

# Fully differentiable Lagrangian convolutional neural network for physics-informed precipitation nowcasting

Peter Pavlík<sup>a,b</sup>,<sup>\*</sup>, Martin Výboh<sup>a</sup>, Anna Bou Ezzeddine<sup>a</sup>, Viera Rozinajová<sup>a</sup>

<sup>a</sup> *Kempelen Institute of Intelligent Technologies, Bottova 7939/2A, Bratislava, 811 09, Slovakia*

<sup>b</sup> *Faculty of Information Technology, Brno University of Technology, Božetěchova 2, Brno, 612 00, Czech Republic*

## ARTICLE INFO

Dataset link: [Maly Javornik Radar \(Original data\)](#)

### Keywords:

Rain  
Forecasting  
Neural networks  
Meteorological radar

## ABSTRACT

This paper presents a convolutional neural network model for precipitation nowcasting that combines data-driven learning with physics-informed domain knowledge. We propose LUPIN, a Lagrangian Double U-Net for Physics-Informed Nowcasting, that draws from existing extrapolation-based nowcasting methods. It consists of a U-Net that dynamically produces mesoscale advection motion fields, a differentiable semi-Lagrangian extrapolation operator, and an advection-free U-Net capturing the growth and decay of precipitation over time. Using our approach, we successfully implement the Lagrangian convolutional neural network for precipitation nowcasting in a fully differentiable and GPU-accelerated manner. This allows for end-to-end training and inference, including the data-driven Lagrangian coordinate system transformation of the data at runtime. We evaluate the model and compare it with other related AI-based models both quantitatively and qualitatively in an extreme event case study. Based on our evaluation, LUPIN matches and even exceeds the performance of the chosen benchmarks, opening the door for other Lagrangian machine learning models.

## 1. Introduction

Precipitation nowcasting is the task of predicting the future intensity of precipitation (rainfall, snowfall, hail, etc.) in the near future with high local detail (WMO, 2017). The nowcasts are valuable for early warning about extreme events and can help mitigate damage and loss of lives. Accurate precipitation nowcasting can therefore increase our resilience to the climate crisis, particularly in the face of increased precipitation variability (Zhang et al., 2024). Radar observations from ground-based stations provide reliable approximations of the current rainfall rate and are therefore the primary data source for precipitation nowcasting. While alternative sources such as satellite sensors or outputs from numerical weather prediction models are sometimes employed (Ionescu et al., 2021; Wang et al., 2025; Hachemi et al., 2023), ground radar remains the standard due to its high spatial and temporal resolution.

Traditionally, the nowcasting is performed using radar extrapolation methods based on the assumption that precipitation at the mesoscale behaves and advects according to the continuity equation. A few past precipitation observations are used to estimate an advection motion field and the last observation is extrapolated forward in time according to this field. This is the basis of many operational methods.

However, in the last few years, convolutional neural networks (CNNs) have become the de facto state-of-the-art models for precipitation nowcasting. The first application of a CNN-based architecture in the domain was the ConvLSTM model (Shi et al., 2015). It leveraged that the problem can be posed as a spatiotemporal sequence forecasting problem similar to next-frame video prediction and therefore is a good fit for processing using CNNs.

One obvious problem with this architecture is that the convolutional recurrence structure in ConvLSTM is location-invariant, therefore not explicitly taking into consideration the advection motion of precipitation over time. This led to the development of the Trajectory GRU (TrajGRU) model (Shi et al., 2017) that can learn the location-variant structure for recurrent connections, therefore internally considering the precipitation advection. It was shown that the TrajGRU is more efficient in capturing the spatiotemporal correlations in the data.

Despite the TrajGRU model seemingly combining the best of both worlds – the physics-informed location-variant structure with data-driven CNN-based processing – location-invariant architectures continued to be developed and refined. In Ayzel et al. (2020), a U-Net architecture was successfully applied for precipitation nowcasting to create RainNet. Despite being a purely data-driven model without

<sup>\*</sup> Corresponding author at: Kempelen Institute of Intelligent Technologies, Bottova 7939/2A, Bratislava, 811 09, Slovakia.

E-mail addresses: [peter.pavlik@kinit.sk](mailto:peter.pavlik@kinit.sk) (P. Pavlík), [martin.vyboh@kinit.sk](mailto:martin.vyboh@kinit.sk) (M. Výboh), [anna.bou.ezzeddine@kinit.sk](mailto:anna.bou.ezzeddine@kinit.sk) (A.B. Ezzeddine), [viera.rozinajova@kinit.sk](mailto:viera.rozinajova@kinit.sk) (V. Rozinajová).

<https://doi.org/10.1016/j.acags.2025.100296>

Received 4 February 2025; Received in revised form 26 September 2025; Accepted 29 September 2025

Available online 1 October 2025

2590-1974/© 2025 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

explicit inclusion of domain knowledge, it performs well and is often viewed as a simple benchmark CNN-based model.

The main challenge for all the mentioned data-driven nowcasting models is the growing smoothness and fading out of the nowcast quantities with growing time horizons. Due to the highly chaotic behaviour of the precipitation in the atmosphere, there is a high degree of uncertainty in the predictions. When training with traditional gridpoint-based loss functions, we face the so-called “double penalty problem”. A forecast of a precipitation feature that is correct in terms of intensity, size, and timing, but incorrect concerning location, results in a very large error (Keil and Craig, 2009). This leads to “mathematically optimal” predictions being visibly blurry and physically inconsistent, limiting their usefulness as nowcasts for early warning of extreme events, as those especially are smoothed out.

In an attempt to solve the blurriness challenge, some models introduced a GAN-based framework to produce realistically looking nowcasts (Tian et al., 2019; Wang et al., 2021). A notable example is the Deep Generative Model of Radar (DGMR) model (Ravuri et al., 2021) that is able to produce accurate and realistic nowcasts up to 90 min ahead. A visual verification by meteorologists was used, and it was selected as the best in most cases by a large margin out of all the compared nowcasting models at the time.

Recently, the GAN approach was combined with the explicit consideration of advection motion fields in NowcastNet (Zhang et al., 2023). It incorporates a differentiable evolution network that generates a motion field and performs a deterministic evolution of the precipitation according to the continuity equation before passing it to the main generative model. This helped the model outperform the DGMR by a large margin in expert evaluation at up to 3-hour ahead lead times.

Another recent model that explicitly takes into consideration the advection of the precipitation is the Lagrangian CNN (L-CNN) introduced in Ritvanen et al. (2023). Here, the advection motion field obtained from an optical flow algorithm is used to transform the input into the Lagrangian coordinate system, where a U-Net processes the temporally differenced advection-free inputs, thus separating the learning of rainfall growth and decay from mesoscale motion. The model outperformed a benchmark U-Net-based RainNet, showing that incorporating insights from rainfall physics can lead to substantial improvements in machine learning nowcasting methods. However, due to the motion field being a product of a non-differentiable optical flow method, the model cannot be trained using a fully end-to-end approach, constraining its utility.

In this paper, we present a new model that takes inspiration from the differentiable evolution network of NowcastNet and applies it to the ingenious Lagrangian advection-free approach of L-CNN to form a single end-to-end CNN-based model. We implement a differentiable semi-Lagrangian extrapolation that allows us to train the model to produce an advection motion field and dynamically map the inputs into the Lagrangian coordinates to provide a better insight for the subsequent learning of precipitation evolution by performing the Lagrangian coordinate transformation and temporal differencing during runtime. We present LUPIN: A Lagrangian Double U-Net for Physics-Informed Nowcasting.

## 2. The Lagrangian view

To understand the theoretical underpinnings of our approach, we first present the assumptions and methods on which the traditional nowcasting methods are based. The formalism is mostly based on the notations of Germann and Zawadzki (2002) and Ritvanen et al. (2023).

