DEEP LEARNING AS A TOOL TO FORECAST HYDROLOGIC RESPONSE FOR LANDSLIDE-PRONE HILLSLOPES

by

ELIJAH ORLAND

A THESIS

Presented to the Department of Earth Sciences
and the Graduate School of the University of Oregon
in partial fulfillment of the requirements
for the degree of
Master of Science

June 2020

THESIS APPROVAL PAGE

Student: Elijah Orland

Title: Deep Learning as a Tool to Forecast Hydrologic Response for Landslide-Prone Hillslopes

This thesis has been accepted and approved in partial fulfillment of the requirements for the Master of Science degree in the Earth Sciences by:

Joshua Roering	Chairperson
Matthew Polizzotto	Member
Diego Melgar	Member
Jake Searcy	Member

and

Kate Mondloch	Interim Vice Provost and Dean of the Graduate School

Original approval signatures are on file with the University of Oregon Graduate School.

Degree awarded June 2020

THESIS ABSTRACT

Elijah Orland

Master of Science

Department of Earth Sciences

June 2020

Title: Deep Learning as a Tool to Forecast Hydrologic Response for Landslide-Prone Hillslopes

Empirical thresholds for landslide warning systems have benefitted from the incorporation of soil-hydrologic monitoring data, but the mechanistic basis for their predictive capabilities is limited. Although physically based hydrologic models can accurately simulate changes in soil moisture and pore pressure that promote landslides, their utility is restricted by high computational costs and non-unique parameterization issues. We construct a Deep Learning model using soil-moisture, pore-pressure, and rainfall monitoring data acquired from landslide-prone hillslopes in Oregon, USA, to predict the timing and magnitude of hydrologic response at multiple soil depths for 36-hour intervals. We find that observation records as short as six months are sufficient for accurate predictions, and our model captures hydrologic response for high-intensity rainfall events even when those storm types are excluded from model training. We conclude that machine learning can provide an accurate, and computationally efficient alternative to empirical methods or physical modeling for landslide hazard warning.

CURRICULUM VITAE

NAME OF AUTHOR: Elijah Orland

GRADUATE AND UNDERGRADUATE SCHOOLS ATTENDED:

University of Oregon, Eugene
Middlebury College, Middlebury

DEGREES AWARDED:

Master of Science, 2020, University of Oregon

Bachelor of Arts, 2018, Middlebury College

PUBLICATIONS:

Orland, E., Roering, J.J., Mirus, B.B., Thomas, M.A., (2020). Deep Learning as a tool to forecast hydrologic response for landslide-prone hillslopes. *Geophysical Research Letters* (Accepted)

ACKNOWLEDGMENTS

This Thesis would not be possible without the support of my advisor Josh Roering, and the careful input of my co-authors, Matt Thomas and Ben Mirus in the corresponding publication. I would also like to thank Johnathan Godt, Joel Smith, Jeff Coe, and Rex Baum for developing and maintaining installations at the Knife Ridge monitoring site, as well as two anonymous reviewers and Francis Rengers for their constructive inputs. I would furthermore like to acknowledge all the support I have received throughout my time at the University, including my roommates (plus cats), lab mates, and office staff.

To my parents, thank you for the ever continuing support of my education and helping to provide the means to pursue my goals. The two of you are a beacon of positivity and encouragement whenever my own world feels dark.

Finally, Perri: three years ago my world changed when we decided that no distance was too far to keep us apart. Being on opposite sides of the country certainly tested that, and I could not be more grateful to see the final result. Thank you for our nightly chats, our video calls, and the occasional visit. I'm glad that time is behind us.

And to each of you that I mentioned: *this work is yours, too*. One that would not exist without your incredible patience and thoughtful input as I made my way through the University. I am grateful for every single one of you.

To my parents: this is dedicated to you.

You have taught me that education is priceless, yet have always given so much to ensure I have it.

TABLE OF CONTENTS

LIST OF FIGURES

Figure Page

"Who really can face the future? All you can do is project from the past, even when the past shows that such projections are often wrong. And who really can forget the past? What else is there to know?" – Robert Pirsig

I. INTRODUCTION

Rainfall-induced shallow landslides pose a threat to communities situated in mountainous terrain. These natural hazards are deadly and can incur up to $500 million (USD, 2020 adjusted) in economic losses from a single event (Turner & Schuster, 1996). Furthermore, there are over 294,000 individual landslides identified in the conterminous United States, with additional mapping still needed to complete a nationwide inventory (Mirus et al., 2020). Assessing the potential for landslide initiation is a crucial step for developing an early warning system for at-risk communities, and relies on predictions of hillslope response to rainfall. This is often facilitated by hydrologic models characterized by varying degrees of complexity, such as empirical thresholds that relate rainfall and/or soil moisture measurements that correspond to past landslide events (e.g., Guzzetti et al., 2008; Wieczorek & Guzzetti, 2000; Mirus et al., 2018), or deterministic, physically based numerical models (e.g., Bellugi et al., 2015; Montgomery & Dietrich, 1994; Simunek et al., 2005). There are unique advantages to each method: empirical relationships between soil moisture, rainfall, and landsliding provide an efficient, regional-scale means of characterizing landslide initiation potential, but have limited predictive accuracy (Bogaard & Greco, 2018; van Natijne et al., 2020). On the other hand, hydrologic modeling of saturated and unsaturated flow can be used to develop thresholds for assessing distributed landslide potential and can be extended beyond the observed record (e.g., Fusco et al., 2019; Thomas et al., 2018a) instead of relying solely on statistical analyses. Hydrologic modeling of subsurface flow is relevant for hazard assessment, as it provides quantitative estimates of pore pressures in soils, which can be used to approximate the corresponding reduction of resisting stresses on hillslopes that can lead to slope failure (Lu & Likos, 2004; Lu & Godt, 2008; Thomas et al., 2018b). That said, models of variably-saturated hydraulic and mechanical hillslope processes require significant effort to parameterize and validate, and the computational burden associated with their solution can limit their utility for operational (real-time) warning systems. Further complications include the natural heterogeneity of soil structure, such as the presence of tree roots which add quantifiable tensile strength (Roering et al., 2003), and macropores that provide preferred channels of non-Darcy flow (McDonnell, 1990). Consequently, physical based modeling may not necessarily capture all the underlying physical processes in natural field settings.

