



HESS Opinions: Incubating deep-learning-powered hydrologic science advances as a community

Chaopeng Shen¹, Eric Laloy², Amin Elshorbagy³, Adrian Albert⁴, Jerad Bales⁵, Fi-John Chang⁶, Sangram Ganguly⁷, Kuo-Lin Hsu⁸, Daniel Kifer⁹, Zheng Fang¹⁰, Kuai Fang¹, Dongfeng Li¹⁰, Xiaodong Li¹¹, and Wen-Ping Tsai¹

¹Civil and Environmental Engineering, Pennsylvania State University, University Park, PA 16802, USA

²Institute for Environment, Health and Safety, Belgian Nuclear Research Centre, Mol, Belgium

³Dept. of Civil, Geological, and Environmental Engineering, University of Saskatchewan, Saskatoon, Canada

⁴National Energy Research Supercomputing Center, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA

⁵Consortium of Universities for the Advancement of Hydrologic Science, Inc. (CUAHSI), Cambridge, MA, USA

⁶Department of Bioenvironmental Systems Engineering, National Taiwan University, Taipei, 10617, Taiwan

⁷NASA Ames Research Center/BAER Institute, Moffett Field, CA 94035, USA

⁸Civil and Environmental Engineering, University of California, Irvine, Irvine, CA 92697, USA

⁹Computer Science and Engineering, Pennsylvania State University, University Park, PA 16802, USA

¹⁰Civil Engineering, University of Texas at Arlington, Arlington, TX 76013, USA

¹¹State Key Laboratory of Hydraulics and Mountain River Engineering, Sichuan University, Sichuan, China

Correspondence: Chaopeng Shen (cshen@enr.psu.edu)

Received: 2 April 2018 – Discussion started: 9 April 2018

Revised: 2 September 2018 – Accepted: 19 September 2018 – Published: 1 November 2018

Abstract. Recently, deep learning (DL) has emerged as a revolutionary and versatile tool transforming industry applications and generating new and improved capabilities for scientific discovery and model building. The adoption of DL in hydrology has so far been gradual, but the field is now ripe for breakthroughs. This paper suggests that DL-based methods can open up a complementary avenue toward knowledge discovery in hydrologic sciences. In the new avenue, machine-learning algorithms present competing hypotheses that are consistent with data. Interrogative methods are then invoked to interpret DL models for scientists to further evaluate. However, hydrology presents many challenges for DL methods, such as data limitations, heterogeneity and co-evolution, and the general inexperience of the hydrologic field with DL. The roadmap toward DL-powered scientific advances will require the coordinated effort from a large community involving scientists and citizens. Integrating process-based models with DL models will help alleviate data limitations. The sharing of data and baseline models will improve the efficiency of the community as a whole. Open competitions could serve as the organizing events to greatly propel growth and nurture data science education in

hydrology, which demands a grassroots collaboration. The area of hydrologic DL presents numerous research opportunities that could, in turn, stimulate advances in machine learning as well.

1 Overview

Deep learning (DL) is a suite of tools centered on artfully designed large-size artificial neural networks. The deep networks at the core of DL are said to have “depth” due to their multi-layered structures, which help deep networks represent abstract concepts about the data (Schmidhuber, 2015). Given input attributes that describe an instance, deep networks can be trained to make predictions of some dependent variables, either continuous or categorical, for this instance. For example, for standard computer vision problems, deep networks can recognize the theme or objects from a picture (Guo et al., 2016; He et al., 2016; Simonyan and Zisserman, 2014) or remotely sensed images (Zhu et al., 2017). For sequential data, DL can associate natural language sequences with commands (Baughman et al., 2014; Hirschberg and Man-

ning, 2015) or predict the action of an actor in the next video frame (Vondrick et al., 2016). DL can also generate (or synthesize) images that carry certain artistic styles (Gatys et al., 2016) or a natural language response to questions (Leviathan and Matias, 2018; Zen and Sak, 2015). With the support of deep architectures, deep networks can automatically engineer relevant concepts and features from large datasets, instead of requiring human experts to define these features (Sect. 2.2.2). As a foundational component of modern artificial intelligence (AI), DL has made substantial strides in recent years and helped to solve problems that have resisted AI for decades (LeCun et al., 2015).

While DL has stimulated exciting advances in many disciplines and has become the method of choice in some areas, hydrology so far has only had a very limited set of DL applications (Shen, 2018) (hereafter referred to as Shen18). Despite scattered reports of promising DL results (Fang et al., 2017; Laloy et al., 2017, 2018; Tao et al., 2016; Vandal et al., 2017; Zhang et al., 2018), hydrologists have not widely adopted these new tools. This collective opinion paper argues that there are many opportunities in hydrological sciences where DL can help provide both stronger predictive capabilities and a complementary avenue toward scientific discovery. We then reflect on why it has been challenging to harness the power of DL and big data in hydrology and explore what we can do as a community to incubate progress. Readers who are less familiar with machine learning or deep learning are referred to a companion review paper (Shen18), which provides a more comprehensive and technical background than this opinion paper. Many details behind the arguments in Sect. 2 are provided in Shen18.

2 The emergence of a complementary research avenue

We are witnessing the growth of three pillars needed for DL to support a research avenue that is complementary to traditional hypothesis-driven research: big hydrologic data, powerful machine learning algorithms, and interrogative methods (such as visualization and techniques) to extract interpretable knowledge from the trained networks. This new avenue starts from data, uses DL methods to generate hypotheses, and applies interrogative methods to help us understand hydrologic system functioning. We discuss these aspects in the following sections.

2.1 With more data, opportunities arise

The fundamental supporting factor for emerging opportunities with DL is the growth of big hydrologic data, with all surface, sub-surface, urban, infrastructure, and ecosystem dimensions. In this paper, hydrology refers to both the complete natural and engineered water cycle and associated processes in the ecosystem and geologic media. There are ever increasing amounts of hydrologic data available through re-

mote sensing (see a summary in Srinivasan, 2013) and data compilations. For example, satellite-based datasets include precipitation, surface soil moisture (Entekhabi, 2010; Jackson et al., 2016; Mecklenburg et al., 2008), vegetation states and indices, e.g., Knyazikhin et al. (1999), derived evapotranspiration products (Mu et al., 2011), terrestrial water storage (Wahr et al., 2006), snow cover (Hall et al., 2006), and a planned mission for estimating streamflows (Pavelsky et al., 2014), etc. On the data compilation side, there are now compilations of geologic (Gleeson et al., 2014) and soil datasets; these include the centralized management of streamflow and groundwater data in the United States, Europe, parts of South America and Asia, or globally for some large rivers (GRDC, 2017), water chemistry, groundwater samples, and other biogeophysical datasets. The Consortium of Universities for the Advancement of Hydrologic Science, Inc. (CUAHSI) operates two systems for the discovery and the archiving of water data: the Hydrologic Information System (CUAHSI, 2018c) for time series and HydroShare for all water data types (Horsburgh et al., 2016). An Internet of Water (Aspen, 2017) has been proposed and is beginning to develop, thereby improving access to these emerging datasets.

Moreover, unconventional data sources are starting to emerge. A high-resolution sensing of Earth will be provided by increasing numbers of CubeSats, unmanned aerial vehicles, balloons, inexpensive photogrammetric sensing, and many other sources (McCabe et al., 2017). These new sources provide new forms and scales of measurements not envisioned before. For example, cell phone signal strength and cell-phone pictures can contribute to high-resolution monitoring of rainfall intensity (Allamano et al., 2015). Inexpensive infrared camera images can detect water levels in complex urban water flows (Hiroi and Kawaguchi, 2016). Internet-of-Things (IoT) sensors embedded in the water infrastructure can transmit data about the states of water in our environment (Zhang et al., 2018). These new sources of information provide unprecedented volumes and multifaceted coverages of the natural and built environment. However, since each new data source has its own characteristics and peculiarities, the identification of the appropriate approaches to fully exploit their value, especially synergistically, creates a significant challenge. In contrast, DL models can be built without significant human expertise and extensive manual labor to rapidly derive useful information from these data.

2.2 DL: a big step forward

2.2.1 Rapid adoption

The field of hydrology has witnessed the flows and ebbs of several generations of machine learning methods in the past few decades, from regularized linear regression (Tibshirani and Tibshirani, 1994) to support vector regression (Drucker et al., 1996), from genetic programming (Koza, 1992) to artificial neural networks (ANNs) (Chang et al., 2014; Chen et

Table 1. Number of papers returned from searches on ISI Web of Science.

year	DL-non-CS	DL-CS	ML-non-CS	ML-CS	DL/ML-CS	DL/ML-non-CS
2011	0	23	1068	1838	1 %	0 %
2012	15	25	1310	1899	1 %	1 %
2013	35	80	1677	2360	3 %	2 %
2014	84	238	2228	3050	8 %	4 %
2015	308	709	3074	4405	16 %	10 %
2016	841	1462	4414	5361	27 %	19 %
2017	2035	2723	6125	5860	46 %	33 %

DL-CS results were obtained by searching for “topic” (TS) = “deep learning”, “and”, “research area” (SU) = “computer science”; ML-CS was obtained the same way as DL-CS, only that “deep learning” was replaced by “machine learning”; DL-non-CS was obtained by TS = “deep learning” *not* SU = “computer science” *not* SU = education. Education was removed because entries in this category were not related to our definition of DL. There were also 19 articles in 2011 where deep learning was about education in disciplines other than SU = Education. Therefore, 19 was used as a blank value and also subtracted from the DL-non-CS column. DL/ML-CS is ratio of DL-CS to ML-CS expressed as a percentage. DL/ML-non-CS was obtained similarly.

al., 2018; Hsu et al., 1995, 1997, 2002), from classification and regression trees to random forest (Ho, 1995), and from a Gaussian process (Snelson and Ghahramani, 2006) to a radial basis function network (Moradkhani et al., 2004). There are especially numerous hydrologic applications of ANNs, including the modeling of rainfall-runoff, groundwater, water quality, salinity, rainfall estimation, and soil parameterization (Govindaraju, 2000; Maier et al., 2010 – see a brief discussion in Appendix A in Shen18). Each approach offered useful solutions to a set of problems, but each also faced its own limitations. As a result, over time, some may have grown dispassionate about progress in machine learning, while some others may have concerns about whether DL represents real progress or is just “hype.”

The progress in AI brought forth by DL to various industries and scientific disciplines is revolutionary (Sect. 4 in Shen18) and can no longer be ignored by the hydrologic community. Major technology firms have rapidly adopted and commercialized DL-powered AI (Evans et al., 2018). For example, Google has re-oriented its research priority from “mobile-first” to “AI-first” (Dignan, 2018). The benefits of these industrial investments can now be felt by ordinary users of their services such as machine translation and digital assistants, which can engage in conversations sounding like humans (Leviathan and Matias, 2018). Moreover, AI patents grew at a 34 % compound annual growth rate between 2013 and 2017, apparently after DL’s breakthroughs in 2012 (Columbus, 2018). Also reported in Columbus (2018), more than 65 % of data professionals responded to a survey indicating AI as their company’s most significant data initiative for next year.