### 2.1. Advection equation and Lagrangian persistence

The precipitation in the near future is highly dependent on the present weather. The basis of all traditional nowcasting methods is therefore the assumption of persistence — the precipitation in the near future will stay roughly the same as observed in the present.

In the Eulerian coordinates (where we look at fluid flow in relation to the space that it moves through), this means the simplest produced nowcast is an observation identical to the last one. However, the performance of this Eulerian persistence “model” falls rather quickly with increasing lead times as the current precipitation moves elsewhere.

The assumption of persistence is more realistic under the Lagrangian specification of the flow field — where we follow fluid parcels as they move through space and time. Lagrangian persistence assumes the overall shape of the precipitation will stay the same, but accounts for the displacement of the precipitation according to the advection equation. The advection equation can be expressed mathematically as follows:

$$\frac{\partial \Psi}{\partial t} + \nabla \cdot (\Psi \mathbf{u}) = 0 \quad (1)$$

where  $\Psi$  denotes the 2D radar precipitation field and  $\mathbf{u}$  the horizontal motion field. It is often assumed that the flow is incompressible. Therefore, the divergence of the motion field is zero ( $\mathbf{u}$  is solenoidal):

$$\nabla \cdot \mathbf{u} = 0 \quad (2)$$

Under this zero divergence assumption, the advection Eq. (1) can be written in the form:

$$\frac{\partial \Psi}{\partial t} + \mathbf{u} \cdot \nabla \Psi = 0 \quad (3)$$

We can transform the precipitation field  $\Psi$  to a Lagrangian coordinate system precipitation field  $\tilde{\Psi}$  by extrapolating it according to the motion field  $\mathbf{u}$  so that  $\mathbf{u} \cdot \nabla \tilde{\Psi} = 0$  holds and the advection equation can finally be written as:

$$\frac{\partial \tilde{\Psi}}{\partial t} = 0 \quad (4)$$

The advection equation now clearly denotes the Lagrangian persistence assumption, expressing that the precipitation field in Lagrangian coordinates will not change over time.

### 2.2. Implementation of the Lagrangian transformation

To transform the input 2D precipitation fields  $\Psi_{1,2,\dots,n}$  to the Lagrangian coordinates  $\tilde{\Psi}_{1,2,\dots,n}$ , all the inputs need to be extrapolated to a single reference field. We define an extrapolation operator  $\xi^t(\Psi, \mathbf{u})$  as an operation of extrapolating a precipitation field  $\Psi$  by  $t$  time steps forward in time along the motion field  $\mathbf{u}$ .

Semi-Lagrangian extrapolation (Sawyer, 1963) is commonly used as an implementation of this operator. For estimating the motion field  $\mathbf{u}$ , various optical flow methods are available, the Lucas-Kanade algorithm (Lucas and Kanade, 1981; Bouguet et al., 2001) being a popular one. Both of these are implemented in the Python language in the pySTEPS nowcasting library (Pulkkinen et al., 2019).

Formally, to transform  $n$  successive input precipitation fields to Lagrangian coordinates (choosing the last one as reference), the extrapolation operator must be applied to each input image the number of times equal to the time step difference between it and the reference image:

$$\tilde{\Psi}_i = \xi^{n-i}(\Psi_i, \mathbf{u}), \quad i = 1, 2, \dots, n-1 \quad (5)$$

where  $n$  is the total number of successive input precipitation fields (and the index of the last one used as reference). After this transformation, we can estimate the actual temporal difference of precipitation field from Eq. (4) at a given time  $t$  as:

$$\frac{\partial \tilde{\Psi}}{\partial t} \approx \Delta \tilde{\Psi}_t = \tilde{\Psi}_t - \tilde{\Psi}_{t-1} \quad (6)$$

### 2.3. The Lagrangian nowcasting

For a more realistic nowcasting of the rapidly-changing evolution of precipitation over time, we need to introduce one more term to the advection equation – a source–sink term  $S_\psi$  which allows for the refinement of the Lagrangian persistence assumption:

$$\frac{\partial \tilde{\Psi}}{\partial t} = S_\psi \quad (7)$$

where the source–sink term  $S_\psi$  represents the changes in the radar precipitation field over time without considering its displacement according to the advection motion field. It can therefore be considered the advection-free change of the field over time. However, predicting the source–sink term  $S_\psi$  and therefore the growth and decay of precipitation systems is very difficult. It is the biggest challenge for the nowcasting systems. The state-of-the-art extrapolation nowcasting methods such as S-PROG (Seed, 2003), LINDA (Pulkkinen et al., 2021), STEPS (Bowler et al., 2006) and ANVIL (Pulkkinen et al., 2020) include an autoregressive model to estimate the source–sink term  $S_\psi$  based either on the input precipitation fields in Lagrangian coordinates  $\tilde{\Psi}$  or the estimated time derivatives  $\Delta \tilde{\Psi}$ .

### 2.4. Incorporating convolutional neural networks

Analogously to the autoregressive approaches, some CNN-based models have taken the Lagrangian approach to estimate the source–sink term  $S_\psi$ .

In the case of TrajGRU (Shi et al., 2017), the model implicitly learns the flow field  $u$  via the dynamically determined neural connections, therefore learning to perform the right extrapolation before estimating the source–sink term.

The L-CNN model (Ritvanen et al., 2023) takes a different approach, transforming the precipitation fields to Lagrangian coordinates before the training of the model (based on Lucas-Kanade advection motion field) and time differencing the transformed fields. Then they train a U-Net to predict the future temporal differences  $\Delta \tilde{\Psi}/S_\psi$ .

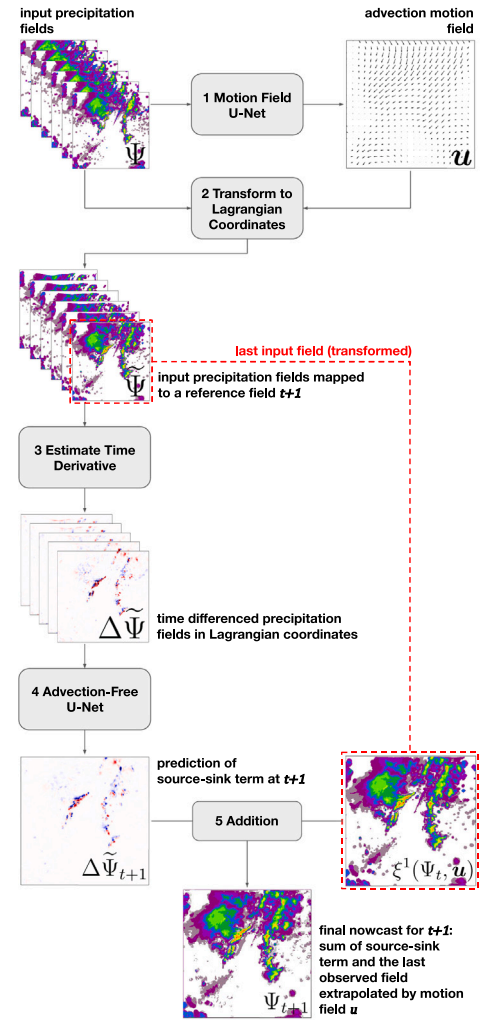
Another model worth mentioning here is the NowcastNet (Zhang et al., 2023), specifically its Evolution Network. While it does not transform the precipitation maps to Lagrangian coordinates, it learns to separate the prediction of the advection motion field from the prediction of changes in precipitation field intensity, analogous to the source–sink term  $S_\psi$ . Also, it introduces a motion-regularization term loosely based on the continuity (advection) Eq. (1) to enforce a smoother motion field. In contrast to L-CNN, the differentiable evolution operator allows this model to be trainable in an end-to-end manner.