To better constrain hydrologic response to rainfall, we propose an alternative method based on Deep Learning (DL)—an approach that requires less physical parameterization and trains faster than existing multi-dimensional flow models. DL-based approaches have already shown significant promise in the field of hydrology (Shen et al., 2018), and here, we propose the use of specific DL-based algorithms that can learn and forecast hydrologic conditions from time series. This method relies on the quality and abundance of in-situ soil moisture or pore-pressure measurements to provide an empirically derived, site-specific hydrologic dataset which implicitly captures the characteristics of multi-dimensional variably saturated flow. We base our DL application on three unique requirements that are essential considerations for integrating a DL model with an early warning system: (1) evaluating how well a calibrated model can perform on unseen data (i.e. general model performance); (2) finding a critical length of input data for sufficient model performance (i.e. setting minimum performance requirements); and, (3) how well a model can predict conditions associated with high-intensity rainfall when the observation data used for model training does not include high-intensity events (i.e., model extrapolation capability).

Here, we develop an explicit, data-driven analysis of these requirements for successful early warning system integration and demonstrate the utility of DL for accurately modeling 3D variably saturated flow from data representative of typical early warning systems. We further show how DL algorithms provide accurate forecasts of pore pressure response up to 36 hours in advance—a relevant timescale for landslide risk mitigation. While we do not suggest that DL can replace the insights gained from physically based models, it is a promising method for forecasting soil hydrologic response in highly heterogenous or otherwise complex landscapes where comprehensive physical modeling would be computationally prohibitive.

II. METHODS

2.1 Data Source

We used data collected from the Knife Ridge monitoring site in the Elliot State Forest, Oregon, USA (Smith et al., 2014) (Figure 1a). The Oregon Coast Range is characterized by steep, soil-mantled hillslopes that are prone to shallow landslides that transition into debris flows, such as those triggered by historic storms in 1964 and 1996 (Coe et al., 2011; Montgomery et al., 2009; Robison et al., 1999; Wiley, 2000). Previous studies have shown that hydrologic response in this landscape is challenging to model as a result of strong hydrologic connectivity between the heterogeneous soil mantle and shallow zones of fractured bedrock (Ebel et al., 2008; Ebel et al., 2007; Mirus et al., 2016). The monitoring site is located within a valley axis with an upslope contributing area of approximately 300 m^2, and features greater than 30° slopes with an average soil depth between 0.9-2.1m (Smith et al., 2014). The instrumental record includes hourly measurements of rainfall, as well as soil moisture and pore pressure (suction) at three depths acquired from 2009 to 2012 (Smith et al., 2014) and has been the subject of prior 3D physically based modeling efforts (Mirus et al., 2016).

2.2 Machine Learning Methods

Artificial Neural Networks (ANNs) have broad application to classification and regression-based problem solving, owing to their ability to approximate a wealth of mathematical relationships between input and output variables (Hornik et al., 1989). ANNs are already utilized in geoscience applications (e.g., DeVries et al., 2017), and have been used to model highly non-linear processes common in hydrology (Abrahart &

See, 2007). Other uses of neural networks or machine learning (ML) in hydrology include time-dependent ANNs capable of filling data gaps in hydrologic times series (Ren et al., 2019), as well as tree-based based machine learning algorithms for estimating soil-hydrologic parameters from physical or textural soil properties (Araya & Ghezzehei, 2019).

Specifically, ANNs consist of one or more hidden layers of neurons that each contain a set of weights that are applied to all input variables, plus an added bias. The weighted sum of a neuron is subsequently non-linearized by an activation function. Hidden layers of neurons connect input and output variables. The various combinations of weights and biases across all neurons are then trained by changing their values such that they minimize a fixed error metric like the root mean squared error (RMSE) or mean absolute error (MAE) to provide the proper non-linear transformation of input variables to a corresponding output variable or variables.

A variation of traditional ANNs is the Recurrent Neural Network (RNN), which is well suited for processing sequential or temporal data, given a RNN's ability to incorporate information from the previous timestep and make predictions for the next timestep(s) based on weighting past and present inputs. A popular variant on the traditional RNN cell is the Long Short Term Memory (LSTM) cell (Hochreiter & Schmidhuber, 1997). By design, a LSTM cell contains an internal state which acts as its memory, and is iteratively adjusted and propagated through time. The internal state allows the cell to handle longer term temporal dependencies that traditional RNN cells are not able to explicitly incorporate within their architecture.

