
Appendix: Learning Similarity Metrics for Numerical Simulations

This supplemental document contains an analysis of the proposed metric design with respect to properties of metrics in general (App. A) and details to the used network architectures (App. B). Afterwards, material that deals with the data sets is provided. It contains examples and failure cases for each of the data domains and analyzes the impact of the data difficulty (App. C and D). Next, the evaluation on real-world data is described in more detail (App. E). Finally, we explore additional metric evaluations (App. F) and give an overview on the used notation (App. G).

The source code for using the trained *LSiM* metric and re-training the model from scratch are available at <https://github.com/tum-pbs/LSiM>. This includes the full data sets and the corresponding data generation scripts for the employed PDE solver.

A. Discussion of Metric Properties

To analyze if the proposed method qualifies as a *metric*, it is split in two functions $m_1 : \mathbb{I} \rightarrow \mathbb{L}$ and $m_2 : \mathbb{L} \times \mathbb{L} \rightarrow [0, \infty)$, which operate on the input space \mathbb{I} and the latent space \mathbb{L} . Through flattening elements from the input or latent space into vectors, $\mathbb{I} \simeq \mathbb{R}^a$ and $\mathbb{L} \simeq \mathbb{R}^b$ where a and b are the dimensions of the input data and all feature maps respectively, and both values have a similar order of magnitude. m_1 describes the non-linear function computed by the base network combined with the following normalization and returns a point in the latent space. m_2 uses two points in the latent space to compute a final distance value, thus it includes the latent space difference and the aggregation along the spatial, layer, and channel dimensions. With the Siamese network architecture, the resulting function for the entire approach is

$$m(\mathbf{x}, \mathbf{y}) = m_2(m_1(\mathbf{x}), m_1(\mathbf{y})).$$

The identity of indiscernibles mainly depends on m_1 because, even if m_2 itself guarantees this property, m_1 could still be non-injective, which means it can map different inputs to the same point in latent space $\tilde{\mathbf{x}} = \tilde{\mathbf{y}}$ for $\mathbf{x} \neq \mathbf{y}$. Due to the complicated nature of m_1 , it is difficult to make accurate predictions about the injectivity of m_1 . Each base network layer of m_1 recursively processes the result of the preceding layer with various feature extracting operations. Here, the intuition is that significant changes in the input should produce different feature map results in one or more layers of the network. As very small changes in the input lead to zero valued distances predicted by the CNN (i.e., an

identical latent space for different inputs), m_1 is in practice not injective. In an additional experiment, the proposed architecture was evaluated on about 3500 random inputs from all our data sets, where the CNN received one unchanged and one slightly modified input. The modification consisted of multiple pixel adjustments by one bit (on 8-bit color images) in random positions and channels. When adjusting only a single pixel in the 224×224 input, the CNN predicts a zero valued distance on about 23% of the inputs, but we never observed an input where seven or more changed pixels resulted in a distance of zero in all experiments.

In this context, the problem of numerical errors is important because even two slightly different latent space representations could lead to a result that seems to be zero if the difference vanishes in the aggregation operations or is smaller than the floating point precision. On the other hand, an automated analysis to find points that have a different input but an identical latent space image is a challenging problem and left as future work.

The evaluation of the base network and the normalization is deterministic, and hence $\forall \mathbf{x} : m_1(\mathbf{x}) = m_1(\mathbf{x})$ holds. Furthermore, we know that $m(\mathbf{x}, \mathbf{x}) = 0$ if m_2 guarantees that $\forall m_1(\mathbf{x}) : m_2(m_1(\mathbf{x}), m_1(\mathbf{x})) = 0$. Thus, the remaining properties, i.e., non-negativity, symmetry, and the triangle inequality, only depend on m_2 since for them the original inputs are not relevant, but their respective images in the latent space. The resulting structure with a relaxed identity of indiscernibles is called a *pseudometric*, where $\forall \tilde{\mathbf{x}}, \tilde{\mathbf{y}}, \tilde{\mathbf{z}} \in \mathbb{L}$:

$$m_2(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) \geq 0 \tag{1}$$

$$m_2(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) = m_2(\tilde{\mathbf{y}}, \tilde{\mathbf{x}}) \tag{2}$$

$$m_2(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) \leq m_2(\tilde{\mathbf{x}}, \tilde{\mathbf{z}}) + m_2(\tilde{\mathbf{z}}, \tilde{\mathbf{y}}) \tag{3}$$

$$m_2(\tilde{\mathbf{x}}, \tilde{\mathbf{x}}) = 0 \tag{4}$$

Notice that m_2 has to fulfill these properties with respect to the latent space but not the input space. If m_2 is carefully constructed, the metric properties still apply, independently of the actual design of the base network or the feature map normalization.

A first observation concerning m_2 is that if all aggregations were sum operations and the element-wise latent space difference was the absolute value of a difference operation, m_2 would be equivalent to computing the L^1 norm of the difference vector in latent space:

$$m_2^{sum}(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) = \sum_{i=1}^b |\tilde{\mathbf{x}}_i - \tilde{\mathbf{y}}_i|.$$

Similarly, adding a square operation to the element-wise distance in the latent space and computing the square root at the very end leads to the L^2 norm of the latent space difference vector. In the same way, it is possible to use any L^p norm with the corresponding operations:

$$m_2^{sum}(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) = \left(\sum_{i=1}^b |\tilde{\mathbf{x}}_i - \tilde{\mathbf{y}}_i|^p \right)^{\frac{1}{p}}.$$

In both cases, this forms the metric induced by the corresponding norm, which by definition has all desired properties (1), (2), (3), and (4). If we change all aggregation methods to a weighted average operation, each term in the sum is multiplied by a weight w_i . This is even possible with learned weights, as they are constant at evaluation time if they are clamped to be positive as described above. Now, w_i can be attributed to both inputs by distributivity, meaning each input is element-wise multiplied with a constant vector before applying the metric, which leaves the metric properties untouched. The reason is that it is possible to define new vectors in the same space, equal to the scaled inputs. This renaming trivially provides the correct properties:

$$m_2^{weighted}(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) = \sum_{i=1}^b w_i |\tilde{\mathbf{x}}_i - \tilde{\mathbf{y}}_i|,$$

$$\stackrel{w_i \geq 0}{=} \sum_{i=1}^b |w_i \tilde{\mathbf{x}}_i - w_i \tilde{\mathbf{y}}_i|.$$

Accordingly, doing the same with the L^p norm idea is possible, and each w_i just needs a suitable adjustment before distributivity can be applied, keeping the metric properties once again:

$$m_2^{weighted}(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) = \left(\sum_{i=1}^b w_i |\tilde{\mathbf{x}}_i - \tilde{\mathbf{y}}_i|^p \right)^{\frac{1}{p}}$$




Figure 1. Proposed base network architecture consisting of five layers with up to 192 feature maps that are decreasing in spatial size. It is similar to the feature extractor from AlexNet as identical spatial dimensions for the feature maps are used, but it reduces the number of feature maps for each layer by 50% to have fewer weights.

