

DISCHARGE ESTIMATION USING PHYSICS-INFORMED NEURAL NETWORKS WITH IMAGE-BASED SURFACE VELOCIMETRY

ESTIMATION DU DÉBIT AU MOYEN DE RÉSEAUX DE NEURONES INFORMÉS PAR LA PHYSIQUE ET DE LA VÉLOCIMÉTRIE DE SURFACE FONDÉE SUR L'IMAGE

Yen-Cheng Lin¹, and Hao-Che Ho²

ABSTRACT

Accurate discharge estimation is fundamental for evaluating water use efficiency and supporting adaptive irrigation management under increasing water resource pressure. Image-based velocimetry techniques, particularly Large-Scale Particle Image Velocimetry (LSPIV), have been increasingly adopted in irrigation channels and natural rivers due to their non-intrusive deployment and high spatial resolution. The integration of convolutional neural networks (CNNs) has led to CNN-LSPIV approaches that enhance surface velocity detection robustness under varying illumination, reflections, and surface texture conditions. Nevertheless, data-driven velocity reconstructions may exhibit locally non-conservative behavior that violates mass continuity, especially in regions influenced by image distortion or insufficient tracers, introducing significant uncertainty in discharge quantification. To address these challenges, a physics-informed discharge estimation framework is developed based on Physics-Informed Neural Networks (PINNs), embedding the shallow water continuity equation directly into the training loss together with observational consistency and hydraulic boundary constraints. Two operational modes are proposed: a Velocity Field Refinement mode (PINN-VR), which refines CNN-LSPIV velocity fields prior to leapfrog-based depth integration, and a Direct Depth Inversion mode (PINN-DDI), which simultaneously reconstructs velocity and water depth fields through global physics-constrained optimization.

The framework is validated through laboratory flume experiments under three bed configurations covering 23 experimental cases, with ADV measurements providing reference validation data. CNN-LSPIV achieves mean velocity error reductions of 44 to 70% relative to conventional LSPIV. PINN-VR performs best under lateral composite bed conditions (average depth RMSE = 0.0135 m), whereas PINN-DDI achieves superior accuracy under longitudinal bed variations (average depth RMSE = 0.0217 m) by reducing sensitivity to directional error accumulation. At the discharge estimation scale, PINN-DDI yields the lowest errors under both composite bed configurations. Field validation using UAV imagery over a natural gravel-bed river reach with ADCP cross-sectional measurements confirms the framework's transferability, with the best-performing configuration reducing average discharge error by approximately 27% relative to LSPIV-based methods. These results demonstrate that embedding physical conservation constraints into image-based flow reconstruction substantially advances non-contact discharge estimation toward hydraulically consistent flow quantification, strengthening the role of image-derived data in irrigation system performance assessment and adaptive water management.

Keywords: Discharge estimation; Physics-Informed Neural Networks; CNN-LSPIV; Image-based velocimetry; Irrigation channel monitoring.

1 PhD student, Department of Civil Engineering, National Taiwan University, No. 1, Sec. 4, Roosevelt Rd., Taipei 106319, Taiwan, Email: d10521010@ntu.edu.tw

2 Associate professor, Department of Civil Engineering, National Taiwan University, No. 1, Sec. 4, Roosevelt Rd., Taipei 106319, Taiwan, Email: haocheho@ntu.edu.tw

RÉSUMÉ

Une estimation précise du débit est fondamentale pour évaluer l'efficacité d'utilisation de l'eau et soutenir une gestion adaptative de l'irrigation dans un contexte de pression croissante sur les ressources en eau. Les techniques de vélocimétrie fondées sur l'image, notamment la vélocimétrie par images de particules à grande échelle (LSPIV), sont de plus en plus employées dans les canaux d'irrigation et les cours d'eau naturels en raison de leur déploiement non intrusif et de leur haute résolution spatiale. L'intégration de réseaux de neurones convolutifs (CNN) a conduit à des approches CNN-LSPIV qui renforcent la robustesse de la détection de la vitesse de surface sous des conditions variables d'éclairage, de réflexion et de texture de surface. Néanmoins, les reconstructions de vitesse fondées sur les données peuvent présenter localement un comportement non conservatif qui contrevient à la continuité de la masse, en particulier dans les zones affectées par la distorsion des images ou par une quantité insuffisante de traceurs, ce qui introduit une incertitude importante dans la quantification du débit. Pour répondre à ces difficultés, un cadre d'estimation du débit informé par la physique est élaboré à partir de réseaux de neurones informés par la physique (PINN), en intégrant directement l'équation de continuité des eaux peu profondes dans la fonction de perte d'apprentissage, conjointement avec la cohérence des observations et les contraintes hydrauliques aux limites. Deux modes de fonctionnement sont proposés : un mode d'affinement du champ de vitesse (PINN-VR), qui affine les champs de vitesse CNN-LSPIV avant l'intégration de la profondeur selon un schéma saute-mouton, et un mode d'inversion directe de la profondeur (PINN-DDI), qui reconstruit simultanément les champs de vitesse et de profondeur d'eau au moyen d'une optimisation globale contrainte par la physique.