### 3. LUPIN – Lagrangian double U-Net for Physics-Informed Nowcasting

This section is dedicated to introducing and describing the architecture and training process of our precipitation nowcasting model, LUPIN. We compare it to previous approaches, explain the differences, and elaborate on the thought process behind them. The main point of comparison will be the Lagrangian CNN (L-CNN) model from Ritvanen et al. (2023), as it served as the main inspiration for our approach, whose shortcomings we aim to address.

#### 3.1. LUPIN architecture

The overall architecture of the LUPIN model is based on the architecture of the L-CNN model (Ritvanen et al., 2023). We took the key idea of the L-CNN – to transform the input data to Lagrangian coordinates and use a CNN to model the time derivatives of rainfall fields (the source–sink term  $S_\psi = \Delta \tilde{\Psi}_{t+1}$ ) – and decided to explore implementing this approach as a fully data-driven model.

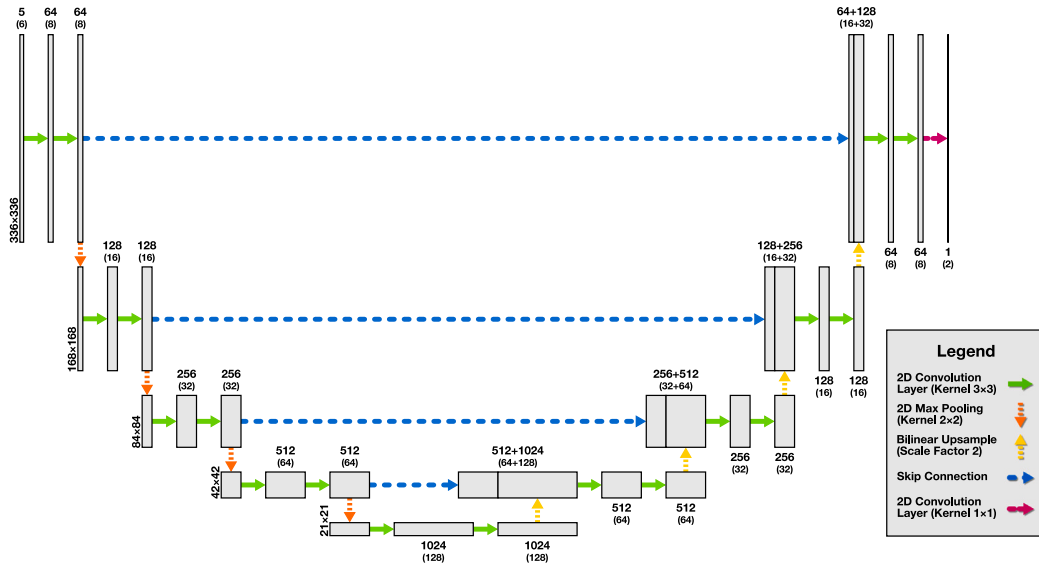


**Fig. 1.** Architecture of the LUPIN model and the inference process of a nowcast one time step ahead. The grey processing blocks are marked 1 to 5 in the order they are called during inference. **1** The first step is the calculation of the advection motion field  $u$ . **2** Then, the  $n$  input fields  $\Psi$  can be transformed into Lagrangian coordinates as described in Eq. (5). **3** The precipitation maps in Lagrangian coordinates are differenced over time to obtain  $n - 1$  time differenced fields  $\Delta \tilde{\Psi}$ . **4** The differenced fields are used as an input to the Advection-Free U-Net that predicts the next differenced field  $\Delta \tilde{\Psi}_{t+1}$  (representing the source–sink term  $S_\psi$  at time  $t + 1$ ). **5** The predicted source–sink term is added to the last input field (already transformed in step 2 to reference field  $t + 1$ ), the product being the final nowcast for the next time step.

The inclusion of the optical flow algorithm in the L-CNN model makes the training process cumbersome. The algorithm is orders of magnitude slower compared to the time needed for GPU-accelerated inference and backpropagation in a CNN. This makes the training feasible only with inputs and outputs transformed to Lagrangian coordinates beforehand. At the same time, the longer inference limits the application of the model for ensemble nowcasting.

Our solution was to implement a differentiable semi-Lagrangian-like extrapolation operation in the PyTorch machine learning framework (Paszke et al., 2019). This allowed us to create a model that learns to produce the motion field in a data-driven manner and then performs the transformation to Lagrangian coordinates dynamically during runtime in a fraction of the time.

The overall LUPIN architecture is in the form of a double U-Net, see Fig. 1. The first U-Net (Motion Field U-Net) produces the motion fields



**Fig. 2.** The figure illustrates the U-Net convolutional neural network architecture employed in LUPIN for both motion field estimation (Step 1 in Fig. 1) and source-sink term prediction (Step 4 in Fig. 1). The figure consists of blocks connected by arrows, where block shapes represent the dimensions of the feature maps at each stage, and arrows denote the operations applied between them. The left half depicts the encoder, which consists of repeated pairs of convolutional layers followed by max pooling, progressively reducing spatial resolution while increasing the number of channels. The right half shows the decoder, which mirrors this structure using pairs of convolutional layers followed by bilinear upsampling to restore spatial resolution and decrease the number of channels. Encoder and decoder components are linked through skip connections, where the output of each encoder stage is concatenated with the input of the corresponding decoder stage. The final decoder layer reduces the number of channels to match the target output. Spatial resolutions at each stage are indicated alongside the leftmost blocks, ranging from  $336 \times 336$  down to  $21 \times 21$  pixels. The number of channels after each convolutional block is shown above or below the feature maps: larger values correspond to the more complex Advection-Free U-Net, while smaller values correspond to the less complex Motion Field U-Net.

and functionally replaces the Lukas-Kanade optical flow algorithm in the L-CNN. The second one (Advection-free U-Net) works with the Lagrangian-transformed and time differenced inputs analogous to the L-CNN specification. Both U-Nets closely follow the RainNet architecture (Ayzel et al., 2020), with the Motion Field U-Net using 8 times fewer channels per layer, as estimating motion fields does not require the same level of model complexity. See Fig. 2 for a diagram of the U-Net architecture used in the LUPIN model.

The model is to be applied iteratively in a rolling window manner. Each inference produces a new motion field and the corresponding nowcast for the next time step. The nowcast then subsequently serves as the last input observation for the next iteration. By choosing the reference time for the Lagrangian mapping to match the one-step ahead nowcast window  $t + 1$ , we have eliminated the need to perform the transformation back to Eulerian coordinates as the final prediction step, simplifying the inference process in comparison to the L-CNN.

### 3.2. Training process

Training the LUPIN is a three-stage process. While it is feasible to train the whole model at once, preliminary experiments have shown this to be unstable. We therefore used the following process:

1. Train the Motion Field U-Net (MF-U-Net) to produce satisfactory motion fields by itself.
2. Train the Advection-Free U-Net (AF-U-Net) with the MF-U-Net's weights frozen.
3. Unfreeze MF-U-Net's weights and continue the training of both U-Nets together to act jointly as a single model.

#### 3.2.1. Training MF-U-Net

Training the MF-U-Net to produce physically consistent advection motion fields  $\mathbf{u}$  is not trivial. To be usable in the Lagrangian framework, the motion fields should be stable over time when produced iteratively and locally smooth so as not to interfere with the subsequent estimation

of temporal differences in Lagrangian coordinates. We want to emulate a simple Lagrangian persistence model in a data-driven way.