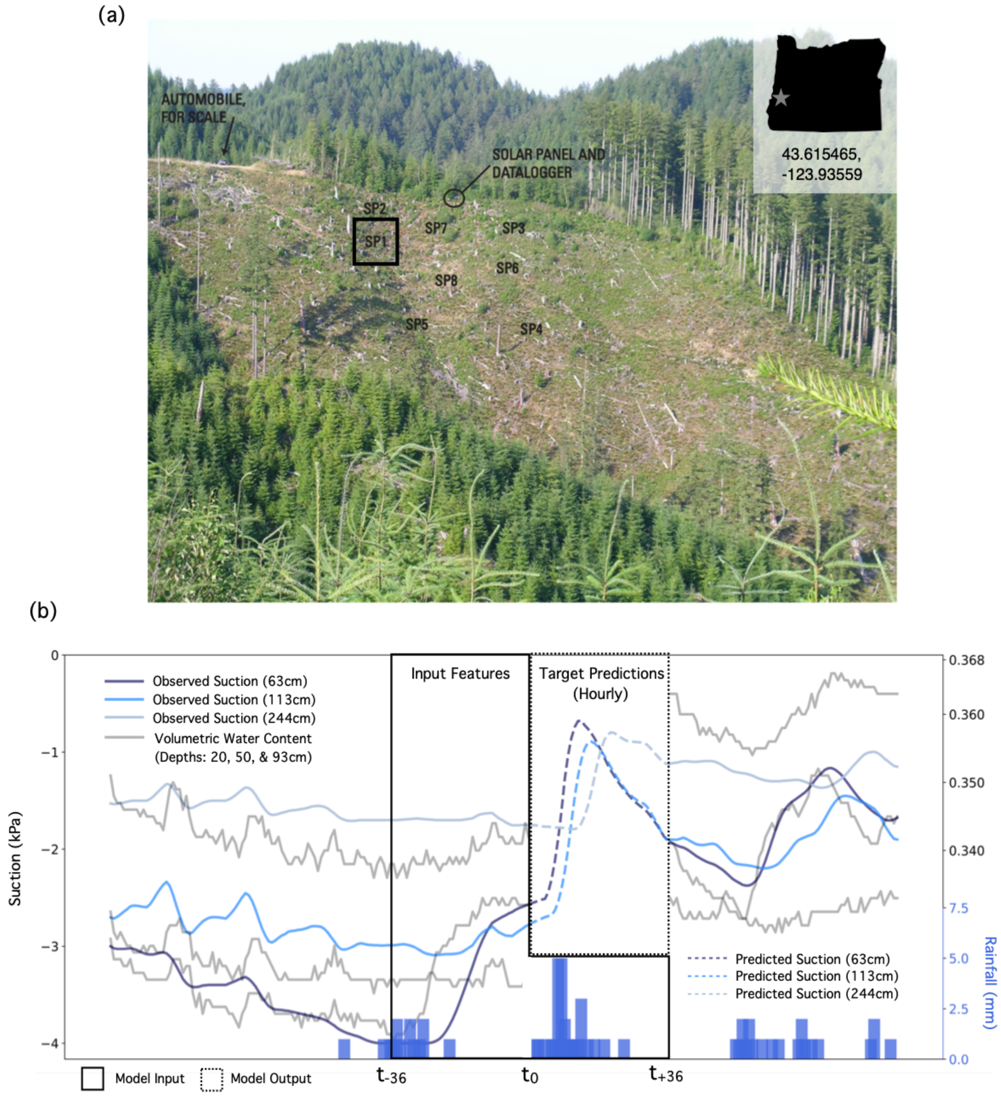

Figure 1. (a) Oblique view of the Knife Ridge monitoring site featuring all instrumented soil pits (SP), including the pit for which we model pore pressures (black box). As the site sits on a slope >30° located within a convergent valley axis, variably saturated flow takes place both laterally and vertically. Adapted from Smith et al. (2014). (b) Conceptual diagram showing model inputs and outputs. During training, a moving window slides across the input data provided to it (solid), and the model adjusts its parameters to produce a sequence of pore pressure values (dashed) from the model inputs that best matches the observed pore pressure sequence for those timesteps (not pictured). This process repeats until the model converges on a set of a parameters that produces the lowest mean squared error measured across all predicted and observed sequences.

Furthermore, LSTM-based neural networks are conceptually better suited than statistically-based autoregressive time series models, given a LSTM's capacity to approximate non-linear relationships between input and output variables, as opposed to assuming linear relationships between lagged endogenous or exogeneous variables in most autoregressive models (Box et al., 2008). This is particularly relevant for the non-linear physical processes that govern variably saturated subsurface flow (Richards, 1931). For a more in depth discussion of LSTM cells, we refer to Olah (2015), and for applications to hydrology, we refer to (Kratzert et al., 2018).

We use an LSTM-based "encoder-decoder" model with a global Luong attention mechanism (Luong et al., 2015)—an architecture we chose through an extensive trial and error-based process considering computational processing time, ease of use, and consistency of results. The encoder-decoder architecture consists of two parts: an "encoder" that reads in and learns an input sequence; and a "decoder" that receives a fixed representation (understanding) of the input sequence from the encoder, and learns the proper non-linear transformations to translate the input to an output sequence of variable length (Sutskever et al., 2014). For temporal data, the encoder learns the functional relationships of the input sequence, and the decoder learns the necessary transformations to translate the previous input sequence to a sequence of predicted values. The decoder does so one step at a time based on the input sequence itself, as well as the previous timestep of the prediction sequence.

Our model consists of 24 unidirectional LSTM units in both the encoder and decoder, which feed to a layer of 75 connected neurons with a 50% dropout regularization rate for further processing, and finally connect to a layer of 3 neurons outputting a vector of 36 hourly predictions at each of the three tensiometer depths. We script our model in the Python programming language, utilizing the Keras frontend API (https://keras.io/) with a Tensorflow backend (https://www.tensorflow.org/), and the Adam optimizer (Kingma & Ba, 2014). We use Scikit-learn (Pedregosa et al., 2011) for data pre- and post-processing. As model weights are set randomly upon initialization, each model run typically converges on a similar, but not necessarily identical solution. However, training for 2000-5000 epochs consistently results in a series of weights and biases that lead to comparable results across model runs. Typical training time takes between 2-10 minutes on a single GPU processor in the Google Collaboratory scripting environment (*colab.research.google.com).*

2.3 Data Preprocessing and Experimental Framework

Modeling steps include preparing input sequences of the previous 36 hours of data at 1hr resolution from all sensors (rainfall, soil moisture, tensiometer), and the measured maximum and cumulative rainfall for the next 36-hr interval. We represent this future 36-hr rainfall input (herein referred to as a *forecast*) through 18 two-hour intervals of max and cumulative rainfall data. These input data pair with the corresponding pore pressure data within the prediction interval at each tensiometer depth, and frames our study as supervised learning regression problem. In doing so, we provide our model with both prior and anticipated hydrologic information as inputs and seek to draw an explicit link between these inputs and the corresponding pore pressure response within the 36-hr rainfall forecast period (Figure 1b). Model training then occurs by finding the best set of non-linearized weights and biases that apply to current and past information which most closely match the observed pore pressure response. For this implementation, we choose the recorded rainfall data as forecasted information because it allows our algorithm to approximate the functional relationships between soil moisture response and rainfall. As such, we do not incorporate uncertainty into rainfall forecasts in this contribution.

