Final Report: West Virginia Statewide Sinkhole Mapping Project

Funder: United States Department of Agriculture (USDA) Natural Resource Conservation Service (NRCS)
Performance Period: October 1, 2023 to September 30, 2025
NRCE Federal Award Identification Number: NR233D47XXXXC005

Overview

The goal of this project was to generate a statewide dataset of sinkhole features within karst regions of West Virginia. To take into account errors in the karst geology boundaries, as defined by the West Virginia Geological and Economic Survey (WVGES) 1:250,000 scale map data, all areas within 0.5 miles (0.80 km) of the karst formations were investigated. The mapped areas are shown in Figure 1 below. We specifically focused on mapping the smallest depressions as opposed to the larger depressions that contain these smaller features. An example of the product is shown in Figure 2.

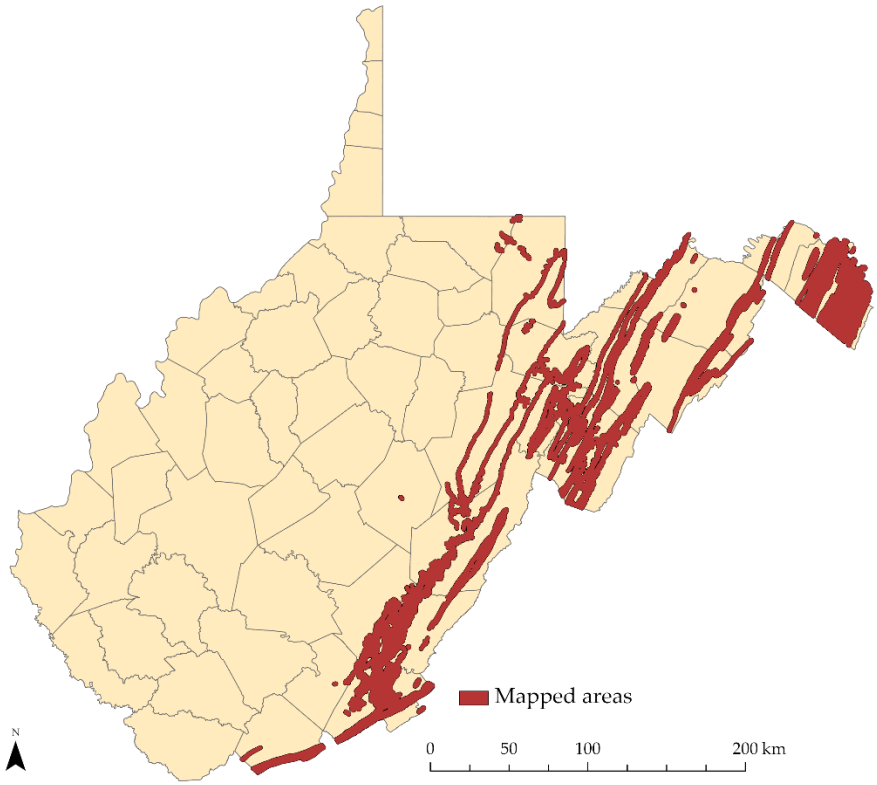

Figure 1. Mapped extent (all areas within 0.5 miles of karst geology in West Virginia).