$$= \left(\sum_{i=1}^b w_i |\tilde{\mathbf{x}}_i - \tilde{\mathbf{y}}_i| |\tilde{\mathbf{x}}_i - \tilde{\mathbf{y}}_i| \dots |\tilde{\mathbf{x}}_i - \tilde{\mathbf{y}}_i| \right)^{\frac{1}{p}}$$

$$= \left(\sum_{i=1}^b w_i^{\frac{1}{p}} |\tilde{\mathbf{x}}_i - \tilde{\mathbf{y}}_i| w_i^{\frac{1}{p}} |\tilde{\mathbf{x}}_i - \tilde{\mathbf{y}}_i| \dots w_i^{\frac{1}{p}} |\tilde{\mathbf{x}}_i - \tilde{\mathbf{y}}_i| \right)^{\frac{1}{p}},$$

$$\stackrel{w_i \geq 0}{=} \left(\sum_{i=1}^b |w_i^{\frac{1}{p}} \tilde{\mathbf{x}}_i - w_i^{\frac{1}{p}} \tilde{\mathbf{y}}_i|^p \right)^{\frac{1}{p}}.$$

With these weighted terms for m_2 , it is possible to describe all used aggregations and latent space difference methods. The proposed method deals with multiple higher order tensors instead of a single vector. Thus, the weights w_i additionally depend on constants such as the direction of the aggregations and their position in the latent space tensors. But it is easy to see that mapping a higher order tensor to a vector and keeping track of additional constants still retains all properties in the same way. As a result, the described architecture by design yields a pseudometric that is suitable for comparing simulation data in a way that corresponds to our intuitive understanding of distances.

B. Architectures

The following sections provide details regarding the architecture of the base network and some experimental design.

B.1. Base Network Design

Fig. 1 shows the architecture of the base network for the $LSiM$ metric. Its purpose is to extract features from both inputs of the Siamese architecture that are useful for the further processing steps. To maximise the usefulness and to avoid feature maps that show overly similar features, the chosen kernel size and stride of the convolutions are important. Starting with larger kernels and strides means the network has a big receptive field and can consider simple, low-level features in large regions of the input. For the two




Figure 2. Analysis of the distributions of learned feature map aggregation weights across the base network layers. Displayed is a base network with pre-trained weights (left) in comparison to our method for fully training the base network (right). Note that the percentage of unused feature maps for most layers of our base network is 0%.

following layers, the large strides are replaced by additional MaxPool operations that serve a similar purpose and reduce the spatial size of the feature maps.

For the three final layers, only small convolution kernels and strides are used, but the number of channels is significantly larger than before. These deep features maps typically contain high-level structures, which are most important to distinguish complex changes in the inputs. Keeping the number of trainable weights as low as possible was an important consideration for this design to prevent overfitting to certain simulations types and increase generality. We explored a weight range by using the same architecture and only scaling the number of feature maps in each layer. The final design shown in Fig. 1 with about 0.62 million weights worked best for our experiments.

In the following, we analyze the contributions of the per-layer features of two different metric networks to highlight differences in terms of how the features are utilized for the distance estimation task. In Fig. 2, our *LSiM* network yields a significantly smaller standard deviation in the learned weights that aggregate feature maps of five layers, compared to a pre-trained base network. This means, all feature maps contribute to establishing the distances similarly, and the aggregation just fine-tunes the relative importance of each feature. In addition, almost all features receive a weight greater than zero, and as a result, more features are contributing to the final distance value.

Employing a fixed pre-trained feature extractor, on the other hand, shows a very different picture: Although the mean across the different network layers is similar, the contributions of different features vary strongly, which is visible in the standard deviation being significantly larger. Furthermore, 2-10% of the feature maps in each layer receive a weight of zero and hence were deemed not useful at all for establishing the distances. This illustrates the usefulness of a targeted network in which all features contribute to the distance inference.

B.2. Feature Map Normalization

In the following, we analyze how the different feature map normalizations discussed in Section 3.2 of the main paper affect the performance of our metric. We compare using no normalization $norm_{none}(\mathbf{G}) = \mathbf{G}$, the unit length normalization via division by the norm of a feature vector $norm_{unit}(\mathbf{G}) = \mathbf{G} / \|\mathbf{G}\|_2$ proposed by Zhang et al., a global unit length normalization $norm_{global}(\mathbf{G}) = \mathbf{G} / \max(\|\mathbf{G}_0\|_2, \|\mathbf{G}_1\|_2, \dots)$ that considers the norm of all feature vectors in the entire training set, and the proposed normalization to a scaled chi distribution

$$norm_{dist}(\mathbf{G}) = \frac{1}{\sqrt{g_c - 1}} \frac{\mathbf{G} - \text{mean}(\mathbf{G}_0, \mathbf{G}_1, \dots)}{\text{std}(\mathbf{G}_0, \mathbf{G}_1, \dots)}.$$

Fig. 3 shows a comparison of these normalization methods on the combined test data. Using no normalization is significantly detrimental to the performance of the metric as succeeding operations cannot reliably compare the features. A unit length normalization of a single sample is already a major improvement since following operations now have a predictable range of values to work with. This corresponds to a cosine distance, which only measures angles of the feature vectors and entirely neglects their length.




Figure 3. Performance on our test data for different feature map normalization approaches.

Using the maximum norm across all training samples (computed in a pre-processing step and fixed for training) introduces additional information as the network can now compare magnitudes as well. However, this comparison is not stable as the maximum norm can be an outlier with respect to the typical content of the corresponding feature.

The proposed normalization forms a chi distribution by individually transforming each component of the feature vector to a standard normal distribution. Afterwards, scaling with the inverse mode of the chi distribution leads to a consistent average magnitude close to one. It results in the best performing metric since both length and angle of the feature vectors can be reliably compared by the following operations.

B.3. Recursive “Meta-Metric”

Since comparing the feature maps is a central operation of the proposed metric calculations, we experimented with replacing it with an existing CNN-based metric. In theory, this would allow for a recursive, arbitrarily deep network that repeatedly invokes itself: first, the extracted representations of inputs are used and then the representations extracted from the previous representations, etc. In practice, however, using more than one recursion step is currently not feasible due to increasing computational requirements in addition to vanishing gradients.

Fig. 4 shows how our computation method can be modified for a CNN-based latent space difference, instead of an element-wise operation. Here we employ *LPIPS* (Zhang et al., 2018). There are two main differences compared to proposed method. First, the *LPIPS* latent space difference creates single distance values for a pair of feature maps instead of a spatial feature difference. As a result, the following aggregation is a single learned average operation and spatial or layer aggregations are no longer necessary. We also performed experiments with a spatial *LPIPS* version here, but due to memory limitations, these were not successful. Second, the convolution operations in *LPIPS* have a lower limit for spatial resolution, and some feature maps of our base network are quite small (see Fig. 1). Hence, we up-scale the feature maps below the required spatial size of 32×32 using nearest neighbor interpolation.

