

This work was written as part of one of the author's official duties as an Employee of the United States Government and is therefore a work of the United States Government. In accordance with 17 U.S.C. 105, no copyright protection is available for such works under U.S. Law.

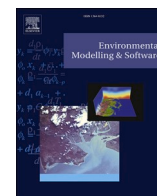
Public Domain Mark 1.0

<https://creativecommons.org/publicdomain/mark/1.0/>

Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

Please provide feedback

Please support the ScholarWorks@UMBC repository by emailing scholarworks-group@umbc.edu and telling us what having access to this work means to you and why it's important to you. Thank you.



Building a landslide hazard indicator with machine learning and land surface models

T.A. Stanley^{a,b,c,*}, D.B. Kirschbaum^c, S. Sobieszczuk^d, M.F. Jasinski^c, J.S. Borak^{c,e}, S. L. Slaughter^f

^a Universities Space Research Association, Columbia, MD, USA

^b Goddard Earth Sciences Technology and Research, Columbia, MD, USA

^c Hydrological Sciences Laboratory, NASA Goddard Space Flight Center, Greenbelt, MD, USA

^d Oregon Water Science Center, United States Geological Survey, Portland, OR, USA

^e Earth System Science Interdisciplinary Center, University of Maryland, College Park, College Park, MD, USA

^f Landslide Hazards Program, United States Geological Survey, Golden, CO, USA

ARTICLE INFO

Keywords:

XGBoost
Washington
Oregon
Land data assimilation system
Gradient boosting machine
National climate assessment

ABSTRACT

The U.S. Pacific Northwest has a history of frequent and occasionally deadly landslides caused by various factors. Using a multivariate, machine-learning approach, we combined a Pacific Northwest Landslide Inventory with a 36-year gridded hydrologic dataset from the National Climate Assessment – Land Data Assimilation System to produce a landslide hazard indicator (LHI) on a daily 0.125-degree grid. The LHI identified where and when landslides were most probable over the years 1979–2016, addressing issues of bias and completeness that muddy the analysis of multi-decadal landslide inventories. The seasonal cycle was strong along the west coast, with a peak in the winter, but weaker east of the Cascade Range. This lagging indicator can fill gaps in the observational record to identify the seasonality of landslides over a large spatiotemporal domain and show how landslide hazard has responded to a changing climate.

1. Introduction

Landslides have caused substantial property losses and fatalities in the Pacific Northwest of the United States. Fortunately, the states of Washington and Oregon have invested in landslide research, as have federal and local agencies. An important fruit of this effort, landslide inventories have already been of great assistance to urban planning, testing of early warning systems, and prioritizing future research. In this manuscript, we show how machine learning can describe the seasonality of landslides over a broad region by fusing landslide inventories and land data assimilation systems.

In the city of Seattle, most landslides happen in winter, especially the month of January (Laprade and Tubbs, 2008). A much shorter record of landslides in the city of Portland, Oregon showed frequent landslide occurrence from December to March (Mirus et al., 2018b). A similar pattern was found in a Pacific Northwest Landslide Inventory (PNLI) that included data from Portland, Seattle, and many other locations (Kirschbaum et al., 2016). However, data on the timing of landslides are rare outside of metropolitan areas. Furthermore, major temporal gaps

can be seen in the PNLI, especially in the years 1979–1995 and 1998–2005 (Stanley et al., 2019). Given the diversity of terrains and climates across the Pacific Northwest, it seems unwise to assume that all locations follow the seasonal pattern established within two well-documented metropolitan regions. Instead, landslide models must fill gaps in the recorded history of the Pacific Northwest.

Previous research has identified atmospheric rivers (long plumes of water vapor in the troposphere) as the principal cause of major floods in Western Washington (Neiman et al., 2011), and landslides along the West Coast (Biasutti et al., 2016; Cordeira et al., 2019). However, atmospheric rivers reaching the coast of Washington or Oregon are most frequent in November (Gershunov et al., 2017), which is quite early, seasonally, relative to the historical landslide climatology. Although the duration, intensity, and frequency of atmospheric rivers along the U.S. West Coast do not follow the same climatology (Gershunov et al., 2017), vertically integrated horizontal vapor transport also peaks early in the season, during December; and atmospheric rivers last longer early in the season. Recent research suggests that “families” (temporal clusters) of atmospheric rivers are substantially more hazardous than lone

* Corresponding author. Universities Space Research Association, Columbia, MD, USA.

E-mail address: thomas.a.stanley@nasa.gov (T.A. Stanley).

<https://doi.org/10.1016/j.envsoft.2020.104692>

Received 4 December 2019; Received in revised form 2 March 2020; Accepted 9 March 2020

Available online 13 March 2020

1364-8152/© 2020 The Authors.

Published by Elsevier Ltd.

This is an open access article under the CC BY-NC-ND license

(<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

atmospheric rivers (Ogle et al., 2018). Taken together, these findings suggest that individual storms are not sufficient to explain patterns of landslide occurrence; antecedent conditions must also be considered.

Researchers have long been aware of the importance of antecedent rainfall to hydrologic response in this region (Istok and Boersma, 1986). In the City of Seattle, landslide occurrence was estimated by combining 3 days of recent and 15 days of antecedent precipitation (Chleborad, 2000). More recently, an exhaustive analysis revealed that 32 days of antecedent precipitation better identified landslide events in much of the Seattle Metro area (Scheevel et al., 2017). Applying a similar recent-antecedent (RA) rainfall threshold across the Pacific Northwest resulted in numerous successfully identified landslides, but the model's false positive rate (FPR) was quite high in some locations, and the monthly climatology overemphasized the beginning of landslide season in November (Stanley et al., 2017). Since antecedent rainfall functions not as the direct cause of landslides but as a proxy for groundwater state, an antecedent wetness index was created to describe the accumulation of water in a slope (Godt et al., 2006). This index was combined with the intensity and duration of rainstorms to predict most major landslide events in Seattle, Washington.

Subsequent research showed that replacing the antecedent rainfall component of a RA threshold with soil saturation measured by subsurface probes reduced the frequency of false positives, especially in the summer and fall (Mirus et al., 2018a). Melting snow can also be a risk factor for landslides (Chleborad, 2000). Snow depth has been incorporated in a prototype landslide early warning system for Seattle (Chleborad et al., 2008). In particular, rain on snow (ROS) can generate massive runoff and landslide events (Chleborad, 2000; Marks et al., 1998). Beyond these well-established variables, other factors, such as air temperature (Chleborad, 2000, 1998, 1997), soil temperature (Tubbs, 1974), and wind (Buma and Johnson, 2015), have been identified as predictors of landslides in mountainous or coastal areas of the western United States. A global analysis of satellite precipitation estimates indicated that landslide behavior of the western United States is not well explained by rainfall alone, which suggests that other factors are important in this region (Jia et al., 2019). Although intense rainfall is an important trigger of landslides in the Pacific Northwest, many other factors help to explain the spatiotemporal patterns of landslide occurrence.

Prior landslide models rarely considered all of these time-varying factors, but the failure to do so can result in outputs that are climatologically incorrect or confined to a specific place. Machine learning offers a way to model landslide occurrence in a consistent, objective manner, without omitting major hydrological variables.

Machine learning has been applied in landslide research for over a decade (Lee et al., 2002; Marjanovic et al., 2009). The vast majority of this research focused on static landslide susceptibility (Korup and Stolle, 2014), but machine learning was also applied to dynamic landslide prediction (Farahmand and AghaKouchak, 2013), prediction of landslide displacement from in situ data (Lian et al., 2014; Zhao et al., 2018), and the remote sensing of landslides (Chen et al., 2014; Ghorbanzadeh et al., 2019; Stumpf and Kerle, 2011). Most scientists reported good to excellent predictive performance (Korup and Stolle, 2014; Lun et al., 2017). However, researchers should take steps to avoid overfitting (Brenning, 2005) and evaluate the geomorphic plausibility of model outputs (Steger et al., 2016). Researchers must also consider the potential effects of incomplete or inaccurate landslide inventories upon both the training and testing of empirical models (Stanley and Kirschbaum, 2017; Steger et al., 2017). Despite these notes of caution, machine learning is a powerful tool for understanding complex climate data and predicting rare events like landslides.

We describe the seasonality of landslides with a landslide hazard indicator (LHI) that scales well across space and time. Here, LHI is defined as the daily probability that a landslide will occur, averaged over a given space-time domain. At coarser temporal resolutions, this lagging indicator is similar—but not identical to—landslide hazard (the

yearly probability of landslide occurrence is greater than the mean of the daily probabilities). LHI offers a few benefits for landslide research. First, it is directly comparable to observed frequencies of a Boolean variable. This means that for a well-calibrated model, the mean of predictions and observations will be equal. Second, the simplicity of calculation leads to a simplicity of interpretation: the values shown are always the mean daily probability over a given area (in this case, 0.125-degree grid cells). In contrast, using the number of predicted landslides as an indicator might result in numerous cases of fractional landslides, while return periods are notorious for misinterpretation (Read and Vogel, 2015; Sayers, 2016; Vogel et al., 2014). Third, LHI enables comparisons across space, time, and even other variables. Even geographic areas of unequal size or time periods of unequal length could be compared. LHI identifies the times and places at the greatest hazard with a consistent quantitative metric that facilitates the search for patterns across space and time.