Le cadre est validé par des expériences en canal de laboratoire menées selon trois configurations de fond et couvrant 23 cas expérimentaux, les mesures ADV servant de données de référence pour la validation. Par rapport à la LSPIV conventionnelle, la CNN-LSPIV réduit de 44 à 70 % l'erreur moyenne sur la vitesse. Le mode PINN-VR donne les meilleurs résultats pour les fonds composites latéraux (RMSE moyenne de la profondeur = 0,0135 m), tandis que le mode PINN-DDI offre une précision supérieure pour les variations longitudinales du fond (RMSE moyenne de la profondeur = 0,0217 m) en réduisant la sensibilité à l'accumulation directionnelle des erreurs. À l'échelle de l'estimation du débit, le PINN-DDI produit les erreurs les plus faibles pour les deux configurations de fond composite. Une validation sur le terrain à l'aide d'images acquises par UAV au-dessus d'un tronçon naturel de rivière à lit graveleux, avec des mesures ADCP en section transversale, confirme la transférabilité du cadre ; la configuration la plus performante réduit l'erreur moyenne de débit d'environ 27 % par rapport aux méthodes fondées sur la LSPIV. Ces résultats montrent que l'intégration des contraintes de conservation physique dans la reconstruction des écoulements à partir d'images fait progresser sensiblement l'estimation sans contact du débit vers une quantification hydrauliquement cohérente des écoulements, renforçant ainsi le rôle des données issues de l'imagerie dans l'évaluation des performances des systèmes d'irrigation et la gestion adaptative de l'eau.

Mots-clés : Estimation du débit; Réseaux neuronaux à base physique; CNN-LSPIV; Vélocimétrie fondées sur l'image; Surveillance des canaux d'irrigation.

1. INTRODUCTION

The growing demand for reliable water resource management under climate change has intensified the need for accurate and operationally feasible discharge monitoring in irrigation channels and natural rivers. Conventional contact-based discharge measurement methods are often difficult or unsafe to perform during high-flow conditions, due to unstable cross-sections, equipment damage risks, and field personnel safety concerns (Rantz, 1982). Non-contact discharge estimation techniques capable of operating under adverse hydraulic conditions have therefore become an important research priority.

LSPIV has been widely adopted for surface velocity measurement because it tracks naturally occurring floating materials or surface textures from image sequences, providing wide spatial coverage without physical intrusion (Fujita and Komura, 1994). However, LSPIV performance is sensitive to image quality: illumination variation, surface reflections, insufficient tracers, and wind-induced disturbances may produce erroneous velocity vectors that propagate into discharge estimation errors. CNN-LSPIV approaches address this by replacing the conventional cross-correlation algorithm with a convolutional neural network to improve feature recognition and velocity estimation stability (Ho et al., 2023). Despite these improvements, CNN-based methods remain fundamentally data-driven, and when image quality deteriorates, the reconstructed velocity field may violate mass conservation, leading to spatially inconsistent and biased discharge estimates.

Physics-Informed Neural Networks (PINNs) address this limitation by embedding governing physical equations directly into the neural network training process, restricting the solution space to physically admissible states (Raissi et al., 2019). The shallow water continuity equation is selected as the governing constraint because mass conservation is the most fundamental kinematic requirement of open-channel flow and involves no parameters requiring field calibration. Reliable water depth information is further essential for discharge computation. The leapfrog numerical scheme offers an effective approach for depth inversion from velocity fields under controlled conditions (Lin et al., 2022), though local velocity errors may accumulate along the stepwise integration path under complex bed topography.

This study develops a physics-informed framework integrating CNN-LSPIV with PINN-based physical constraints, proposes two water depth inversion strategies suited to different hydraulic environments, and evaluates the complete framework through systematic laboratory experiments and UAV-based field observations.