DL is gaining adoption in a wide range of scientific disciplines and, in some areas, has started to substantially transform those disciplines. The fast growth is clearly witnessed from literature searches. Since 2011, the number of entries with DL as a topic increased almost exponentially, showing around a 100 % compound annual growth rate through

2017 (Table 1). DL evolved from occupying less than 1 % of machine learning (ML) entries in computer science (CS) in 2011 to 46 % in 2017. This change showcases a massive conversion from traditional machine learning to DL within computer science. Other disciplines lagged slightly behind, but also experienced an exponential increase, with the DL/ML ratio jumping from 0 % in 2011 to 33 % in 2017. As reviewed in Shen18, DL has enhanced the statistical power of data in high-energy physics, and the use of DL accounts for a 25 % increase in the experimental dataset (Baldi et al., 2015). In biology, DL has been used to predict potential pathological implications from genetic sequences (Angermueller et al., 2016). DL models fed with raw-level data have been shown to outperform those with expert-defined features when they predict high-level outcomes, e.g., toxicity, from molecular compositions (Goh et al., 2017). Just like other methods, DL may eventually be replaced by newer ones, but that is not a reason to hold out when progress is possible.

Many of the above-mentioned advances were driven by DL’s domination in AI competitions:

- The ImageNet Challenge is an open competition to evaluate algorithms for object detection and image classification (Russakovsky et al., 2014). Topics change during each contest, and a dataset of ~ 14 million tagged images and videos were cumulatively compiled, with convenient and uniform data access provided by the organizers. The 2010 Challenge was won by a large-scale support vector machine (SVM). A deep convolutional neural network first won this contest in 2012 (Krizhevsky et al., 2012). This victory heralded the exponential growth of DL in popularity. Since then, and until 2017 (the last contest), the vast majority of entrants and all contest winners used CNNs, which are preferred to over other methods by large margins (Schmidhuber, 2015).
- The IJCNN traffic sign recognition contest, which is composed of 50 000 images (48 pixels × 48 pixels), wit-

nessed visual recognition performance (that exceeded human capability) using CNN-based methods (Stalrkamp et al., 2011). CNNs also performed better than humans in the recognition of cancers from medical images (Yu et al., 2016).

- The TIMIT speech corpus is a dataset that holds the recordings from 630 English speakers. Models based on long short-term memory (LSTM) significantly outperformed the hidden Markov model (HMM) results (Graves et al., 2013) in recognizing the speeches. Similarly, LSTM-based methods significantly outperformed all statistical approaches in keyword spotting (Indermuhle et al., 2012), optical character recognition, language identification, text-to-speech synthesis, social signal classification, machine translation, and Chinese handwriting recognition.
- An LSTM-based speech recognition system has achieved “human parity” in conversational speech recognition on the Switchboard corpus (Xiong et al., 2016). A parallel version achieved best-known pixel-wise brain image segmentation results on the MR-BrainS13 dataset (Stollenga et al., 2015).
- A time-series forecasting contest, the Computational Intelligence in Forecasting Competition, was won by a combination of fuzzy and exponential models in 2015 when no LSTM was present, but LSTM won the contest in 2016 (CIF, 2016).

In contrast, only a handful of applications of big data DL could be found in hydrology, but they have already demonstrated great promise. Vision DL has been employed to retrieve precipitation from satellite images, where it demonstrated a more materially superior performance than earlier-generation neural networks (Tao et al., 2017, 2018). A generative adversarial network (GAN) was used to imitate and generate the scanning of images of geologic media (Laloy et al., 2018), where the authors showed realistic replication of training image patterns. Time-series DL was employed to temporally extend satellite-sensed soil moisture observations (Fang et al., 2017) and was found to be more reliable than simpler methods. Time series DL rainfall-runoff models that are confined to certain geographic regions have been created (Kratzert et al., 2018). There are also DL studies, based on smaller datasets, to help predict water flows in the urban environment (Assem et al., 2017) and in the water infrastructure (Zhang et al., 2018). In addition to utilizing big data, DL was able to create valuable, big datasets that could not have otherwise been possible. For example, utilizing DL, researchers were able to generate new datasets for tropical cyclones, atmospheric rivers, and weather fronts (Liu et al., 2016; Matsuoka et al., 2017) by tracking them. Machine learning has also been harnessed to tackle the convection parameterization issue in climate modeling (Gentine et al., 2018).

2.2.2 Technical advances

Underpinning the powerful performance of DL are its technical advances. The deep architectures have several distinctive advantages: (1) Deep networks are designed with the capacity to represent extremely complex functions. (2) After training, the intermediate layers can perform modular functions, which can be migrated to other tasks in a process called transfer learning, and can extend the value of the training data. (3) The hidden layer structures have been designed to automatically extract features, which helps in dramatically reducing labor, expertise, and the trial and error time needed for feature engineering. (4) Compared to earlier models like classification trees, most of the deep networks are differentiable, meaning that we can calculate derivatives of outputs with respect to inputs or the parameters in the network. This feature enables highly efficient training algorithms that exploit these derivatives. Moreover, the differentiability of neural networks enables the querying of DL models for the sensitivity analysis of output to input parameters, a task of key importance in hydrology.

Metaphorically, the intermediate (or hidden) layers in DL algorithms can be understood as workbenches or placeholders for tools that are to be built by deep networks themselves. These hidden layers are trained to calculate certain features from the data, which are then used by downstream layers to predict the dependent variables. For example, Yosinski (2015) showed that some intermediate layers in a deep vision-recognition network are responsible for identifying the location of human or animal faces; Karpathy et al. (2015) showed that some hidden cells in a text prediction network act as line-length counters, while some others keep track of whether the text is in quotes or not. These functionalities were not bestowed by the network designers, but they emerged by themselves after network training. Earlier network architectures either did not have the needed depth or were not designed in an artful way, such that the intermediate layers could be effectively trained. For more technical details, refer to an introduction in Schmidhuber (2015) and Shen (2018).

Given that deep networks can identify features without a human guide, it follows that they may extract features that the algorithm designers were unaware of or did not intentionally encode the network to do. If we could believe that there is latent knowledge about the hydrologic system that humans are not yet aware of but can be determined from data, the automatic extraction of features leads to a potential pathway toward knowledge discovery. For example, deep networks recently showed that grid-like neuron response structures automatically emerge at intermediate network layers for a network trained to imitate how mammals perform navigation, providing strong support to a Nobel-winning neuroscience theory about the functioning of these structures (Banino et al., 2018).

Deep networks may be more robust than simpler models despite their large size if they are regularized properly (regularization techniques apply a penalty to model complexity to make the model more robust) and are chosen based on validation errors in a two-stage approach (Kawaguchi et al., 2017). Effective regularization techniques include (i) early stopping, which monitors the training progress on a separate validation set and stop the training once validation metrics start to deteriorate and/or (ii) novel regularization techniques such as dropout (Srivastava et al., 2014). DL architectures addressed issues like the vanishing gradient (Hochreiter, 1998). The training of the large networks today was computationally implausible until scientists started to exploit the parallel processing power of graphical processing units (GPUs). Nowadays, new application-specific integrated circuits have also been created to specifically tackle DL, although DL architectures are still evolving.

Primary types of successful deep learning architectures include convolutional neural networks (CNN) for image recognition (Krizhevsky et al., 2012; Ranzato et al., 2006), the long short-term memory (LSTM) (Greff et al., 2015; Hochreiter and Schmidhuber, 1997) networks for time-series modeling, variational autoencoders (VAE) (Kingma and Welling, 2013), and deep belief networks for pattern recognition and data (typically image but also text, sound, etc.) generation (Sect. 3.2 in Shen18). Besides these new architectures, a novel generative model concept called the generative adversarial network (GAN) has become an active area of research. The key characteristic of GANs is that they are learned by creating a competition between the actual generative model or “generator” and a discriminator in a zero-sum game framework (Goodfellow et al., 2014) in which these components are learned jointly. Compared to other generative models, GANs potentially offer much greater flexibility in the patterns to be generated. The power of GANs has been recognized recently in the geoscientific community, especially in the machine learning research inspired by physics, where GANs have been used to generate certain complicated physical, environmental, and socio-economic systems (Albert et al., 2018; Laloy et al., 2018).

While showing many advantages, DL models will require a substantial amount of computing expertise. The tuning of hyper-parameters, e.g., network size, learning rate, batch size, etc., often requires a priori experiences and trial and error. The computational paradigm, e.g., computing on graphical processing units, is also substantially different from typical hydrologists’ educational background. The fundamental theories on why DL generalizes so well have not been maturely developed (Sect. 2.7 in Shen18). In the ongoing debates, some argue that a large part of DL’s power comes from memorization while others counter that DL prioritizes learning simple patterns (Arpit et al., 2017; Krueger et al., 2017), and a two-stage procedure (training and testing) also helps (Kawaguchi et al., 2017). Despite these explanations, it has been found in vision DL that deep networks can be fooled by

adversarial examples where small, unperceivable perturbations to input images sometimes cause large changes in predictions, leading to incorrect outcomes (Goodfellow et al., 2015; Szegedy et al., 2013). It remains to be seen whether such adversarial examples exist for hydrologic DL applications. If we can recreate adversarial examples, they can be added into the training dataset to improve the robustness of the model (Ororbia et al., 2016).

2.3 Network interrogative methods to enable knowledge extraction from deep networks

Conventionally, neural networks were primarily used to approximate mappings between inputs and outputs. The focus was put on improving predictive accuracy. In terms of the use of neural networks in scientific research, then, there has been a major concern that DL and more generally machine learning are referred to as black boxes that cannot be understood by humans and, thus, cannot serve to advance scientific understanding. At the same time, data-driven research may lack clearly stated hypotheses, which is in contrast to traditional hypothesis-driven scientific methods. There has been significant pressure from both inside and outside the DL community to make the network decisions more explainable. For example, new (as of January 2018) European data privacy laws dictate that automated individual decision-making, which significantly influences the algorithm’s users, must provide a “right to explanation” where a user can ask for an explanation of an algorithmic decision (Goodman and Flaxman, 2016).

Some recent progress in DL research focused on addressing these concerns. Notably, a new sub-discipline, known as “AI neuroscience” has produced useful interrogative techniques to help scientists interpret the DL model (see literature in Sect. 3.2 in Shen18). The main classes of interpretive methods include (i) the reverse engineering of the hidden layers, which attributes deep network decisions to input features or a subset of inputs; (ii) transferring knowledge from deep networks to interpretable, reduced-order models; and (iii) visualization of network activations. Many scientists have also devised case-by-case ad hoc methods, e.g., to investigate the correlation between inputs and cell activations (Shen, 2018; Voosen, 2017).

So far, interpretive DL methods have not been employed in hydrology or even the geosciences. However, to give some examples from other domains, in medical image diagnosis, some researchers used reverse engineering methods to show which pixel on an image led the network to make its decision about anatomy classifications (Kumar and Menkovski, 2016). They found that the network traced its decisions to image landmarks mostly often used by human experts. In more recent research, AI researchers trained their network not only to classify an image but also to didactically explain why the decision was made and why an image is one class instead of another (Fig. 1). Extending this idea to the pre-




Figure 1. Authors trained a joint classification and explanation network for image classification. The bolded text is a “class-relevant” attribute (a distinguishing attribute for the class) in the explanation. Their classification network extracts visual features (regions on the image) responsible for the decision. Then, the explanation network links these regions to distinguishing words in a dictionary to form an explanation that explains the reason for the classification and why it does not belong to other classes. This level of explanation may be difficult to achieve for hydrologic problems due to limited supervising data (annotated dictionary for classes), but it is possible to borrow the idea of associating features in the input data with some descriptive words (Hendricks et al., 2016).

precipitation retrieval problem in hydrology (Tao et al., 2017, 2018), we could let DL inform us of which features on the satellite cloud image are helpful for reducing bias in precipitation retrieval.