2. Data

2.1. National Climate Assessment – land data assimilation system (NCA-LDAS)

The NCA-LDAS is a terrestrial-hydrology satellite data-assimilation system that was created to support the National Climate Assessment (NCA) (Jasinski et al., 2019; Kumar et al., 2018; Xia et al., 2012). Developed within the NASA Land Information System (LIS) modeling framework (Kumar et al., 2006), NCA-LDAS integrates input data from multiple sources, including a 36-year record of satellite observations, to provide gridded, daily time-series over the continental U.S. of land surface variables, including precipitation, runoff, latent and sensible heat fluxes, soil moisture at multiple depths, and snow-water equivalent. NCA-LDAS Version 2.0 employs the Noah Version 3.3 land surface model (Ek et al., 2003) and is driven by atmospheric forcing data from the North American Land Data Assimilation System-Phase 2 (NLDAS-2) (Xia et al., 2012). NCA-LDAS data products are provided at 0.125-degree resolution for the satellite era – water years 1980–2015 (Jasinski et al., 2018). The NLDAS-2 precipitation forcing field is a product of a temporal disaggregation of a gauge-only Climate Prediction Center analysis of daily precipitation (Chen et al., 2008; Higgins et al., 2000) and includes an orographic adjustment (Daly et al., 1994). The rest of the forcing fields used by NCA-LDAS, including surface temperature, radiation, wind, and humidity were derived from the analysis fields of the National Centers for Climate Prediction North American Regional Reanalysis. LIS generated a variety of hydrological variables from these forcings, including soil moisture and snow depth. LIS performed simultaneous assimilation of satellite data with a 1-D ensemble Kalman filter. This data assimilation process systematically improved the accuracy of soil moisture and snow depth, especially when data from newer sensors were available (Kumar et al., 2018). Analysis of these data revealed predominantly declining trends across the Pacific Northwest in several variables relevant to landslide activity, including heavy precipitation, snow cover, and runoff (Jasinski et al., 2019).

Because precipitation thresholds for landslide warning have been shown to vary widely over space (Baum and Godt, 2009), we rescaled rainfall and all other NCA-LDAS variables in proportion to each grid cell's full range: $\text{rescaled grid cell value} = (\text{grid cell value} - \text{grid cell minimum}) / (\text{grid cell maximum} - \text{grid cell minimum} + 10^{-20})$. (A small value was added to the denominator to avoid divide-by-zero errors in rare cases.) For some machine-learning techniques, rescaling or normalizing data greatly improves performance, but this is not required for the tree-based method presented below. This preprocessing provides a strong localization effect that respects both the effects of rainfall frequency and maximum intensity, and it is easily reversed, interpreted, and replicated. Nevertheless, we do not claim that this is the optimal method for model localization; future research may show that a

more sophisticated approach delivers a benefit to model performance.

2.2. 3DEP

Although very high-resolution digital elevation models (DEM) have been derived from LIDAR over much of the Pacific Northwest, these data did not provide complete coverage. The coarse resolution of the NCA-LDAS forcing variables implied that a relatively coarse DEM with a high level of consistency would be more appropriate. The 3DEP dataset with a resolution of 0.333 arcseconds (U.S. Geological Survey, 2016) combines a high level of consistency and trust with a resolution that is still relevant to landslide behavior. In order to represent terrain at the modeling resolution of 0.125° , we calculated a relief value by subtracting the lowest from the highest elevation within each 0.125° -degree grid cell. This range of elevations does not capture the geomorphological nuances useful for landslide modeling at the local scale, but it does provide a simple topographic index to modulate the response to rainfall and other triggers. As our primary goal was to examine the role of climate in landslide activity, additional topographic variables were not added to the model.

2.3. Ecoregions

In order to summarize and disaggregate landslide seasonality over the Pacific Northwest, we adopted the Level I Ecoregions dataset (U.S. Environmental Protection Agency, 2010). The Commission for Environmental Cooperation developed this ecological classification to address environmental issues shared by Canada, Mexico, and the United States. Level I is the coarsest classification, with 15 ecoregions that cover the entirety of North America. Only three of these zones occur in the states of Oregon and Washington: the Marine West Coast Forest, the Northwestern Forested Mountains, and the North American Deserts (Fig. 1).

2.4. PNLI

A landslide inventory for the Pacific Northwest was created by

merging data from several existing inventories (Kirschbaum et al., 2016), as well as collecting new landslide reports from news media and other sources. Most of the landslides in the PNLI originated in the Global Landslide Catalog (Kirschbaum et al., 2015, 2010), the Statewide Landslide Information Database for Oregon (Burns, 2014), and landslide inventories from the Washington Department of Natural Resources (Sarikhani et al., 2009; Sarikhani et al., 2008; Washington Division of Geology and Earth Resources, 2016). Over 85% of the inventoried landslides were considered shallow landslides or debris flows. In order to determine relationships between historical landslide occurrence and possible explanatory variables from NCA-LDAS, we converted the 14,495 landslides with known dates to a gridded dataset with the same temporal and spatial resolution as NCA-LDAS. Unfortunately, this process required omission of most data in the source inventories, because these—often prehistoric—landslides do not have specific known dates. In this paper, we shall refer to this gridded record of landslide occurrence as the PNLI, although “PNLI” has previously been used to describe the underlying vector data. The PNLI is affected by a reporting bias towards populated areas and highways, as the date of landslide events is much more likely to be known when humans are impacted (Kirschbaum et al., 2016). This problem is most salient in Washington’s Olympic Peninsula, where the PNLI does not reflect the large number of landslides with known locations but unknown dates of occurrence (Fig. 2). In addition, the PNLI is systematically biased towards the present since data was collected more consistently over the past decade. Thus, the PNLI alone is inadequate to provide a record of the climate’s effect on landslide behavior over multiple decades. We sought to address this problem by developing a landslide model that responds to changes in a multivariate, spatially consistent, gridded climate record.

Despite these limitations, some general patterns can be observed in the PNLI. Landslides have been reported much more frequently to the west of the Cascade Mountains (Fig. 2). Furthermore, landslide activity followed a clear seasonal cycle (Fig. 3), with the most frequent landslide activity in the winter months of December, January, and February. Landslides were also common in March and November. The North American Deserts did not follow this cycle, but very few landslides were reported in this ecoregion. Landslides were most frequent in the Marine




Fig. 1. The states of Washington and Oregon contain three Level I Ecoregions: the Marine West Coast Forest, the Northwestern Forested Mountains, and the North American Deserts.




Fig. 2. This map shows the number of reported landslide events per year from 2007 to 2015 over the 0.125-degree grid cells. At most locations, landslides have been reported less than once a year. Landslides were reported much more frequently in the city of Portland, which we believe is primarily caused by a reporting bias towards urban areas.

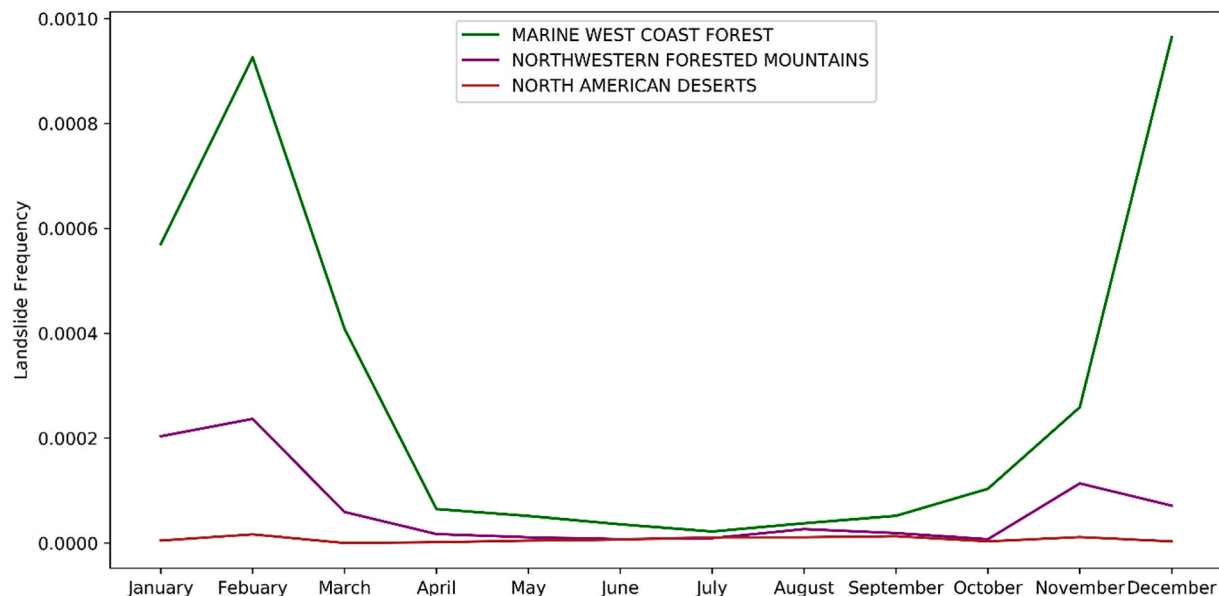


Fig. 3. The historical landslide frequency was much higher in the Marine West Coast Forest ecoregion (green), but the gap shrank during the summer (1996–2015). Landslides were rarely reported in the North American Deserts (brown); this ecoregion did not show an obvious seasonal cycle. Overall, landslides were reported most frequently during the winter months. The dual peaks of activity in December and February were caused by the presence of two major events in the Pacific Northwest Landslide Inventory on February 6, 1996 and December 3, 2007. (For interpretation of the references to color in this figure legend, the reader is referred to the Web version of this article.)

West Coast Forest ecoregion, although the difference was less pronounced from April to September.