2. METHODOLOGY

2.1 CNN-Enhanced LSPIV

CNN-LSPIV replaces the conventional Direct Cross-Correlation (DCC) algorithm with a convolutional neural network for image feature extraction and matching discrimination (Ho et al., 2023). The network architecture comprises convolutional and pooling layers (3×3 kernel, 2×2 max pooling) followed by fully connected layers and a softmax classifier. ReLU activation and Batch Normalization are applied throughout to enhance training stability. The Particle Average Movement (PAM) method is applied at inference time for sub-pixel correction, maintaining high matching stability and velocity field continuity under texture-deficient and variable-quality imaging conditions.

2.2 PINN Framework

The PINN adopts a Multilayer Perceptron (MLP) taking spatial coordinates (x, y) as input and mapping them to surface velocity components and water depth. Six hidden layers with 128 neurons per layer and tanh activation are adopted, ensuring continuous partial derivatives required for PDE residual computation. The Softplus function is applied as the output activation for water depth to enforce strict positivity throughout the domain.

The total training loss combines three terms (Wang et al., 2021):

$$L_{total} = \lambda_{data} \cdot L_{data} + \lambda_{pde} \cdot L_{pde} + \lambda_{bc} \cdot L_{bc} \quad (1)$$

The data loss anchors network outputs to CNN-LSPIV observations. The PDE loss enforces the two-dimensional steady shallow water continuity equation:

$$L_{pde} = \frac{1}{N} \sum_{i=1}^N \left[\frac{\partial(\hat{H}^{(i)} \cdot \alpha \hat{u}_s^{(i)})}{\partial x} + \frac{\partial(\hat{H}^{(i)} \cdot \alpha \hat{v}_s^{(i)})}{\partial y} \right]^2 \quad (2)$$

where $\alpha = 0.85$ is the surface-to-depth-averaged velocity ratio, a value widely adopted and validated under turbulent open-channel flow conditions (Le Coz et al., 2010; Welber et al., 2016). The boundary condition loss constrains water depth at the upstream and downstream cross-sections, providing the absolute magnitude calibration reference essential for the PINN-DDI mode. Partial derivatives are computed via automatic differentiation, avoiding truncation errors associated with finite difference schemes. Network parameters are updated using the Adam optimization algorithm (Kingma and Ba, 2015), which incorporates adaptive momentum estimation to accelerate convergence and reduce oscillations during training. A warm-up training strategy reduces the effective PDE loss weight to $0.1 \lambda_{pde}$ during the first 20% of training epochs, allowing the network to first establish a reasonable velocity distribution before the full physical constraint is applied (Wang et al., 2021).




Figure 1. Network architecture of the PINN.

2.3 Two Inversion Modes

PINN-VR (Velocity Field Refinement) refines the CNN-LSPIV velocity field under mass conservation constraints while preserving the observed magnitude. Regions with large continuity-equation residuals receive stronger corrections, while physically consistent regions are minimally disturbed. The refined surface velocity is converted to depth-averaged velocity using $\alpha = 0.85$ and subsequently input to the leapfrog scheme on an Arakawa C-grid (Mesinger and Arakawa, 1976) for finite-difference integration along the streamwise direction to reconstruct the water depth field.

PINN-DDI (Direct Depth Inversion) incorporates water depth directly as an inversion target within the PINN, solved simultaneously with the velocity field under global continuity constraints (Raissi et al., 2019). This formulation eliminates the leapfrog integration step entirely, and because both velocity and depth are optimized jointly within the same physical constraint framework, the solution is less sensitive to strong streamwise topographic variations.

3. EXPERIMENTAL DESIGN

3.1 Laboratory Flume Experiments

Laboratory experiments were conducted at the Ujigawa Open Laboratory, Disaster Prevention Research Institute, Kyoto University, Japan. The recirculating flume measures 21.0 m in length, 0.50

m in width, and 0.50 m in depth. A digital camera was mounted vertically at approximately 4 m above the water surface, recording at 2704×1520 pixels and 60 fps, with 20 ground control points used for orthorectification. White polystyrene spheres (diameter approximately 3 mm, density approximately 0.02 g/cm^3) served as tracer particles. A 16-MHz Acoustic Doppler Velocimeter (ADV) and a supplementary electromagnetic current meter provided reference velocity measurements at 15 points across three cross-sections.