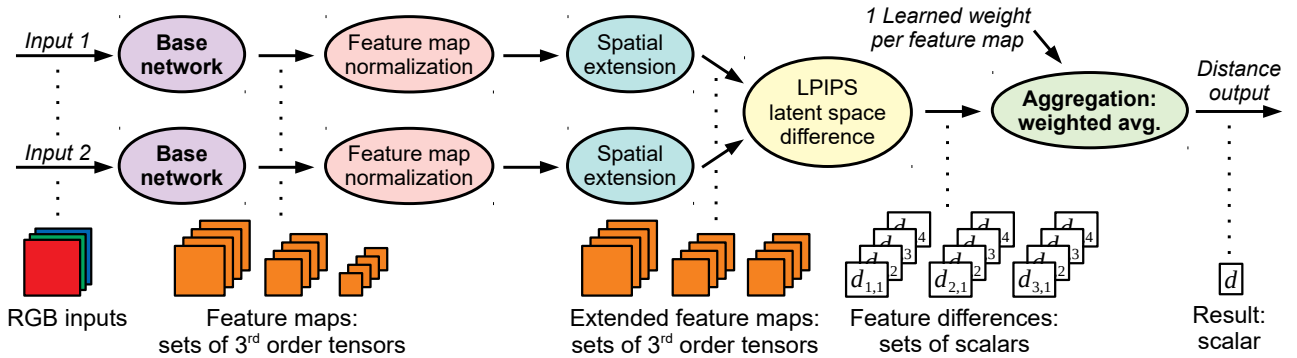


Figure 4. Adjusted distance computation for a *LPIPS*-based latent space difference. To provide sufficiently large inputs for *LPIPS*, small feature maps are spatially enlarged with nearest neighbor interpolation. In addition, *LPIPS* creates scalar instead of spatial differences leading to a simplified aggregation.

On our combined test data, such a metric with a fully trained base network achieves a performance comparable to *AlexNet_{random}* or *AlexNet_{frozen}*.

B.4. Optical Flow Metric

In the following, we describe our approach to compute a metric via optical flow (OF). For an efficient OF evaluation, we employed a pre-trained network (Ilg et al., 2016). From an OF network $f : \mathbb{I} \times \mathbb{I} \rightarrow \mathbb{R}^{i_{max} \times j_{max} \times 2}$ with two input data fields $\mathbf{x}, \mathbf{y} \in I$, we get the flow vector field $f^{\mathbf{x}\mathbf{y}}(i, j) = (f_1^{\mathbf{x}\mathbf{y}}(i, j), f_2^{\mathbf{x}\mathbf{y}}(i, j))^T$, where i and j denote the locations, and f_1 and f_2 denote the components of the flow vectors. In addition, we have a second flow field $f^{\mathbf{y}\mathbf{x}}(i, j)$ computed from the reversed input ordering. We can now define a function $m : \mathbb{I} \times \mathbb{I} \rightarrow [0, \infty)$:

$$m(\mathbf{x}, \mathbf{y}) = \sum_{i=0}^{i_{max}} \sum_{j=0}^{j_{max}} \sqrt{(f_1^{\mathbf{x}\mathbf{y}}(i, j))^2 + (f_2^{\mathbf{x}\mathbf{y}}(i, j))^2} + \sqrt{(f_1^{\mathbf{y}\mathbf{x}}(i, j))^2 + (f_2^{\mathbf{y}\mathbf{x}}(i, j))^2}.$$

Intuitively, this function computes the sum over the magnitudes of all flow vectors in both vector fields. With this definition, it is obvious that $m(\mathbf{x}, \mathbf{y})$ fulfills the metric properties of non-negativity and symmetry (see Eq. (1) and (2)). Under the assumption that identical inputs create a zero flow field, a relaxed identity of indiscernibles holds as well (see Eq. (4)). Compared to the proposed approach, there is no guarantee for the triangle inequality though, thus $m(\mathbf{x}, \mathbf{y})$ only qualifies as a pseudo-semimetric.

Fig. 5 shows flow visualizations on data examples produced by FlowNet2. The metric works relatively well for inputs that are similar to the training data from FlowNet2 such as the shape data example in the top row. For data that provides some outline, e.g., the smoke simulation example in the middle row or also liquid data, the metric does not work

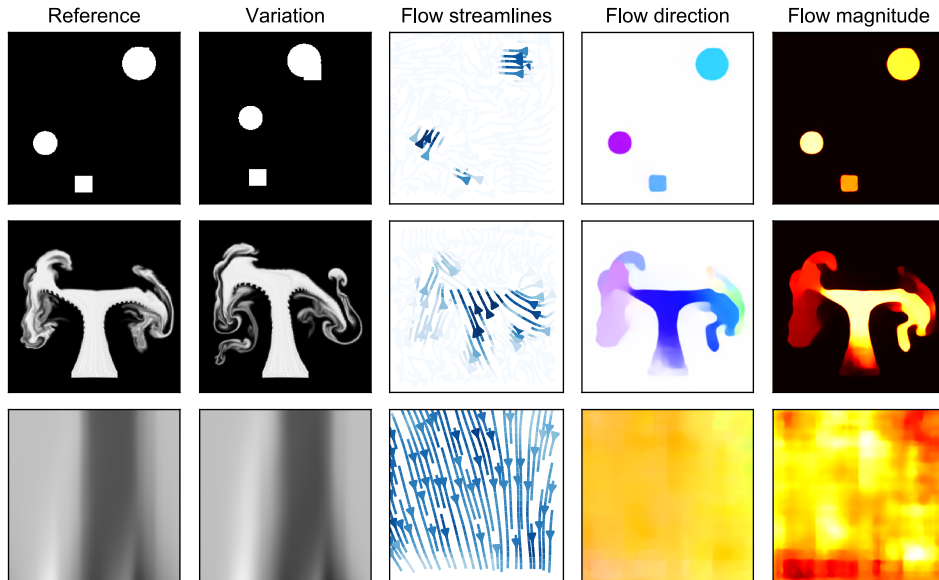


Figure 5. Outputs from FlowNet2 on data examples. The flow streamlines are sparse visualization of the resulting flow field and indicate the direction of the flow by their orientation and its magnitude by their color (darker being larger). The two visualizations on the right show the dense flow field and are color-coded to show the flow direction (blue/yellow: vertical, green/red: horizontal) and the flow magnitude (brighter being larger).




Figure 6. Non-Siamese network architecture with the same feature extractor used in Fig. 1. It uses both stacked inputs and directly predicts the final distance value from the last set of feature maps with several fully connected layers.

as well but still provides a reasonable flow field. However, for full spatial examples such as the Burger’s or Advection-Diffusion cases (see bottom row), the network is no longer able to produce meaningful flow fields. The results are often a very uniform flow with similar magnitude and direction.

B.5. Non-Siamese Architecture

To compute a metric without the Siamese architecture outlined above, we use a network structure with a single output as shown in Fig. 6. Thus, instead of having two identically feature extractors and combining the feature maps, here the distance is directly predicted from the stacked inputs with a single network with about 1.24 million weights. After using the same feature extractor as described in Section B.1, the final set of feature maps is spatially reduced with an adaptive MaxPool operation. Next, the result is flattened, and

three consecutive fully connected layers process the data to form the final prediction. Here, the last activation function is a sigmoid instead of ReLU. The reason is that a ReLU would clamp every negative intermediate value to a zero distance, while a sigmoid compresses the intermediate value to a small distance that is more meaningful than directly clamping it.

In terms of metric properties, this architecture only provides non-negativity (see Eq. (1)) due to the final sigmoid function. All other properties cannot be guaranteed without further constraints. This is the main disadvantage of a non-Siamese network. These issues could be alleviated with specialized training data or by manually adding constraints to the model, e.g., to have some amount of symmetry (see Eq. (2)) and at least a weakened identity of indiscernibles (see Eq. (4)). However, compared to a Siamese network that guarantees




Figure 7. Network architecture with skip connections for better information transport between feature maps. Transposed convolutions are used to upscale the feature maps in the second half of the network to match the spatial size of earlier layers for the skip connections.

them by design, these extensions are clearly sub-optimal. As a result of the missing properties, this network has significant problems with generalization. While it performs well on the training data, the performance noticeably deteriorates for several of the test data sets.