The distribution of landslides was highly skewed (Fig. 4). Just eight days over the years 1996–2015 accounted for 87% of landslide grid cells, as well as almost 90% of historical landslides from 1996 to 2016 (Stanley et al., 2017). No landslides were reported on the vast majority of days (6,820) during this period. The dominance of a handful of landslide events suggests the importance of long records to a good understanding of landslide hazard.

3. Methods

3.1. Gradient boosting

From the wide variety of machine-learning software, we selected XGBoost, a form of decision tree model that uses gradient boosting to achieve a good empirical fit. XGBoost has become the first choice of many data scientists due to its well-documented speed and accuracy (Chen and Guestrin, 2016). Furthermore, landslide scientists have applied it to both susceptibility mapping (Chakraborty et al., 2019) and displacement monitoring (Zhao et al., 2018). XGBoost is a gradient boosting machine (GBM), a class of models that optimizes classification performance by combining many weakly predictive models into a single, highly accurate ensemble. The algorithm relies upon multiple iterations, each of which corrects to the model produced in the previous step, a process known as “boosting”. A GBM optimizes performance by following the steepest gradient along a differentiable loss function (Friedman, 2001). With this method, scientists can fit a model that is highly predictive of a training dataset. However, overfitting (Fig. 5) can reduce the performance of the model on new data. This is commonly addressed by random subsampling of the training data, a technique which can greatly improve the ability of a GBM to generalize over new data (Friedman, 2002). In addition, XGBoost addresses this issue by regularizing the objective function (Chen and Guestrin, 2016); i.e. it adds a term that represents model complexity to the loss function that describes classification error:

$$L(\phi) = \sum_i l(y_i, \hat{y}_i) + \sum_k \Omega(f_k)$$

Here, l is the loss function that compares the predicted value, \hat{y}_i (in this case, the probability that a landslide has occurred), to the observed value, y_i (1 if a landslide was recorded, otherwise 0), and Ω represents the complexity of each tree added to the model, f_k . In XGBoost, this regularization is primarily applied through the γ parameter, which




Fig. 4. The geographic reach of a landslide event was related to its frequency. The x-axis shows the number of 0.125-degree grid cells in which a landslide was recorded in the Pacific Northwest Landslide Inventory for a given day. The y-axis shows the number of days on which a landslide event with a given geographic reach occurred. A small number of landslide events dominated the historical record (1996–2015).




Fig. 5. An example of overfitting from a preliminary model that incorporated mean sediment thickness. The red dot shows the modeled probability that a landslide occurred near Mukilteo, Washington on December 19, 2012. The blue dots show the modeled probability that a landslide would occur, given other sediment thicknesses and holding all other variables constant. The model has clearly identified one historical landslide in the training data, rather than learned a general rule about landslide behavior. Sediment thickness was not included in the final model. (For interpretation of the references to color in this figure legend, the reader is referred to the Web version of this article.)

requires each new split to improve the existing tree by a minimum amount. Although we utilized both subsampling and objective regularization when training the model, these were inadequate to obtain a stable and generally applicable prediction of landslide activity. Fortunately, XGBoost offers several additional parameters.

The rate at which machine learning proceeds can be directly controlled with the η parameter, also known as shrinkage or learning rate. This number represents the ratio by which the output of each iteration will be reduced. Thus, increasing η will speed up machine learning and result in a model with a small number of strongly predictive trees. In contrast, decreasing η will slow learning and produce more trees, each of which is less fitted to the data but can have improved prediction as a whole. In practice, larger ensembles of trees produce more accurate and reliable outputs. The default η of 0.3 is excellent for exploratory work, but data scientists often lower the value to achieve better generalization. We set η to 0.01 for the final model in our application.

Although XGBoost offers several parameters to mitigate overfitting, monotonicity constraints were especially effective. Although this parameter has been used in cybersecurity research (Incer et al., 2018), we believe that landslide scientists have not applied it previously. This may be because it nullifies an important advantage of machine learning over linear regression: the ability to model non-monotonic relationships. This loss may reduce performance on existing data, but it can increase the model's reliability for new conditions (Incer et al., 2018). Fortunately, researchers can control the presence and direction of monotonic constraints for every input variable. Even very conservative values for γ and other parameters failed to eliminate overfitting, so we applied monotonic constraints to all predictors for which a direction could be assigned with confidence from prior knowledge (Table 1). Other variables were not included in the model. We also made a simplification of soil temperature, which is available for each soil layer in the NCA-LDAS record. Rather than include four related variables that might not have a monotonic relationship with landslide probability, we checked if the temperature of any layer was below 0 °C. Although variable selection and monotonicity introduced a certain amount of subjectivity into the modeling process, the increase in reliability made it the best choice for obtaining a regional landslide climatology.

Table 1

These variables may predict landslides. We applied constraints to the GBM that forced it to eliminate any splits in the model's trees where the output does not increase or decrease monotonically. Although these constraints prevent the landslide model from achieving a perfect fit to the training data, enforcing monotonicity greatly increased its physical realism and generalization capacity. All predictors have a 0.125-degree, daily resolution.

Predictor	Monotonicity	Justification
Antecedent conditions		
Relief (derived from 3DEP elevations)	Increasing	Mountainous regions generally experience more landslides than flat plains.
CanopInt (Plant canopy surface water)	Increasing	Water stored in the canopy of forests may travel downwards and infiltrate or run off.
SnowDepth	Increasing	If melted, snow provides a large source of additional water.
Frozen (derived from soil temperatures)	Decreasing	Frozen ground may prevent the movement of both water and soil.
SoilMoist0_10 cm	Increasing	Soil moisture can serve as a proxy for groundwater conditions near the surface, and it may be correlated with conditions at depth.
SoilMoist10_40 cm	Increasing	
SoilMoist40_100 cm	Increasing	
SoilMoist100_200 cm	Increasing	
Forcings (Triggers)		
Rainf	Increasing	Rainfall is a well-known trigger of landslides.
Snowf	Increasing	Falling snow may accumulate mass on a slope or melt and introduce additional liquid.
Wind_f (Wind speed)	Increasing	Wind can topple trees, weakening root mass strength or directly triggering landslides. Wind can increase the rate of snowmelt. Wind may also serve as an indicator of storm intensity not captured by daily rainfall totals.
Qs (Storm surface runoff)	Increasing	Excessive runoff may initiate debris flows or undercut stream banks.
Qsb (Base flow-groundwater runoff)	Increasing	Water infiltrating from surficial soils may increase pore pressures at depth.
Qsm (Snow melt)	Increasing	Water from melting snow may infiltrate into ground water or run off across the ground surface.

The difference between the number of grid cells with and without recorded landslides, known as class imbalance, is an important challenge in landslide research (Stumpf and Kerle, 2011; Van Den Eeckhaut et al., 2009), as well as many other domains. Dynamic landslide modelers must contend with a much greater imbalance than is typically encountered in susceptibility mapping, because exact dates are not available for most inventoried landslides. Thus, the inventories of dated landslides are much smaller than the undated landslide inventories used for susceptibility mapping. Furthermore, the limited stockpile of landslide events must be divided across thousands of time steps. In this case, the ratio of landslides to other pixels is 7×10^{-5} . The strength of this class imbalance destabilized gradient descent, even with a low η . Fortunately, XGBoost provides two tools to address class imbalance. The first, `scale_pos_weight`, can raise the importance of rare events during model training. This can increase predictive ability (Song et al., 2018), but will produce incorrect probability estimates (xgboost developers, 2016). In order to avoid this problem, we focused on the second parameter, `max_delta_step`, which limits the strength of updates performed in each iteration of training. This differs from η in that it applies a hard limit to each update, rather than rescaling the value. A value of 10 produced a much more stable series of updates when training this landslide model. Table 2 shows the full list of parameter choices used to generate a landslide climatology of the Pacific Northwest.

In order to achieve the most accurate model, we limited the training data to the most recent section of the PNLI: water years 2011–2015. We withheld water years 1996–2010 to evaluate the performance of the trained model. This time period contains two very large landslide events and relatively few small landslide events. Since these events are outside

Table 2

These parameters control how XGBoost fits model to data. Many other options are available but were not used in the final version of this landslide model.

Parameter	Value
Tree method	"hist"
Objective	"binary:logistic"
Number of rounds	500
η (shrinkage)	0.01
γ (minimum loss reduction)	0.01
Maximum delta step	10
Subsample	0.8
Monotone constraints	See Table 1
Other parameters	Default values

the range of data for 2011–2015, they represent a difficult test for the model to overcome. Most of these early events were recorded in Oregon. The remainder of the PNLI is not a complete record of water year 2016, so it was not included in either the training or validation datasets. The entirety of the states of Washington and Oregon were utilized, but the coarse resolution of the NCA-LDAS dataset meant that a few landslides were located within water pixels and did not contribute to either training or testing the model.

3.2. Landslide climatology

While the process of building machine-learning models can, in itself, be informative, we produced a landslide hazard indicator in order to understand the historical patterns of landslide behavior across the Pacific Northwest. At the most basic level, we averaged the LHI over the full time series (1979–2016) at each location to produce a map that shows where landslides would be expected more frequently. We applied a similar analysis over time by averaging LHI for each month across the whole study area. We also produced a climatology for each ecoregion by determining the long-term mean LHI within each zone and month. These methods could be repeated for any subset of the study domain. In order to understand the climatology of landslides in more detail, we calculated a seasonality index (SI) that rates the temporal concentration of the LHI at each grid cell. Although this can be produced on a monthly basis (Markham, 1970), it can also be calculated from daily data (Burn, 1997). The SI is determined by treating each unit of time as an angle on the unit circle and assigning a length equal to the property of interest—in this case, LHI. Next, the resultant of these vectors is divided by the annual total to produce the SI value. Thus, a uniform distribution of LHI across the year would produce a SI of zero, while a perfect concentration of all landslides in a single month or day would produce a SI of one. In addition, the mean direction of the vectors can be determined and converted back to the relevant unit of time. Landslide activity does not necessarily peak on the resulting date; rather it is the center of the landslide season as a whole. Mean direction is more reliable for sites with high SI values; low values may indicate a lack of consistent seasonality and, thus, the average date might be uninformative for these sites. This range of climatological measures shows where and when landslides were likely to happen.

