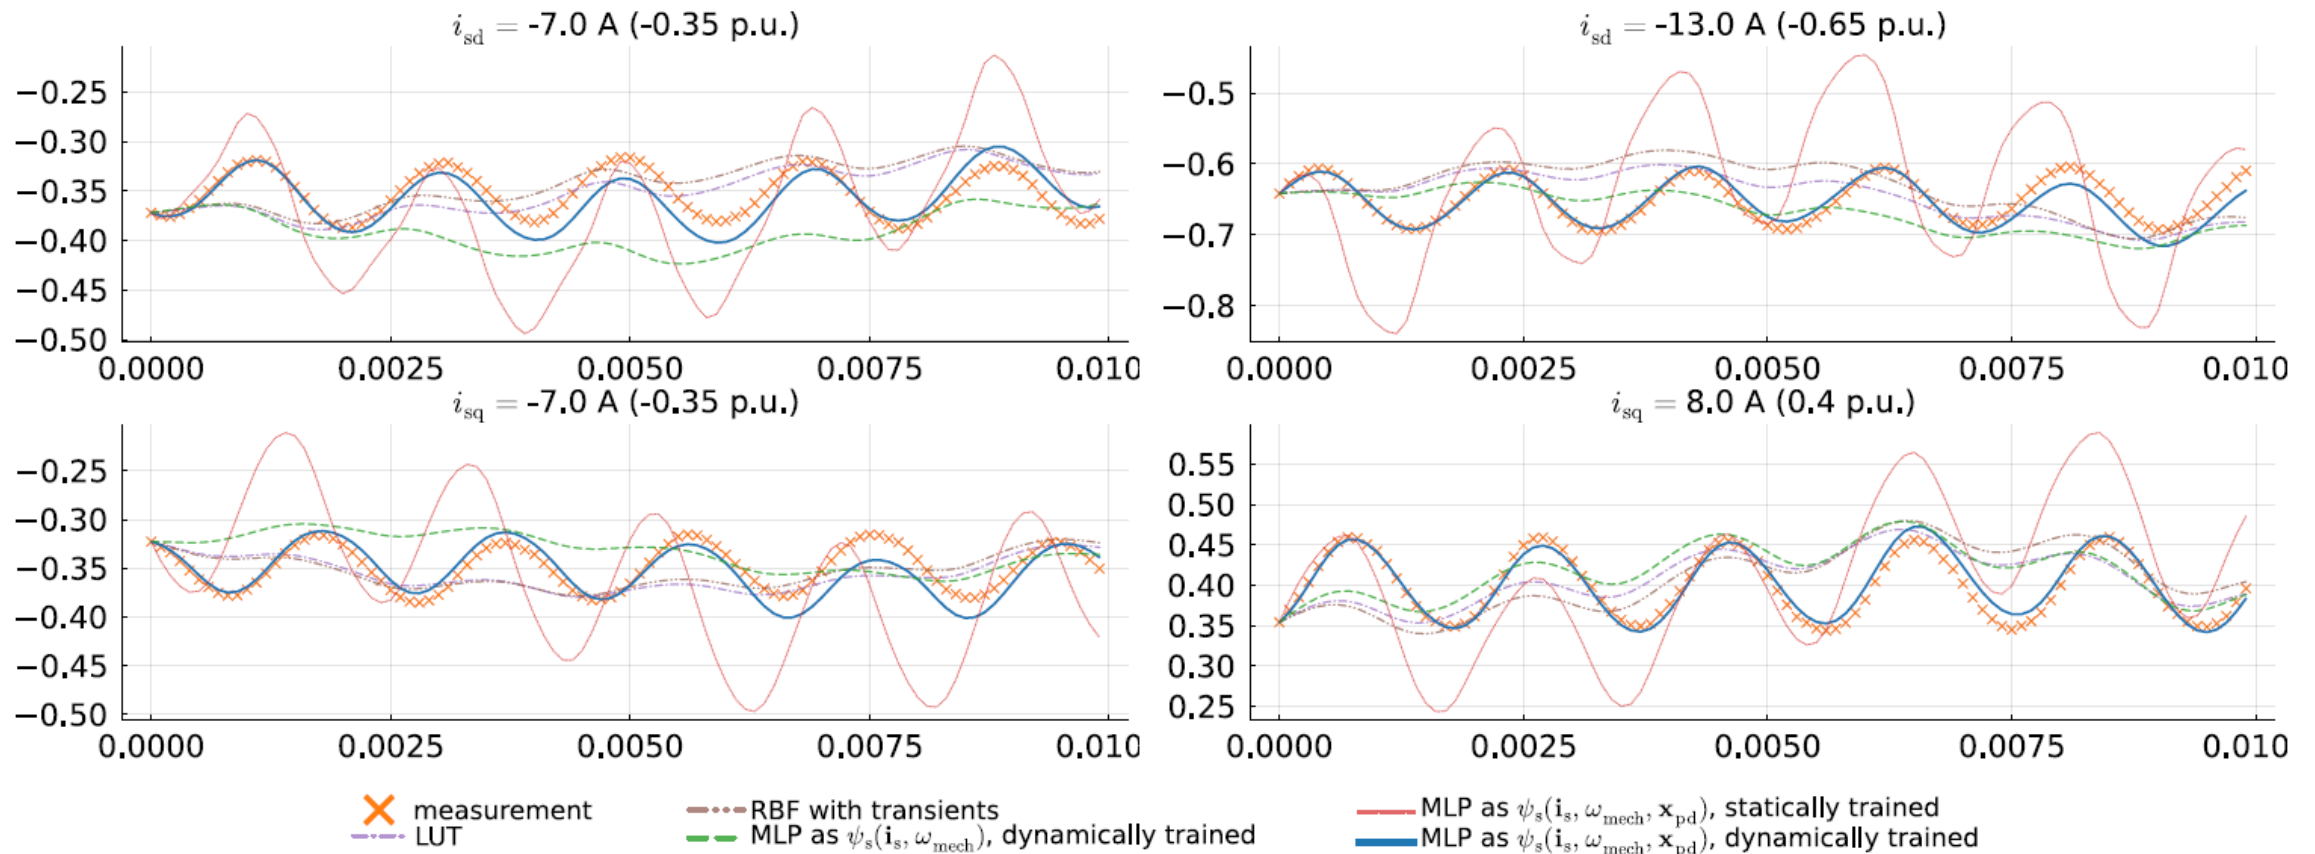


Fig. 5. Example of 100-step ahead predictions of stator currents i_{sd} (top row) and i_{sq} (bottom row) using trained flux linkage model of the IPMSM. The measured data has been acquired at constant speed -1300 RPM.

LITERATURE

- Chen, R. T., Rubanova, Y., Bettencourt, J., & Duvenaud, D. K. (2018). Neural ordinary differential equations. *Advances in neural information processing systems*, 31.
- Dupont, E., Doucet, A., & Teh, Y. W. (2019). Augmented neural odes. *Advances in neural information processing systems*, 32.
- Zhang, H., Gao, X., Unterman, J., & Arodz, T. (2020). Approximation capabilities of neural ODEs and invertible residual networks. In *International Conference on Machine Learning* (pp. 11086-11095).
- Kidger, P. (2022). On neural differential equations. *Doctoral thesis, University of Oxford. arXiv:2202.02435.*
- Ortombina, L., Pasqualotto, D., Tinazzi, F., & Zigliotto, M. (2020). Magnetic model identification of synchronous motors considering speed and load transients. *IEEE Transactions on Industry Applications*, 56(5), 4945-4954.
- Rackauckas, C., Ma, Y., Martensen, J., Warner, C., Zubov, K., Supekar, R., ... & Edelman, A. (2020). Universal differential equations for scientific machine learning. *arXiv preprint arXiv:2001.04385.*

NEURAL ORDINARY DIFFERENTIAL EQUATIONS

Thank you for Your attention

Jakub Ševčík

377 634 180

jsevcik@fel.zcu.cz

fel.zcu.cz



FACULTY OF ELECTRICAL
ENGINEERING
UNIVERSITY OF WEST BOHEMIA

RICE