To assess the efficacy of our LSTM-based model, we explicitly address the three requirements necessary for application in early warning systems. Specifically, we ask the following: (1) Given a model trained on two years of data, how well can it forecast pore pressure for the following year?; (2) To what extent does the length of the training dataset affect our ability to make accurate predictions?; and (3) If our training dataset does not include high intensity events, how well can we predict pore pressure for high intensity rainfall outside of the model's calibrated range? For this exclusion process, we omit rainfall events with intensities >4mm/hr across all water years—a relatively low intensity value that corresponds to the mean hourly intensities of two 1996 storms that caused mass wasting across the Oregon Coast Range. The first question demonstrates the algorithm's ability to learn from a variety of soil moisture response examples and tests anticipated general performance; the second question tests how much information our model needs to provide accurate forecasts, thus establishing a minimum requirement for available training data; and finally, the third question tests our algorithm's ability to extrapolate response governed by highly non-linear processes and potentially predict response to extreme, never-recorded events.

III. RESULTS

Our results demonstrate that the LTSM-based model can yield highly accurate predictions (RMSE < 1 kPa for positive pore pressures) for Water Year 3 (Oct 1, 2011 – Sept 31, 2012) in as little as two minutes (computational time) of training on Water Years 1 & 2 using a single GPU processor, thus addressing our model's general performance capacity. As an example, Figure 2a shows predictions for pore pressure (suction) at 36 hour intervals from January 14 – February 14, 2012 (a subsection of Year 3) at all tensiometer depths. This one-month illustration includes many 36-hr prediction intervals that show a rise and decay of pore pressure in response to rainfall events. While 36-hr prediction intervals generally agree with the corresponding observed intervals, our model sometimes fails to predict periods of wetting or drying during intervals that feature continuous rainfall over several days. For instance, in the long sequence of rainfall events between January 17-21, our model incorrectly predicts drying conditions with a decrease in pore pressure. Nonetheless, these types of errors are infrequent and unrepresentative of general model performance, as the correlation coefficients between observed and simulated values exceed 0.96 (Supplemental Figure 1).

To establish minimum performance requirements, we progressively resampled our training dataset into shorter and shorter subsets and then tested its performance for the third year of our record. Figure 2b shows the drop in model performance is most prominent when we reduce sampling to less than six months of observations, which we posit as a minimum sampling threshold for the dataset. While this number will vary based on the dataset provided, having a baseline value for training data provides valuable insight for similar endeavors at other field sites and demonstrates an important minimum observation requirement for accurate predictions.