3.3. Co-occurrence of rain on snow and landslides

Since rain falling on snow has been identified as a major cause of landslides in this region, we also analyzed the historical frequency of these events. Drawing on a previously published model (Chleborad et al., 2008), we defined ROS as days in every grid cell where greater than 25.4 mm of rain fell on a layer of snow deeper than 150 mm. The historical frequency of landslides (3.0×10^{-3}) in these events is roughly two orders of magnitude higher than the overall baseline for the study area (8.5×10^{-5}). It is also slightly higher than the frequency of landslides when rainfall is greater than 25.4 mm and snow depth is not

greater than 150 mm (2.8×10^{-3}). In addition, a threshold of 25 mm recent rainfall has been identified as a rough minimum trigger for landslide activity across this region (Stanley et al., 2017). We computed the long-term frequency of ROS days at each location.

3.4. Trends in rain on snow and landslide seasonality

In addition to the primary analyses, we also calculated the trends over time in ROS, the LHI, and the SI derived from LHI. The rate of change was calculated with the Theil-Sen method for slope estimation (Theil, 1992), with Mann-Kendall significance testing (Kendall, 1948; Mann, 1945).

4. Results

4.1. Gradient boosting model evaluation

The model produced maps of LHI with a daily, 0.125-degree resolution. Typically, the western third of the study area showed the highest probability of landslide occurrence (Fig. 6), but this was not always the case (Fig. 7). This geographic pattern broadly matches the historical record of landslides (Fig. 3). Standard quantitative measures also indicate success in reproducing the observations in the PNLI. Binary predictions, including landslide maps, are often judged with the curve from the receiver operating characteristic (Petchko et al., 2014). Although the shape of the curve can convey more complex information, the area under this curve (AUC) serves as a standard measure of performance. Results of the model comparison have an AUC of 0.91 for the model training period (2011–2015), and an AUC of 0.95 for the evaluation period (1996–2010), suggesting excellent overall performance.

The process of model development is iterative, and we tested multiple models against the validation data. In order to ensure complete

independence, we did not test predictions for the year 2016 during the model development process. Therefore, we obtained two small inventories of dated landslides: NASA's Global Landslide Catalog (GLC) and Recent Landslides in Washington State (Washington DNR, 2019). Rock falls were removed from both inventories, and landslides with a spatial accuracy worse than 10 km were removed from the GLC. We benchmarked the new XGBoost model against a simple recent-antecedent (RA) rainfall threshold that was fitted to this study area (Stanley et al., 2017). The RA threshold predicted 37% of the GLC landslides and 31% of the Recent Landslides in Washington State, with a FPR of 6%. In order to compare these binary results to the probabilistic outputs of XGBoost, we applied a probability threshold of 0.00033126, which generates a binary output with the same FPR as the RA threshold, but a different spatiotemporal structure. With this binary output, XGBoost correctly predicted 55% of the GLC and 69% of the Washington landslides.

Since quantitative metrics can give false confidence in empirical models (Steger et al., 2016), we also sought to examine the plausibility of the relationships between model inputs and outputs. XGBoost provides a few default measures of variable importance; the most useful is “gain”, the average error reduction made by introducing each variable to the model (Chen and Guestrin, 2016). This measure shows that rainfall is the most useful predictor of landslides, while snowfall is the least (Fig. 8). Soil moisture and temperature also play important roles in modulating the output.

4.2. Landslide climatology

LHI provides an overview of landslide seasonality across the whole of the Pacific Northwest. As shown in Fig. 6, landslides were more probable along the west coast. Landslide-triggering conditions also occur frequently on the west side of the Cascades. The temporal pattern of




Fig. 6. Mean LHI by grid cell for the period 1979–2016. Light color indicates low probability of a landslide occurring at a location, and dark red indicates a higher probability. (For interpretation of the references to color in this figure legend, the reader is referred to the Web version of this article.)




Fig. 7. LHI on August 1, 2009. Light color indicates low probability of a landslide occurring on this day, and dark red indicates a higher probability. On this summer day, landslides were relatively improbable (one order of magnitude lower than the average rate), and the region of greatest hazard had shifted from the Coast Range to the North Cascades. (For interpretation of the references to color in this figure legend, the reader is referred to the Web version of this article.)




Fig. 8. This chart shows the information gained from each time a variable was considered by the model. Since each variable can show up many times, the error reduction was averaged across all of the instances of that variable in the model. In this case, the most informative predictors were rainfall, soil moisture, and soil temperature.

landslide behavior was revealed by calculating the monthly mean LHI of all grid cells in the study area (Fig. 9). The highest monthly LHI from 1979 to 2016 was recorded for February 1996, which was also the peak of landslide activity recorded in the PNLI. (The PNLI does not contain any landslides with precise dates reported before 1996.) Other notable

peaks in LHI can be seen in December 1981, February 1986, December 1996, February 1999, and January 2006. LHI typically peaked in winter, fell to very low levels in the summer, and did not rise until late in the fall. We also analyzed the monthly climatology within each major ecoregion (Fig. 10). LHI was highest in the “Marine West Coast Forest” ecoregion




Fig. 9. LHI was averaged across the study area by month. Both observations from the Pacific Northwest Landslide Inventory (blue) and the modeled LHI (orange) show a consistent seasonal cycle with a peak in February 1996, the first month for which dated landslides were available. The observed and modeled LHI magnitudes match closely during the training period (gray box) and the first half of water year 2016 but diverge for many of the earlier years. In most cases, gaps in the historical record probably caused the difference between the low frequency of observed landslides and the winter peaks in modeled LHI. (For interpretation of the references to color in this figure legend, the reader is referred to the Web version of this article.)




Fig. 10. This monthly landslide climatology was produced by determining the mean LHI for each month and EPA level 1 ecoregion. LHI was substantially higher in the Marine West Coast Forest ecoregion, although the gap narrows during the summer months. The Northwestern Forested Mountains experienced a flatter distribution of LHI than the coast, but the seasonal cycle still peaked in the winter. LHI was lowest in the North American Deserts ecoregion at all times of year. Unlike the Pacific Northwest as a whole, LHI peaked during March in the desert region, not December and January.

and lowest in the “North American Deserts” ecoregion. LHI peaked in the “Marine West Coast Forest” and “Northwestern Forest Mountains” during December and January, but March was the peak month in “North American Deserts”. Calculation of SI and mean direction of LHI at each grid cell corroborated this climatology (Fig. 11). The average date of the local landslide season was typically in the winter months of December, January, and February for locations west of the Cascades. Landslide season was less consistent east of the Cascades; the average dates were spread over fall, winter, or spring. However, LHI did not peak during summer anywhere in the study area. SI was very high along the west coast, indicating a strong seasonal pattern to LHI in this area (Fig. 12). In

some locations, SI exceeded 0.7. East of the Cascades, SI was more variable. In some of these areas, SI exceeded 0.6, while $SI < 0.1$ at many locations. The largest area without a consistent landslide season was located in central Washington; SI was also low over much of eastern Oregon. Overall, this climatological analysis showed that landslide hazard was most severe in winter, that LHI was declining but highly variable, and that the seasonal cycle was more consistent along the west coast.

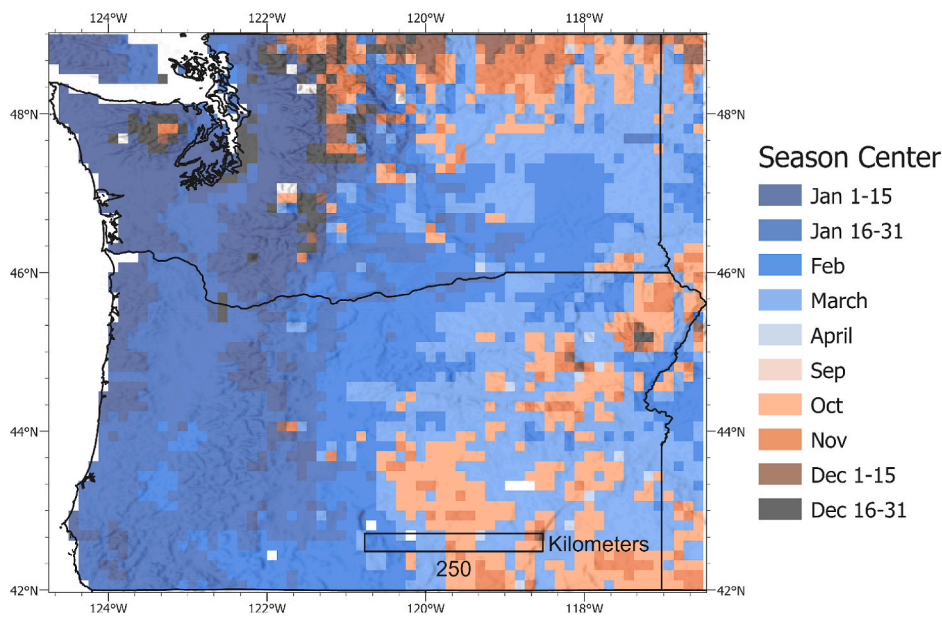


Fig. 11. The day of the year around which the typical landslide season is centered (known as “mean direction”). This date is not necessarily the same as the date of maximum LHI. It is calculated by treating each day of the year and its corresponding climatological mean LHI as a vector, then determining the resultant. West of the Cascade range, landslide season occurs during the months of December, January, and February. East of the Cascades, the seasonal mean direction can also occur in the Spring or the Fall.