Three bed configurations were tested to evaluate framework performance under progressively complex hydraulic conditions. The smooth bed provides a uniform baseline for model calibration. The lateral composite bed features one side raised by 0.10 m over a 0.35 m sloped transition (total model length 5.00 m), simulating the asymmetric cross-section of a point bar and generating pronounced lateral velocity gradients. The longitudinal composite bed consists of a 0.10 m elevated platform (total length 3.50 m) with 1:7.5 sloped transitions, simulating bedforms or sills that induce sequential flow acceleration, local disturbance, and wake reattachment. A total of 23 experimental cases were conducted with Froude numbers ranging from 0.086 to 0.214, consistent with subcritical flow throughout




Figure 2. (a) Lateral composite bed. (b) Longitudinal composite bed.

Table 1. Summary of experimental cases under three bed configurations.

Bed configuration	Cases	Q (m ³ /s)	H (m)	Fr
Smooth	S1 to S8 (8 cases)	0.022-0.035	0.200-0.350	0.108-0.214
Lateral composite	C1 to C8 (8 cases)	0.022-0.035	0.200-0.350	0.086-0.214
Longitudinal composite	B1 to B7 (7 cases)	0.022-0.035	0.250-0.360	0.086-0.153

3.2 Field Experiment

Field experiments were conducted at a gravel-bed river reach in northern Taiwan, approximately 20 m in width, with naturally occurring surface textures suitable for image velocimetry. A RTK-equipped UAV

hovered at 70 m altitude, recording at 4K resolution (30 fps) with a field of view of approximately 100 m × 55 m (spatial resolution approximately 2.45 cm/pixel). Inter-frame image stabilization was performed using the Scale-Invariant Feature Transform (SIFT) algorithm (Lowe, 2004) to correct UAV platform motion induced by wind and GPS drift. An Acoustic Doppler Current Profiler (ADCP) was deployed at four cross-sections to provide reference water depth and discharge measurements.

4. RESULTS AND DISCUSSION

4.1 Surface Velocity Measurement

Under smooth bed conditions, CNN-LSPIV substantially outperforms conventional LSPIV across all flow depths. As shown in Table 2, CNN-LSPIV achieves mean percentage errors (E_{avg}) of 5.6%, 5.8%, and 5.4% at $H = 0.20$, 0.25 , and 0.30 m respectively, remaining consistently within 5 to 6% regardless of depth, compared to 19.3%, 10.3%, and 14.8% for LSPIV. RMSE reductions of 70%, 44%, and 62% are achieved across the three depth conditions. The markedly fluctuating error of LSPIV with changing water depth contrasts with the stable performance of CNN-LSPIV, demonstrating the robustness of learned convolutional feature representation under varying hydraulic and imaging conditions.


Table 2. Surface velocity error comparison between LSPIV and CNN-LSPIV (smooth bed).

Method	$H = 0.20$ m RMSE (m/s)	E_{avg} (%)	$H = 0.25$ m RMSE (m/s)	E_{avg} (%)	$H = 0.30$ m RMSE (m/s)	E_{avg} (%)
LSPIV	0.0738	19.3	0.0308	10.3	0.0351	14.8
CNN-LSPIV	0.0221	5.6	0.0174	5.8	0.0133	5.4

4.2 PINN Architecture and Loss Function

Architecture calibration using the smooth bed baseline confirms that 6 hidden layers achieves the best balance between depth inversion accuracy and mass conservation improvement, with PINN-VR reaching $\Delta R = 95.2$ to 97.4% and PINN-DDI reaching $\Delta R = 95.0$ to 96.9% . Increasing to 10 layers marginally improves RMSE but reduces the stability of the continuity equation residual, suggesting a trade-off between data fitting emphasis and consistent physical constraint enforcement.

The loss function ablation study confirms the indispensable role of each term. Removing the PDE loss reduces ΔR substantially to 69.6 to 87.6%, confirming its central role as the mass conservation constraint. Removing the boundary condition loss causes PINN-DDI depth RMSE to increase by up to one order of magnitude, from 0.0063 to 0.0091 m to 0.0318 to 0.1135 m, as the network loses its absolute depth magnitude anchor and adjusts the depth scale freely to minimize the PDE residual, producing physically unrealistic results despite numerically elevated ΔR values. PINN-VR is less sensitive to this removal because the leapfrog upstream boundary condition already provides a magnitude reference. Only the combined action of all three terms yields solutions that are simultaneously quantitatively accurate and physically plausible.






Figure 3. Surface velocity distribution under different flow depths: (a) LSPIV ($H=0.20\text{m}$). (b) LSPIV ($H=0.30\text{m}$). (c) CNN-LSPIV ($H=0.20\text{m}$). (d) CNN-LSPIV ($H=0.30\text{m}$).