B.6. Skip Connections in Base Network

As explained above, our base network primarily serves as a feature extractor to produce activations that are employed to evaluate a learned metric. In many state-of-the-art methods, networks with skip connections are employed (Ronneberger et al., 2015; He et al., 2016; Huang et al., 2017), as experiments have shown that these connections help to preserve information from the inputs. In our case, the classification “output” of a network such as the AlexNet plays no actual role. Rather, the features extracted along the way are crucial. Hence, skip connections should not improve the inference task for our metrics.

To verify that this is the case, we have included tests with a base network (see Fig. 7) similar to the popular UNet architecture (Ronneberger et al., 2015). For our experiments, we kept the early layers closely in line with the feature extractors that worked well for the base network (see Section B.1). Only the layers in the decoder part have an increased spatial feature map size to accommodate the skip connections. As expected, this network can be used to compute reliable metrics for the input data without negatively affecting the performance. However, as expected, the improvements of skip connections for regular inference tasks do not translate into improvements for the metric calculations.

C. Impact of Data Difficulty




Figure 8. Impact of increasing data difficulty for a reduced training data set. Evaluations on training data for L^2 and $LPIPS$, and the test performance of models trained with the different reduced data sets ($LSiM_{reduced}$) are shown.

We shed more light on the aspect of noise levels and data difficulty via six reduced data sets that consist of a smaller amount of Smoke and Advection-Diffusion data with differently scaled noise strength values. Results are shown in Fig. 8. Increasing the noise level creates more difficult data as shown by the dotted and dashed plots representing the performance of the L^2 and the $LPIPS$ metric on each data set. Both roughly follow an exponentially decreasing function. Each point on the solid line plot is the test result of a reduced $LSiM$ model trained on the data set with the corresponding noise level. Apart from the data, the entire training

setup was identical. This shows that the training process is very robust to the noise, as the result on the test data only slowly decreases for very high noise levels. Furthermore, small amounts of noise improve the generalization compared to the model that was trained without any noise. This is somewhat expected, as a model that never saw noisy data during training cannot learn to extract features which are robust with respect to noise.

D. Data Set Details

In the following sections, the generation of each used data set is described. For each figure showing data samples (consisting of a reference simulation and several variants with a single changing initial parameter), the leftmost image is the reference and the images to the right show the variants in order of increasing parameter change. For the figures 9, 10, 11, and 12, the first subfigure (a) demonstrates that medium and large scale characteristics behave very non-chaotic for simulations without any added noise. They are only included for illustrative purposes and are not used for training. The second and third subfigure (b) and (c) in each case show the training data of *LSiM*, where the large majority of data falls into the category (b) of normal samples that follow the generation ordering, even with more varying behaviour. Category (c) is a small fraction of the training data, and the shown examples are specifically picked to show how the chaotic behaviour can sometimes override the ordering intended by the data generation in the worst case. Occasionally, category (d) is included to show how normal data samples from the test set differ from the training data.

D.1. Navier-Stokes Equations

These equations describe the general behaviour of fluids with respect to advection, viscosity, pressure, and mass conservation. Eq. (5) defines the conservation of momentum, and Eq. (6) constraints the conservation of mass:

$$\frac{\partial u}{\partial t} + (u \cdot \nabla)u = -\frac{\nabla P}{\rho} + \nu \nabla^2 u + g, \quad (5)$$

$$\nabla \cdot u = 0. \quad (6)$$

In this context, u is the velocity, P is the pressure the fluid exerts, ρ is the density of the fluid (usually assumed to be constant), ν is the kinematic viscosity coefficient that indicates the thickness of the fluid, and g denotes the acceleration due to gravity. With this PDE, three data sets were created using a smoke and a liquid solver. For all data, 2D simulations were run until a certain step, and useful data fields were exported afterwards.

SMOKE

For the smoke data, a standard Eulerian fluid solver using a preconditioned pressure solver based on the conjugate

gradient method and Semi-Lagrangian advection scheme was employed.

The general setup for every smoke simulation consists of a rectangular smoke source at the bottom with a fixed additive noise pattern to provide smoke plumes with more details. Additionally, there is a downwards directed, spherical force field area above the source, which divides the smoke in two major streams along it. We chose this solution over an actual obstacle in the simulation in order to avoid overfitting to a clearly defined black obstacle area inside the smoke data. Once the simulation reaches a predefined time step, the density, pressure, and velocity fields (separated by dimension) are exported and stored. Some example sequences can be found in Fig. 9. With this setup, the following initial conditions were varied in isolation:

- Smoke buoyancy in x- and y-direction
- Strength of noise added to the velocity field
- Amount of force in x- and y-direction provided by the force field
- Orientation and size of the force field
- Position of the force field in x- and y-direction
- Position of the smoke source in x- and y-direction


Overall, 768 individual smoke sequences were used for training, and the validation set contains 192 sequences with different initialization seeds.

LIQUID


For the liquid data, a solver based on the fluid implicit particle (FLIP) method (Zhu & Bridson, 2005) was employed. It is a hybrid Eulerian-Lagrangian approach that replaces the Semi-Lagrangian advection scheme with particle based advection to reduce numerical dissipation. Still, this method is not optimal as we experienced problems such as mass loss, especially for larger noise values.

The simulation setup consists of a large breaking dam and several smaller liquid areas for more detailed splashes. After the dam hits the simulation boundary, a large, single drop of liquid is created in the middle of the domain that hits the already moving liquid surface. Then, the extrapolated level set values, binary indicator flags, and the velocity fields (separated by dimension) are saved. Some examples are shown in Fig. 10. The list of varied parameters include:


- Radius of the liquid drop
- Position of the drop in x- and y-direction
- Amount of additional gravity force in x- and y-direction
- Strength of noise added to the velocity field



(a) Data samples generated without noise: tiny output changes following generation ordering




(b) Normal training data samples with noise: larger output changes but ordering still applies




(c) Outlier data samples: noise can override the generation ordering by chance


Figure 9. Various smoke simulation examples using one component of the velocity (top rows), the density (middle rows), and the pressure field (bottom rows).




(a) Data samples generated without noise: tiny output changes following generation ordering



(b) Normal training data samples with noise: larger output changes but ordering still applies




(c) Outlier data samples: noise can override the generation ordering by chance




(d) Data samples from test set with additional background noise


Figure 10. Several liquid simulation examples using the binary indicator flags (top rows), the extrapolated level set values (middle rows), and one component of the velocity field (bottom rows) for the training data and only the velocity field for the test data.




(a) Data samples generated without noise: tiny output changes following generation ordering



(b) Normal training data samples with noise: larger output changes but ordering still applies




(c) Outlier data samples: noise can override the generation ordering by chance




(d) Data samples from test set with additional background noise


Figure 11. Various examples from the Advection-Diffusion equation using the density field.



