

1 We thank all the reviewers for their positive feedback and constructive comments. Below, we first discuss two issues
2 raised by multiple reviewers and then answer and clarify reviewer-specific comments / questions.

3 **Inflation factor α :** **R1** and **R4** asked for clarification of the inflation factor α . It is simply the increase in feature
4 dimensions for the UpDown model. I.e., if the dimension of q_1 is d then the dimension of q_2 is αd . We will clarify this.

5 **Practical implications:** **R3** and **R4** were curious whether our experiments demonstrate practical benefits of our
6 approach. We show state-of-the-art results for rotated MNIST (Fig. 4). However, we have not tackled large-scale
7 learning tasks yet. Our primary contribution is, as acknowledged by **R4**, to “enrich the understanding and discussion
8 about continuous time deep estimators” and to offer “a different viewing angle on the parametrization of neural ODEs.”
9 We are in the process of extending our approach to convolutional networks and to ResNet-like structures with striding
10 and feature-dimension changes. This will allow us to explore the performance of our approach for larger-scale tasks.

11 Response to **R1**

12 **Paper too dense:** This is a fair point. Our aim was to make the main manuscript as self-contained as possible. Hence,
13 we moved many details to the supplementary material. We will, as suggested, add to Fig. 1 an algorithmic description.

14 **Code complexity:** We will add a more detailed README file and will supplement the existing code with Jupyter
15 notebooks so that a reader can more easily follow. We will also add a simple PyTorch example illustrating the approach
16 in parallel with the new algorithmic description and independent of the more complex implementation we submitted.

17 Response to **R3**

18 **“Line 5: It’s not clear to me why having constant parameters throughout makes a model not continuous-depth”**
19 We will rephrase our sentence. The static parameters setting is indeed continuous-depth but obviously does not leverage
20 time varying parameters. **“I don’t understand the statement in line 106 with 3d images vs \mathbb{R}^d .”** In the LDDMM
21 community, the deformation map operates on 2D or 3D spaces, whereas in the DL setting, the map is usually defined in
22 much higher dimensions, e.g. the number of pixels of an image or for us the state-space of our UpDown model. **Line**
23 **120: Is $R(\Theta(t))$ some corresponding norm?** In our current setting, $R(\theta(t))$ is the Frobenius norm. However, $R(\theta(t))$
24 can be more general (see, e.g., Appendix A.2 on the Barron norm). We will clarify these points in the manuscript.

25 Response to **R4**

26 **“... a convolutional layer formulation is not included”** We had, in fact, already implemented shooting for convo-
27 lutional layers (it is part of the submitted code), but did not include this material and associated experiments to
28 avoid making the paper even more dense. **“It might not be straight forward to apply the shooting formulation to**
29 **different neural network layers ...”** We agree that there is a trade-off between the ease with which [10] can be
30 deployed for θ constant in time and the increased complexity of our approach for handling dynamic θ . In the “Rotated
31 MNIST” experiment, e.g., our formulation clearly shows benefits over its non-dynamic (`stat. direct` in Fig. 4)
32 counterpart. Hence, increased complexity can improve performance and can therefore be beneficial. Note also that our
33 implementation allows us to automatically derive the shooting equations (Appendix C and implementation) if desired.

34 **“How is this work connected to kernel machines? In particular, as A.2 mentions the connection to reproducing**
35 **kernel Hilbert spaces.”** Since our space of vector fields is finite dimensional, it is a proper RKHS when endowed with
36 the Frobenius norm. This connection still needs further exploration using other perspectives present in the DL literature.
37 **“Can a time-dependent θ also be simulated by increasing the state space for static direct, which is closest to the**
38 **original work of neural ODEs [10]?”** This is an excellent question. Unfortunately, we do not have a theoretically
39 conclusive answer at this point. **“How much more expensive in run-time is the proposed method? Although,**
40 **only a forward ODE solver run is required, the method needs to track each particle along the way.”** Currently,
41 depending on the number of particles, the method is more expensive in terms of run-time compared to [10]. However,
42 our primary focus was on functionality at the current stage, leaving substantial room for optimization and making
43 runtime comparisons difficult at this point. E.g., the matrices in the UpDown model are obtained via outer-products.
44 Hence, explicitly forming these matrices could be avoided. Further, for convolutional formulations the particles do not
45 need to have the same dimension as the input images, but could be much smaller patches. For the experiments in the
46 manuscript the `static direct` approach (which does not use any particles) indeed runs faster, but not dramatically so.

47 **“... [the supplementary material] does not include all the experiments ...”** We decided to omit any code that would
48 require reviewers to download additional data (e.g., for the “bouncing balls” experiment). Code for *all* experiments
49 will be included in the final version. **“... something wrong with the second part of equation 3.9.”** It is actually a
50 renaming of the variables – we overloaded the notations. We will fix this. **“Equations in the appendix seem not to**
51 **be numbered correctly”** We apologize that equation numbers in the appendix can easily be confused with section
52 numbers. We will fix this in the final version. **“... more discussion [of broader impacts] is needed.”** We will add
53 more discussion related to the societal impacts of the algorithms that may result from our work.