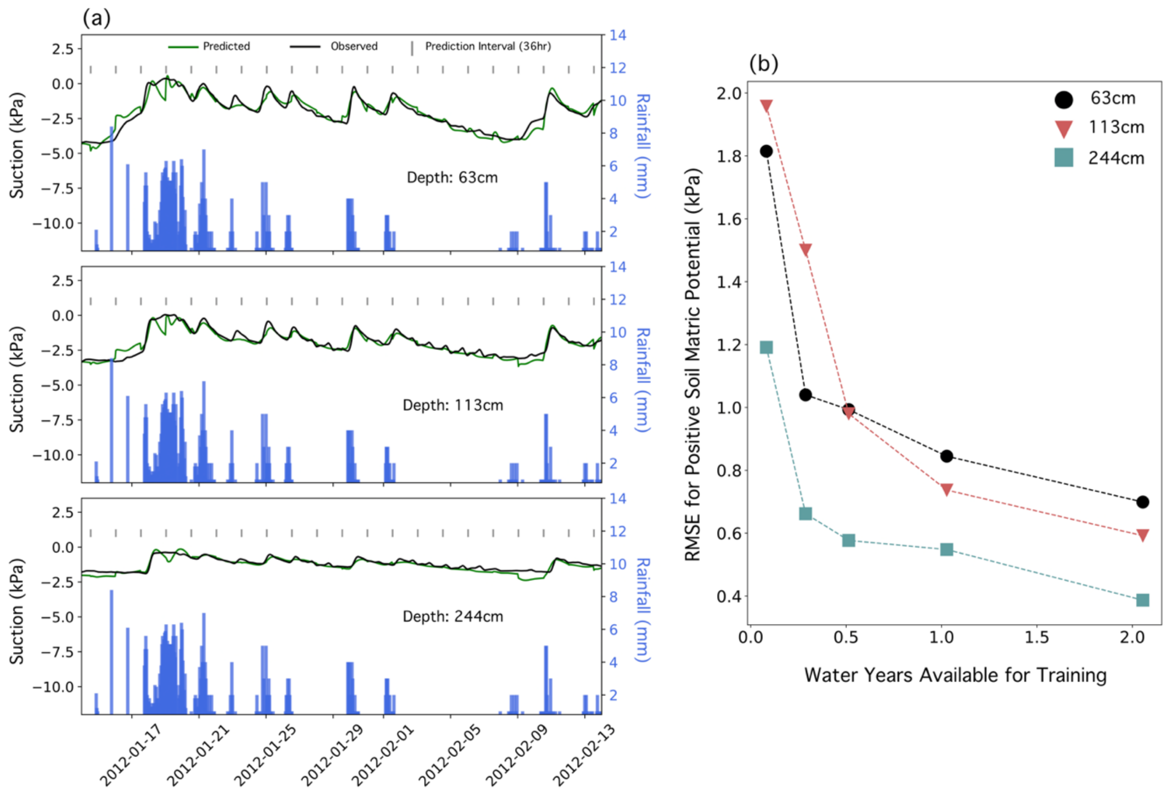

Figure 2. (a) Predicted and observed pore pressure response from January 14 - February 14, 2012, delineated by each 36-hr prediction interval. (b) RMSE of positive pore pressure predictions as a function of available training data. Data subsampling occurred by randomly selecting training samples and their corresponding targets without replacement from a total population of training data from Water Years 1 & 2. This introduces possible random sampling error, but eliminates any bias from choosing only samples from dry or wet months. With a limited number of training samples, however, we suggest selecting data that most closely approximates the intended conditions one chooses to model.

Finally, to demonstrate model extrapolation capacity, Figure 3a shows 36-hr predictions for intervals with rainfall intensities above the 4mm/hr training threshold. Overall, our intensity-limited model demonstrates reasonable accuracy in estimating rising, peak, and falling limbs of individual prediction intervals, but shows infrequent areas of high inaccuracy, particularly for prolonged periods of rainfall over the 4mm/hr threshold. Figure 3b shows absolute prediction error between our original and intensity-limited models at each tensiometer depth. While the intensity-limited model features the largest error values, overall error approximates the mean error values of our original model, both of which lie within or near the ± 0.5 kPa margin of instrumental tensiometer error (Smith et al., 2014). Thus, even constraining the model training to non-anomalous rainfall events still allows for accurate—if not always consistent—predictions in response to higher intensity rainfall events.

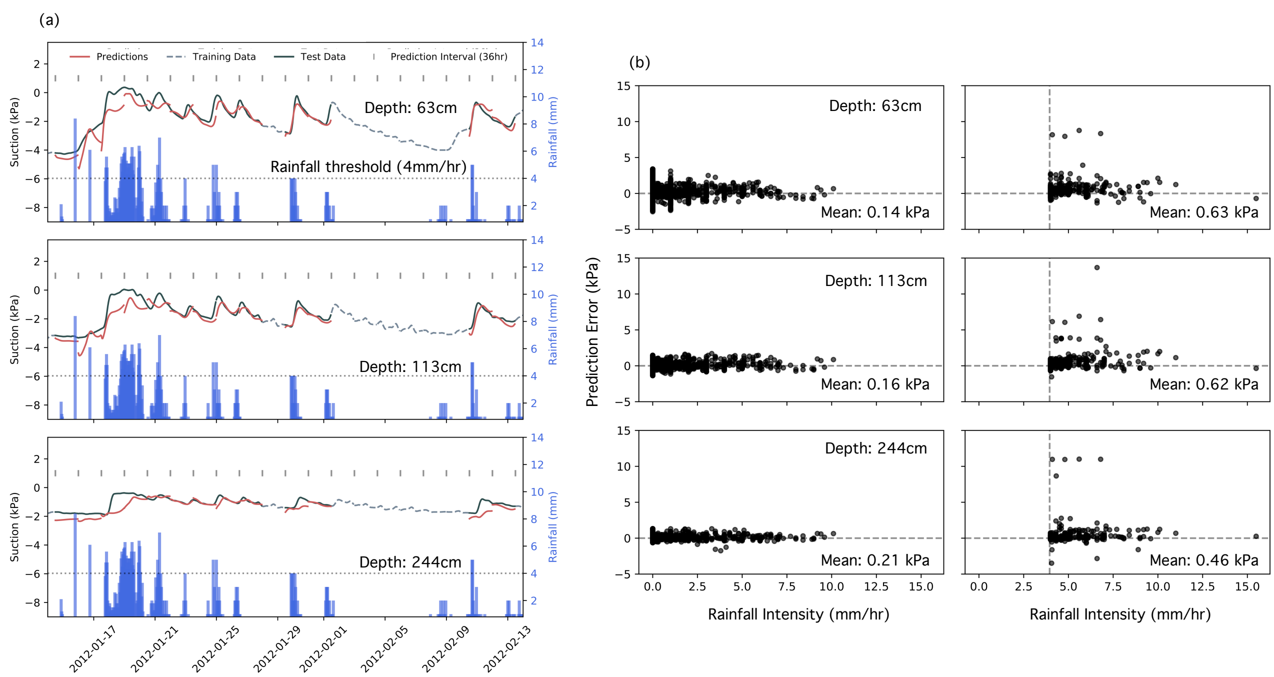

Figure 3. (a) Intensity-limited pore pressure response from January 14 - February 14, 2012. (b) Comparison of error values based on rainfall intensity. Model absolute error on testing data after training on all rainfall intensities featured in Water Years 1 & 2 (left). Intensity limited model absolute error from all rainfall intensities at or above 4mm/hr across all water years (right).

IV. DISCUSSION

The heterogeneity of landslide-prone hillslopes poses challenges to the processes encoded in hydrologic models. For example, complicating factors such as tree roots, macropores, soil horizons and otherwise heterogeneous soil conditions limit the ability for a simple, single parametrization to properly characterize field conditions. Mirus et al. (2016) simulate pore pressure values on Knife Ridge through 3D numerical subsurface flow modeling. Their model parameterization includes topographic data, soil depth information, and five hydraulic property parameters for each of the three hydrogeologic units included in their study. As a result, their work provides an important physical analog and gives valuable insight into the controls of hydrologic response, but does not reach the same level of accuracy with pore pressure predictions in comparison to our DL model. An increase in the number of modeling scenarios such as those utilized by Thomas et al. (2018a) can begin to account for inaccuracies behind a single physical model parameterization, yet the extended processing time even for 1D infiltration scenarios is problematic, and this limitation will only increase with additional dimensions. The ML-based modeling presented here outperforms the processing time of the HYDRUS-1D (Simunek et al., 2005) model utilized by Thomas et al. (2018a) by as many as three orders of magnitude (128hrs vs 2-10 minutes). Thus, with the naturaand uncertainty of soil conditions for field sites, a well calibrated ML-based model has the capability to learn and account for site-specific characteristics without imposing a computational limitation.