4.3 Water Depth Inversion

Four approaches are compared: LSPIV-LP (LSPIV with leapfrog), CNN-LSPIV-LP, PINN-VR, and PINN-DDI. Average RMSE results are summarized in Table 3.

Table 3. Average RMSE (m) of water depth inversion under three bed configurations.

Method	Smooth bed	Lateral composite	Longitudinal composite
LSPIV-LP	0.0834	0.0435	0.1016
CNN-LSPIV-LP	0.0050	0.0154	0.0536
PINN-VR	0.0046	0.0135	0.0528
PINN-DDI	0.0086	0.0187	0.0217




Figure 4. Three-dimensional reconstructions of water depth under smooth bed (S4), lateral composite bed (C1), and longitudinal composite bed (B2) conditions.

LSPIV-LP exhibits the largest errors under all configurations, increasing substantially with bed complexity. Under longitudinal composite bed conditions, RMSE reaches 0.2041 m in individual cases, reflecting systematic error accumulation along the stepwise integration path when velocity measurements are spatially inconsistent.

Replacing LSPIV with CNN-LSPIV as the leapfrog input reduces average RMSE by over 90% under smooth bed conditions. However, under longitudinal composite bed conditions, CNN-LSPIV-LP average RMSE remains at 0.0536 m, indicating that improved velocity accuracy alone cannot overcome the inherent directional error accumulation of leapfrog integration under strong streamwise topographic variation.

The two PINN-based strategies exhibit a clear bed-dependent performance pattern. Under lateral composite bed conditions, PINN-VR achieves the lowest average RMSE (0.0135 m). The physics-refined velocity field provides a spatially more coherent input to the leapfrog scheme, enabling more effective tracking of the pronounced lateral depth gradient between the main channel and the bar region. Under longitudinal composite bed conditions, PINN-DDI achieves the best performance with an average RMSE of 0.0217 m, compared to 0.0528 m for PINN-VR. Global optimization under continuous constraints avoids the directional error accumulation that plagues stepwise integration under strong streamwise bed variations.

4.4 Discharge Estimation

Table 4. Average discharge error E_{avg} (%) under three bed configurations.

Method	Smooth bed	Lateral composite	Longitudinal composite
LSPIV-LP	15.24	25.82	41.20
CNN-LSPIV-LP	8.94	15.50	22.45
PINN-VR	8.76	10.96	20.81
PINN-DDI	9.48	10.59	18.63

Under smooth bed conditions, CNN-LSPIV-LP, PINN-VR, and PINN-DDI all achieve average discharge errors below 10%, with minor differences among the three methods. As bed complexity increases, the advantages of the physics-informed framework become more pronounced. Under lateral composite bed conditions, both PINN strategies reduce average discharge error to approximately 11%, compared to 15.50% for CNN-LSPIV-LP. Under longitudinal composite bed conditions, PINN-DDI achieves the best accuracy at 18.63%, outperforming CNN-LSPIV-LP (22.45%) and PINN-VR (20.81%), consistent with its superior depth inversion performance under strong streamwise bed variations.




Figure 5. Comparison of discharge estimation errors under different bed configurations.

4.5 Field Experiment Results

At the field site, CNN-LSPIV reduces the average surface velocity mean percentage error from 42.18% to 28.72%, a 31.9% relative improvement over conventional LSPIV. The larger absolute errors compared to laboratory conditions reflect non-uniform tracer distributions, time-varying water surface illumination, and complex three-dimensional flow structures induced by channel planform curvature and natural bed topography. Water depth inversion results confirm the trends observed in laboratory experiments. PINN-DDI achieves the best performance with an average E_{avg} of 16.88%, compared to 27.17% for LSPIV-LP, 18.39% for CNN-LSPIV-LP, and 18.32% for PINN-VR. The advantage of global optimization becomes more prominent under the complex three-dimensional bed topography of the natural channel. For discharge estimation, PINN-DDI achieves an average error of 24.11%, representing approximately 27% reduction relative to LSPIV-LP (33.04%). Results are summarized in Table 5.

Table 5. Summary of average field experiment results (four cross-sections).

Method	Velocity E_{avg} (%)	Depth E_{avg} (%)	Discharge E_{avg} (%)
LSPIV-LP	42.18	27.17	33.04
CNN-LSPIV-LP	28.72	18.39	24.93
PINN-VR	-	18.32	24.81
PINN-DDI	-	16.88	24.11

CONCLUSIONS

This study develops a physics-informed framework for non-contact river discharge estimation by integrating CNN-LSPIV surface velocity measurement with PINN-based physical constraints under the governing shallow water continuity equation. The following principal conclusions are drawn.