(a) Data samples generated without noise: tiny output changes following generation ordering



(b) Normal training data samples with noise: larger output changes but ordering still applies



(c) Outlier data samples: noise can override the generation ordering by chance

Figure 12. Different simulation examples from the Burger’s equation using the velocity field.

The liquid training set consists of 792 sequences and the validation set of 198 sequences with different random seeds. For the liquid test set, additional background noise was added to the velocity field of the simulations as displayed in Fig. 10(d). Because this only alters the velocity field, the extrapolated level set values and binary indicator flags are not used for this data set, leading to 132 sequences.

D.2. Advection-Diffusion and Burger’s Equation

For these PDEs, our solvers only discretize and solve the corresponding equation in 1D. Afterwards, the different time steps of the solution process are concatenated along a new dimension to form 2D data with one spatial and one time dimension.

ADVECTION-DIFFUSION EQUATION

This equation describes how a passive quantity is transported inside a velocity field due to the processes of advection and diffusion. Eq. (7) is the simplified Advection-Diffusion equation with constant diffusivity and no sources or sinks.

$$\frac{\partial d}{\partial t} = \nu \nabla^2 d - u \cdot \nabla d, \quad (7)$$

where d denotes the density, u is the velocity, and ν is the kinematic viscosity (also known as diffusion coefficient) that determines the strength of the diffusion. Our solver employed a simple implicit time integration and a diffusion solver based on conjugate gradient without preconditioning. The initialization for the 1D fields of the simulations was created by overlaying multiple parameterized sine curves with random frequencies and magnitudes.

In addition, continuous forcing controlled by further parameterized sine curves was included in the simulations over time. In this case, the only initial conditions to vary are the forcing and initialization parameters of the sine curves and the strength of the added noise. From this PDE, only the passive density field was used as shown in Fig. 11. Overall, 798 sequences are included in the training set and 190 sequences with a different random initialization in the validation set.

For the Advection-Diffusion test set, the noise was instead added directly to the passive density field of the simulations. This results in 190 sequences with more small scale details as shown in Fig. 11(d).

BURGER'S EQUATION

This equation is very similar to the Advection-Diffusion equation and describes how the velocity field itself changes due to diffusion and advection:

$$\frac{\partial u}{\partial t} = \nu \nabla^2 u - u \cdot \nabla u. \quad (8)$$

Eq. (8) is known as the viscous form of the Burger's equation that can develop shock waves, and again u is the velocity and ν denotes the kinematic viscosity. Our solver for this PDE used a slightly different implicit time integration scheme, but the same diffusion solver as used for the Advection-Diffusion equation.

The simulation setup and parameters were also the same; the only difference is that the velocity field instead of the density is exported. As a consequence, the data in Fig. 12 looks relatively similar to those from the Advection-Diffusion equation. The training set features 782 sequences, and the validation set contains 204 sequences with different random seeds.

D.3. Other Data-Sets

The remaining data sets are not based on PDEs and thus not generated with the proposed method. The data is only used to test the generalization of the discussed metrics and not for training or validation. The Shapes test set contains 160 sequences, the Video test set consists 131 sequences, and the TID test set features 216 sequences.

SHAPES

This data set tests if the metrics are able to track simple, moving geometric shapes. To create it, a straight path between two random points inside the domain is generated

and a random shape is moved along this path in steps of equal distance. The size of the used shape depends on the distance between the start and end point such that a significant fraction of the shape overlaps between two consecutive steps. It is also ensured that no part of the shape leaves the domain at any step by using a sufficiently big boundary area when generating the path.

With this method, multiple random shapes for a single data sample are produced, and their paths can overlap such that they occlude each other to provide an additional challenge. All shapes are moved in their parametric representation, and only when exporting the data, they are discretized onto a fixed binary grid. To add more variations to this simple approach, we also apply them in a non-binary way with smoothed edges and include additive Gaussian noise over the entire domain. Examples are shown in Fig. 13.

VIDEO

For this data set, different publicly available video recordings were acquired and processed in three steps. First, videos with abrupt cuts, scene transitions, or camera movements were discarded, and afterwards the footage was broken down into single frames. Then, each frame was resized to match the spatial size of our other data by linear interpolation. Since directly using consecutive frames is no challenge for any analyzed metric and all of them recovered the ordering almost perfectly, we achieved a more meaningful data set by skipping several intermediate frames. For the final data set, we defined the first frame of every video as the reference and collected subsequent frames in an interval step of ten frames as the increasingly different variations. Some data examples can be found in Fig. 14.




Figure 13. Examples from the shapes data set using a field with only binary shape values (first row), shape values with additional noise (second row), smoothed shape values (third row), and smoothed values with additional noise (fourth row).




Figure 14. Multiple examples from the video data set.



Figure 15. Examples from the TID2013 data set proposed by Ponomarenko et al.. Displayed are a change of contrast, three types of noise, denoising, jpg2000 compression, and two color quantizations (from left to right and top to bottom).

TID2013

This data set was created by Ponomarenko et al. and used without any further modifications. It consists of 25 reference images with 24 distortion types in five levels. As a result, it is not directly comparable to our data sets; thus, it is excluded from the test set aggregations. The distortions focus on various types of noise, image compression, and color changes. Fig. 15 contains examples from the data set.

D.4. Hardware

Data generation, training, and metric evaluations were performed on a machine with an Intel i7-6850 (3.60Ghz) CPU and an NVIDIA GeForce GTX 1080 Ti GPU.

E. Real-World Data

Below, we give details of the three data sets used for the evaluation in Section 6.3 of the main paper.

E.1. ScalarFlow

The *ScalarFlow* data set (Eckert et al., 2019) contains 3D velocities of real-world scalar transport flows reconstructed from multiple camera perspectives. For our evaluation, we cropped the volumetric $100 \times 178 \times 100$ grids to $100 \times 160 \times 100$ such that they only contain the area of interest and convert them to 2D with two variants: either by using the center slice or by computing the mean along the z-dimension. Afterwards, the velocity vectors are split by channels, linearly interpolated to 256×256 , and then normalized. Variations for each reconstructed plume are acquired by using frames in equal temporal intervals. We employed the velocity field reconstructions from 30 plumes (with simulation IDs 0 – 29) for both compression methods. Fig. 16 shows some example sequences.

E.2. Johns Hopkins Turbulence Database

The Johns Hopkins Turbulence Database (*JHTDB*) (Perlman et al., 2007) features various data sets of 3D turbu-




Figure 16. Four different smoke plume examples of the processed *ScalarFlow* data set using one of the three velocity components. The two top rows show the center slice, and the two bottom rows show the mean along the *z*-dimension. The temporal interval between each image is ten simulation time steps.




Figure 17. Data samples extracted from the Johns Hopkins Turbulence Database with a spatial or temporal interval of ten using one of the three velocity components. From top to bottom: *mhd1024* and *isotropic1024coarse* (varied time step), *isotropic4096* and *rotstrat4096* (varied *z*-position), *channel* and *channel5200* (varied *x*-position).




Figure 18. Examples of the processed *WeatherBench* data: high-res temperature data *1.40625deg/temperature* (upper two rows) and low-res geopotential data *5.625deg/geopotential_500* (lower two rows). The temporal interval spacing between the images is twenty hours.

lent flow fields created with direct numerical simulations (DNS). Here, we used three forced isotropic turbulence data sets with different resolutions (*isotropic1024coarse*, *isotropic1024fine*, and *isotropic4096*), two channel flows with different Reynolds numbers (*channel* and *channel-5200*), the forced magneto-hydrodynamic isotropic turbulence data set (*mhd1024*), and the rotating stratified turbulence data set (*rotstrat4096*).