As our sensitivity tests demonstrate, the limits of our model's accuracy are determined by the quantity and quality of the training data, which as a pre-requisite must capture the underlying physical relationships we seek to model. Because our study area is characterized by a Mediterranean climate, training data is required during fall and winter months when rainfall events that promote landslides are most common. Furthermore, our results suggest that if training data happen to be acquired during a period lacking rainfall events of significant intensity, the model can approximate soil hydrologic response to high rainfall events, albeit with reduced accuracy. This issue of extrapolation beyond conditions within the period of record is common to many environmental and hazards research pursuits.

By closely approximating the timing and magnitude of rises in pore pressure along a variety of wetting/drying conditions, our LSTM-based model demonstrates a fast, accurate, and computationally efficient alternative to physically based modeling of landslide hazard. We anticipate that this approach can provide accurate predictions for integration into landslide early warning systems. Our results suggest that it is possible to train an algorithm that does not require the same sets of parameters as physical models, but uses site-specific data to capture the complexity of in situ measurements, and replicates the physical relationships behind 3D flow for variably saturated subsurface conditions.

We present our model as a proof-of-concept for how machine learning can help forecast conditions for use in early warning systems. There are several modifications that could generate improved results, such as an increase in the length of observation data, the potential for additional insight from multi-head convolutional layers such as those utilized by Zheng et al. (2014), or adapting more advanced attention-based models such as the Transformer (Vaswani et al., 2017). The latter represents a promising next step. As the physical significance of a neural network's weights are challenging to interpret, a more in-depth study of attention in future network architectures may provide further insight into the influence of each physical input at each timestep.

V. CONCLUSIONS

We demonstrate how an LSTM-based DL algorithm can accurately predict the timing and magnitude of pore pressure response to rainfall in 36-hour intervals. Our model is computationally efficient and can be used to provide real-time forecasts of hydrologic response conditions on landslide-prone hillslopes for early warning. The quantity and quality of the training data used for the model serves as a first-order limit on the model's predictive capability. We further show that when trained on exclusively low intensity rainfall, our model can approximate pore pressure response to high intensity events. With previous work exemplifying the mathematical capabilities of ANNs to learn and understand physical processes in hydrology and geosciences, we present this model as a next step to more accurately constrain near-surface hydrologic response in naturally heterogeneous areas where physically based modeling is computationally prohibitive.

APPENDIX: SUPPLEMENTAL FIGURE 1

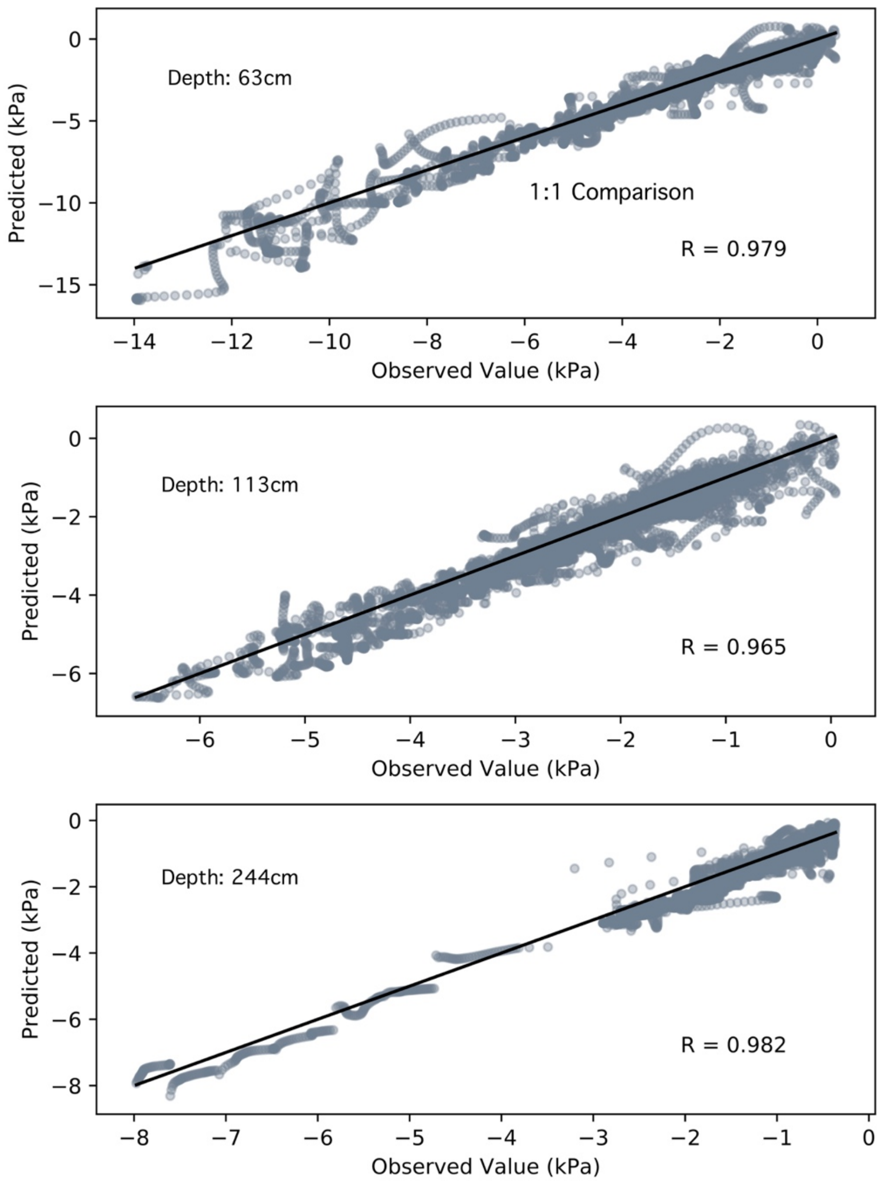

Supplemental Figure 1. Comparison of predicted and observed soil matric potential (suction) values. Reference line demonstrates 1:1 fit. These data are inherently temporally autocorrelated; we do not suggest causality between the two individually and separately autocorrelated datasets, but rather highlight the strong agreement between observed and predicted values across all possible observations.