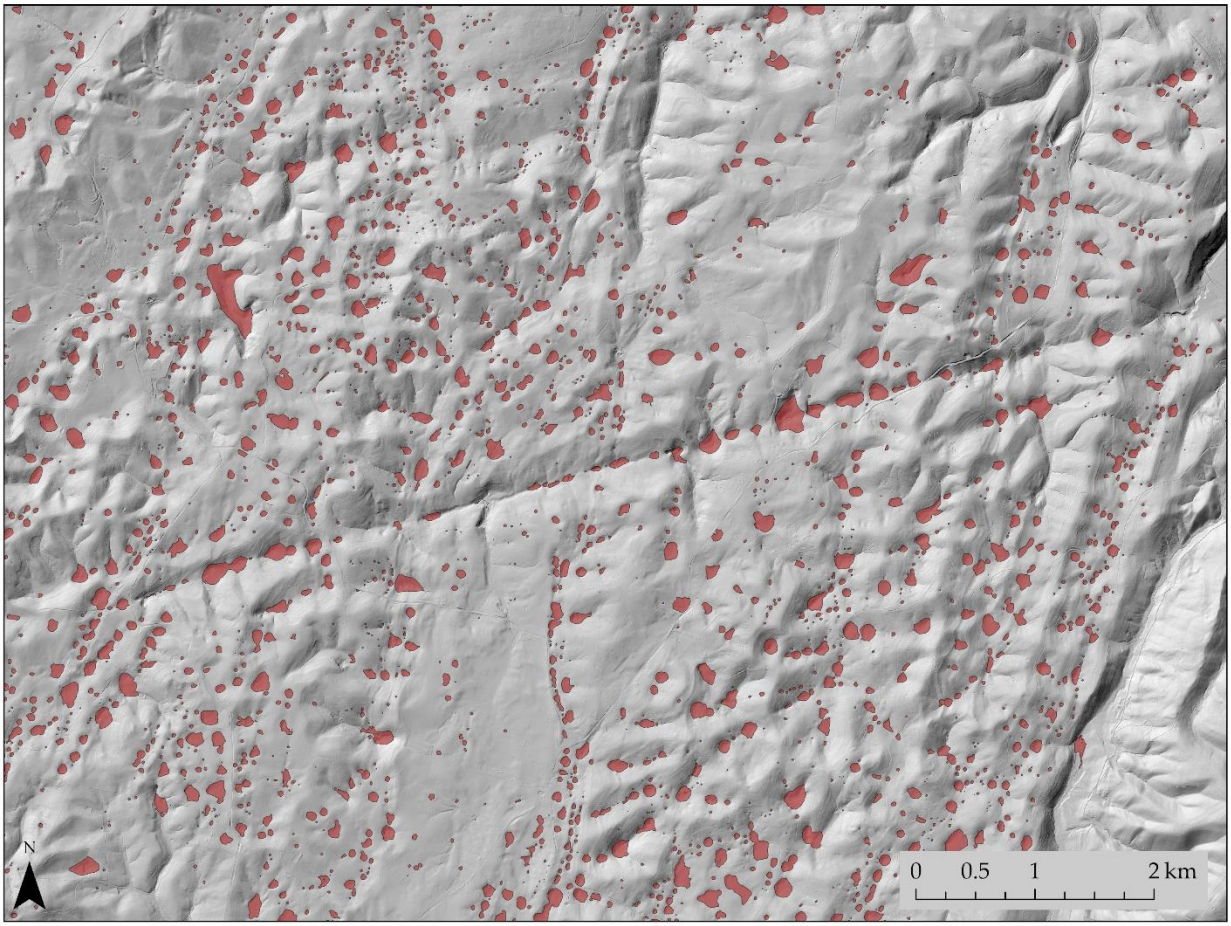

Figure 2. Example of mapped sinkholes.

The mapping team is named below. The PIs for the project were Drs. Dorothy Vesper and Aaron Maxwell. One graduate research assistant (GRA) was also funded: Marissa Loftus. All other students were funded hourly. One salaried analyst at the Werst Virginia GIS Technical Center (WVGISTC), Shannon Maynard, also participated.