2.4 The complementary research avenue

As the interrogative methods grow further, a research avenue complementary to the traditional hypothesis-driven one emerges toward attaining knowledge (Fig. 2). The data-driven research avenue can be divided into four steps: (i) hypotheses are generated by machine learning algorithms from data; (ii) the validation step is where data is withheld from training and, differently from training, are employed to evaluate the machine-learning-generated hypotheses; (iii) interpretive methods are employed to extract data-consistent and human-understandable hypotheses (Mount et al., 2016; described in Sect. 2.3); (iv) the retained hypotheses are presented to scientists for analysis and further data collection; and the process iterates.

The classical avenue, especially when applied to modeling studies, attracts non-uniqueness and subjectivity. To give a concrete example, consider a classical problem of rainfall-runoff modeling. Suppose a hydrologist found that hydrologic responses in several nearby basins are different. Some basins produce flashier peaks while others have

smaller peaks in summer, large seasonal fluctuation, and large peak streamflows only in winter. Taking a modeling approach, the hydrologist might invoke a conceptual hydrologic model, e.g., Topmodel (Beven, 1997), only to find that the model results do not adequately describe the observed heterogeneity in the rainfall-runoff response. The hydrologist might hypothesize that the different behaviours are due to heterogeneity in soil texture, which is not well represented in the model. Subsequently, the hydrologist incorporates processes that represent soil spatial heterogeneity, such as refined soil pedo-transfer functions that can differentiate between the soil types in different regions. Perhaps with some parameter adjustments, this model can provide streamflow predictions that are qualitatively similar to the observations. This procedure then increases the hydrologist’s confidence that the heterogeneity in soil hydraulic parameters is indeed responsible for their different hydrologic responses. However, this improvement is not conclusive due to process equifinality; there can be alternative processes that can also result in similar outcomes, e.g., the influence of soil thickness, karst geology, terrain, or drainage density. The identification of potential improvement might be dependent on the hydrologist’s intuition or preconceptions, which are nonetheless important but are potentially biased. While the intention of a process-based model may be deductive (Beven, 1989), the example




Figure 2. Comparing two alternative avenues toward gaining knowledge from data. In the classical avenue, scientists compile and interpret data, form hypotheses, and (optionally) build models to describe data and hypotheses (the green pathway). Then the model results with data to affirm or reject the hypotheses and the feedbacks (the yellow pathway) allow the scientist to revise the model and iterate. In the data-driven avenue, scientists collect data and define the target variables of DL models (the green path). Then interpretive methods are invoked to extract data-consistent and human-understandable hypotheses (the yellow path). There must be a hypotheses validation step where data withheld from training is used to evaluate or reject the hypotheses.

process given above is, in fact, abductive reasoning (Josephson and Josephson, 1994), as it seeks a plausible but not exhaustive thorough explanation of the phenomenon. Furthermore, incorporating all the physics into the model may prove technically challenging, computationally impractical, or too time-consuming.

Compared to the classical avenue, the inductive data-driven approach allows us to more efficiently explore a larger set of hypotheses. Although it cannot be said that the machine learning algorithms present no human bias (because inputs are human-defined and some hyper-parameters are empirically adjusted), the larger set of hypotheses presented will at least greatly reduce that risk. First, let us examine a data-driven approach based on classification and regression tree algorithms (CART-based – Fang and Shen, 2017). We could start with physiographic data for many basins in this region, including terrain, soil type, soil thickness, etc. We can use CART to model the process-based model’s errors, which allows us to separate out the conditions under which these errors occur more frequently. We let the pattern emerge out of data without enforcing a strong human preconceived hypothesis. Attention must be paid to the robustness of the data mining and utilize holdout datasets or cross-validation to verify the generality of the conclusion. Data may suggest that soil thickness is the main reason for the error. Or if data do not prefer one hypothesis over the other, then all hypotheses are equally possible and cannot be ruled out. This advantage of DL can be summarized in a short phrase, “an algorithm has no ego”. On a practical level, this approach can more efficiently and simultaneously examine multiple competing hypotheses.

One example of such analyses was carried out by Fang and Shen (2017), where differences in basin storage–streamflow correlations were explained by physical factors using CART, an earlier-generation data mining method (Fig. 3). The data

mining analysis allowed patterns to emerge, which inspired hypotheses about how key factors that control the hydrologic functioning of different systems, such as soil thickness and soil bulk density, are important controls of streamflow–storage relationships. In another example, the data-mining analysis showed that drought recovery time is associated with temperature and precipitation, while biodiversity has only secondary importance (Schwalm et al., 2017). Scientists need to define the predictors and general model types, but they do not pose strongly constraining hypotheses about the controlling factors and instead “let the data speak”. The key to this approach is a large amount of data from which patterns emerge.

However, in working with DL models, we need to further resort to interrogative methods to interpret the results (Fig. 2 right panel). For example, we can construct DL models to predict the errors of the process-based model and then use visualization techniques to see which variable, under which condition, lead to the error. Because DL can absorb a large amount of data, it can find commonality among data as well as identify differences. Whereas CART models are limited by the amount of data and face stability problems in lower branches (data are exponentially less at lower branches), DL models may produce a more robust interpretation.

The machine learning paradigm lends us to finding “unrecognized linkages” (Wagener et al., 2010) or complex patterns in the data that humans could not easily realize or capture. Owing to the strong capability of DL, it can better approximate the “best achievable model” (BAM) for the mapping of relations between inputs and output. As such, it lends support to measuring the information content contained in the inputs about the output. Nearing et al. (2016) utilized Gaussian process regression to approximate the BAM. DL can play similar roles and can also allow for modeling, perhaps in a more




Figure 3. We calculated storage–streamflow correlation patterns over the contiguous United States (CONUS) and divided small or mesoscale basins into multiple classes. We studied which physical factors most cleanly separate different correlation patterns. Panel (a) shows the distribution of basin classes on and to the south of the Appalachian ranges. Panel (b) shows a one-level CART model that finds the most effective criterion to split basins into clusters with different distances to class #1 (storage and streamflow are highly correlated across all flow regimes) correlation patterns. This CART model does not achieve complete separation of class #1 basins, because it is based on criteria that could be established using input attributes, and there is also resolution limitation with terrestrial water storage observations, which are from the GRACE satellite. However, on a large spatial scale, what separates the blue class and the green class (low correlation between storage and high flow bands) turned out to be soil thickness. It suggests the blue basins in the south have high correlation *because* they have thick soils, which facilitates infiltration, water storage, and increases the importance of groundwater-contributed streamflow (Fang and Shen, 2017).

thorough way. The simplicity of building DL models and altering inputs makes them an ideal test bed for new ideas.

Outputs from the hidden layers of deep networks can now be visualized to gain insights about the transformations performed on the input data by the network (Samek et al., 2017). For image recognition tasks, one can invert the DL model to find out the parts of the inputs that led the network to make a certain decision (Mahendran and Vedaldi, 2015). There are also means to visualize outputs from recurrent networks, e.g., showing the conditions under which certain cells are activated (Karpathy et al., 2015). These visualizations can illustrate the relationships that the data-driven model has identified.

Considering the above potential benefits, the data-driven avenue should at least be given an opportunity in hydrological sciences discovery. However, this avenue may be uncomfortable for some. In the classical avenue, the scientist must originate the hypotheses before constructing models; in the data-driven one, the data mining and knowledge discovery process is a precursor step to the main hypotheses formation – hypotheses cannot be generated before the data mining analysis (Mount et al., 2016). Specifically, hypotheses can no longer be clearly stated during the proposal stage of research.

Granted, the interrogative methods as a whole are new, and time is required for them to grow. We need to note that the nascent “DL neuroscience” literature did not exist until 2015. However, if we outright reject the complementary avenue based on the habitual thinking that neural networks are black boxes, we may deny ourselves an opportunity for breakthroughs.

3 Challenges and opportunities for DL in hydrology

The field of hydrology has a unique set of challenges that are also research opportunities for DL. Many of these science challenges have, to date, not been effectively addressed using traditional methods and cannot be sufficiently tackled by individual research groups. Some challenges for which DL approaches might be exploited are presented below.

Observations in hydrology and water science are generally regionally and temporally imbalanced. For example, while streamflow observations are relatively dense in the United States, such data are sparse in many other parts of the world, because measurements have either not been made or are not made accessible. There is often a dearth of observations that can be used as comprehensive training datasets for DL algorithms. Few hydrologic applications have as much data as available to standard AI research applications such as image recognition or natural language processing. Remote sensing of hydrologic variables also has limitations, including the effects of canopies and clouds, which can limit observations, the temporal density of observations because of orbital paths, and observation footprints, which create challenges when trying to validate satellite observations with field point measurements. A body of literature studying this problem across different geographic regions can be loosely summarized under the topic of “prediction in ungauged basins” (PUB – Hrachowitz et al., 2013). PUB problems pose a significant challenge to data-driven methods.

Global change is altering the hydrologic and related cycles, and hydrologists must now make predictions in anticipation of changes beyond previously observed ranges (Wa-

gener et al., 2010). More frequent extremes have been observed for many parts of the world, and such extremes have been projected to occur even more frequently in the future (Stocker et al., 2013). Data-driven methods must demonstrate their capability to make reasonable predictions when applied out of the range of the training dataset.

Observations of the water cycle tend to focus on one aspect of the water cycle and seldom offer a complete description. For example, we can estimate total terrestrial water storage (Wahr, 2004) or the top 5 cm surface soil moisture via multiple satellite missions. It is difficult, however, to directly combine such observations of components of the water cycle into a complete picture of the water cycle. Challenges, then, are merging distinct observations with all their space–time discontinuities to aid predictions, validating models, and providing a more complete understanding of the global water cycle.

Hydrologic data are accompanied by a large amount of strongly heterogeneous (Blöschl, 2006) “contextual variables” such as land use, climate, geology, and soil properties. The proper scale at which to represent heterogeneity in natural systems is a vexing problem (Archfield et al., 2015). For example, representing the micro-scale heterogeneity in soil moisture and texture is not computationally realistic in hydrologic models. The scale at which heterogeneity should be represented varies with the setting and elements of the water cycle (Ajami et al., 2016). Moreover, while we recognize that heterogeneity exists in contextual features, many of these features, such as soil properties and hydrogeology, are poorly characterized across landscapes, but both features play an important role in controlling water movement. Heterogeneity needs to be adequately represented without radically bloating the parameter space of the models and thus data demand.

Furthermore, the heterogeneous physiographic factors co-evolve and covary (Troch et al., 2013) with complicated causal and non-causal connections (Faghmous and Kumar, 2014). The relationships of soil, terrain, and vegetation are further conditioned by geologic and climate history and often do not transfer to other regions (Thompson et al., 2006). Consequently, training with insufficient data may result in overfitted data-driven models or many alternative DL models that cannot be rejected. On the flip side, such complexity due to co-evolution also precludes a reductionist approach where all or most of these relationships are clearly described in fundamental laws.

Hydrologic problems fit poorly into the template of problems for which standard network structures (Sect. 3.2 in Shen18) are designed, i.e., pure image recognition, time-series prediction problems, or a mixture of both. For example, catchment hydrologic problems are characterized by both spatially heterogeneous but temporally static attributes (topography and hydrogeology) and temporal (atmospheric forcing) dimensions. Such input dimensions are not effi-

ciently represented by typical input dimensions of LSTM or CNN.