Fig. 12. A seasonality index (SI) indicates the concentration (purple) or dispersion (brown) of LHI over a typical year, where a value of one would indicate that only one day contains the entire annual hazard and a value of 0 indicates a uniform distribution across all days. The areas west of the Cascade Range show a strong seasonal pattern in LHI, except for the Olympic Mountains. SI ranges from very high to very low values to the east of the Cascade Range, with a fairly consistent seasonal pattern along the Columbia and Snake Rivers. (For interpretation of the references to color in this figure legend, the reader is referred to the Web version of this article.)




Fig. 13. Monthly mean frequency of rain on snow across the study area for 1979–2016 is highest in the months of December, January, and February.

4.3. Co-occurrence of rain on snow and landslides

Rain falls on snow most frequently during the winter months of December, January, and February (Fig. 13). The months of March and November represent a transition between summer and winter conditions, with a lesser number of ROS days. This season matches the landslide cycle quite well (Fig. 3).

ROS was most common in the Olympic and Cascade Ranges, particularly the North Cascades (Fig. 14). The remainder of the study area experiences ROS only on rare occasions. This geographic distribution does not match that of historically recorded landslides (Fig. 2). Specifically, no dated landslides were recorded in the Olympic Mountains, even though this area experiences ROS frequently. In addition,

large numbers of landslides have been recorded in the Coast Range and Cascades of Oregon; ROS occurs rarely in most of these areas.

4.4. Trends

We found statistically significant declining trends in the annual number of ROS days at many locations (Fig. 15a). The rate of change was no more than a half day per year over the period 1979–2016; in most locations, it was much slower. Very few locations experienced increased ROS. Statistical significance was largely confined to the areas with frequent ROS. The overall picture is one of broadly declining numbers of ROS days at higher elevations and infrequent ROS at lower elevations and in drier regions. Although statistically significant trends were not




Fig. 14. Frequency of rain-on-snow days by location over the time period 1979–2016. Rain falls on snow more often in the Olympic and Cascade Ranges. These events are infrequent over most of the study area.

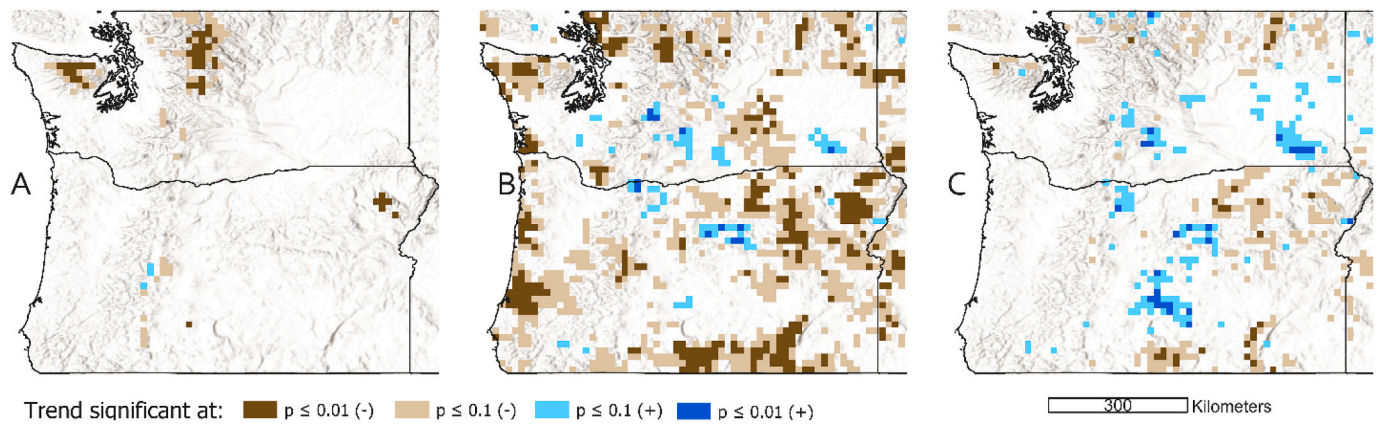


Fig. 15. Trends in (a) the number of rain-on-snow (ROS) days, (b) the landslide hazard indicator (LHI), (c) the seasonality index (SI) were calculated with the Theil-Sen method for slope estimation. Upward trends are shown in blue and downward trends in brown. Mann-Kendall tests showed that most locations have not experienced statistically significant changes (no color). Trends significant at the $p < 0.1$ level are shown with light colors, while trends significant at the $p < 0.01$ level are shown with dark colors. ROS declined at higher elevations, with minor upward trends observed in only a few locations (a). Only statistically significant trends ($p > 0.1$) are shown, but these tend to be found at the sites with the most ROS days. LHI decreased throughout the study area (b), despite the diversity of terrain and climate. The geographic distribution of declines in ROS and LHI do not match, which implies that the change in LHI is not dominated by changes in the frequency of ROS. Few locations west of the Cascades showed a significant trend in the SI, and the trends were not clustered into a few large regions (c). The combination of low statistical significance, roughly equal numbers of increases and declines, and lack of association between trends and major climate zones suggests that no major shift in the SI took place over the years 1979–2016. (For interpretation of the references to color in this figure legend, the reader is referred to the Web version of this article.)

found at most grid cells, LHI declined at most locations on both sides of the Cascade Range and increased in relatively few locations (Fig. 15b). Statistically significant trends in SI were relatively uncommon, even at the 0.1 level (Fig. 15c). SI increased and decreased in roughly equal numbers of locations. Nearly all of these were located east of the Cascades. However, neither the increases nor the declines were clustered in a specific area.

5. Discussion

While we implicitly or explicitly modeled a wide range of triggers, pre-conditioning factors, and landslide types, some topics were not analyzed. Rock falls were explicitly omitted from the PNLI and wave action on coastal bluffs was not modeled. We did not consider the cascading impacts from volcanoes, which means the results do not show potentially substantial, if infrequent, lahar hazards from the Cascades. While wildfires are probably not a major destabilizing factor in this region, the omission of burned areas as a predictor might have a confounding effect on model performance in some locations. Perhaps the most important omission was the absence of anthropogenic terrain modification from the list of predictors. Human activities can have a profoundly destabilizing effect, and they have played a major role in landslide activity from local (Chleborad et al., 2008) to global scales (Froude and Petley, 2018). The challenge here was differentiating the true effects of human activity from the apparent effects that are due to inventory bias. Landslides were much more commonly recorded near highways in the PNLI (Kirschbaum et al., 2016), but this probably reflects the chance that a landslide was recorded more than the chance that a landslide occurred. Including any anthropogenic activity index, such as distance to road, would probably have acted as a strong confounding factor on the analysis of hazard. Nevertheless, we acknowledge that it would be desirable to consider anthropogenic slope modification as soon as improved inventories are available.

The incomplete nature of the PNLI implies some uncertainty in the LHI values. The relatively low modeled probabilities associated with historical landslide events might be due not only to the low frequency of historical landslides, but also the uncertainty due to false negative labels in the training data. It also implies that the probabilities of landslide occurrence represent a minimum level of hazard; this is one reason that

the LHI cannot be interpreted as a “hazard map” of the Pacific Northwest. Another reason is that the dominance of a handful of major events (Fig. 4) suggests that risks from the extreme tail of the distribution are very important. This implies that the 20-year inventory may be too short to accurately estimate landslide hazard. For now, models are a preferable way to assess hazard, although the current model does underestimate the impact of two major events.

Although ROS is rare (Fig. 6) in the areas with the most landslide reports (Fig. 2), this does not mean that ROS plays no role in triggering landslides. Since major landslide events are rare (Fig. 4), infrequent ROS could still be a critical factor in slope failure.

Although we have used the PNLI to evaluate the model, a more holistic comparison is often warranted. For example, Fig. 1 shows a surprising lack of dated landslides in the Olympic Peninsula of Washington. Frequent ROS events imply a high landslide rate (Fig. 14), as does the rugged terrain. In fact, numerous landslides have been catalogued in this area (Kirschbaum et al., 2016), but the record rarely included dates. By one interpretation, the absence of human populations in the Olympic Mountains has reduced the number of landslides reported by the news media, a major source of dated records. This seems probable, but the spatial pattern of the LHI suggests that other factors were in play (Fig. 6). Landslide-triggering conditions occurred far less often in the northeast quadrant of the peninsula than along the Pacific coastline. It seems possible that part of the explanation for the gap in Fig. 1 is that landslides in this area are simply infrequent. Another interpretation is that the absence of human activity on the Olympic Peninsula has reduced the number of destabilized slopes relative to more populous regions of the Pacific Northwest. In this hypothesis, much of the reported landslide activity in the PNLI has been affected by anthropogenic influence. Finally, the model explicitly omitted rock falls, which are a major process at higher elevations. If these were included, LHI in rugged mountains like the Olympic Range might have been higher. We cannot resolve the issue without a more detailed investigation, but we are inclined to attribute most of the gap to an inventory bias towards populated areas.

The divergence between model and observations from 1999 to 2005 could be interpreted as a failure of the model to generalize outside of the training period. In this view, the model should have been trained with a randomized train-test data split instead of the division by water years. However, that approach would have introduced numerous false

negatives from years that probably experienced numerous landslides. We attribute the gaps in the inventory prior to 2007 to lack of reporting, rather than any physical phenomenon. If this assumption is correct, the current train-test split should produce a more accurate model by including fewer labeling errors in the training data. Of course, a much larger and more complete inventory of dated landslides would obviate this approach and allow a conventional randomized training data split.