REFERENCES CITED

Abrahart, R. J., & See, L. M. (2007). Neural network modelling of non-linear hydrological relationships. *Hydrology and Earth System Sciences*, *11*(5), 1563–1579. https://doi.org/10.5194/hess-11-1563-2007

Araya, S. N., & Ghezzehei, T. A. (2019). Using Machine Learning for Prediction of Saturated Hydraulic Conductivity and Its Sensitivity to Soil Structural Perturbations. *Water Resources Research*, *55*(7), 5715–5737. https://doi.org/10.1029/2018WR024357

Bellugi, D., Milledge, D. G., Dietrich, W. E., Perron, J. T., & McKean, J. (2015). Predicting shallow landslide size and location across a natural landscape: Application of a spectral clustering search algorithm. *Journal of Geophysical Research: Earth Surface*, *120*(12), 2552–2585. https://doi.org/10.1002/2015JF003520

Bogaard, T., & Greco, R. (2018). Invited perspectives: Hydrological perspectives on precipitation intensity-duration thresholds for landslide initiation: proposing hydro-meteorological thresholds. *Natural Hazards and Earth System Sciences*, *18*(1), 31–39. https://doi.org/10.5194/nhess-18-31-2018

Box, G. E. P., Jenkins, G. M., & Reinsel, G. C. (2008). *Time Series Analysis*. *Time Series Analysis: Forecasting and Control: Fourth Edition*. Wiley. https://doi.org/10.1002/9781118619193

Coe, J. A., Michael, J. A., & Burgos, M. M. (2011). *Map of debris flows caused by rainfall during 1996 in parts of the Reedsport and Deer Head Point quadrangles, Douglas County, southern Coast Range, Oregon Open-File Report 2011-1150*. Retrieved from http://pubs.usgs.gov/of/2011/1150/

DeVries, P. M. R., Thompson, T. Ben, & Meade, B. J. (2017). Enabling large-scale viscoelastic calculations via neural network acceleration. *Geophysical Research Letters*, *44*(6), 2662–2669. https://doi.org/10.1002/2017GL072716

Ebel, B. A., Loague, K., Montgomery, D. R., & Dietrich, W. E. (2008). Physics-based continuous simulation of long-term near-surface hydrologic response for the Coos Bay experimental catchment. *Water Resources Research*. https://doi.org/10.1029/2007WR006442

Ebel, B. A., Loague, K., Vanderkwaak, J. E., Dietrich, W. E., Montgomery, D. R., Torres, R., & Anderson, S. P. (2007). Near-surface hydrologic response for a steep, unchanneled catchment near Coos Bay, Oregon: 2. Physics-based simulations. *American Journal of Science*. https://doi.org/10.2475/04.2007.03

Fusco, De Vita, Mirus, Baum, Allocca, Tufano, … Calcaterra. (2019). Physically Based Estimation of Rainfall Thresholds Triggering Shallow Landslides in Volcanic Slopes of Southern Italy. *Water*, *11*(9), 1915. https://doi.org/10.3390/w11091915

Guzzetti, F., Peruccacci, S., Rossi, M., & Stark, C. P. (2008). The rainfall intensity-duration control of shallow landslides and debris flows: An update. *Landslides*, *5*(1), 3–17. https://doi.org/10.1007/s10346-007-0112-1

Hochreiter, S., & Schmidhuber, J. (1997). Long Short-Term Memory. *Neural Computation*. https://doi.org/10.1162/neco.1997.9.8.1735

Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks*, *2*(5), 359–366. https://doi.org/10.1016/0893-6080(89)90020-8

Kingma, D. P., & Ba, J. (2014). Adam: A Method for Stochastic Optimization, 1–15. Retrieved from http://arxiv.org/abs/1412.6980

Kratzert, F., Klotz, D., Brenner, C., Schulz, K., & Herrnegger, M. (2018). Rainfall-runoff modelling using Long Short-Term Memory (LSTM) networks. *Hydrology and Earth System Sciences*. https://doi.org/10.5194/hess-22-6005-2018

Lu, N., & Likos, W. J. (2004). Rate of Capillary Rise in Soil. *Journal of Geotechnical and Geoenvironmental Engineering*, *130*(6), 646–650. https://doi.org/10.1061/(ASCE)1090-0241(2004)130:6(646)

Lu, Ning, & Godt, J. (2008). Infinite slope stability under steady unsaturated seepage conditions. *Water Resources Research*. https://doi.org/10.1029/2008WR006976

Luong, M. T., Pham, H., & Manning, C. D. (2015). Effective approaches to attention-based neural machine translation. In *Conference Proceedings - EMNLP 2015: Conference on Empirical Methods in Natural Language Processing*. https://doi.org/10.18653/v1/d15-1166

McDonnell, J. J. (1990). The influence of macropores on debris flow initiation. *Quarterly Journal of Engineering Geology*. https://doi.org/10.1144/GSL.QJEG.1990.023.04.06