Because large and diverse datasets are needed for DL application, access to properly pre-processed and formatted datasets presents practical challenges. The steps of data compilation, pre-processing, and formatting often occupy too much time for researchers. Many of the processing tasks for images cannot be handled by individual research groups. Compared to the DL communities in AI and chemistry, etc., the DL learning community in hydrology is not sufficiently coordinated, resulting in the significant waste of effort and “reinvention of wheels”.

Deep generative models such as GANs can be used for the stochastic generation of natural textures. This capability has recently led to methodological advances in subsurface hydrology (Laloy et al., 2017, 2018; Mosser et al., 2017) where the ability to efficiently and accurately simulate complicated geologic structures with given (non-Gaussian) geostatistical properties is of paramount importance for the uncertainty quantification of subsurface flow and transport models. However, amongst other directions for future research, more work is needed to (i) generate the complete range of structural complexity observed in geologic layers, (ii) deal efficiently with large 3-D domains, and (iii) account for various types of direct (e.g., observed geologic facies at a given location, mean property value over a specific area, etc.) and indirect (e.g., measured hydrologic state variables to be used within an inverse modeling procedure) conditioning data in the simulation.

4 A community roadmap toward DL-powered scientific advances in hydrology

Despite the challenges articulated above, here we propose a community roadmap for advancing hydrologic sciences using DL (Fig. 4). A well-coordinated community is much more efficient in resolving problems, as we have seen in other scientific endeavours. As Montanari et al. (2013) noted, “future science must be based on an interdisciplinary approach” and “the research challenges in hydrology for the next 10 years should be tackled through a collective effort”. We see that several steps are crucial in this roadmap: the integration of physical knowledge, process-based models (PBMs), and DL; using DL to infer unknown quantities; utilizing community approaches in sharing and accessing data; open and transparent model competitions; baseline models and visualization packages; and an education program that introduces data-driven methods at various levels.

4.1 Integrating physical knowledge, process-based models, and DL models

To address the challenge of data limitations (data quantity), we envision that a critical and necessary step is to more




Figure 4. A roadmap toward DL-powered scientific discovery in hydrologic science. Data availability can be increased by (green arrows) collecting and compiling existing data, incorporating novel data sources such as those collected by citizen scientists, remote sensing, and modeling datasets. DL can be employed to predict data that are currently difficult to observe. The modeling competitions and the integration between PBM and DL will build an important shared computing and analytic infrastructure, which together with data sources, supports a wide range of hydrologic applications. Interpretive methods should be attempted to extract knowledge from trained deep networks (orange arrows). Underpinning these activities is the enhanced, community-based educational program for machine learning in hydrology (purple arrows). However, these activities, especially the modeling competitions, might in turn feed back into the educational activity.

organically integrate hydrologic knowledge, process-based models, and DL. Process-based models, as they are derived from underlying physics, require less data for calibration than data-driven models. They can provide estimates for spatial and temporal data gaps and unobserved hydrologic processes. Well-constructed PBMs should also be able to represent temporal changes and trends. However, because data-driven models directly target observations, these models may have better performance in locations and periods where data are available. Also, as discussed earlier, data-driven models are less prone to a priori model structural error than PBMs. We should maximally utilize the best features of both types of models.

There will be a diversity of approaches with which PBMs and data-driven models could be combined. Karpatne et al. (2017) compiled a list of approaches of what they collectively call “theory-guided data science”, which include (i) using knowledge to design data-driven models; (ii) using knowledge to initialize network states; (iii) using physical knowledge to construct priors to constrain the data-driven models; (iv) using knowledge-based constrained optimization (although this may be difficult to implement in practice); (v) using theory as regularization terms for the data-driven model, which will force the model to respect these constraints; and (vi) learning hybrid models, where the data-driven method is used as surrogate for certain parts of the

physical model. One may also impose multiple learning objectives based on the knowledge of the problem.

There are a multitude of potential approaches, and this list can be further expanded to accommodate various objectives. First, we can focus on PBM errors (the difference between PBM simulation and observations). Non-deep machine learning has already shown promise in correcting PBM errors. Abramowitz et al. (2006) developed an ANN to predict the error in net ecosystem exchange from a land surface model and achieved a 95 % reduction in annual error. More importantly, an ANN trained to correct the error at one biome corrects the PBM in another biome with a different temperature regime (Abramowitz et al., 2007). In the context of weather forecasts, machine learning methods were used to learn the patterns from past forecasting errors (Delle Monache et al., 2011, 2013), leading to a 20 % improvement in performance for events of similar characteristics (Junk et al., 2015). Their results suggest that PBMs make structural errors that are independent of the state-variable regimes, although there is a lack of theories to guide the separation of error types. We envision that PBMs can better resolve the impacts of regime changes, while DL can better capture state-independent error patterns and carry out mild state-dependent extrapolations. A co-benefit of modeling the PBM error is that insights are gained about the PBM. Using interrogative methods to reverse engineer what DL has learned about PBM error provides possible avenues for improving the underlying PBM processes.

Second, PBMs can augment input data for DL models. PBMs can be used to increase supervising data for DLs, for example, for climate or land-use scenarios that have not existed presently and to augment existing data. Given model structural uncertainty (uncertainty with hydrologic processes), frameworks like the Structure for Unifying Multiple Modeling Alternatives (Clark et al., 2015) and automated model building (Marshall, 2017) could be employed to generate a range of outputs. Furthermore, with limited data, we may not be able to reject alternative DL models that could generate unphysical or unrealistic outputs. Providing PBM simulations as either training data or regularization terms can help nudge DL models to generate physically meaningful outputs. The extent to which errors in PBM model results affect DL outcomes remains to be explored. A theoretical framework is lacking for separately estimating aleatory uncertainty (resulting from data noise), epistemic uncertainty (resulting from PBM error and training data paucity), and uncertainty due to a regime shift. The advantages and disadvantages of various approaches could be systematically and efficiently evaluated in a community-coordinated fashion.

4.2 Multifaceted, community-coordinated hydrologic modeling competitions

There are many possible approaches and many alternative model structures for using DL to make hydrologic predic-

tions and to provide insight into hydrologic processes. In light of these challenges, we argue that open, fast, and standardized competitions are an effective way of accelerating the progress. The competitions can evaluate the models not only in terms of predictive performance but also in the attainment of understanding.

The impacts of competitions are best evidenced in the community-coordinated AI challenges, which use a standardized set of problems. These competitions have strongly propelled the advances in AI. Some have argued that the contributions of the ImageNet dataset and the competition may be more significant than the winning algorithms arising from these contests (Gershgorin, 2017). New methods can be evaluated objectively and disseminated rapidly through competitions. Because the problems are standardized, they remove biases due to data sources and pre-processing. The community can quickly learn the advantages and disadvantages of alternative model design through these competitions, which also encourages reproducibility.

We envision multifaceted hydrologic modeling competitions where various models ranging from process-based ones to DL ones are evaluated and compared. The coordinators can, for example, provide a set of standard atmospheric forcings, landscape characteristics, and observed variables and provide targeted questions that participants must address. Importantly, the evaluation criteria should include not only performance-type criteria such as model efficiency coefficients and bias but also *qualitative* and *explanatory* ones, such as explanations for control variables and model errors. Overly simplified or poorly constructed models may provide more accessible explanations, but they might be misleading because the models may be overfitted to a given situation. Their simplicity may also constrain their ability to digest large datasets as a way of reducing uncertainty. Multifaceted competitions allow us to also identify the “Pareto front” of interpretability and performance, and they help rule out false explanations. The objective of the competition is not only to seek the best simulation performance but also those methods that offer deeper insights into hydrologic processes.

Another important value of competitions is that organizers will provide a standard input dataset and well-defined tasks, which the entire community can leverage. A substantial amount of effort is required to establish such a dataset, which may only be possible under a specifically designed project. Moreover, open competitions in the computer science field have produced well-known models such as AlexNet (Krizhevsky et al., 2012), GoogLeNet (Szegedy et al., 2015), etc. These models serve as benchmarks and quick entry points for others. They can greatly improve reproducibility and the effectiveness of comparisons.

4.3 Community-shared resources and broader involvement

A useful approach in addressing the major obstacle of data limitation is to increase our data repositories and open access to existing data. The value of data can be greatly enhanced by centralized data compilations, a task many institutions are already undertaking. For example, the Consortium of Universities for the Advancement of Hydrologic Science, Inc. (CUAHSI) provides access to large amounts of hydrologic data (CUAHSI, 2018a). In another example, in 2015, a project called Collaborative Research Actions (Endo et al., 2015) was proposed in the Belmont Forum, which is a group of the world’s major and emerging funders of global environmental change research. Many scientists from different countries join the project and focus on the same issue, the food–energy–water nexus. They shared their data (heterogeneous data) and research results from different regions.

ML has already been used to create useful hydrologic datasets such as soil properties (Chaney et al., 2016; Schaap et al., 2001), land cover (Helber et al., 2017; Zhu et al., 2017), and cyclones (Liu et al., 2016). We envision that there will continue to be significant progress in this regard, and the key to success will be the availability of ground-truth datasets. Using data sharing standards will advance data sharing across domains (WaterML2, 2018). Providing access to data through web services such as those used by CUAHSI negates the problem of storing data in a single location and enhances the ability for them to be discovered. Data brokers also provide more channels to share experiences, scholarly discussions, and debates along with the generation of data.

An important area where DL is expected to deliver significant value is the analysis of big and sub-research-quality data such as those collected by citizen scientists. Many aspects of the water cycle are directly accessible by everyone. Citizen scientists already gather data about precipitation (CoCoRaHS, 2018), temperature, humidity, soil moisture, river stage (CrowdHydrology, 2018), and potentially groundwater levels. These quantities can be measured using inexpensive instruments such as cameras, pressure gauges, and moisture sensors. Volunteer scientists can also be solicited for data in places where such data can best reduce the uncertainty of the DL model, as in a framework called active learning (Settles, 2012). Social data have been used to help monitor flood inundation (Sadler et al., 2018; Wang et al., 2018). Crowdsourced data have played roles in DL research, where a large but noisy dataset was argued to be more useful than a much smaller but well-curated dataset (Huang et al., 2016; Izadinia et al., 2015). Even though there are problems related to data quality, these can be overcome using AI approaches. An important co-benefit of involving citizen scientists is education and outreach to the public. The active engagement is much more effective when the public has a stake in the research outcomes.

4.4 Education

A major barrier to realizing the benefits of data science and DL lies in our undergraduate and graduate curricula. Little in hydrologists' standard curriculum prepares them for a future with substantially more data-driven science. Statistical courses often do not cover machine learning basics, while data mining courses offered by computer science departments lack the connections to the water discipline. Given the interdisciplinary nature of hydrology, it has been long recognized that it takes a community to raise a hydrologist (Merwade and Ruddell, 2012; Wagener et al., 2012). We propose a concerted effort by current hydrologic machine learning researchers, along with participation from computer scientists, to pool and share educational content. Such efforts will form the basis of a hydrologic data mining curriculum and leverage the wit of the community. Collaborations may form through either grassroots collaborations or institutionally supported education projects (e.g., CUAHSI, 2018b). The open competitions would be a great source of education materials. A diversity of models that have been evaluated and contrasted help clarify the pros and cons of different methods. Shared datasets, DL algorithms, and data pre-processing software can be leveraged in classrooms.