Dividing the data by water year also has the benefit of challenging the model's ability to function well under more extreme circumstances than those in the training data. Splitting the data at 2011 leaves the largest landslide events (Fig. 4) out of the training sample. If these extreme conditions were handled well, that would indicate a robust model that is capable of extrapolation outside of its training data. For the most part, this has been achieved. The model shows temporal peaks in 1996, 2006, 2008, and 2009 that correspond to months with more observed landslides. It also identifies the unseasonably low level of landslide activity in the winter of 2009–2010. It did underestimate the peaks of landslide activity in the winters of 1995–1996 and 2007–2008. However, it succeeded in identifying the month of February 1996 as the most hazardous in the entire time series. The results include good predictions of most out-of-sample major events, which should reinforce confidence in the predictions made for years with no historical landslide data.

XGBoost is a powerful tool for prediction, but it requires careful supervision to avoid spurious associations between landslides and predictors. In the course of this research, a preliminary model accurately predicted every grid cell for the training period 2011–2015. Unfortunately, it was severely overfitted (Fig. 5), so we were forced to revise both the model parameters and input variables. Nevertheless, the preliminary GBM performed well against independent data ($AUC = 0.95$). Our experience suggests that landslide scientists should be very cautious when evaluating the results of empirical modeling, particularly when numerous predictors are available. Furthermore, it demonstrates that the need for common sense in model evaluation is not limited to landslide susceptibility (Steger et al., 2016) but also applies to dynamic landslide models. While we did not attempt to develop an early warning system for the Pacific Northwest, machine learning with multivariate inputs appears likely to advance that field of research as well. Uncertainty and sensitivity analyses of machine-learning models could also help identify what climatic variables are most valuable for understanding landslide behavior; in turn, this can inform decisions about future climate research.

The downwards trend in ROS (Fig. 15a) can be attributed to declines in both snow cover and very heavy (>20 mm) precipitation over the Pacific Northwest (Jasinski et al., 2019). ROS appears to follow the same temporal distribution as landslides (Fig. 2), but its spatial distribution suggests that it accounts for only a small portion of landslide hazard (Fig. 14). Since ROS primarily affects mountainous areas, changes in snowfall have the highest impact on these locations.

When evaluated across the entire record, LHI showed a declining trend over time (Fig. 15b), which paralleled consistent decreasing trends in the underlying hydrologic variables such as extreme rainfall (Jasinski et al., 2019). As with ROS (Fig. 15a), LHI decreased across the Pacific Northwest. Unlike ROS, these declines occurred in all major regions, not just one climate zone.

Although Mann-Kendall analysis showed statistically significant trends in SI (Fig. 15c), many false trend detections would occur through random chance. Combined with the fact that roughly equal numbers of decreasing and increasing trends were observed and the fact that no association between climatic zones and SI trends was obvious (other than the absence of trends along the coast), we suspect that the trends in SI are spurious and do not reflect a major regime shift over the past 4 decades.

Although there has been a decreasing trend in ROS and LHI, it is possible that anthropogenic disturbances have increased both landslide hazard and risk in the states of Washington and Oregon. Population has




Fig. 16. The population of the Pacific Northwest has increased steadily, especially in the most hazardous areas (CIESIN, 2018; CIESIN, 2005).

been growing steadily in the Pacific Northwest for decades (CIESIN, 2018; CIESIN, 2005). This growth has occurred broadly, but the rate of growth is fastest in the most hazardous quarter of the study area (Fig. 16). Although this increase in exposed population does not guarantee an increase in risk (we did not analyze vulnerability), it suggests the need to enforce construction standards that can mitigate or avoid landslide damages and injuries. Fortunately, Washington and Oregon have embraced landslide research and disaster preparedness more generally.

Analyzing the NCA-LDAS dataset produced a multi-decadal overview of landslide hazard across the Pacific Northwest, but this lagging indicator is not without limitations. The LHI emphasizes the geographic or temporal extent of landslide-triggering events, rather than the number, arrangement, speed, or volume of landslides at each location. This emphasis, as well as the coarse resolution, limits its utility for deriving landslide risk, and the LHI probably has little relevance to determining sediment budgets. The LHI is not appropriate for calculating a landslide return period, nor is it a substitute for site-specific investigations prior to any construction project. However, LHI does serve an important purpose by enabling the comparison of landslide hazard across broad spatial and temporal scales. It can also inform the inventory mapping process, by identifying dates and places to look for historical landslides. In the supplement, we have included a listing of grid cells that could form the basis for inventory expansion through historical imagery analysis or searches of media archives. LHI also serves to correct or inform the historical record by providing a view into the conditions at the time of failure. For example, many landslides not considered in this analysis were reported without a known trigger; a high LHI value for that date and location implies that the slide was triggered by one of the predictors from the NCA-LDAS data, and a detailed examination of the model's response to inputs could identify which factors were most important for that event. LHI can probably be extended to projections of future hazard and new regions, although these may require additional predictors.

6. Conclusions

Antecedent conditions help describe landslide behavior not predicted solely by rainfall intensity and duration or recorded in historical inventories. Machine learning is an effective method for incorporating rainfall intensity, atmospheric rivers, antecedent soil moisture, and melting snow from land data assimilation systems into a unified indicator of rainfall-triggered landslide hazard.

Multivariate modeling reveals that landslides follow general geographic and temporal patterns, but with substantial interannual and local variability. West of the Cascade Range, landslide frequency was high; and the season was well defined, with activity centered on the month of January. East of the Cascades, landslide season was not as strongly defined, and the frequency of hazardous conditions was much lower. Although some landslides do occur in summer, this is not the

dominant period for any location in the Pacific Northwest. The LHI serves well as a guide to landslide seasonality because it acts as a consistent measure across time and space that can be easily compared to the frequency of historical reports.

Software accessibility

We used free, open source software to perform this analysis. We have also provided the model dump file to ensure full replicability of this study's results.

XGBoost has been documented thoroughly in an academic paper (Chen and Guestrin, 2016), as well as an online manual at <https://xgboost.readthedocs.io/>. Although XGBoost is written in C++, a wide variety of interfaces are now available. We accessed its functions through scripts written in the Python language on a linux virtual machine with 10 central processing units.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This research was made possible thanks to the contributions of those at NASA Goddard Space Flight Center, USGS Oregon Water Science Center, DOGAMI, ODOT, WADNR and the other federal, state, and local groups responsible for rigorous efforts in comprehensively inventorying landslides in the Pacific Northwest. We would like to specifically acknowledge those individuals whose insights, efforts, and participation greatly improved our research, including Jordan Psaltakis (NASA) for assembling the PNLI, Bill Burns (DOGAMI) for providing insight into historical landslides in Oregon, and Benjamin Mirus (USGS) for providing critical feedback on the aims, modeling approach, and limitations of this research.

This work was supported by the NASA National Climate Assessment Project, grant # NNH14ZDA001N-INCA "Climate Indicators and Data Products for Future National Climate Assessments", and by the NASA Science Mission Directorate Earth Science Division support of the US Global Change Research Program (USGCRP) National Climate Assessment. NCA-LDAS daily data products were obtained from https://disc.gsfc.nasa.gov/datasets/NCALDAS_NOAH0125_D_2.0/summary. NCA-LDAS trends products were obtained from https://disc.gsfc.nasa.gov/datasets/NCALDAS_NOAH0125_Trends_2.0/summary.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.envsoft.2020.104692>.