The first intuitive approach was to just minimize the error between the last input observation extrapolated by one time step and the ground truth. We denote the error criterion  $C$  and the target precipitation field  $\hat{\Psi}$ :

$$\mathcal{L}_{naive} = C(\xi^1(\Psi_t, \mathbf{u}), \hat{\Psi}_{t+1}) \quad (8)$$

However, regardless of the criterion  $C$ , this loss function produces highly unrealistic motion fields full of sinks and sources, functionally interfering with the future task of the AF-U-Net. It goes against the assumption described by Eq. (2) of the divergence-free motion field for the Lagrangian persistence.

To mitigate this, we introduce two modifications to the training process. Firstly, we produce a pooled optimal motion field by minimizing the error of extrapolating the precipitation fields by one-time step over the whole input sequence, not just the last observation and the one following it:

$$\mathcal{L}_{MF} = \sum_{i=1}^n C(\xi^1(\Psi_i, \mathbf{u}), \hat{\Psi}_{i+1}) \quad (9)$$

The loss function above iterates over all  $n$  input observations and sums the error of extrapolating them one step forward in time. This pushes the model to produce more stable motion fields better fit for transforming the inputs into the Lagrangian coordinate system.

Secondly, we introduce an additional physics-informed regularization loss term to directly penalize breaking Eq. (2) – the assumption of incompressible, divergence-free flow – in the produced motion field. By enforcing it, we can push the model towards creating solutions consistent with the continuity equation derived from conservation laws. The assumption is also one of the Navier-Stokes partial differential equations that govern fluid dynamics (Batchelor, 1967). Since our data captures the movement of air in the atmosphere, it should behave according to these equations. The regularization will therefore make the motion fields more physically consistent, resulting in better generalization of the model.

The divergence  $\nabla \cdot \mathbf{u}$  is calculated using partial derivatives of the motion field:

$$\nabla \cdot \mathbf{u} = \frac{\partial u_x}{\partial x} + \frac{\partial u_y}{\partial y} \quad (10)$$

We can compute discrete approximations of motion field partial derivatives using 2D convolutions with a  $3 \times 3$  Sobel filter (Kanopoulos et al., 1988):

a)		
-1	0	1
-2	0	2
-1	0	1

b)		
-1	-2	-1
0	0	0
1	2	1

The filter (a) is used to estimate  $\frac{\partial u_x}{\partial x}$  and the filter (b) estimates  $\frac{\partial u_y}{\partial y}$ . Using these values, we define the physics-informed continuity-consistency loss function at every point of the motion field as:

$$\mathcal{L}_{PI} = |\nabla \cdot \mathbf{u}| \quad (11)$$

In the vocabulary of physics-informed machine learning, this is a soft constraint or learning bias, meaning we allow the network to break the constraint, but penalize it as it diverges (Karniadakis et al., 2021).

The final loss function that is minimized is simply the weighted sum of the data-driven and physics-informed losses:

$$\mathcal{L} = (1 - \beta)\mathcal{L}_{MF} + \beta\mathcal{L}_{PI} \quad \beta \in (0, 1) \quad (12)$$

### 3.2.2. Training AF-U-Net

The second training stage starts with taking the pretrained MF-U-Net and freezing the weights, so as not to unlearn the motion field generation from unstable gradients being backpropagated from the downstream AF-U-Net. Because the MF-U-Net has been trained to effectively perform the Lagrangian persistence nowcasting, we can train the AF-U-Net to post-process its predictions by learning to produce a source–sink term  $S_\psi$  capturing the growth and decay of precipitation fields over time.

As shown in the L-CNN paper, the Lagrangian-transformed and time-differenced fields can allow the Advection-Free U-Net to better grasp the precipitation growth and decay in the data, resulting in more accurate predictions with less blurring as lead time increases. Therefore, the input fields to our Advection-Free U-Net are also being transformed and time-differenced. The Lagrangian transformation is performed dynamically during training using our implementation of the semi-Lagrangian extrapolation.

The prediction task the Advection-Free U-Net performs is in essence auto-regressive, as the target output (source–sink term  $S_{\psi_{t+1}} = \Delta \tilde{\psi}_{t+1}$ ) is the expected difference of the last input extrapolated by one time step ( $\xi^1(\Psi_t, \mathbf{u})$ ) from the ground truth observation ( $\hat{\psi}_{t+1}$ ). To obtain the nowcast for the next time step, a simple addition of these fields is performed:

$$\Psi_{t+1} = S_{\psi_{t+1}} + \xi^1(\Psi_t, \mathbf{u}) \quad (13)$$

The error of this nowcast in relation to the observed precipitation field target forms the loss function during this training stage:

$$\mathcal{L}_{AF} = C(\Psi_{t+1}, \hat{\psi}_{t+1}) \quad (14)$$

In a way, this approach can be viewed as a form of residual learning, with the output of the AF-U-Net serving as a residual being added to a single skip-connected input channel.

### 3.2.3. Training LUPIN

By training the models only separately, an issue of out-of-distribution samples arises. The AF-U-Net produces nowcasts that get blurrier as lead time increases due to uncertainty. This means that, eventually, the outputs do not look like the input images from the train set. This poses a problem for the MF-U-Net in later iterations when it is given input data that does not look like the data it was trained on, resulting in unstable advection motion fields.

Therefore, for these models to cooperate, a final stage of training is needed. This fine-tuning stage is mostly similar to the previous one. The only change is that we unfreeze the weights of the MF-U-Net and add its loss functions to the total loss function being minimized.

In this training stage, the MF-U-Net parameters are updated not only by the gradients backpropagated from the loss function described in Eq. (12), but also the gradients from the AF-U-Net due to the differentiable implementation of the whole LUPIN model. We hypothesize that this might allow the model to produce more useful advection motion fields than the general optical flow algorithms, as the model is trained specifically to produce motion fields for the optimal Lagrangian coordinate transformation of the precipitation fields.

The total weighted loss optimized in this stage, consisting of the three separate weighted loss terms, is specified as:

$$\mathcal{L} = (1 - \beta)((1 - \gamma)\mathcal{L}_{AF} + \gamma\mathcal{L}_{MF}) + \beta\mathcal{L}_{PI} \quad \beta, \gamma \in (0, 1) \quad (15)$$

## 4. The precipitation dataset

For training the LUPIN model, we use a dataset provided by the Slovak Hydrometeorological Institute (SHMU). The dataset consists of radar reflectivity observations from the 1st of January 2016 until the 30th of June 2019, collected periodically in 5 min intervals. Data originates from the Malý Javorník weather radar station, located in the Carpathian Mountains in Western Slovakia. The dataset contains 355761 unique observations in ODIM HDF5 format.

The radar station captures precipitation by sending out radar beams and measuring the echo returned after bouncing off the precipitation particles in the atmosphere. This helps to determine the volume or type of precipitation present in the air. Intensity of the returned echo is called radar reflectivity ( $Z$ ). It is measured in logarithmic dimensionless units called decibels related to  $Z$  (dBZ). The dataset consists of reflectivity observations at multiple data points captured by radar, called gates or bins. Reflectivity gates are measured at multiple horizontal and vertical angles relative to the station, resulting in reflectivity data encoded in polar coordinates.

The data is processed from raw radar reflectivity observations to Cartesian coordinates using the Py-ART Python library (Helmus and Collis, 2016). We aggregate the 3-dimensional precipitation observations into a single two-dimensional reflectivity map by a method called CMAX. It groups data by returning the maximum value of reflectivity in the vertical column. Resulting reflectivity map is a  $340\text{km} \times 340\text{km}$  sized centre slice with  $1\text{km} \times 1\text{km}$  resolution. The radar reflectivity factor  $Z$  is converted to rainfall rate  $R$  (in mm/h) using the Marshall–Palmer formula (Marshall and Palmer, 1948), which is given by

$$Z = 200 R^{1.6}, \quad (16)$$

where the constants 200 and 1.6 are empirically derived constants.