For the evaluation, five 256×256 reference slices in the x/y -plane from each of the seven data sets are used. The spatial and temporal position of each slice is randomized within the bounds of the corresponding simulation domain. We normalize the value range and split the velocity vectors by component for an individual evaluation. Variants for each reference are created by gradually varying the x - and z -position of the slice in equal intervals. The temporal position of each slice is varied as well if a sufficient amount of temporally resolved data is available (for *isotropic1024coarse*, *isotropic1024fine*, *channel*, and *mhd1024*). This leads to 216 sequences in total. Fig. 17 shows examples from six of the *JHTDB* data sets.

E.3. WeatherBench

The *WeatherBench* repository (Rasp et al., 2020) represents a collection of various weather measurements of different atmospheric quantities such as precipitation, cloud coverage, wind velocities, geopotential, and temperature. The data ranges from 1979 to 2018 with a fine temporal resolution and is stored on a Cartesian latitude-longitude grid of the earth. In certain subsets of the data, an additional dimension such as altitude or pressure levels is available. As all measurements are available as scalar fields, only a linear interpolation to the correct input size and a normalization was necessary in order to prepare the data. We used the low-

resolution geopotential data set at 500hPa (i.e., at around 5.5km height) with a size of 32×64 yielding smoothly changing features when upsampling the data. In addition, the high-res temperature data with a size of 128×256 for small scale details was used. For the temperature field, we used the middle atmospheric pressure level at 850hPa corresponding to an altitude of 1.5km in our experiments.

To create sequences with variations for a single time step of the weather data, we used frames in equal time intervals, similar to the *ScalarFlow* data. Due to the very fine temporal discretization of the data, we only use a temporal interval of two hours as the smallest interval step of one in Fig. 19. We sampled three random starting points in time from each of the 40 years of measurements, resulting in 120 temperature and geopotential sequences overall. Fig. 18 shows a collection of example sequences.

E.4. Detailed Results

For each of the variants explained in the previous sections, we create test sets with six different spatial and temporal intervals. Fig. 19 shows the combined Spearman correlation of the sequences for different interval spacings when evaluating various metrics. For the results in Fig. 7 in the main paper, all correlation values shown here are aggregated by data source via mean and standard deviation.

While our metric reliably recovers the increasing distances within the data sets, the individual measurements exhibit interesting differences in terms of their behavior for varying distances. As *JHTDB* and *WeatherBench* contain relatively uniform phenomena, a larger step interval creates more difficult data as the simulated and measured states contain changes that are more and more difficult to analyze along a sequence. For *ScalarFlow*, on the other hand, the diffi-




Figure 19. Detailed breakdown of the results when evaluating $LSiM$ on the individual data sets of *ScalarFlow* (30 sequences each), *JHTDB* (90 sequences each), and *WeatherBench* (120 sequences each) with different step intervals.

culty decreases for larger intervals due to the large-scale motion of the reconstructed plumes. As a result of buoyancy forces, the observed smoke rises upwards into areas where no smoke has been before. For the network, this makes predictions relatively easy as the large-scale translations are indicative of the temporal progression, and small scale turbulence effects can be largely ignored. For this data set, smaller intervals are more difficult as the overall shape of the plume barely changes while the complex evolution of small scale features becomes more important.

Overall, the $LSiM$ metric recovers the ground truth ordering of the sequences very well as indicated by the consistently high correlation values in Fig. 19. The other metrics come close to these results on certain sub-datasets but are significantly less consistent. $SSiM$ struggles on *JHTDB* across all interval sizes, and $LPIPS$ cannot keep up on *WeatherBench*, especially for larger intervals. L^2 is more stable overall, but consistently stays below the correlation achieved by $LSiM$.

F. Additional Evaluations

In the following, we demonstrate other ways to compare the performance of the analyzed metrics on our data sets. In

Tab. 1, the Pearson correlation coefficient is used instead of Spearman’s rank correlation coefficient. While Spearman’s correlation measures monotonic relationships by using ranking variables, it directly measures linear relationships.

The results in Tab. 1 match very closely to the values computed with Spearman’s rank correlation coefficient. The best performing metrics in both tables are identical; only the numbers slightly vary. Since a linear and a monotonic relation describes the results of the metrics similarly well, there are no apparent non-linear dependencies that cannot be captured using the Pearson correlation.

In the Tables 2 and 3, we employ a different, more intuitive approach to determine combined correlation values for each data set using the Pearson correlation. We are no longer analyzing the entire predicted distance distribution and the ground truth distribution at once as done above. Instead, we individually compute the correlation between the ground truth and the predicted distances for the single data samples of the data set. From the single correlation values, we compute the mean and standard deviations shown in the tables. Note that this approach potentially produces less accurate comparison results, as small errors in the individual

Table 1. Performance comparison on validation and test data sets measured in terms of the Pearson correlation coefficient of ground truth against predicted distances. **Bold+underlined** values show the best performing metric for each data set, **bold** values are within a 0.01 error margin of the best performing, and *italic* values are 0.2 or more below the best performing. On the right a visualization of the combined test data results is shown for selected models.

Metric	Validation data sets				Test data sets					
	Smo	Liq	Adv	Bur	TID	LiqN	AdvD	Sha	Vid	All
L^2	0.66	0.80	0.72	0.60	0.82	0.73	0.55	<i>0.66</i>	0.79	0.60
SSIM	0.69	0.74	0.76	0.70	0.78	<i>0.26</i>	0.69	<i>0.49</i>	0.73	0.53
LPIPS v0.1	0.63	0.68	0.66	0.71	0.85	<i>0.49</i>	0.61	0.84	0.83	0.65
AlexNet _{random}	0.63	0.69	0.67	0.65	0.83	0.64	0.63	0.74	0.81	0.65
AlexNet _{frozen}	0.66	0.69	0.68	0.71	0.85	<i>0.39</i>	0.61	0.86	0.83	0.64
Optical flow	0.63	<i>0.56</i>	<i>0.37</i>	<i>0.39</i>	<i>0.49</i>	<i>0.45</i>	<i>0.28</i>	<i>0.61</i>	0.74	<i>0.48</i>
Non-Siamese	0.77	0.84	0.78	0.74	0.67	0.81	0.64	<i>0.27</i>	0.79	0.60
Skip _{from scratch}	0.79	0.83	0.80	0.73	0.85	0.78	0.61	0.79	0.84	0.71
LSiM _{noiseless}	0.77	0.77	0.76	0.72	0.86	0.62	0.58	0.84	0.83	0.68
LSiM _{strong noise}	0.65	<i>0.64</i>	0.66	0.68	0.81	<i>0.39</i>	0.53	0.90	0.82	0.64
LSiM (ours)	0.78	0.82	0.79	0.74	0.86	0.79	0.58	0.87	0.82	0.72

Table 2. Performance comparison on validation data sets measured by computing mean and standard deviation (in brackets) of Pearson correlation coefficients (ground truth against predicted distances) from individual data samples. **Bold+underlined** values show the best performing metric for each data set, **bold** values are within a 0.01 error margin of the best performing, and *italic* values are 0.2 or more below the best performing. On the right a visualization of the combined test data results is shown for selected models.