References

- Baum, R.L., Godt, J.W., 2009. Early warning of rainfall-induced shallow landslides and debris flows in the USA. *Landslides* 7, 259–272. <https://doi.org/10.1007/s10346-009-0177-0>.
- Biasutti, M., Seager, R., Kirschbaum, D.B., 2016. Landslides in West Coast metropolitan areas: the role of extreme weather events. *Weather Clim. Extrem.* 14, 67–79. <https://doi.org/10.1016/j.wace.2016.11.004>.
- Brenning, A., 2005. Spatial prediction models for landslide hazards: review, comparison and evaluation. *Nat. Hazards Earth Syst. Sci.* 5, 853–862.
- Buma, B., Johnson, A.C., 2015. The role of windstorm exposure and yellow cedar decline on landslide susceptibility in southeast Alaskan temperate rainforests. *Geomorphology* 228, 504–511. <https://doi.org/10.1016/j.geomorph.2014.10.014>.
- Burn, D.H., 1997. Catchment similarity for regional flood frequency analysis using seasonality measures. *J. Hydrol.* 202, 212–230. [https://doi.org/10.1016/S0022-1694\(97\)00068-1](https://doi.org/10.1016/S0022-1694(97)00068-1).
- Burns, W.J., 2014. Statewide landslide information database for Oregon, release 3.2. Center for International Earth Science Information Network - CIESIN - Columbia University, 2018. Gridded population of the World, Version 4 (GPWv4): Population Count, Revision 11. <https://doi.org/10.7927/H4JW8BX5>.
- Center for International Earth Science Information Network - CIESIN, Columbia University, United Nations Food and Agriculture Programme - FAO, Centro Internacional de Agricultura Tropical - CIAT, 2005. Gridded Population of the World, Version 3 (GPWv3): Population Count Grid. <https://doi.org/10.7927/H4639MPP>.
- Chakraborty, T., Alam, M.S., Islam, M.D., 2019. Landslide susceptibility mapping using XGBoost model in Chittagong District, Bangladesh. In: *International Conference on Disaster Risk Management*. Dhaka, Bangladesh, pp. 431–434.
- Chen, M., Shi, W., Xie, P., Silva, V.B.S., Kousky, V.E., Wayne Higgins, R., Janowiak, J.E., 2008. Assessing objective techniques for gauge-based analyses of global daily precipitation. *J. Geophys. Res.* 113, D04110. <https://doi.org/10.1029/2007JD009132>.
- Chen, T., Guestrin, C., 2016. XGBoost. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '16*. ACM Press, New York, New York, USA, pp. 785–794. <https://doi.org/10.1145/2939672.2939785>.
- Chen, W., Li, X., Wang, Y., Chen, G., Liu, S., 2014. Forested landslide detection using LiDAR data and the random forest algorithm: a case study of the Three Gorges, China. *Remote Sens. Environ.* 152, 291–301. <https://doi.org/10.1016/j.rse.2014.07.004>.
- Chleborad, A.F., 2000. Preliminary Method for Anticipating the Occurrence of Precipitation-Induced Landslides in Seattle, Washington, Open-File Report 2000-469.
- Chleborad, A.F., 1998. Use of air temperature data to anticipate the onset of snowmelt-season landslides, Open-File Report. <https://doi.org/10.3133/OFR98124>.
- Chleborad, A.F., 1997. Temperature, snowmelt, and the onset of spring season landslides in the central Rocky Mountains, Open-File Report. <https://doi.org/10.3133/OFR9727>.
- Chleborad, A.F., Baum, R.L., Godt, J.W., Powers, P.S., 2008. A prototype system for forecasting landslides in the Seattle, Washington, area. In: Baum, R.L., Godt, J.W., Highland, L.M. (Eds.), *Reviews in Engineering Geology*. Geological Society of America, Boulder, Colorado USA, pp. 103–120. <https://doi.org/10.1130/2008.4020.06>.
- Cordeira, J.M., Stock, J., Dettinger, M.D., Young, A.M., Kalansky, J.F., Ralph, F.M., Cordeira, J.M., Stock, J., Dettinger, M.D., Young, A.M., Kalansky, J.F., Ralph, F.M., 2019. A 142-year climatology of northern California landslides and atmospheric rivers. *Bull. Am. Meteorol. Soc.* <https://doi.org/10.1175/BAMS-D-18-0158.1>.
- Daly, C., Neilson, R.P., Phillips, D.L., 1994. A statistical-topographic model for mapping climatological precipitation over mountainous terrain. *J. Appl. Meteorol.* 33, 140–158. [https://doi.org/10.1175/1520-0450\(1994\)033<0140:ASTMFM>2.0.CO;2](https://doi.org/10.1175/1520-0450(1994)033<0140:ASTMFM>2.0.CO;2).
- Ek, M.B., Mitchell, K.E., Lin, Y., Rogers, E., Grunmann, P., Koren, V., Gayno, G., Tarpley, J.D., 2003. Implementation of Noah land surface model advances in the National Centers for Environmental Prediction operational mesoscale Eta model. *J. Geophys. Res. Atmos.* 108. <https://doi.org/10.1029/2002JD003296>.
- Farahmand, A., AghaKouchak, A., 2013. A satellite-based global landslide model. *Nat. Hazards Earth Syst. Sci.* 13, 1259–1267. <https://doi.org/10.5194/nhess-13-1259-2013>.
- Friedman, J.H., 2002. Stochastic gradient boosting. *Comput. Stat. Data Anal.* 38, 367–378. [https://doi.org/10.1016/S0167-9473\(01\)00065-2](https://doi.org/10.1016/S0167-9473(01)00065-2).
- Friedman, J.H., 2001. Greedy function approximation: a gradient boosting machine. *Ann. Stat.* 29, 1189–1232. <https://doi.org/10.1214/aos/1013203451>.
- Froude, M.J., Petley, D.N., 2018. Global fatal landslide occurrence from 2004 to 2016. *Nat. Hazards Earth Syst. Sci.* 18, 2161–2181. <https://doi.org/10.5194/nhess-18-2161-2018>.
- Gershunov, A., Shulgina, T., Ralph, F.M., Lavers, D.A., Rutz, J.J., 2017. Assessing the climate-scale variability of atmospheric rivers affecting western North America. *Geophys. Res. Lett.* 44, 7900–7908. <https://doi.org/10.1002/2017GL074175>.
- Ghorbanzadeh, O., Blaschke, T., Gholamnia, K., Meena, S., Tiede, D., Aryal, J., 2019. Evaluation of different machine learning methods and deep-learning convolutional neural networks for landslide detection. *Rem. Sens.* 11, 196. <https://doi.org/10.3390/rs11020196>.
- Godt, J.W., Baum, R.L., Chleborad, A.F., 2006. Rainfall characteristics for shallow landsliding in Seattle, Washington, USA. *Earth Surf. Process. Landforms* 31, 97–110. <https://doi.org/10.1002/esp.1237>.
- Higgins, R.W., Shi, W., Yarosh, E., Joyce, R., 2000. Improved United States Precipitation Quality Control System and Analysis [WWW Document]. ATLAS No. 7. http://www.cpc.ncep.noaa.gov/research_papers/ncep_cpc_atlas/7/toc.html. (Accessed 7 September 2019).
- Incer, I., Theodorides, M., Afroz, S., Wagner, D., 2018. Adversarially robust malware detection using monotonic classification. In: *Proceedings of the Fourth ACM International Workshop on Security and Privacy Analytics - IWSPA '18*. ACM Press, New York, New York, USA, pp. 54–63. <https://doi.org/10.1145/3180445.3180449>.
- Istok, J.D., Boersma, L., 1986. Effect of antecedent rainfall on runoff during low-intensity rainfall. *J. Hydrol.* [https://doi.org/10.1016/0022-1694\(86\)90098-3](https://doi.org/10.1016/0022-1694(86)90098-3).
- Jasinski, M.F., Borak, J.S., Kumar, S.V., Mocko, D., Peters-Lidard, C.D., Rodell, M., Rui, H., Beaudoin, H.K., Vollmer, B.E., Arsenault, K.R., Li, B., Bolten, J.D., Tangdamrongsub, N., 2019. NCA-LDAS: overview and analysis of hydrologic trends for the National climate assessment. *J. Hydrometeorol.* <https://doi.org/10.1175/JHM-D-17-0234.1>.