As with the design of any education effort, it is important to consider inclusiveness and diversity. Especially for hydrologic DL, the source field of AI appears to have an extremely poor track record of gender balance (Simonite, 2018). The reason for such an imbalance could be rooted in introductory computer science classes, undergraduate curricula, and social stereotypes. Research has found that the introductory computer science classes, especially those taken by non-majors, are instrumental in developing a desire to stay in the field (Lehman, 2017). In addition, the portrayal of gender stereotypes regarding computing and the increase in weed-out courses (Aspray, 2016) have discouraged women students in computer science (Sax et al., 2017). To counter such negative impacts, the introductory courses in the curriculum need to assume little prior programming experiences. Special attention must be paid by the educators to shatter the stereotypes. On the other hand, the richness of natural beauty in hydrology and the connection between data and the real world may be employed to help bridge the gender gap.

5 Concluding remarks

In this opinion paper, we argue that hydrologic scientists ought to give consideration to a research avenue that complements traditional approaches wherein DL-powered data mining is used to generate hypotheses, predictions, and insights. Although there may have been strong reservations toward black-box approaches in the past, recent efforts and advances have been made in the interpretation and understanding of deep learning networks. The black-box percep-

tion of ML in hydrology is perhaps a self-reinforcing curse on this juvenile field, as a rejection of the research avenue based on this perception will, in turn, jeopardize the development of more transparent algorithms. Progress in hydrology and other disciplines show that there is substantial promise in incorporating DL into hydrologists' toolbox. However, challenges such as data limitation and model variability demand a community-coordinated approach.

We have also argued for open hydrologic competitions that emphasize both performance and the ability to explain. These competitions will greatly improve the growth of the field as a whole. They serve as valuable "organizing events", where different threads in algorithm development, model evaluations and comparisons, reproducibility tests, dataset compilation, resource sharing, and community organization all come to a convergence to spur growth in the field.

Data availability. No data sets were used in this article.

Author contributions. CS initiated the discussion and the Opinion paper. All co-authors contributed to the writing of the paper.

Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. We thank Matthew McCabe, Keith Sawicz, and an anonymous reviewer for their valuable comments, which helped to improve the paper. We thank the editor for handling the manuscript. The discussion for this opinion paper was supported by U.S. Department of Energy under contract DE-SC0016605. The funding to support the publication of this article was provided by the U.S. National Science Foundation (NSF) grant EAR-1832294 to CS, Canadian NSERC-DG 403047 to AE, NSF grant EAR-1338606 to JB, Key R&D projects of the Science and Technology department in Sichuan Province grant 2018SZ0343 and the open fund of State Key Laboratory of Hydraulics and Mountain River Engineering, Sichuan University to XL, Belgian Nuclear Research Centre to EL, and NSF grant CCF-1317560 to DK.

Edited by: Louise Slater

Reviewed by: Matthew McCabe, Keith Sawicz, and one anonymous referee

References

- Abramowitz, G., Gupta, H., Pitman, A., Wang, Y., Leuning, R., Cleugh, H., Hsu, K., Abramowitz, G., Gupta, H., Pitman, A., Wang, Y., Leuning, R., Cleugh, H., and Hsu, K.: Neural Error Regression Diagnosis (NERD): A Tool for Model Bias Identification and Prognostic Data Assimilation, *J. Hydrometeorol.*, 7, 160–177, <https://doi.org/10.1175/JHM479.1>, 2006.

- Abramowitz, G., Pitman, A., Gupta, H., Kowalczyk, E., Wang, Y., Abramowitz, G., Pitman, A., Gupta, H., Kowalczyk, E., and Wang, Y.: Systematic Bias in Land Surface Models, *J. Hydrometeorol.*, 8, 989–1001, <https://doi.org/10.1175/JHM628.1>, 2007.
- Ajami, H., Khan, U., Tuteja, N. K., and Sharma, A.: Development of a computationally efficient semi-distributed hydrologic modeling application for soil moisture, lateral flow and runoff simulation, *Environ. Model. Softw.*, 85, 319–331, <https://doi.org/10.1016/J.ENVSOF.2016.09.002>, 2016.
- Albert, A., Strano, E., Kaur, J., and Gonzalez, M.: Modeling urbanization patterns with generative adversarial networks, arXiv:1801.02710, available at: <http://arxiv.org/abs/1801.02710>, last access: 24 March 2018.
- Allamano, P., Croci, A., and Laio, F.: Toward the camera rain gauge, *Water Resour. Res.*, 51, 1744–1757, <https://doi.org/10.1002/2014WR016298>, 2015.
- Angermueller, C., Pärnamaa, T., Parts, L., and Stegle, O.: Deep learning for computational biology, *Mol. Syst. Biol.*, 12, 878, <https://doi.org/10.15252/MSB.20156651>, 2016.
- Archfield, S. A., Clark, M., Arheimer, B., Hay, L. E., McMillan, H., Kiang, J. E., Seibert, J., Hakala, K., Bock, A., Wagener, T., Farmer, W. H., Andréassian, V., Attinger, S., Viglione, A., Knight, R., Markstrom, S., and Over, T.: Accelerating advances in continental domain hydrologic modeling, *Water Resour. Res.*, 51, 10078–10091, <https://doi.org/10.1002/2015WR017498>, 2015.
- Arpit, D., Jastrzebski, S., Ballas, N., Krueger, D., Bengio, E., Kanwal, M. S., Maharaj, T., Fischer, A., Courville, A., Bengio, Y., and Lacoste-Julien, S.: A Closer Look at Memorization in Deep Networks, in *Proceedings of the 34 th International Conference on Machine Learning*, Sydney, Australia, PMLR 70, available at: <https://arxiv.org/abs/1706.05394>, last access: 19 November 2017.
- Aspen: Internet of Water: Sharing and Integrating Water Data for Sustainability, A Rep. from Aspen Inst. Dialogue Ser. Water Data, available at: <https://www.aspeninstitute.org/publications/internet-of-water/> (last access: 27 August 2018), 2017.
- Aspray, W.: Women and underrepresented minorities in computing: a historical and social study, Springer, Basel, Switzerland, 2016.
- Assem, H., Ghariba, S., Makrai, G., Johnston, P., Gill, L., and Pilla, F.: Urban Water Flow and Water Level Prediction Based on Deep Learning, in *ECML PKDD 2017: Machine Learning and Knowledge Discovery in Databases*, 317–329, Springer, Cham., 2017.
- Baldi, P., Sadowski, P., and Whiteson, D.: Enhanced Higgs Boson to $\tau^+\tau^-$ Search with Deep Learning, *Phys. Rev. Lett.*, 114, 111801, <https://doi.org/10.1103/PhysRevLett.114.111801>, 2015.
- Banino, A., Barry, C., Uria, B., Blundell, C., Lillicrap, T., Mirowski, P., Pritzel, A., Chadwick, M. J., Degris, T., Modayil, J., Wayne, G., Soyer, H., Viola, F., Zhang, B., Goroshin, R., Rabinowitz, N., Pascanu, R., Beattie, C., Petersen, S., Sadik, A., Gaffney, S., King, H., Kavukcuoglu, K., Hassabis, D., Hassell, R., and Kumaran, D.: Vector-based navigation using grid-like representations in artificial agents, *Nature*, 557, 429–433, <https://doi.org/10.1038/s41586-018-0102-6>, 2018.
- Baughman, A. K., Chuang, W., Dixon, K. R., Benz, Z., and Basilico, J.: DeepQA Jeopardy! Gamification: A Machine-Learning Perspective, *IEEE Trans. Comput. Intell. AI Games*, 6, 55–66, <https://doi.org/10.1109/TCIAIG.2013.2285651>, 2014.
- Beven, K.: Changing ideas in hydrology – The case of physically-based models, *J. Hydrol.*, 105, 157–172, [https://doi.org/10.1016/0022-1694\(89\)90101-7](https://doi.org/10.1016/0022-1694(89)90101-7), 1989.
- Beven, K.: Topmodel?: A Critique, *Hydrol. Process.*, 11, 1069–1085, 1997.
- Blöschl, G.: Hydrologic synthesis: Across processes, places, and scales, *Water Resour. Res.*, 42, W03S02, <https://doi.org/10.1029/2005WR004319>, 2006.
- Chaney, N. W., Herman, J. D., Ek, M. B., and Wood, E. F.: Deriving Global Parameter Estimates for the Noah Land Surface Model using FLUXNET and Machine Learning, *J. Geophys. Res. Atmos.*, 121, 13218–13235, <https://doi.org/10.1002/2016JD024821>, 2016.
- Chang, L.-C., Shen, H.-Y., and Chang, F.-J.: Regional flood inundation nowcast using hybrid SOM and dynamic neural networks, *J. Hydrol.*, 519, 476–489, <https://doi.org/10.1016/J.JHYDROL.2014.07.036>, 2014.
- Chen, I.-T., Chang, L.-C., and Chang, F.-J.: Exploring the spatio-temporal interrelation between groundwater and surface water by using the self-organizing maps, *J. Hydrol.*, 556, 131–142, <https://doi.org/10.1016/J.JHYDROL.2017.10.015>, 2018.
- CIF: Results, Int. Time Ser. Forecast. Compet. – Comput. Intell. Forecast., available at: <http://irafm.osu.cz/cif/main.php?c=Static&page=results> (last access: 24 March 2018), 2016.
- Clark, M. P., Nijssen, B., Lundquist, J. D., Kavetski, D., Rupp, D. E., Woods, R. A., Freer, J. E., Gutmann, E. D., Wood, A. W., Brekke, L. D., Arnold, J. R., Gochis, D. J., and Rasmussen, R. M.: A unified approach for process-based hydrologic modeling: 1. Modeling concept, *Water Resour. Res.*, 51, 2498–2514, <https://doi.org/10.1002/2015WR017198>, 2015.
- CoCoRaHS: Community Collaborative Rain, Hail and Snow Network (CoCoRaHS), “Volunteers Work. together to Meas. Precip. across nations”, available at: <https://www.cocorahs.org/>, last access: 23 August 2018.
- Columbus, L.: Roundup Of Machine Learning Forecasts And Market Estimates, 2018, *Forbes Contrib.*, available at: <https://www.forbes.com/sites/louisacolumbus/2018/02/18/roundup-of-machine-learning-forecasts-and-market-estimates-2018/#7d3c5bd62225>, last access: 30 July 2018.
- CrowdHydrology: CrowdHydrology, available at: <http://crowdhydrology.geology.buffalo.edu/>, last access: 23 August 2018.
- CUAHSI: Consortium of Universities Allied for Water Research, Inc (CUASI) Cyberseminars, available at: <https://www.cuahsi.org/education/cyberseminars/>, last access: 8 October 2018a.
- CUAHSI: Data Driven Education, available at: <https://www.cuahsi.org/education/data-driven-education/>, last access: 26 September 2018b.
- CUAHSI: HydroClient, available at: <http://data.cuahsi.org/>, last access: 19 August 2018c.
- Delle Monache, L., Nipen, T., Liu, Y., Roux, G., and Stull, R.: Kalman Filter and Analog Schemes to Postprocess Numerical Weather Predictions, *Mon. Weather Rev.*, 139, 3554–3570, <https://doi.org/10.1175/2011MWR3653.1>, 2011.
- Delle Monache, L., Eckel, F. A., Rife, D. L., Nagarajan, B., and Searight, K.: Probabilistic Weather Prediction with an Analog Ensemble, *Mon. Weather Rev.*, 141, 3498–3516, <https://doi.org/10.1175/MWR-D-12-00281.1>, 2013.