Metric	Validation data sets			
	Smo	Liq	Adv	Bur
L^2	0.68 (0.27)	0.82 (0.18)	0.74 (0.24)	0.63 (0.33)
SSIM	0.71 (0.23)	0.75 (0.23)	0.79 (0.21)	0.73 (0.33)
LPIPS v0.1	0.66 (0.29)	0.71 (0.24)	0.70 (0.29)	0.75 (0.28)
AlexNet _{random}	0.65 (0.28)	0.71 (0.29)	0.71 (0.27)	0.68 (0.31)
AlexNet _{frozen}	0.69 (0.27)	0.72 (0.25)	0.71 (0.27)	0.74 (0.29)
Optical flow	0.66 (0.38)	<i>0.59 (0.47)</i>	<i>0.38 (0.52)</i>	<i>0.41 (0.49)</i>
Non-Siamese	0.80 (0.19)	0.87 (0.14)	0.81 (0.20)	0.76 (0.32)
Skip _{from scratch}	0.81 (0.19)	0.85 (0.16)	0.82 (0.19)	0.77 (0.30)
LSiM _{noiseless}	0.79 (0.21)	0.79 (0.20)	0.79 (0.23)	0.76 (0.29)
LSiM _{strong noise}	0.67 (0.28)	<i>0.66 (0.29)</i>	0.68 (0.30)	0.70 (0.32)
LSiM (ours)	0.81 (0.20)	0.84 (0.16)	0.81 (0.19)	0.78 (0.28)

Table 3. Performance comparison on test data sets measured by computing mean and std. dev. (in brackets) of Pearson correlation coefficients (ground truth against predicted distances) from individual data samples. **Bold+underlined** values show the best performing metric for each data set, **bold** values are within a 0.01 error margin of the best performing, and *italic* values are 0.2 or more below the best performing.

Metric	Test data sets					
	TID	LiqN	AdvD	Sha	Vid	All
L^2	0.84 (0.08)	0.75 (0.18)	0.57 (0.38)	<i>0.67 (0.18)</i>	0.84 (0.27)	0.69 (0.29)
SSIM	0.81 (0.20)	<i>0.26 (0.38)</i>	0.71 (0.31)	<i>0.53 (0.32)</i>	0.77 (0.28)	0.58 (0.38)
LPIPS v0.1	0.87 (0.11)	<i>0.51 (0.34)</i>	0.63 (0.34)	0.85 (0.14)	0.87 (0.22)	0.71 (0.31)
AlexNet _{random}	0.84 (0.10)	0.67 (0.24)	0.65 (0.33)	0.74 (0.18)	0.85 (0.26)	0.72 (0.28)
AlexNet _{frozen}	0.86 (0.11)	<i>0.41 (0.37)</i>	0.64 (0.34)	0.87 (0.14)	0.87 (0.22)	0.70 (0.34)
Optical flow	0.74 (0.67)	<i>0.50 (0.34)</i>	<i>0.32 (0.53)</i>	<i>0.63 (0.45)</i>	0.78 (0.45)	<i>0.53 (0.49)</i>
Non-Siamese	0.87 (0.12)	0.84 (0.12)	0.66 (0.34)	<i>0.31 (0.45)</i>	0.83 (0.26)	0.64 (0.39)
Skip _{from scratch}	0.87 (0.12)	0.80 (0.16)	0.63 (0.37)	0.80 (0.17)	0.87 (0.20)	0.76 (0.27)
LSiM _{noiseless}	0.87 (0.11)	<i>0.64 (0.29)</i>	0.60 (0.38)	0.86 (0.15)	0.86 (0.22)	0.73 (0.31)
LSiM _{strong noise}	0.83 (0.12)	<i>0.39 (0.38)</i>	0.55 (0.36)	0.91 (0.17)	0.86 (0.25)	0.67 (0.37)
LSiM (ours)	0.88 (0.10)	0.81 (0.15)	0.60 (0.37)	0.88 (0.16)	0.85 (0.23)	0.77 (0.28)

computations can accumulate to larger deviations in mean and standard deviation. Still, both tables lead to very similar conclusions: The best performing metrics are almost the same, and low combined correlation values match with results that have a high standard deviation and a low mean.

Fig. 20 shows a visualization of predicted distances c against ground truth distances d for different metrics on every sample from the test sets. Each plot contains over 6700 individual data points to illustrate the global distance distributions created by the metrics, without focusing on single cases. A theoretical optimal metric would recover a perfectly narrow distribution along the line $c = d$, while worse metrics recover broader, more curved distributions. Overall, the sample distribution of an L^2 metric is very wide. $LPIPS$ manages to follow the optimal diagonal a lot better, but our approach approximates it with the smallest deviations, as also shown in the tables above. The L^2 metric performs very poorly on the shape data indicated by the too steeply increasing blue lines that flatten after a ground truth distance of 0.3. $LPIPS$ already significantly reduces this problem, but $LSiM$ still works slightly better.

A similar issue is visible for the Advection-Diffusion data, where for L^2 a larger number of red samples is below the

optimal $c = d$ line, than for the other metrics. $LPIPS$ has the worst overall performance for liquid test set, indicated by the large number of fairly chaotic green lines in the plot. On the video data, all three metrics perform similarly well.

A fine-grained distance evaluation in 200 steps of L^2 and our $LSiM$ metric via the mean and standard deviation of different data samples is shown in Fig. 21. Similar to Fig. 20, the mean of an optimal metric would follow the ground truth line with a standard deviation of zero, while the mean of worse metrics deviates around the line with a high standard deviation. The plot on the left combines eight samples with different seeds from the *Sha* data set, where only a single shape is used. Similarly, the center plot aggregates eight samples from *Sha* with more than one shape. The right plot shows six data samples from the *LiqN* test set that vary by the amount of noise that was injected into the simulation.

The task of only tracking a single shape in the example on the left is the easiest of the three shown cases. Both metrics have no problem to recover the position change until a variation of 0.4, where L^2 can no longer distinguish between the different samples. Our metric recovers distances with a continuously rising mean and a very low standard deviation. The task in the middle is already harder, as multiple shapes




Figure 20. Distribution evaluation of ground truth distances against normalized predicted distances for L^2 , $LPIPS$ and $LSiM$ on all test data (color coded).