- Principal Investigators
    - Dorothy Vesper (Professor, West Virginia University Department of Geology and Geography)
    - Aaron Maxwell (Associate Professor/West Virginia GIS Technical Center Faculty Director, West Virginia University Department of Geology and Geography)
- Mappers
    - Marissa Loftus (graduate research assistant funded on project; lead mapper)
    - Charlie Alavi (undergraduate student)
    - William Beadling (MS student)
    - Conner Channels (undergraduate student)
    - Gabe Jones (undergraduate student)
    - Rowan Keifer (undergraduate student)
    - Oliver Kelly (undergraduate student)

- Shannon Maynard (GIS analyst at WVGISTC)
- Kennedy Ochieng (PhD student)
- Sam Stockton (MA student)

Data Layers

All generated data layers are provided as both standalone shapefiles and feature classes within a file geodatabase (*wvSinkholes.gdb*). All layers are projected into the NAD83 UTM Zone 17N projection with horizontal units in meters.

Primary Deliverable

1. *sinkholesWVGISTC_WVDEP*: final edited sinkhole datasets including data generated as part of this project and a subset of those generated by the WVDEP. Only WVDEP features that do not overlap with the WVGISTC features were included. The original WVDEP data are available at https://tagis.dep.wv.gov/home/gis%20data. The associated attributes are described in a later section.
    - Features from deep learning: 37,072
    - Features from WVGISTC mappers: 27,788
    - Features ingested from WVDEP dataset: 21,131
    - **Total**: 85,991
2. *sinkholesWVGISTC_WVDEP_lowest_point*: final edited sinkholes dataset represented as point features where the point coordinate corresponds to the center of the elevation cell with the lowest elevation of all of the cells that intersect the sinkhole object.
3. *kd500.tif*: kernel density of sinkhole points (from *sinkholesWVGISTC_WVDEP_lowest_point*) calculated using kernel density with a 500 m radius.
4. *kd500_area.tif*: kernel density of sinkhole points (from *sinkholesWVGISTC_WVDEP_lowest_point*) calculated using kernel density with a 500 m radius using estimated sinkhole area as a weighting.
5. *kd500_vol.tif*: kernel density of sinkhole points (from *sinkholesWVGISTC_WVDEP_lowest_point*) calculated using kernel density with a 500 m radius using estimated sinkhole volume as a weighting.

Ancillary Products

1. *deepLearningOuput*: unedited output from the deep learning process (n = 48,022)
2. *checkGrid*: grid used to organize checking process (total of 1,458 tiles; tileID = unique ID)
3. *sinkholesFromWVGISTC*: edited sinkhole dataset not including those from the WVDEP project (n = 64,939)

Input Data

1. *mlra*: NRCS major land resource areas (MLRAs) (https://www.nrcs.usda.gov/resources/data-and-reports/major-land-resource-area-mlra)
2. *huc12Watersheds*: hydrologic unit code (HUC 12) watershed boundaries

3. gelogy250k: 1:250,000 scale geology data (https://wvgis.wvu.edu/data/dataset.php?ID=43)
4. *geology250kKarst*: just karst formations/groups from 1:250,00 scale geology data
5. *ecoregions*: EPA ecoregions (https://www.epa.gov/eco-research/ecoregions)
6. *counties*: county boundaries
7. *soilUnitsSSURGO*: SSURGO soil unit boundaries

Mapping Process

Deep learning

Sinkhole detection is generally undertaken using one of the following processes:

1. Field surveys
2. Manual mapping based on interpretation of digital elevation data and aerial orthophotography via heads-up, on-screen digitizing
3. Pit filling/detection algorithms with subsequent automated and manual clean-up to remove false detections
4. Supervised learning from labelled data

For this project, we chose to use supervised learning followed by manual clean-up via interpretation of digital elevation data and aerial orthophotography. Initial results using the supervised learning and pit filling methods suggested that both techniques would require manual clean-up; however, it was gauged that less manual clean-up would be required for the supervised learning-based output.

A variety of supervised learning architectures are available. We chose to make use of the UNet deep learning architecture (Ronneberger et al., 2015), which is based on convolutional neural networks (CNNs) and allows for capturing spatial patterns at different spatial resolutions. It is conceptualized in Figure 3. The purpose of the encoder component, which uses a combination of convolutional layers, batch normalization, activation functions, and max pooling, is to learn spatial patterns at different spatial scales. The decoder, which uses convolutional layers, batch normalization, activation functions, and transpose convolution, then restores the data to its original spatial resolution to allow for cell-level predictions (Ronneberger et al., 2015).