- Dignan, L.: Google Research becomes Google AI to reflect AI-first ambitions, *zdnet.com*, available at: <https://www.zdnet.com/article/google-research-becomes-google-ai-to-reflect-ai-first-ambitions/>, last access: 20 August 2018.
- Drucker, H., Burges, C. J. C., Kaufman, L., Smola, A., and Vapnik, V.: Support vector regression machines, *Proc. 9th Int. Conf. Neural Inf. Process. Syst.*, 155–161, available at: <https://dl.acm.org/citation.cfm?id=2999003> (last access: 5 January 2018), 1996.
- Endo, A., Burnett, K., Orenco, P., Kumazawa, T., Wada, C., Ishii, A., Tsurita, I., and Taniguchi, M.: Methods of the Water-Energy-Food Nexus, *Water*, 7, 5806–5830, <https://doi.org/10.3390/w7105806>, 2015.
- Entekhabi, D.: The Soil Moisture Active Passive (SMAP) mission, *Proc. IEEE*, 98, 704–716, <https://doi.org/10.1109/JPROC.2010.2043918>, 2010.
- Evans, H., Gervet, E., Kuchembuck, R., and Hu, M.: Will You Embrace AI Fast Enough? ATKearney Operations & Performance Transformation report, available at: <https://www.atkearney.com/operations-performance-transformation/article/?a/will-you-embrace-ai-fast-enough>, last access 10 August 2018.
- Faghmous, J. H. and Kumar, V.: A Big Data Guide to Understanding Climate Change: The Case for Theory-Guided Data Science, *Big Data*, 2, 155–163, <https://doi.org/10.1089/big.2014.0026>, 2014.
- Fang, K. and Shen, C.: Full-flow-regime storage-streamflow correlation patterns provide insights into hydrologic functioning over the continental US, *Water Resour. Res.*, 53, 8064–8083, <https://doi.org/10.1002/2016WR020283>, 2017.
- Fang, K., Shen, C., Kifer, D., and Yang, X.: Prolongation of SMAP to Spatio-temporally Seamless Coverage of Continental US Using a Deep Learning Neural Network, *Geophys. Res. Lett.*, 44, 11030–11039, <https://doi.org/10.1002/2017GL075619>, 2017.
- Gatys, L. A., Ecker, A. S., and Bethge, M.: Image Style Transfer Using Convolutional Neural Networks, 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2414–2423, 2016.
- Gentine, P., Pritchard, M., Rasp, S., Reinaudi, G., and Yacalis, G.: Could Machine Learning Break the Convection Parameterization Deadlock?, *Geophys. Res. Lett.*, 45, 5742–5751, <https://doi.org/10.1029/2018GL078202>, 2018.
- Gershgorin, D.: The data that transformed AI research-and possibly the world, *Quartz*, available at: <https://qz.com/1034972/the-data-that-changed-the-direction-of-ai-research-and-possibly-the-world/>, last access: 8 October 2018, 2017.
- Gleeson, T., Moosdorf, N., Hartmann, J., and van Beek, L. P. H.: A glimpse beneath earth's surface: GLobal HYdrogeology MaPS (GLHYMPS) of permeability and porosity, *Geophys. Res. Lett.*, 41, 3891–3898, <https://doi.org/10.1002/2014GL059856>, 2014.
- Goh, G. B., Hodas, N. O., and Vishnu, A.: Deep learning for computational chemistry, *J. Comput. Chem.*, 38, 1291–1307, <https://doi.org/10.1002/jcc.24764>, 2017.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y.: Generative Adversarial Networks, in *Proceedings of the 27th International Conference on Neural Information Processing Systems (NIPS'14)*, available at: <http://arxiv.org/abs/1406.2661> (last access: 25 February 2017), 2014.
- Goodfellow, I., Shlens, J., and Szegedy, C.: Explaining and Harnessing Adversarial Examples, *International Conference on Learning Representations*, available at: <http://arxiv.org/abs/1412.6572> (last access: 25 February 2017), 2015.
- Goodman, B. and Flaxman, S.: European Union regulations on algorithmic decision-making and a “right to explanation”, *arXiv:1606.08813*, available at: <http://arxiv.org/abs/1606.08813> (last access: 7 February 2018), 2016.
- Govindaraju, R. S.: Artificial Neural Networks in Hydrology. II: Hydrologic Applications, *J. Hydrol. Eng.*, 5, 124–137, [https://doi.org/10.1061/\(ASCE\)1084-0699\(2000\)5:2\(124\)](https://doi.org/10.1061/(ASCE)1084-0699(2000)5:2(124)), 2000.
- Graves, A., Mohamed, A., and Hinton, G.: Speech recognition with deep recurrent neural networks, in: 2013 IEEE International Conference on Acoustics, Speech and Signal Processing, 6645–6649, IEEE, Vancouver, Canada, 2013.
- GRDC: River Discharge Data, Glob. Runoff Data Cent., available at: http://www.bafg.de/GRDC/EN/02_srvcs/21_tmsrs/riverdischarge_node.html, last access: 28 July 2017.
- Greff, K., Srivastava, R. K., Koutník, J., Steunebrink, B. R., and Schmidhuber, J.: LSTM: A Search Space Odyssey, available at: <http://arxiv.org/abs/1503.04069> (last access: 18 July 2016), 2015.
- Guo, Y., Liu, Y., Oerlemans, A., Lao, S., Wu, S., and Lew, M. S.: Deep learning for visual understanding: A review, *Neurocomputing*, 187, 27–48, <https://doi.org/10.1016/J.NEUCOM.2015.09.116>, 2016.
- Hall, D. K., Riggs, G. A., and Salomonson, V. V.: MODIS/Terra Snow Cover 5-Min L2 Swath 500 m, Version 5, Boulder, Colorado USA, 2006.
- He, K., Zhang, X., Ren, S., and Sun, J.: Deep Residual Learning for Image Recognition, 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 770–778, 2016.
- Helber, P., Bischke, B., Dengel, A., and Borth, D.: EuroSAT: A Novel Dataset and Deep Learning Benchmark for Land Use and Land Cover Classification, available at: <http://arxiv.org/abs/1709.00029> (last access: 2 September 2018), 2017.
- Hendricks, L. A., Akata, Z., Rohrbach, M., Donahue, J., Schiele, B., and Darrell, T.: Generating Visual Explanations, in: *Computer Vision – ECCV 2016*, 3–19, available at: https://doi.org/10.1007/978-3-319-46493-0_1, Springer, Cham, Amsterdam, The Netherlands, 2016.
- Hiroi, K. and Kawaguchi, N.: FloodEye: Real-time flash flood prediction system for urban complex water flow, 2016 IEEE SENSORS, 1–3, 2016.
- Hirschberg, J. and Manning, C. D.: Advances in natural language processing, *Science*, 349, 261–266, <https://doi.org/10.1126/science.aaa8685>, 2015.
- Ho, T. K.: Random decision forests, in *Proceedings of 3rd International Conference on Document Analysis and Recognition*, IEEE Comput. Soc. Press., 1, 278–282, 1995.
- Hochreiter, S.: The Vanishing Gradient Problem During Learning Recurrent Neural Nets and Problem Solutions, *Int. J. Uncertainty, Fuzziness Knowledge-Based Syst.*, 6, 107–116, <https://doi.org/10.1142/S0218488598000094>, 1998.
- Hochreiter, S. and Schmidhuber, J.: Long Short-Term Memory, *Neural Comput.*, 9, 1735–1780, <https://doi.org/10.1162/neco.1997.9.8.1735>, 1997.
- Horsburgh, J. S., Morsy, M. M., Castronova, A. M., Goodall, J. L., Gan, T., Yi, H., Stealey, M. J., and Tarboton, D. G.: HydroShare:

- Sharing Diverse Environmental Data Types and Models as Social Objects with Application to the Hydrology Domain, *J. Am. Water Resour. As.*, 52, 873–889, <https://doi.org/10.1111/1752-1688.12363>, 2016.
- Hrachowitz, M., Savenije, H. H. G., Blöschl, G., McDonnell, J. J., Sivapalan, M., Pomeroy, J. W., Arheimer, B., Blume, T., Clark, M. P., Ehret, U., Fenicia, F., Freer, J. E., Gelfan, A., Gupta, H. V., Hughes, D. A., Hut, R. W., Montanari, A., Pande, S., Tetzlaff, D., Troch, P. A., Uhlenbrook, S., Wagener, T., Winsemius, H. C., Woods, R. A., Zehe, E., and Cudennec, C.: A decade of Predictions in Ungauged Basins (PUB) – a review, *Hydrol. Sci. J.*, 58, 1198–1255, <https://doi.org/10.1080/02626667.2013.803183>, 2013.
- Hsu, K., Gupta, H. V., and Sorooshian, S.: Artificial Neural Network Modeling of the Rainfall-Runoff Process, *Water Resour. Res.*, 31, 2517–2530, <https://doi.org/10.1029/95WR01955>, 1995.
- Hsu, K., Gao, X., Sorooshian, S., Gupta, H. V., Hsu, K., Gao, X., Sorooshian, S., and Gupta, H. V.: Precipitation Estimation from Remotely Sensed Information Using Artificial Neural Networks, *J. Appl. Meteorol.*, 36, 1176–1190, [https://doi.org/10.1175/1520-0450\(1997\)036<1176:PEFRSI>2.0.CO;2](https://doi.org/10.1175/1520-0450(1997)036<1176:PEFRSI>2.0.CO;2), 1997.
- Hsu, K., Gupta, H. V., Gao, X., Sorooshian, S., and Imam, B.: Self-organizing linear output map (SOLO): An artificial neural network suitable for hydrologic modeling and analysis, *Water Resour. Res.*, 38, 38-1–38-17, <https://doi.org/10.1029/2001WR000795>, 2002.
- Huang, W., He, D., Yang, X., Zhou, Z., Kifer, D., and Giles, C. L.: Detecting Arbitrary Oriented Text in the Wild with a Visual Attention Model, in *Proceedings of the 2016 ACM on Multimedia Conference – MM '16*, 551–555, ACM Press, New York, USA, 2016.
- Indermuhle, E., Frinken, V., and Bunke, H.: Mode Detection in Online Handwritten Documents Using BLSTM Neural Networks, *2012 International Conference on Frontiers in Handwriting Recognition*, 302–307, IEEE, 2012.
- Izadinia, H., Russell, B. C., Farhadi, A., Hoffman, M. D., and Hertzmann, A.: Deep Classifiers from Image Tags in the Wild, *Proceedings of the 2015 Workshop on Community-Organized Multimodal Mining: Opportunities for Novel Solutions*, 13–18, ACM, 2015.
- Jackson, T., O'Neill, P., Njoku, E., Chan, S., Bindlish, R., Colliander, A., Chen, F., Burgin, M., Dunbar, S., Piepmeier, J., Cosh, M., Caldwell, T., Walker, J., Wu, X., Berg, A., Rowlandson, T., Pacheco, A., McNairn, H., Thibeault, M., Martínez-Fernández, J., González-Zamora, Á., Seyfried, M., Bosch, D., Starks, P., Goodrich, D., Prueger, J., Su, Z., van der Velde, R., Asanuma, J., Palecki, M., Small, E., Zreda, M., Calvet, J., Crow, W., Kerr, Y., Yueh, S., and Entekhabi, D.: Soil Moisture Active Passive (SMAP) Project Calibration and Validation for the L2/3_SM_P Version 3 Data Products, *SMAP Proj. JPL D-93720*, available at: <https://nsidc.org/data/smap/technical-references> (last access: 27 July 2017), 2016.
- Josephson, J. R. and Josephson, S. G.: *Abductive inference?: computation, philosophy, technology*, Cambridge University Press, 1994.
- Junk, C., Delle Monache, L., Alessandrini, S., Cervone, G., and von Bremen, L.: Predictor-weighting strategies for probabilistic wind power forecasting with an analog ensemble, *Meteorol. Z.*, 24, 361–379, <https://doi.org/10.1127/metz/2015/0659>, 2015.
- Karpathy, A., Johnson, J., and Fei-Fei, L.: Visualizing and Understanding Recurrent Networks, *ICLR 2016 Workshop*, available at: <http://arxiv.org/abs/1506.02078> (last access: 7 November 2016), 2015.
- Karpatne, A., Atluri, G., Faghmous, J. H., Steinbach, M., Banerjee, A., Ganguly, A., Shekhar, S., Samatova, N., and Kumar, V.: Theory-Guided Data Science: A New Paradigm for Scientific Discovery from Data, *IEEE Trans. Knowl. Data Eng.*, 29, 2318–2331, <https://doi.org/10.1109/TKDE.2017.2720168>, 2017.
- Kawaguchi, K., Kaelbling, L. P., and Bengio, Y.: Generalization in Deep Learning, *arXiv:1710.05468*, available at: <http://arxiv.org/abs/1710.05468> (last access: 12 March 2018), 2017.
- Kingma, D. P. and Ba, J.: Adam: A Method for Stochastic Optimization, in *3rd International Conference for Learning Representations*, San Diego, CA, available at: <http://arxiv.org/abs/1412.6980> (last access: 30 March 2018), 2014.
- Kingma, D. P. and Welling, M.: Auto-Encoding Variational Bayes, in *Proceedings of the 2014 International Conference on Learning Representations (ICLR)*, available at: <http://arxiv.org/abs/1312.6114> (last access: 24 March 2018), 2013.
- Knyazikhin, Y., Glassy, J., Privette, J. L., Tian, Y., Lotsch, A., Zhang, Y., Wang, Y., Morisette, J. T., P. Votava, Myneni, R. B., Nemani, R. R., and Running, S. W.: MODIS Leaf Area Index (LAI) and Fraction of Photosynthetically Active Radiation Absorbed by Vegetation (FPAR) Product (MOD15) Algorithm Theoretical Basis Document, https://modis.gsfc.nasa.gov/data/atbd/atbd_mod15.pdf (last access: 8 October 2018), 1999.
- Koza, J. R.: *Genetic Programming: on the Programming of Computers by Means of Natural Selection*, MIT Press, Cambridge, MA, USA, 1992.
- Kratzert, F., Klotz, D., Brenner, C., Schulz, K., and Hernegger, M.: Rainfall-Runoff modelling using Long-Short-Term-Memory (LSTM) networks, *Hydrol. Earth Syst. Sci. Discuss.*, <https://doi.org/10.5194/hess-2018-247>, in review, 2018.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E.: ImageNet Classification with Deep Convolutional Neural Networks, *Advances in Neural Information Processing Systems* 25, 1097–1105, available at: <https://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-> (last access: 30 March 2018), 2012.
- Krueger, D., Ballas, N., Jastrzebski, S., Arpit, D., Kanwal, M. S., Maharaj, T., Bengio, E., Erraqui, A., Fischer, A., Lacoste-Julien, S., and Courville, A.: Deep nets don't learn via memorization, *ICLR 2017*, available at: <https://openreview.net/pdf?id=rJv6ZgHYg>, 2017.
- Kumar, D. and Menkovski, V.: Understanding Anatomy Classification Through Visualization, *30th NIPS Machine learning for Health Workshop*, available at: <https://arxiv.org/abs/1611.06284> (last access: 24 November 2017), 2016.
- Laloy, E., Hérault, R., Lee, J., Jacques, D., and Linde, N.: Inversion using a new low-dimensional representation of complex binary geological media based on a deep neural network, *Adv. Water Resour.*, 110, 387–405, <https://doi.org/10.1016/J.ADVWATRES.2017.09.029>, 2017.
- Laloy, E., Hérault, R., Jacques, D., and Linde, N.: Training-Image Based Geostatistical Inversion Using a Spatial Generative Ad-

- versarial Neural Network, *Water Resour. Res.*, 54, 381–406, <https://doi.org/10.1002/2017WR022148>, 2018.
- LeCun, Y., Bengio, Y., and Hinton, G.: Deep learning, *Nature*, 521, 436–444, <https://doi.org/10.1038/nature14539>, 2015.
- Lehman, K. J.: Courting the Uncommitted: A Mixed-Methods Study of Undecided Students in Introductory Computer Science Courses, UCLA Electron. Theses Diss., available at: <https://escholarship.org/uc/item/94k326xs> (last access: 31 July 2018), 2017.
- Leviathan, Y. and Matias, Y.: Google Duplex: An AI System for Accomplishing Real-World Tasks Over the Phone, Google AI Blog, available at: <https://ai.googleblog.com/2018/05/duplex-ai-system-for-natural-conversation.html>, last access: 20 August 2018.
- Liu, Y., Racah, E., Prabhat, Correa, J., Khosrowshahi, A., Lavers, D., Kunkel, K., Wehner, M., and Collins, W.: Application of Deep Convolutional Neural Networks for Detecting Extreme Weather in Climate Datasets, ACM SIGKDD 2016 Conference on Knowledge Discovery & Data Mining, available at: <http://arxiv.org/abs/1605.01156>, last access: 21 October 2016.
- Mahendran, A. and Vedaldi, A.: Understanding deep image representations by inverting them, 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 5188–5196, 2015.
- Maier, H. R., Jain, A., Dandy, G. C., and Sudheer, K. P.: Methods used for the development of neural networks for the prediction of water resource variables in river systems: Current status and future directions, *Environ. Model. Softw.*, 25, 891–909, <https://doi.org/10.1016/j.envsoft.2010.02.003>, 2010.
- Marshall, L.: Creativity, Uncertainty, and Automated Model Building, *Groundwater*, 55, 693–697, <https://doi.org/10.1111/gwat.12552>, 2017.
- Matsuoka, D., Nakano, M., Sugiyama, D., and Uchida, S.: Detecting Precursors of Tropical Cyclone using Deep Neural Networks, The 7th International Workshop on Climate Informatics, CI, 2017.
- McCabe, M. F., Rodell, M., Alsdorf, D. E., Miralles, D. G., Uijlenhoet, R., Wagner, W., Lucieer, A., Houborg, R., Verhoest, N. E. C., Franz, T. E., Shi, J., Gao, H., and Wood, E. F.: The future of Earth observation in hydrology, *Hydrol. Earth Syst. Sci.*, 21, 3879–3914, <https://doi.org/10.5194/hess-21-3879-2017>, 2017.
- Mecklenburg, S., Kerr, Y., Font, J., and Hahne, A.: The Soil Moisture and Ocean Salinity Mission – An Overview, IGARSS 2008, IEEE International Geoscience and Remote Sensing Symposium, IV-938–IV-941, 2008.
- Merwade, V. and Ruddell, B. L.: Moving university hydrology education forward with community-based geoinformatics, data and modeling resources, *Hydrol. Earth Syst. Sci.*, 16, 2393–2404, <https://doi.org/10.5194/hess-16-2393-2012>, 2012.
- Montanari, A., Young, G., Savenije, H. H. G., Hughes, D., Wagener, T., Ren, L. L., Koutsoyiannis, D., Cudennec, C., Toth, E., Grimaldi, S., Blöschl, G., Sivapalan, M., Beven, K., Gupta, H., Hipsey, M., Schaeffli, B., Arheimer, B., Boegh, E., Schymanski, S. J., Baldassarre, G. Di, Yu, B., Hubert, P., Huang, Y., Schumann, A., Post, D. A., Srinivasan, V., Harman, C., Thompson, S., Rogger, M., Viglione, A., McMillan, H., Characklis, G., Pang, Z., and Belyaev, V.: “Panta Rhei – Everything Flows”: Change in hydrology and society – The IAHS Scientific Decade 2013–2022, *Hydrol. Sci. J.*, 58, <https://doi.org/10.1080/02626667.2013.809088>, 2013.
- Moradkhani, H., Hsu, K., Gupta, H. V., and Sorooshian, S.: Improved streamflow forecasting using self-organizing radial basis function artificial neural networks, *J. Hydrol.*, 295, 246–262, <https://doi.org/10.1016/j.jhydrol.2004.03.027>, 2004.
- Mosser, L., Dubrule, O., and Blunt, M. J.: Reconstruction of three-dimensional porous media using generative adversarial neural networks, *Phys. Rev. E*, 96, 43309, <https://doi.org/10.1103/PhysRevE.96.043309>, 2017.
- Mount, N. J., Maier, H. R., Toth, E., Elshorbagy, A., Solomatine, D., Chang, F.-J., and Abraham, R. J.: Data-driven modelling approaches for socio-hydrology: opportunities and challenges within the Panta Rhei Science Plan, *Hydrol. Sci. J.*, 61, 1–17, <https://doi.org/10.1080/02626667.2016.1159683>, 2016.
- Mu, Q., Zhao, M., and Running, S. W.: Improvements to a MODIS global terrestrial evapotranspiration algorithm, *Remote Sens. Environ.*, 115, 1781–1800, <https://doi.org/10.1016/j.rse.2011.02.019>, 2011.
- Nearing, G. S., Mocko, D. M., Peters-Lidard, C. D., Kumar, S. V., and Xia, Y.: Benchmarking NLDAS-2 Soil Moisture and Evapotranspiration to Separate Uncertainty Contributions, *J. Hydrometeorol.*, 17, 745–759, <https://doi.org/10.1175/JHM-D-15-0063.1>, 2016.
- Ororbia, A. G., Giles, C. L., and Kifer, D.: Unifying Adversarial Training Algorithms with Flexible Deep Data Gradient Regularization, *Neural Computation*, MIT Press, available at: <http://arxiv.org/abs/1601.07213> (last access: 24 February 2017), 2016.
- Pavelsky, T. M., Durand, M. T., Andreadis, K. M., Beighley, R. E., Paiva, R. C. D., Allen, G. H., and Miller, Z. F.: Assessing the potential global extent of SWOT river discharge observations, *J. Hydrol.*, 519, 1516–1525, <https://doi.org/10.1016/j.jhydrol.2014.08.044>, 2014.
- Ranzato, M., Poultney, C., Chopra, S., and LeCun, Y.: Efficient learning of sparse representations with an energy-based model, *Proc. 19th Int. Conf. Neural Inf. Process. Syst.*, 1137–1144, available at: <https://dl.acm.org/citation.cfm?id=2976599> (last access: 30 March 2018), 2006.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., and Fei-Fei, L.: ImageNet Large Scale Visual Recognition Challenge, *arXiv:1409.0575*, available at: <http://arxiv.org/abs/1409.0575> (last access: 24 March 2018), 2014.
- Sadler, J. M., Goodall, J. L., Morsy, M. M., and Spencer, K.: Modeling urban coastal flood severity from crowd-sourced flood reports using Poisson regression and Random Forest, *J. Hydrol.*, 559, 43–55, <https://doi.org/10.1016/j.jhydrol.2018.01.044>, 2018.
- Samek, W., Binder, A., Montavon, G., Lapuschkin, S., and Müller, K.-R.: Evaluating the Visualization of What a Deep Neural Network Has Learned, *IEEE Trans. Neural Networks Learn. Syst.*, 28, 2660–2673, <https://doi.org/10.1109/TNNLS.2016.2599820>, 2017.
- Sax, L. J., Lehman, K. J., Jacobs, J. A., Kanny, M. A., Lim, G., Monje-Paulson, L., and Zimmerman, H. B.: Anatomy of an Enduring Gender Gap: The Evolution of Women’s Participation in Computer Science, *J. Higher Educ.*, 88, 258–293, <https://doi.org/10.1080/00221546.2016.1257306>, 2017.
- Schaap, M. G., Leij, F. J., and van Genuchten, M. T.: Rosetta: a Computer Program for Estimating Soil Hydraulic Parameters With Hierarchical Pedotransfer Functions, *J. Hydrol.*, 251, 163–176, doi:10.1016/S0022-1694(01)00466-8, 2001.