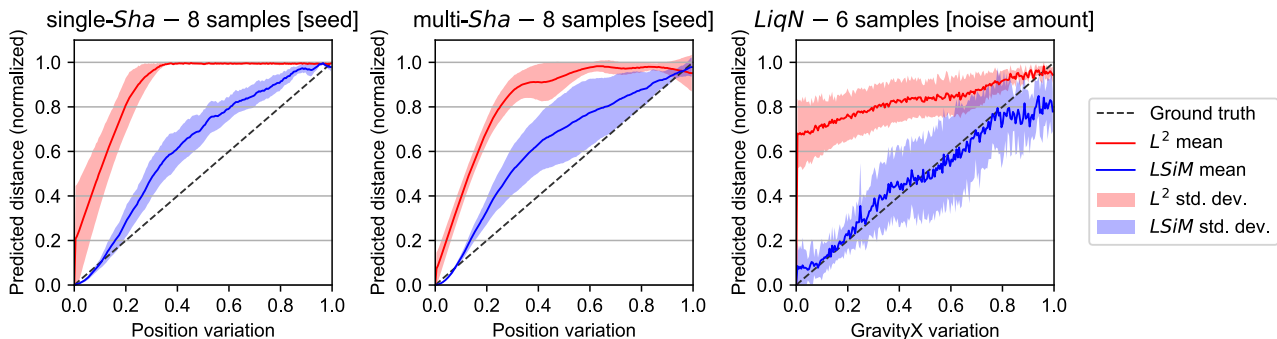


Figure 21. Mean and standard deviation of normalized distances over multiple data samples for L^2 and $LSiM$. The samples differ by the quantity displayed in brackets. Each data sample uses 200 parameter variation steps instead of 10 like the others in our data sets. For the shape data the position of the shape varies and for the liquid data the gravity in x-direction is adjusted.

can occlude each other during the position changes. Starting at a position variation of 0.4, both metrics have a quite high standard deviation, but the proposed method stays closer to the ground truth line. L^2 shows a similar issue as before because it flattens relatively fast. The plot on the right features the hardest task. Here, both metrics perform similar as each has a different problem in addition to an unstable mean. Our metric stays close to the ground truth, but has a quite high standard deviation starting at about a variation of 0.4. The standard deviation of L^2 is lower, but instead it starts off with a big jump from the first few data points. To some degree, this is caused by the normalization of the plots, but it still overestimates the relative distances for small variations in the simulation parameter.

These findings also match with the distance distribution evaluations in Fig. 20 and the tables above: Our method has a significant advantage over shallow metrics on shape data, while the differences of both metrics become much smaller for the liquid test set.

G. Notation

In this work, we follow the notation suggested by Goodfellow et al.. Vector quantities are displayed in bold, and tensors use a sans-serif font. Double-barred letters indicate sets or vector spaces. The following symbols are used:

\mathbb{R}	Real numbers	o_0, o_1, \dots, o_n	Series of outputs of a simulation with increasing ground truth distance to o_0
i, j	Indexing in different contexts	Δ	Amount of change in a single simulation parameter
\mathbb{I}	Input space of the metric, i.e., color images/field data of size $224 \times 224 \times 3$	t_1, t_2, \dots, t_t	Time steps of a numerical simulation
a	Dimension of the input space \mathbb{I} when flattened to a single vector	v	Variance of the noise added to a simulation
$\mathbf{x}, \mathbf{y}, \mathbf{z}$	Elements in the input space \mathbb{I}	\mathbf{c}	Ground truth distance distribution, determined by the data generation via Δ
\mathbb{L}	Latent space of the metric, i.e., sets of 3 rd order feature map tensors	\mathbf{d}	Predicted distance distribution (supposed to match the corresponding \mathbf{c})
b	Dimension of the latent space \mathbb{L} when flattened to a single vector	$\bar{\mathbf{c}}, \bar{\mathbf{d}}$	Mean of the distributions \mathbf{c} and \mathbf{d}
$\tilde{\mathbf{x}}, \tilde{\mathbf{y}}, \tilde{\mathbf{z}}$	Elements in the latent space \mathbb{L} , corresponding to $\mathbf{x}, \mathbf{y}, \mathbf{z}$	$\ \cdot\cdot\cdot\ _2$	Euclidean norm of a vector
w	Weights for the learned average aggregation (1 per feature map)	$m(\mathbf{x}, \mathbf{y})$	Entire function computed by our metric
p_0, p_1, \dots	Initial conditions / parameters of a numerical simulation	$m_1(\mathbf{x}, \mathbf{y})$	First part of $m(\mathbf{x}, \mathbf{y})$, i.e., base network and feature map normalization
n	Number of variations of a simulation parameter, thus determines length of the network input sequence	$m_2(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$	Second part of $m(\mathbf{x}, \mathbf{y})$, i.e., latent space difference and the aggregations
		\mathbf{G}	3 rd order feature tensor from one layer of the base network
		g_b, g_c, g_x, g_y	Batch (g_b), channel (g_c), and spatial dimensions (g_x, g_y) of \mathbf{G}
		f	Optical flow network
		f^{xy}, f^{yx}	Flow fields computed by an optical flow network f from two inputs in \mathbb{I}
		f_1^{xy}, f_2^{xy}	Components of the flow field f^{xy}
		∇, ∇^2	Gradient (∇) and Laplace operator (∇^2)
		∂	Partial derivative operator
		t	Time in our PDEs
		u	Velocity in our PDEs
		ν	Kinematic viscosity / diffusion coefficient in our PDEs
		d, ρ	Density in our PDEs
		P	Pressure in the Navier-Stokes Equations
		g	Gravity in the Navier-Stokes Equations

References

- Eckert, M.-L., Um, K., and Thuerey, N. Scalarflow: A large-scale volumetric data set of real-world scalar transport flows for computer animation and machine learning. *ACM Transactions on Graphics*, 38(6), 2019. doi:10.1145/3355089.3356545.
- Goodfellow, I., Bengio, Y., and Courville, A. *Deep Learning*. MIT Press, 2016. URL <http://www.deeplearningbook.org>.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016. doi:10.1109/CVPR.2016.90.
- Huang, G., Liu, Z., Van Der Maaten, L., and Weinberger, K. Q. Densely connected convolutional networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2261–2269, 2017. doi:10.1109/CVPR.2017.243.
- Ilg, E., Mayer, N., Saikia, T., Keuper, M., Dosovitskiy, A., and Brox, T. FlowNet 2.0: Evolution of optical flow estimation with deep networks. *CoRR*, abs/1612.01925, 2016. URL <http://arxiv.org/abs/1612.01925>.
- Perlman, E., Burns, R., Li, Y., and Meneveau, C. Data exploration of turbulence simulations using a database cluster. In *SC '07: Proceedings of the 2007 ACM/IEEE Conference on Supercomputing*, pp. 1–11, 2007. doi:10.1145/1362622.1362654.
- Ponomarenko, N., Jin, L., Ieremeiev, O., Lukin, V., Egiazarian, K., Astola, J., Vozel, B., Chehdi, K., Carli, M., Battisti, F., and Kuo, C. C. J. Image database TID2013: Peculiarities, results and perspectives. *Signal Processing-Image Communication*, 30:57–77, 2015. doi:10.1016/j.image.2014.10.009.
- Rasp, S., Dueben, P., Scher, S., Weyn, J., Mouatadid, S., and Thuerey, N. Weatherbench: A benchmark dataset for data-driven weather forecasting. *CoRR*, abs/2002.00469, 2020. URL <http://arxiv.org/abs/2002.00469>.
- Ronneberger, O., Fischer, P., and Brox, T. U-net: Convolutional networks for biomedical image segmentation. *CoRR*, abs/1505.04597, 2015. URL <http://arxiv.org/abs/1505.04597>.
- Zhang, R., Isola, P., Efros, A. A., Shechtman, E., and Wang, O. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *2018 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 586–595, 2018. doi:10.1109/CVPR.2018.00068.
- Zhu, Y. and Bridson, R. Animating sand as a fluid. In *ACM SIGGRAPH 2005 Papers*, pp. 965–972, New York, NY, USA, 2005. doi:10.1145/1186822.1073298.