### 4.1. Dataset filtering

To avoid training the neural network on a biased dataset dominated by non-precipitation cases, it is necessary to subset the full dataset. Including large numbers of clear-sky observations provides little to no information for learning, unnecessarily slows down training, and can even harm performance — the model may overfit to non-precipitating conditions, leading to systematic underestimation of precipitation in its forecasts.

**Table 1**

The observation counts of the full dataset, the subset selected for training according to the dataset filtering described in Section 4.1, and the sizes of train, test, and validation splits.

Set	№ obs.	Full %	Target %
Full Dataset	355761	100.00	–
Target & Lead Obs.	63007	17.71	–
Target Obs.	54721	15.38	100.00
Target Obs. Train.	35767	10.05	65.37
Target Obs. Valid.	7998	2.25	14.61
Test Target Obs.	10956	3.08	20.02

To filter the images, we convert CMAX reflectivity maps to rainfall rate, using the Marshall-Palmer formula. Then, we compute the ratio of rainy to clear pixels in the map with the threshold of 0.05mm/5min or 0.6mmh<sup>-1</sup>. Finally, if the rainfall rate map contains at least 5% of rainy pixels and 11 previous observations are available (last hour), it is added to the target observations set.

For training LUPIN, we used six input observations to predict six future observations. Consequently, predicting the target observation requires that eleven lead observations be available in the dataset.

The dataset is split into train, test, and validation subsets. The last 20% of target observations are selected for the test set. This set contains rainfall maps from the 2nd of September 2018 until the 30th of June 2019. The first 80% of the data is chosen for training, out of which 15% is split further into a validation set.

The training set is split randomly into training and validation subsets, ensuring that there are no overlapping observations in these sets (e.g., there are no target observations in the training set that are lead observations for the validation set and vice versa). The exact observation counts can be observed in Table 1.

## 5. Motion field comparison

Before we move to the evaluation of the model as a whole, it is worth comparing only the motion fields generated by the Motion Field U-Net with the aforementioned Lukas-Kanade algorithm (Lucas and Kanade, 1981) it replaces, more specifically, the pyramidal implementation available in the pySTEPS library (Pulkkinen et al., 2019) based on Bouquet et al. (2001). We will test the hypothesis that our model can produce motion fields that are better suited for the Lagrangian transformation task.

We will also use this section to compare the effect of the physics-informed regularization of the motion field we apply during training. Consequently, we will compare three models — the Lukas-Kanade algorithm, the MF-U-Net trained without the physics-informed regularization, and the one trained with it. To visually compare the motion fields from a sample extreme event from the test set, see Fig. 3 (for more information about the event, see Section 6.2).

While the general direction of all the motion fields is the same, notable differences can be observed visually, especially between the motion fields of Lukas-Kanade and the MF-U-Nets. First, it is worth mentioning that Lukas-Kanade produces uniform motion fields across the entire domain (Fig. 3(a)). In contrast, MF-U-Net generates only negligible motion vectors in areas with no precipitation to extrapolate (Figs. 3(b), 3(c)). However, this discrepancy does not impact the Lagrangian transformation as extrapolating these zero-valued areas is of no consequence.

On closer look, we can see more variability in the motion fields generated by MF-U-Nets. In the output of the non-regularized model (Fig. 3(b)), we can see divergent vectors close by in multiple areas. The regularization seems to partly eliminate this behaviour as intended (Fig. 3(c)).

### 5.1. Fitness metrics

To compare the models objectively, we will evaluate the motion fields generated by the models over the whole test set. However, choosing the right metrics to assess their fitness for Lagrangian transformation is not straightforward.

Ideally, the generated motion fields should be locally smooth, without any apparent sources and sinks. This is desirable not only to satisfy the continuity equation but, in the context of the LUPIN and L-CNN models, to preserve the precipitation field intensities during the extrapolation along the generated advection motion field, since it is the role of the following advection-free U-Net to predict the temporal differences of the precipitation field in Lagrangian coordinates. Because of this, we will use the absolute value of divergence as a measure of the fitness of the motion field — same as the divergence regularization loss from Eq. (11). The lower the divergence, the better the motion field.

On the other hand, we do not want the generated motion field to be so homogeneous that it cannot accurately capture the movement of precipitation. In an extreme case, a zero-only motion field is perfectly divergence-free, but useless for performing the transformation to Lagrangian coordinates. The motion fields also need to accurately extrapolate the given input precipitation fields to the successive observations. As there is no ground truth to compare with, we will simply measure how well the last input precipitation field extrapolated by the generated motion fields matches the target observations — measure the error of a simple Lagrangian persistence nowcast using the generated motion fields. The second metric for measuring the error will be the mean square error of the extrapolation.

The quantitative results are presented grouped by lead time to get a better idea about the models' behaviour with increasing lead times. For each lead time, the last input observation is extrapolated by a set number of time steps. For Lukas-Kanade, this is straightforward, as we only consider a singular motion field for every time step. The MF-U-Net models, on the other hand, work iteratively and generate a new motion field for each time step in a rolling window fashion. See Fig. 4 for the results.

From the divergence comparison, we can see that without the physics-informed regularization that penalizes highly divergent outputs, the divergence is higher, which is not surprising. However, it is good to see that the regularization pushes the model to generate motion fields on par with Lukas-Kanade in terms of divergence. Also, both MF-U-Net models exhibit slightly decreasing motion field divergence over growing lead time.

The second graph shows a clear distinction between the models in terms of the accuracy of precipitation field extrapolation using the generated motion fields. The simple nowcasts using MF-U-Net outputs more closely match the target precipitation fields in comparison to Lukas-Kanade. In addition, the regularized MF-U-Net outperforms the unregularized one, further proving the value of the physics-informed regularization.

In conclusion, while it is hard to draw a definitive conclusion about which model generates the best motion fields for the task of Lagrangian transformation, the MF-U-Net with physics-informed regularization definitely performed well in comparison to the benchmark — the Lukas-Kanade algorithm. With the MF-U-Net outperforming it in the extrapolation MSE metric, it could even be argued that it is better suited for the task.

## 6. Nowcasting ability evaluation

This section presents qualitative and quantitative results of the LUPIN training in comparison to selected reference models. As mentioned in Section 1, a visual qualitative comparison of produced nowcasts in an expert study is often more valuable than a comprehensive quantitative evaluation, as there is no clear numerical metric that measures the fitness of nowcasts.

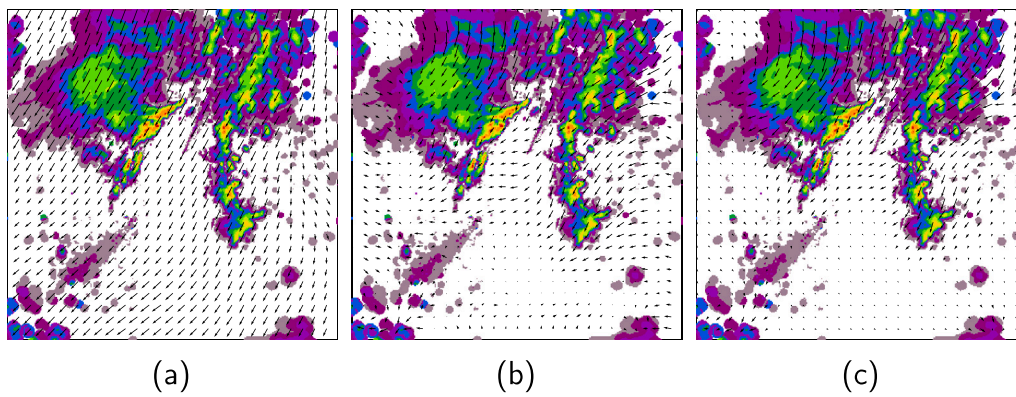


Fig. 3. The motion fields generated from three models for a sample extreme event from the test set. 3(a): Lukas-Kanade algorithm (the basis of L-CNN's Lagrangian transformation), 3(b): MF-U-Net from LUPIN trained without the physics-informed regularization, 3(c): MF-U-Net from LUPIN trained with the physics-informed regularization.