- Jasinski, M.F., Kumar, S.V., Borak, J.S., Mocko, D.M., Peters-Lidard, C.D., Rodell, M., Rui, H., Beaudoin, H.K., Arsenault, K.R., Li, B., Bolten, J.D., 2018. "NCA-LDAS Noah-3.3 Land Surface Model L4 Daily 0.125 X 0.125 Degree V2.0, Daily Data." Goddard Earth Sciences Data and Information Services Center (GES DISC), Greenbelt, Maryland, USA, Version 2.0: [10.5067/7V3N5D004MAS].
- Jia, G., Tang, Q., Xu, X., 2019. Evaluating the performances of satellite-based rainfall data for global rainfall-induced landslide warnings. *Landslides* 17, 283–299. <https://doi.org/10.1007/s10346-019-01277-6>.
- Kendall, M.G., 1948. *Rank Correlation Methods*. Griffin, London.
- Kirschbaum, D.B., Adler, R.F., Hong, Y., Hill, S., Lerner-Lam, A., 2010. A global landslide catalog for hazard applications: method, results, and limitations. *Nat. Hazards* 52, 561–575. <https://doi.org/10.1007/s11069-009-9401-4>.
- Kirschbaum, D.B., Psaltakis, J., Stanley, T.A., 2016. Spatiotemporal properties of landslides in the Pacific Northwest. In: Abstracts with Programs. Geological Society of America, Denver, Colorado, USA. <https://doi.org/10.1130/abs/2016AM-279225>.
- Kirschbaum, D.B., Stanley, T.A., Zhou, Y., 2015. Spatial and temporal analysis of a global landslide catalog. *Geomorphology* 249, 4–15. <https://doi.org/10.1016/j.geomorph.2015.03.016>.
- Korup, O., Stolle, A., 2014. Landslide prediction from machine learning. *Geol. Today* 30, 26–33. <https://doi.org/10.1111/gto.12034>.
- Kumar, S.V., Peters-Lidard, C.D., Tian, Y., Houser, P.R., Geiger, J., Olden, S., Lighty, L., Eastman, J.L., Doty, B., Dirmeyer, P., Adams, J., Mitchell, K., Wood, E.F., Sheffield, J., 2006. Land information system: an interoperable framework for high resolution land surface modeling. *Environ. Model. Software* 21, 1402–1415. <https://doi.org/10.1016/j.envsoft.2005.07.004>.
- Kumar, S.V., Jasinski, M.F., Mocko, D., Rodell, M., Borak, J., Li, B., Kato Beaudoin, H., Peters-Lidard, C.D., Kumar, S.V., Jasinski, M., Mocko, D., Rodell, M., Borak, J., Li, B., Beaudoin, H.K., Peters-Lidard, C.D., 2018. NCA-LDAS land analysis: development and performance of a multisensor, multivariate land data assimilation system for the National Climate Assessment. *J. Hydrometeorol.* <https://doi.org/10.1175/JHM-D-17-0125.1>.
- Laprade, W.T., Tubbs, D.W., 2008. Landslide mapping in Seattle, Washington, in: *landslides and engineering Geology of the Seattle, Washington, area*. *Geol. Soc. Am.* [https://doi.org/10.1130/2008.4020\(02](https://doi.org/10.1130/2008.4020(02).
- Lee, S., Ryu, J., Min, K., Won, J., 2002. Development of two artificial neural network methods for landslide susceptibility analysis. In: *IGARSS 2001. Scanning the Present and Resolving the Future. Proceedings. IEEE 2001 International Geoscience and Remote Sensing Symposium (Cat. No.01CH37217)*. IEEE, pp. 2364–2366. <https://doi.org/10.1109/IGARSS.2001.978003>.
- Lian, C., Zeng, Z., Yao, W., Tang, H., 2014. Extreme learning machine for the displacement prediction of landslide under rainfall and reservoir level. *Stoch. Environ. Res. Risk Assess.* 28, 1957–1972. <https://doi.org/10.1007/s00477-014-0875-6>.
- Lun, N.K., Liew, M.S., Matori, A.N., Zawawi, N.A.W.A., 2017. Recent developments in machine learning applications in landslide susceptibility mapping. In: *AIP Conference Proceedings*, p. 6. <https://doi.org/10.1063/1.5012210>.
- Mann, H.B., 1945. Nonparametric tests against trend. *Econometrica* 13, 245. <https://doi.org/10.2307/1907187>.
- Marjanovic, M., Bajat, B., Kovacevic, M., 2009. Landslide susceptibility assessment with machine learning algorithms. In: *2009 International Conference on Intelligent Networking and Collaborative Systems*. IEEE, pp. 273–278. <https://doi.org/10.1109/INCOS.2009.25>.
- Markham, C.G., 1970. Seasonality of precipitation in the United States. *Ann. Assoc. Am. Geogr.* 60, 593–597. <https://doi.org/10.1111/j.1467-8306.1970.tb00743.x>.
- Marks, D., Kimball, J., Tingey, D., Link, T., 1998. The sensitivity of snowmelt processes to climate conditions and forest cover during rain-on-snow: a case study of the 1996 Pacific Northwest flood. *Hydrol. Process.* 12, 1569–1587. [https://doi.org/10.1002/\(SICI\)1099-1085\(199808/09\)12:10<1569::AID-HYP682>3.0.CO;2-L](https://doi.org/10.1002/(SICI)1099-1085(199808/09)12:10<1569::AID-HYP682>3.0.CO;2-L).
- Mirus, B.B., Becker, R.E., Baum, R.L., Smith, J.B., 2018a. Integrating real-time subsurface hydrologic monitoring with empirical rainfall thresholds to improve landslide early warning. *Landslides* 15, 1909–1919. <https://doi.org/10.1007/s10346-018-0995-z>.
- Mirus, B.B., Morphew, M., Smith, J., Mirus, B.B., Morphew, M.D., Smith, J.B., 2018b. Developing hydro-meteorological thresholds for shallow landslide initiation and early warning. *Water* 10, 1274. <https://doi.org/10.3390/w10091274>.
- Neiman, P.J., Schick, L.J., Ralph, F.M., Hughes, M., Wick, G.A., Neiman, P.J., Schick, L. J., Ralph, F.M., Hughes, M., Wick, G.A., 2011. Flooding in western Washington: the connection to atmospheric rivers. *J. Hydrometeorol.* 12, 1337–1358. <https://doi.org/10.1175/2011JHM1358.1>.
- Ogle, S., Fish, M.A., Wilson, A.M., Oakley, N., Slaughter, S., Ralph, M., 2018. *Atmospheric River Families and Their Relationship to Landslides in Washington State*. AGU, Washington, DC.
- Petschko, H., Brenning, A., Bell, R., Goetz, J., Glade, T., 2014. Assessing the quality of landslide susceptibility maps – case study Lower Austria. *Nat. Hazards Earth Syst. Sci.* 14, 95–118. <https://doi.org/10.5194/nhess-14-95-2014>.
- Read, L.K., Vogel, R.M., 2015. Reliability, return periods, and risk under nonstationarity. *Water Resour. Res.* 51, 6381–6398. <https://doi.org/10.1002/2015WR017089>.
- Sarikhan, Isabelle, Thompson, Elizabeth, Heinitz, A., 2009. Quick Report: Landslide Reconnaissance during the January 7, 2009, Storm. Olympia, WA.
- Sarikhan, I.Y., Stanton, K.D., Contreras, T.A., Polenz, M., Powell, J., Walsh, T.J., Logan, R.L., 2008. Landslide Reconnaissance Following the Storm Event of December 1–3, 2007, in Western Washington.
- Sayers, P., 2016. *Communicating the Chance of a Flood: the Use and Abuse of Probability, Frequency and Return Period*.
- Schevel, C.R., Baum, R.L., Mirus, B.B., Smith, J.B., 2017. *Precipitation Thresholds for Landslide Occurrence Near Seattle, Mukilteo, and Everett, Washington*. U.S. Geological Survey Open-File, Reston, VA. <https://doi.org/10.3133/ofr20171039>.
- Song, Y., Niu, R., Xu, S., Ye, R., Peng, L., Guo, T., Li, S., Chen, T., 2018. Landslide susceptibility mapping based on weighted gradient boosting decision tree in Wanzhou section of the three Gorges reservoir area (China). *ISPRS Int. J. Geo-Inf.* 8, 4. <https://doi.org/10.3390/ijgi8010004>.
- Stanley, T.A., Kirschbaum, D.B., 2017. Effects of inventory bias on landslide susceptibility calculations. In: De Graff, J.V., Shakoar, A. (Eds.), *Landslides: Putting Experience, Knowledge and Emerging Technologies into Practice*, AEG Special Publication No. 27. Association of Environmental and Engineering Geologists, pp. 794–806.
- Stanley, T., Kirschbaum, D., Sobieszczyk, S., Borak, J., Slaughter, S., 2019. A landslide climate indicator from machine learning. In: *2019 Fall AGU Meeting*. AGU, Washington, DC, United States.
- Stanley, T., Kirschbaum, D., Sobieszczyk, S., Jasinski, M.F., Borak, J., Yatheendradas, S., 2017. Assessment of rainfall thresholds for landslide triggering in the Pacific Northwest: extreme short-term rainfall and long-term trends. In: *American Geophysical Union, Fall Meeting 2017*.
- Steger, S., Brenning, A., Bell, R., Glade, T., 2017. The influence of systematically incomplete shallow landslide inventories on statistical susceptibility models and suggestions for improvements. *Landslides* 1–15. <https://doi.org/10.1007/s10346-017-0820-0>.
- Steger, S., Brenning, A., Bell, R., Petschko, H., Glade, T., 2016. Exploring discrepancies between quantitative validation results and the geomorphic plausibility of statistical landslide susceptibility maps. *Geomorphology* 262, 8–23. <https://doi.org/10.1016/j.geomorph.2016.03.015>.
- Stumpf, A., Kerle, N., 2011. Object-oriented mapping of landslides using Random Forests. *Remote Sens. Environ.* 115, 2564–2577. <https://doi.org/10.1016/j.rse.2011.05.013>.
- Theil, H., 1992. A Rank-Invariant Method of Linear and Polynomial Regression Analysis. Springer, Dordrecht, pp. 345–381. https://doi.org/10.1007/978-94-011-2546-8_20.
- Tubbs, D.W., 1974. *Landslides in Seattle*. Olympia.
- U.S. Environmental Protection Agency, 2010. *Level I ecoregions of North America*. Environ. Protect.
- U.S. Geological Survey, 2016. The National map [WWW Document]. In: *3DEP Prod. Serv. Natl. Map, 3D Elev. Progr.*. Web page. http://nationalmap.gov/3DEP/3d_ep_prodserv.html. (Accessed 9 September 2016).
- Van Den Eeckhaut, M., Reichenbach, P., Guzzetti, F., Rossi, M., Poesen, J., 2009. Combined landslide inventory and susceptibility assessment based on different mapping units: an example from the Flemish Ardennes, Belgium. *Nat. Hazards Earth Syst. Sci.* 9, 507–521. <https://doi.org/10.5194/nhess-9-507-2009>.
- Vogel, R., Hecht, J., Read, L., 2014. *Nonstationary Approaches to Hydrologic Design*. Geophysical Research Abstracts.
- Washington Division of Geology and Earth Resources, 2016. *Landslides and Landforms—GIS Data*, February, 2016, Washington Division of Geology and Earth Resources Digital Data Series 12, Version 4.1, Previously Released August 2015.
- Washington DNR, 2019. *Recent Landslides in Washington State* [WWW Document]. Google My Maps. https://www.google.com/maps/d/viewer?mid=1VwV_u07xWo nt7Sb3xRlChBb0PYc&ll=46.707969399301014%2C-119.56221039062495&z=5. (Accessed 6 May 2019).
- xgboost developers, 2016. Notes on Parameter Tuning — Xgboost 0.83.Dev0 Documentation [WWW Document]. https://xgboost.readthedocs.io/en/latest/tutorials/param_tuning.html. (Accessed 19 April 2019).
- Xia, Y., Mitchell, K., Ek, M., Sheffield, J., Cosgrove, B., Wood, E., Luo, L., Alonge, C., Wei, H., Meng, J., Livneh, B., Lettenmaier, D., Koren, V., Duan, Q., Mo, K., Fan, Y., Mocko, D., 2012. Continental-scale water and energy flux analysis and validation for the North American Land Data Assimilation System project phase 2 (NLDAS-2): 1. Intercomparison and application of model products. *J. Geophys. Res. Atmos.* 117. <https://doi.org/10.1029/2011JD016048> n/a-n/a.
- Zhao, J., Liu, Y., Hu, M., 2018. Optimisation algorithm for decision trees and the prediction of horizon displacement of landslides monitoring. *J. Eng.* 1698–1703. <https://doi.org/10.1049/joe.2018.8305>, 2018.