- Schmidhuber, J.: Deep learning in neural networks: An overview, *Neural Networks*, 61, 85–117, <https://doi.org/10.1016/j.neunet.2014.09.003>, 2015.
- Schwalm, C. R., Anderegg, W. R. L., Michalak, A. M., Fisher, J. B., Biondi, F., Koch, G., Litvak, M., Ogle, K., Shaw, J. D., Wolf, A., Huntzinger, D. N., Schaefer, K., Cook, R., Wei, Y., Fang, Y., Hayes, D., Huang, M., Jain, A., and Tian, H.: Global patterns of drought recovery, *Nature*, 548, 202–205, <https://doi.org/10.1038/nature23021>, 2017.
- Settles, B.: Active Learning, in: *Synthesis lectures on artificial intelligence and machine learning*, edited by: Brachman, R. J., Cohen, W. W., and Dietterich, T. G., Morgan & Claypool, Williston, VT, USA, 2012.
- Shen, C.: A trans-disciplinary review of deep learning research and its relevance for water resources scientists, *Water Resour. Res.*, <https://doi.org/10.1029/2018WR022643>, 2018.
- Simonite, T.: AI is the future – but where are the women, *Wired*, available at: <https://www.wired.com/story/artificial-intelligence-researchers-gender-imbalance/>, last access: 1 September 2018.
- Simonyan, K. and Zisserman, A.: Very Deep Convolutional Networks for Large-Scale Image Recognition, *ICLR 2015*, available at: <http://arxiv.org/abs/1409.1556> (last access: 25 August 2018), 2014.
- Snelson, E. and Ghahramani, Z.: Sparse Gaussian Processes using Pseudo-inputs, *Adv. Neural Inf. Process. Syst.*, 18, 1257–1264, 2006.
- Srinivasan, M.: Hydrology from space: NASA's satellites supporting water resources applications, *Water Forum III Droughts Other Extrem. Weather Events*, available at: http://www.jsg.utexas.edu/ciess/files/Srinivasanetal_TWF_Oct14_Final.pdf (last access: 12 July 2016), 2013.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R.: Dropout: A Simple Way to Prevent Neural Networks from Overfitting, *J. Mach. Learn. Res.*, 15, 1929–1958, 2014.
- Stallkamp, J., Schlipsing, M., Salmen, J., and Igel, C.: The German Traffic Sign Recognition Benchmark: A multi-class classification competition, *The 2011 International Joint Conference on Neural Networks*, 1453–1460, 2011.
- Stocker, T. F., Qin, D., Plattner, G.-K., Alexander, L. V., Allen, S. K., Bindoff, N. L., Bréon, F.-M., Church, J. A., Cubasch, U., Emori, S., Forster, P., Friedlingstein, P., Gillett, N., Gregory, J. M., Hartmann, D. L., Jansen, E., Kirtman, B., Knutti, R., Kumar, K. K., Lemke, P., Marotzke, J., Masson-Delmotte, V., Meehl, G. A., Mokhov, I. I., Piao, S., Ramaswamy, V., Randall, D., Rhein, M., Rojas, M., Sabine, C., Shindell, D., Talley, L. D., Vaughan, D. G., and Xie, S.-P.: Technical summary, *Climate Change 2013*, in: *The Physical Science Basis*, edited by: Stocker, T. F., Qin, D., Plattner, G.-K., Tignor, M., Allen, S. K., Boschung, J., Nauels, A., Xia, Y., Bex, V., and Midgley, P. M., *Contrib. Work. Gr. I to Fifth Assess. Rep. Intergov. Panel Clim. Chang.*, available at: http://www.ipcc.ch/pdf/assessment-report/ar5/wg1/WG1AR5_TS_FINAL.pdf (last access: 8 October 2018), 2013.
- Stollenga, M. F., Byeon, W., Liwicki, M., and Schmidhuber, J.: Parallel multi-dimensional LSTM, with application to fast biomedical volumetric image segmentation, *Proc. 28th Int. Conf. Neural Inf. Process. Syst.*, 2, 2998–3006, 2015.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R.: Intriguing properties of neural networks, available at: <http://arxiv.org/abs/1312.6199> (last access: 8 June 2016), 2013.
- Szegedy, C., Wei Liu, Yangqing Jia, Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A.: Going deeper with convolutions, *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1–9, 2015.
- Tao, Y., Gao, X., Hsu, K., Sorooshian, S., and Ihler, A.: A Deep Neural Network Modeling Framework to Reduce Bias in Satellite Precipitation Products, *J. Hydrometeorol.*, 17, 931–945, <https://doi.org/10.1175/JHM-D-15-0075.1>, 2016.
- Tao, Y., Gao, X., Ihler, A., Sorooshian, S., Hsu, K., Tao, Y., Gao, X., Ihler, A., Sorooshian, S., and Hsu, K.: Precipitation Identification with Bispectral Satellite Information Using Deep Learning Approaches, *J. Hydrometeorol.*, 18, 1271–1283, <https://doi.org/10.1175/JHM-D-16-0176.1>, 2017.
- Tao, Y., Hsu, K., Ihler, A., Gao, X., Sorooshian, S., Tao, Y., Hsu, K., Ihler, A., Gao, X., and Sorooshian, S.: A Two-Stage Deep Neural Network Framework for Precipitation Estimation from Bispectral Satellite Information, *J. Hydrometeorol.*, 19, 393–408, <https://doi.org/10.1175/JHM-D-17-0077.1>, 2018.
- Thompson, J. A., Pena-Yewtukhiw, E. M., and Grove, J. H.: Soil-landscape modeling across a physiographic region: Topographic patterns and model transportability, *Geoderma*, 133, 57–70, <https://doi.org/10.1016/J.GEODERMA.2006.03.037>, 2006.
- Tibshirani, R. and Tibshirani, R.: Regression Shrinkage and Selection Via the Lasso, *J. R. Stat. Soc. Ser. B*, 58, 267–288, 1994.
- Troch, P. A., Carrillo, G., Sivapalan, M., Wagener, T., and Sawicz, K.: Climate-vegetation-soil interactions and long-term hydrologic partitioning: signatures of catchment co-evolution, *Hydrol. Earth Syst. Sci.*, 17, 2209–2217, <https://doi.org/10.5194/hess-17-2209-2013>, 2013.
- Vandal, T., Kodra, E., Ganguly, S., Michaelis, A., Nemani, R., and Ganguly, A. R.: DeepSD: Generating High Resolution Climate Change Projections through Single Image Super-Resolution, *23rd ACM SIGKDD Conference on Knowledge Discovery and Data Mining.*, available at: <http://arxiv.org/abs/1703.03126>, last access: 2 December 2017.
- Vondrick, C., Pirsaviash, H., and Torralba, A.: Anticipating Visual Representations from Unlabeled Video, *CVPR 2016*, available at: <http://arxiv.org/abs/1504.08023> (last access: 20 August 2018), 2016.
- Voosen, P.: The AI detectives, *Science*, 80, 357, <https://doi.org/10.1126/science.357.6346.22>, 2017.
- Wagener, T., Sivapalan, M., Troch, P. A., McGlynn, B. L., Harman, C. J., Gupta, H. V., Kumar, P., Rao, P. S. C., Basu, N. B., and Wilson, J. S.: The future of hydrology: An evolving science for a changing world, *Water Resour. Res.*, 46, 1–10, <https://doi.org/10.1029/2009WR008906>, 2010.
- Wagener, T., Kelleher, C., Weiler, M., McGlynn, B., Gooseff, M., Marshall, L., Meixner, T., McGuire, K., Gregg, S., Sharma, P., and Zappe, S.: It takes a community to raise a hydrologist: the Modular Curriculum for Hydrologic Advancement (MOCHA), *Hydrol. Earth Syst. Sci.*, 16, 3405–3418, <https://doi.org/10.5194/hess-16-3405-2012>, 2012.
- Wahr, J.: Time-variable gravity from GRACE: First results, *Geophys. Res. Lett.*, 31, L11501, <https://doi.org/10.1029/2004GL019779>, 2004.

- Wahr, J., Swenson, S., and Velicogna, I.: Accuracy of GRACE mass estimates, *Geophys. Res. Lett.*, 33, L06401, <https://doi.org/10.1029/2005GL025305>, 2006.
- Wang, H., Skau, E., Krim, H., and Cervone, G.: Fusing Heterogeneous Data: A Case for Remote Sensing and Social Media, *IEEE Trans. Geosci. Remote Sens.*, 1–13, <https://doi.org/10.1109/TGRS.2018.2846199>, 2018.
- WaterML2: WaterML2, A global standard for hydrological time series, available at: <http://www.waterml2.org/>, last access: 23 August 2018.
- Xiong, W., Droppo, J., Huang, X., Seide, F., Seltzer, M., Stolcke, A., Yu, D., and Zweig, G.: Achieving Human Parity in Conversational Speech Recognition, Microsoft Res. Tech. Rep. MSR-TR-2016-71. arXiv1610.05256, available at: <http://arxiv.org/abs/1610.05256> (last access: 24 March 2018), 2016.
- Yosinski, J., Clune, J., Nguyen, A., Fuchs, T., and Lipson, H.: Understanding Neural Networks Through Deep Visualization, in *Deep Learning Workshop*, 31 st International Conference on Machine Learning, Lille, France, available at: <http://arxiv.org/abs/1506.06579> (last access: 19 November 2017), 2015.
- Yu, K.-H., Zhang, C., Berry, G. J., Altman, R. B., Ré, C., Rubin, D. L., and Snyder, M.: Predicting non-small cell lung cancer prognosis by fully automated microscopic pathology image features, *Nat. Commun.*, 7, 12474, <https://doi.org/10.1038/ncomms12474>, 2016.
- Zen, H. and Sak, H.: Unidirectional long short-term memory recurrent neural network with recurrent output layer for low-latency speech synthesis, *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 4470–4474, 2015.
- Zhang, D., Lindholm, G., and Ratnaweera, H.: Use long short-term memory to enhance Internet of Things for combined sewer overflow monitoring, *J. Hydrol.*, 556, 409–418, <https://doi.org/10.1016/J.JHYDROL.2017.11.018>, 2018.
- Zhu, X. X., Tuia, D., Mou, L., Xia, G.-S., Zhang, L., Xu, F., and Fraundorfer, F.: Deep Learning in Remote Sensing: A Comprehensive Review and List of Resources, *IEEE Geosci. Remote Sens. Mag.*, 5, 8–36, <https://doi.org/10.1109/MGRS.2017.2762307>, 2017.