Fig. 4. Average motion field divergence (left) and extrapolation MSE (right) for the whole test set, evaluated for lead times up to 30 min ahead. As the Lukas-Kanade produces a singular motion field that is used at every time step, the divergence is constant.

### 6.1. Reference models

The main model to compare against is naturally L-CNN, the model LUPIN was based on. The performance of LUPIN needs to at least match the performance of the L-CNN, or, if our data-driven Lagrangian coordinate transformation is better, exceed it. The U-Net architecture employed in the L-CNN is equivalent to the Advection-Free U-Net in LUPIN, as described in Fig. 2.

The second reference model chosen is RainNet, a data-driven U-Net-based nowcasting model. For a fair evaluation, the U-Net parameters were set to match those of the U-Nets in LUPIN and L-CNN with respect to both architecture and channel dimensions, illustrated in Fig. 2. However, it is a pure machine learning model with no Lagrangian assumptions. It processes the input radar data in the usual – Euler – coordinates. RainNet is therefore a fitting benchmark, as any improved performance can be directly attributed to the Lagrangian approach. It also serves as a kind of ablation study of the designed Lagrangian framework.

The task of all three models is to take 6 successive precipitation fields and predict the next 6 fields. Each model employs the same back-end structure – a U-Net with identical complexity, comprising approximately 31.38 million trainable parameters. Any differences observed are therefore the consequence of the different training regimes and data transformations.

All models were trained and evaluated on a single NVIDIA GeForce RTX 3090 GPU with 24 GB of VRAM. On this hardware, single-sample inference requires around a tenth of a second per model. In the case of L-CNN this is dominated by the additional runtime of the Lukas-Kanade algorithm for motion field estimation, which operates in the order of seconds.

### 6.2. Extreme event case study

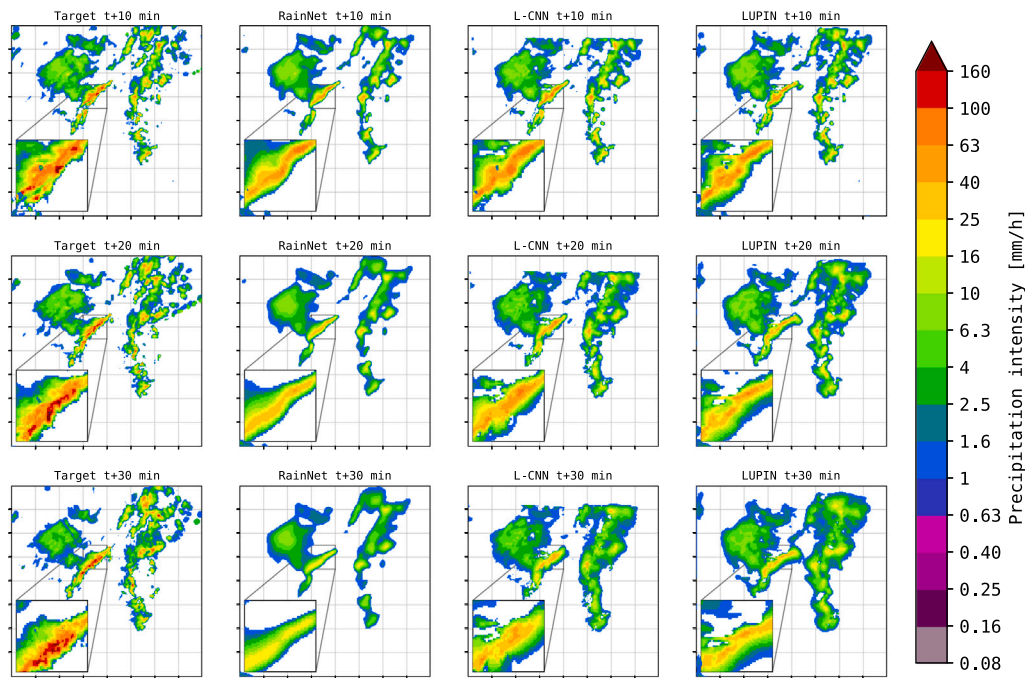
On September 2nd, 2018, southwest Slovakia was hit by a severe convective storm. It flooded underpasses, streets, and basements of buildings and caused public transport to collapse in the capital of Slovakia (ESSL, 2023). The convective storms are complex and hard to predict. They are the most relevant cause of systematic errors in weather and climate models (Reynolds et al., 2020). This rainfall event from the test set was chosen to compare the performance of selected nowcasting models. See Fig. 5 for the nowcasts produced by the models.

The RainNet model exhibits very strong blurring of the output with growing lead time. In comparison, both the L-CNN and LUPIN produce less blurry outputs of similar local detail that better preserve the expected high precipitation intensity, more closely resembling the target precipitation field. The benefits of the Lagrangian approach are clearly visible.


### 6.3. Quantitative comparison

To compare the models quantitatively, their performance is evaluated using multiple metrics over the whole test set. The quantitative results are presented grouped by lead time to get a better idea about the models' behaviour with increasing lead times. First, for the computed mean square error (MSE) and mean bias error (ME), see Fig. 6.

The L-CNN starts best in terms of MSE, though with growing lead times, LUPIN catches up and outperforms it. Both L-CNN and LUPIN exhibit relatively small bias compared to RainNet. The L-CNN model accumulates a small positive bias with growing lead time, and LUPIN, on the other hand, accumulates negative bias.



**Fig. 5.** Extreme event example from the evening of September 2nd, 2018. The leftmost column shows radar rainfall fields of a growing convective rain cell, with the maximum rain rate region zoomed in. Precipitation below 1 mm/h was filtered out. The other columns show nowcasts of the respective target fields at three lead times:  $t + 10$ ,  $t + 20$ , and  $t + 30$  minutes. All the models are able to capture the general southwestern movement of the precipitation field. However, all the models show blurring of the output over time due to uncertainty. The most prominent blurring occurs in the RainNet nowcasts, but L-CNN and LUPIN also exhibit noticeable blurring.



**Fig. 6.** Plots show the test set MSE (top) and ME (bottom) of the compared models in relation to nowcast lead time.

To better evaluate the performance of the models at predicting precipitation events of various levels, we also calculate three binary classification metrics over binary rainfall fields with thresholds of 1, 5, and 10 mm/h, corresponding to light, moderate, and heavy rain, respectively. We calculate precision and recall of the models along with equitable threat score, a common indicator of warning skill in weather forecasting (Schaefer, 1990). See Fig. 7 for comparison.

Precision – a complement of false alarm ratio (FAR) – marks the probability that a positive prediction (an alarm) is really a positive in the target observation. In this metric, the RainNet naturally leads due to its blurring and consequent underestimation. When it produces an alarm, there is a high chance of it being true. However, pure precision is not really of the highest priority for a warning system.