CNN-LSPIV consistently improves velocity measurement accuracy relative to conventional LSPIV under both laboratory and field conditions, achieving RMSE reductions of 44 to 70% across flume depth conditions and a 31.9% reduction in mean percentage error at the field site. The robustness of learned convolutional feature representation under varying illumination and surface texture conditions makes it a reliable foundation for subsequent hydraulic inversion.

The PINN loss function ablation analysis confirms the indispensable role of each constraint term. The PDE residual loss enforces mass conservation throughout the reconstructed flow field. The boundary condition loss provides an absolute magnitude anchor for water depth that is particularly critical in PINN-DDI mode. Only through the combined action of all three terms can the framework simultaneously achieve quantitative accuracy and physical plausibility.

The relative performance of PINN-VR and PINN-DDI is strongly bed-geometry dependent. PINN-VR performs best under lateral composite bed conditions, where physics-refined velocity inputs improve leapfrog integration accuracy along pronounced lateral depth gradients. PINN-DDI performs best under longitudinal composite bed conditions, where global physics-constrained optimization reduces directional error accumulation inherent in stepwise integration. At the discharge estimation scale, PINN-DDI yields the lowest errors under both composite bed configurations.

Field experiments confirm the framework's transferability to natural river environments, with the best-performing configuration reducing average discharge error by approximately 27% relative to conventional LSPIV-based methods. With continuing advances in UAV imaging resolution and real-time processing capability, physics-informed image-based approaches offer a promising pathway for non-contact discharge monitoring in irrigation channels and natural rivers under challenging field conditions.

REFERENCES

- Fujita, I., Komura, S., 1994. Application of video image analysis for measurements of river-surface flows. *Proceedings of Hydraulic Engineering* 38, 733–738. <https://doi.org/10.2208/prohe.38.733>
- Ho, H.C., Chiu, Y.W., Chen, T.Y., Lin, Y.C., 2023. Flow measurement in open channels using imaging techniques in conjunction with a convolutional neural network. *Journal of Hydrology* 618, 129183. <https://doi.org/10.1016/j.jhydrol.2023.129183>
- Kingma, D.P., Ba, J., 2015. Adam: A method for stochastic optimization. *International Conference on Learning Representations (ICLR 2015)*. <https://doi.org/10.48550/arXiv.1412.6980>
- Le Coz, J., Hauet, A., Pierrefeu, G., Dramais, G., Camenen, B., 2010. Performance of image-based velocimetry (LSPIV) applied to flash-flood discharge measurements in Mediterranean rivers. *Journal of Hydrology* 394(1-2), 42–52. <https://doi.org/10.1016/j.jhydrol.2010.05.049>
- Lin, Y.C., Ho, H.C., Lee, T.A., Chen, H.Y., 2022. Application of image technique to obtain surface velocity and bed elevation in open-channel flow. *Water* 14(12), 1895. <https://doi.org/10.3390/w14121895>
- Lowe, D.G., 2004. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* 60(2), 91–110. <https://doi.org/10.1023/B:VISI.0000029664.99615.94>
- Mesinger, F., Arakawa, A., 1976. *Numerical Methods Used in Atmospheric Models*. GARP Publications Series No. 17, Vol. 1. World Meteorological Organization, Geneva.
- Raissi, M., Perdikaris, P., Karniadakis, G.E., 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics* 378, 686–707. <https://doi.org/10.1016/j.jcp.2018.10.045>
- Rantz, S.E., 1982. Measurement and computation of streamflow. U.S. Geological Survey Water-Supply Paper 2175(2). <https://doi.org/10.3133/wsp2175>
- Wang, S., Teng, Y., Perdikaris, P., 2021. Understanding and mitigating gradient flow pathologies in physics-informed neural networks. *SIAM Journal on Scientific Computing* 43(5), A3055–A3081. <https://doi.org/10.1137/20M1318043>
- Welber, M., Le Coz, J., Laronne, J.B., Zolezzi, G., Zamler, D., Dramais, G., Hauet, A., Salvaro, M., 2016. Field assessment of noncontact stream gauging using portable surface velocity radars (SVR). *Water Resources Research* 52(2), 1108–1126. <https://doi.org/10.1002/2015WR017906>