Mirus, B. B., Jones, E. S., Baum, R. L., Godt, J. W., Slaughter, S., Crawford, M. M., … McCoy, K. M. (2020). Landslides across the USA: occurrence, susceptibility, and data limitations. *Landslides*. https://doi.org/10.1007/s10346-020-01424-4

Mirus, B.B., Smith, J., Godt, J., Baum, R., & Coe, J. (2016). Simulated effect of topography and soil properties on hydrologic response and landslide potential under variable rainfall conditions in the Oregon Coast Range, USA. In *Landslides and Engineered Slopes. Experience, Theory and Practice* (pp. 1431–1439). CRC Press. https://doi.org/10.1201/b21520-176

Mirus, B.B, Morphew, M., & Smith, J. (2018). Developing Hydro-Meteorological Thresholds for Shallow Landslide Initiation and Early Warning. *Water*, *10*(9), 1274. https://doi.org/10.3390/w10091274

Montgomery, D. R., & Dietrich, W. E. (1994). A physically based model for the topographic control on shallow landsliding. *Water Resources Research*, *30*(4), 1153–1171. https://doi.org/10.1029/93WR02979

Montgomery, D. R., Schmidt, K. M., Dietrich, W. E., & McKean, J. (2009). Instrumental record of debris flow initiation during natural rainfall: Implications for modeling slope stability. *Journal of Geophysical Research: Earth Surface*. https://doi.org/10.1029/2008JF001078

Olah, C. (2015). Understanding LSTM Networks. Retrieved from http://colah.github.io/posts/2015-08-Understanding-LSTMs/

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., … Duchesnay, É. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*.

Ren, H., Cromwell, E., Kravitz, B., & Chen, X. (2019). Using Deep Learning to Fill Spatio-Temporal Data Gaps in Hydrological Monitoring Networks. *Hydrology and Earth System Sciences Discussions*, (May), 1–20. https://doi.org/10.5194/hess-2019-196

Richards, L. A. (1931). Capillary Conduction of Liquids Trough Porous Mediums. *Physics*, *1*(5), 318–333. https://doi.org/10.1063/1.1745010

Robison, E. G., Mills, K. A., Paul, J., Dent, L., & Skaugset, A. (1999). *Oregon Department of Forestry Storm Impacts and Landslides of 1996: Final Report*. Retrieved from https://www.waterboards.ca.gov/water_issues/programs/tmdl/records/region_1/2003/ref1785.pdf

Roering, J. J., Schmidt, K. M., Stock, J. D., Dietrich, W. E., & Montgomery, D. R. (2003). Shallow landsliding, root reinforcement, and the spatial distribution of trees in the Oregon Coast Range. *Canadian Geotechnical Journal*. https://doi.org/10.1139/t02-113

Shen, C., Laloy, E., Elshorbagy, A., Albert, A., Bales, J., Chang, F.-J., … Tsai, W.-P. (2018). HESS Opinions: Incubating deep-learning-powered hydrologic science advances as a community. *Hydrology and Earth System Sciences*, *22*(11), 5639–5656. https://doi.org/10.5194/hess-22-5639-2018

Simunek, J., Genuchten, M. Van, & Sejna, M. (2005). The HYDRUS-1D software package for simulating the one-dimensional movement of water, heat, and multiple solutes in variably-saturated media. *HYDRUS Software Ser*.

Smith, B. J. B., Godt, J. W., Baum, R. L., Coe, J. A., Burns, W. J., Lu, N., … Jewell, S. (2014). *Hydrologic Monitoring of a Landslide-Prone Hillslope in the Elliott State Forest , Southern Coast Range , Oregon , 2009 – 2012*. *U.S. Geological Survey Open-file Report 2013-1283*. Retrieved from https://pubs.usgs.gov/of/2013/1283/pdf/of13-1283.pdf

Sutskever, I., Vinyals, O., & Le, Q. V. (2014). Sequence to sequence learning with neural networks. In *Advances in Neural Information Processing Systems*.

Thomas, M. A., Mirus, B. B., & Collins, B. D. (2018a). Identifying Physics-Based Thresholds for Rainfall-Induced Landsliding. *Geophysical Research Letters*, *45*(18), 9651–9661. https://doi.org/10.1029/2018GL079662

Thomas, M. A., Mirus, B. B., Collins, B. D., Lu, N., & Godt, J. W. (2018b). Variability in soil-water retention properties and implications for physics-based simulation of landslide early warning criteria. *Landslides*. https://doi.org/10.1007/s10346-018-0950-z

Turner, A., & Schuster, R. L. (1996). *Socioeconomic Significance of Landslides*. *Landslides: Investigation and Mitigation*. Retrieved from http://onlinepubs.trb.org/Onlinepubs/sr/sr247/sr247.pdf

van Natijne, A. L., Lindenbergh, R. C., & Bogaard, T. A. (2020). Machine Learning: New Potential for Local and Regional Deep-Seated Landslide Nowcasting. *Sensors*, *20*(5), 1425. https://doi.org/10.3390/s20051425

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., … Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems*.

Wieczorek, G., & Guzzetti, F. (2000). A review of rainfall thresholds for triggering landslides. *Mediterranean Storms, Proceedings of the EGS Plinius Conference '99*, *8*(January), 404–414. Retrieved from http://www.idrologia.polito.it/~claps/pliniusonline/pdf_proceedings/Plinius/Wieczorek/WIECZOREK.pdf

Wiley, T. J. (2000). Relationship Between Rainfall and Debris Flows in Western Oregon. *Oregon Geology*, *62*(2), 27–43. Retrieved from https://people.wou.edu/~taylors/andrews_forest/refs/wiley_2000.pdf

Zheng, Y., Liu, Q., Chen, E., Ge, Y., & Zhao, J. L. (2014). Time series classification using multi-channels deep convolutional neural networks. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. https://doi.org/10.1007/978-3-319-08010-9_33