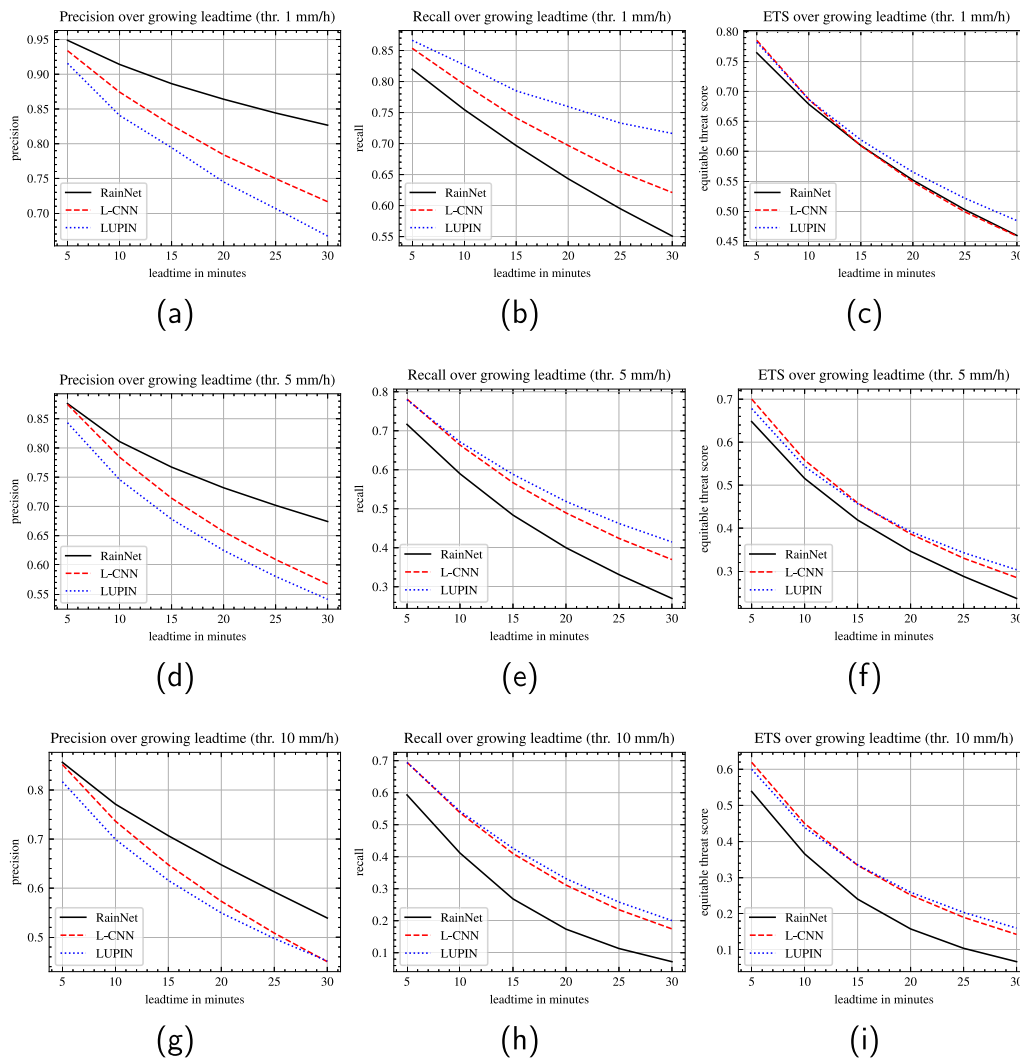
On the other hand, recall – also called the probability of detection (POD) – signifies the probability that a positive (an extreme event) will, in fact, trigger a positive prediction (an alarm). In the context of a warning system, this metric is of high priority to minimize the amount of missed events. Here, L-CNN and LUPIN both perform much better

than RainNet. With increasing lead time, LUPIN achieves better recall compared to L-CNN, with more than 14% better recall at the  $t + 30$  minutes lead time for the 10 mm/h threshold.

Finally, the equitable threat score (ETS) – based on the frequently used critical success index (CSI) – considers both the false alarms and the missed events, while also taking into account hits expected by chance. Here, the RainNet also falls far behind the Lagrangian models. In this metric, the L-CNN performs better at the start, but LUPIN consistently outperforms it with increasing lead times, still ending up with 12% better performance at the final lead time for the 10 mm/h threshold.

#### 6.4. Summary

Both the quantitative and qualitative evaluation have shown that the Lagrangian approach can provide a big improvement in the performance of a nowcasting model in contrast to a purely data-driven RainNet model. The L-CNN and LUPIN perform similarly, proving our



**Fig. 7.** Plots show test set precision (7(a), 7(d), 7(g)), recall (7(b), 7(e), 7(h)), and equitable threat score (7(c), 7(f), 7(i)) of the compared models in relation to nowcast lead time. Each row plots these three metrics for a different rain rate threshold – 1 mm/h (7(a), 7(b), 7(c)), 5 mm/h (7(d), 7(e), 7(f)), and 10 mm/h (7(g), 7(h), 7(i)).

framework is able to match the performance of the L-CNN while allowing much faster inference and a convenient training routine without transforming the whole dataset to Lagrangian coordinates in advance. In quantitative comparison, our LUPIN model consistently outperforms the other two models in the two key warning metrics – recall and equitable threat score – at all rain rate thresholds.

## 7. Conclusion

By implementing the temporal differencing applied to rainfall fields in Lagrangian coordinates introduced by the L-CNN in a fully differentiable manner, we have eliminated most of the drawbacks of the original model, while matching – even exceeding – its performance in key metrics.

Thanks to our implementation of the semi-Lagrangian extrapolation and training the data-driven MF-U-Net to produce physically consistent advection motion fields dynamically during runtime, we mitigate the higher computational cost at the inference of the L-CNN model. The MF-U-Net also mitigates the dependence on the quality of the motion field generated by traditional optical flow methods, which can be impacted by edge effects or gaps in radar coverage in the precipitation fields.

The quality of LUPIN's transformation to Lagrangian coordinates compared to L-CNN is hard to quantify. However, we have shown that

the generated advection motion fields can match the precipitation at successive observations with lower error while exhibiting similarly low divergence when trained with the physics-informed regularization loss. Combined with better overall nowcasting performance of LUPIN with increasing lead time observed in MSE, recall, and ETS metrics compared to the L-CNN, we believe this shows the motion fields learned by LUPIN are a better fit for the task of Lagrangian coordinate transformation of input data than the ones produced by the Lukas-Kanade algorithm.

The LUPIN framework, as defined in this paper, is flexible and not tied to any specific neural network architecture. With the generative models achieving great results in the precipitation nowcasting domain in recent years, our work opens the possibility of them leveraging the Lagrangian approach efficiently. We believe that including domain knowledge and physical laws in the machine learning models is necessary to build more trustworthy and reliable models of the future.

## CRedit authorship contribution statement

**Peter Pavlík:** Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Martin Výboš:** Software, Data curation. **Anna Bou Ezzeddine:** Writing – review & editing, Supervision. **Viera Rozina-jová:** Writing – review & editing, Supervision, Project administration, Funding acquisition.

## Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work the authors used ChatGPT in order to improve the clarity and readability of selected paragraphs. After using this service, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

## Code availability

The source code for all the models used in the experiments, along with the training, evaluation, and visualization scripts, plus animated figures, is available at: <https://github.com/kinit-sk/LUPIN>

## Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Peter Pavlík reports financial support was provided by SOFTEC, spol. s r.o. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgements

This research was partially supported by V4Grid, an Interreg Central Europe Programme project co-funded by the European Union, project No. CE0200803 and in cooperation with Softec, Ltd. and Slovak Hydrometeorological Institute.

## Data availability

The original data is available at [MalyJavornikRadar\(Originaldata\)](#).

## References

- Ayzel, G., Scheffer, T., Heistermann, M., 2020. RainNet v1. 0: a convolutional neural network for radar-based precipitation nowcasting. *Geosci. Model. Dev.* 13 (6), 2631–2644.
- Batchelor, G.K., 1967. *An Introduction to Fluid Dynamics*. Cambridge University Press.
- Bouguet, J.-Y., et al., 2001. Pyramidal implementation of the affine lucas kanade feature tracker description of the algorithm. *Intel Corp.* 5 (1–10), 4.
- Bowler, N.E., Pierce, C.E., Seed, A.W., 2006. STEPS: A probabilistic precipitation forecasting scheme which merges an extrapolation nowcast with downscaled NWP. *Q. J. R. Meteorol. Soc.: A J. Atmospheric Sci. Appl. Meteorol. Phys. Ocean.* 132 (620), 2127–2155.
- ESSL, 2023. European severe weather database. URL <https://eswd.eu/queries/21365.html>.
- Germann, U., Zawadzki, I., 2002. Scale-dependence of the predictability of precipitation from continental radar images. Part I: Description of the methodology. *Mon. Weather Rev.* 130 (12), 2859–2873.
- Hachemi, H., Hamdad, L., Douba, O., et al., 2023. Precipitation nowcasting using CNN-RNN - use case: North of Algeria. <http://dx.doi.org/10.21203/rs.3.rs-2677640/v1>, Preprint, Version 1, Research Square.
- Helmus, J.J., Collis, S.M., 2016. The python ARM radar toolkit (py-ART), a library for working with weather radar data in the python programming language. *J. Open Res. Softw.* 4 (1), e25. <http://dx.doi.org/10.5334/jors.119>.
- Ionescu, V.-S., Czibula, G., Mihuleț, E., 2021. DeePS at: A deep learning model for prediction of satellite images for nowcasting purposes. *Procedia Comput. Sci.* 192, 622–631. <http://dx.doi.org/10.1016/j.procs.2021.08.064>, Knowledge-Based and Intelligent Information & Engineering Systems: Proceedings of the 25th International Conference KES2021. URL <https://www.sciencedirect.com/science/article/pii/S1877050921015519>.
- Kanopoulos, N., Vasanthavada, N., Baker, R.L., 1988. Design of an image edge detection filter using the Sobel operator. *IEEE J. Solid-State Circuits* 23 (2), 358–367.
- Karniadakis, G.E., Kevrekidis, I.G., Lu, L., Perdikaris, P., Wang, S., Yang, L., 2021. Physics-informed machine learning. *Nat. Rev. Phys.* 3 (6), 422–440.
- Keil, C., Craig, G.C., 2009. A displacement and amplitude score employing an optical flow technique. *Weather. Forecast.* 24 (5), 1297–1308. <http://dx.doi.org/10.1175/2009WAF2222247.1>, URL [https://journals.ametsoc.org/view/journals/wefo/24/5/2009waf2222247\\_1.xml](https://journals.ametsoc.org/view/journals/wefo/24/5/2009waf2222247_1.xml).
- Lucas, B.D., Kanade, T., 1981. An iterative image registration technique with an application to stereo vision. In: *IJCAI'81: 7th International Joint Conference on Artificial Intelligence*, vol. 2, pp. 674–679.
- Marshall, J.S., Palmer, W.M.K., 1948. The distribution of raindrops with size. *J. Atmos. Sci.* 5 (4), 165–166. [http://dx.doi.org/10.1175/1520-0469\(1948\)005<0165:TDORWS>2.0.CO;2](http://dx.doi.org/10.1175/1520-0469(1948)005<0165:TDORWS>2.0.CO;2).
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al., 2019. Pytorch: An imperative style, high-performance deep learning library. *Adv. Neural Inf. Process. Syst.* 32.
- Pulkkinen, S., Chandrasekar, V., von Lerber, A., Harri, A.-M., 2020. Nowcasting of convective rainfall using volumetric radar observations. *IEEE Trans. Geosci. Remote Sens.* 58 (11), 7845–7859.
- Pulkkinen, S., Chandrasekar, V., Niemi, T., 2021. Lagrangian integro-difference equation model for precipitation nowcasting. *J. Atmos. Ocean. Technol.* 38 (12), 2125–2145.
- Pulkkinen, S., Nerini, D., Pérez Hortal, A.A., Velasco-Forero, C., Seed, A., Germann, U., Foresti, L., 2019. Pysteps: an open-source python library for probabilistic precipitation nowcasting (v1.0). *Geosci. Model. Dev.* 12 (10), 4185–4219. <http://dx.doi.org/10.5194/gmd-12-4185-2019>, URL <https://gmd.copernicus.org/articles/12/4185/2019/>.
- Ravuri, S., Lenc, K., Willson, M., Kangin, D., Lam, R., Mirowski, P., Fitzsimons, M., Athanassiadou, M., Kashem, S., Madge, S., et al., 2021. Skilful precipitation nowcasting using deep generative models of radar. *Nature* 597 (7878), 672–677.
- Reynolds, C., Williams, K.D., Zadra, A., 2020. WGN system error survey results summary. In: *100th American Meteorological Society Annual Meeting*. AMS.
- Ritvanen, J., Harnist, B., Aldana, M., Mäkinen, T., Pulkkinen, S., 2023. Advection-free convolutional neural network for convective rainfall nowcasting. *IEEE J. Sel. Top. Appl. Earth Obs. Remote. Sens.* 16, 1654–1667. <http://dx.doi.org/10.1109/JSTARS.2023.3238016>.
- Sawyer, J.S., 1963. A semi-Lagrangian method of solving the vorticity advection equation. *Tellus* 15 (4), 336–342.
- Schaefer, J.T., 1990. The critical success index as an indicator of warning skill. *Weather. Forecast.* 5 (4), 570–575.
- Seed, A., 2003. A dynamic and spatial scaling approach to advection forecasting. *J. Appl. Meteorol. Clim.* 42 (3), 381–388.
- Shi, X., Chen, Z., Wang, H., Yeung, D.-Y., Wong, W.-K., Woo, W.-c., 2015. Convolutional LSTM network: A machine learning approach for precipitation nowcasting. *Adv. Neural Inf. Process. Syst.* 28.
- Shi, X., Gao, Z., Lausen, L., Wang, H., Yeung, D.-Y., Wong, W.-k., Woo, W.-c., 2017. Deep learning for precipitation nowcasting: A benchmark and a new model. *Adv. Neural Inf. Process. Syst.* 30.
- Tian, L., Li, X., Ye, Y., Xie, P., Li, Y., 2019. A generative adversarial gated recurrent unit model for precipitation nowcasting. *IEEE Geosci. Remote. Sens. Lett.* 17 (4), 601–605.
- Wang, C., Wang, P., Wang, P., Xue, B., Wang, D., 2021. Using conditional generative adversarial 3-D convolutional neural network for precise radar extrapolation. *IEEE J. Sel. Top. Appl. Earth Obs. Remote. Sens.* 14, 5735–5749.
- Wang, H., Yang, R., He, J., Zeng, Q., Xiong, T., Liu, Z., Jin, H., 2025. Enhancing precipitation nowcasting through dual-attention RNN: Integrating satellite infrared and radar VIL data. *Remote. Sens.* 17 (2), <http://dx.doi.org/10.3390/rs17020238>, URL <https://www.mdpi.com/2072-4292/17/2/238>.
- WMO, 2017. Guidelines for Nowcasting Techniques. (978-92-63-11198-2), World Meteorological Organization, 7 bis, avenue de la Paix, P.O. Box 2300, CH-1211 Geneva 2, Switzerland.
- Zhang, Y., Long, M., Chen, K., Xing, L., Jin, R., Jordan, M.I., Wang, J., 2023. Skilful nowcasting of extreme precipitation with NowcastNet. *Nature* 619 (7970), 526–532.
- Zhang, W., Zhou, T., Wu, P., 2024. Anthropogenic amplification of precipitation variability over the past century. *Science* 385 (6707), 427–432. <http://dx.doi.org/10.1126/science.adp0212>, URL <https://www.science.org/doi/abs/10.1126/science.adp0212>.