The algorithm was implemented using the geodl package (Maxwell et al., 2024) in the R language (R Core Team, 2022). This package was built by PI Maxwell and his research group and supports geospatial semantic segmentation (i.e., pixel-level classification) in R using the torch (Falbel and Luraschi, 2023), terra (Hijmans, 2023), and luz (Falbel, 2023) packages.

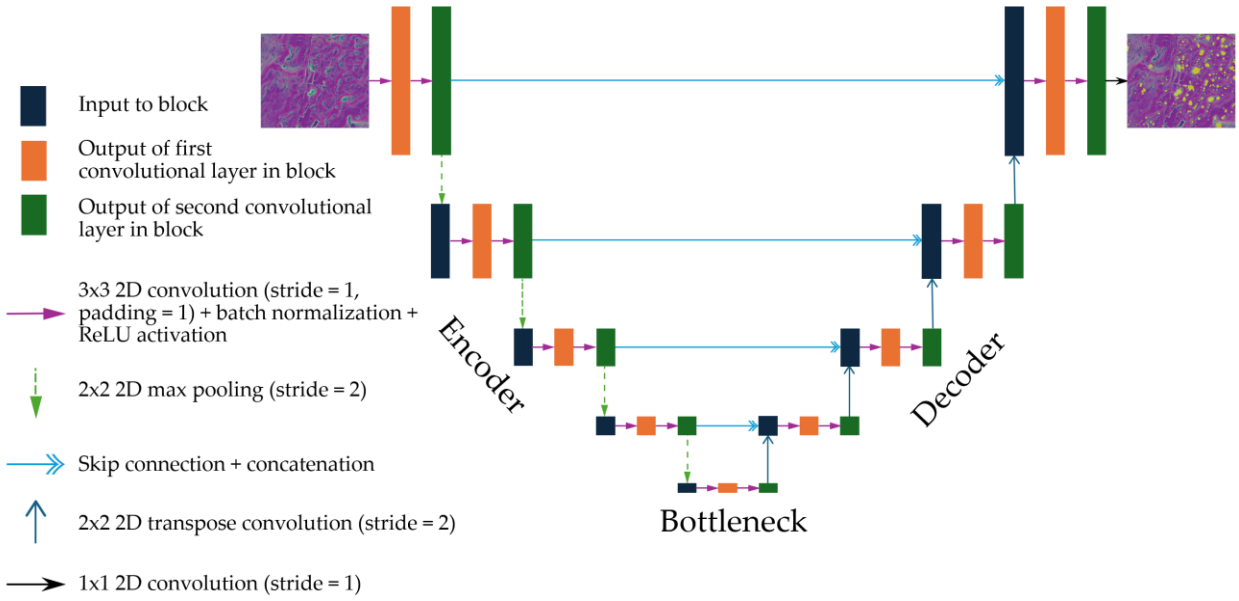

Figure 3. Conceptualization of UNet deep learning semantic segmentation architecture.

Training data

Supervised learning algorithms require input predictor variables and associated labels. As a result, and in order to implement this method, we had to obtain example sinkhole features. We made use of sinkhole data provided by the United States Geological Survey (USGS).

- Monroe County, West Virginia: https://pubs.usgs.gov/publication/sir20235121/full
- Winchester 1:100,000 scale topographic quadrangle: https://www.usgs.gov/maps/sinkholes-and-karst-related-features-shenandoah-valley-winchester-30-x-60-quadrangle-virginia

These data are pictured in Figure 4 as point features due to the scale of the maps. However, each point feature represents a sinkhole polygon. The Monroe County data contained 7,810 features while the Winchester data contained 5,883 features. Both datasets were generated and provided by Dr. Dan Doctor's group at the USGS Florence Bascom Geoscience Center.

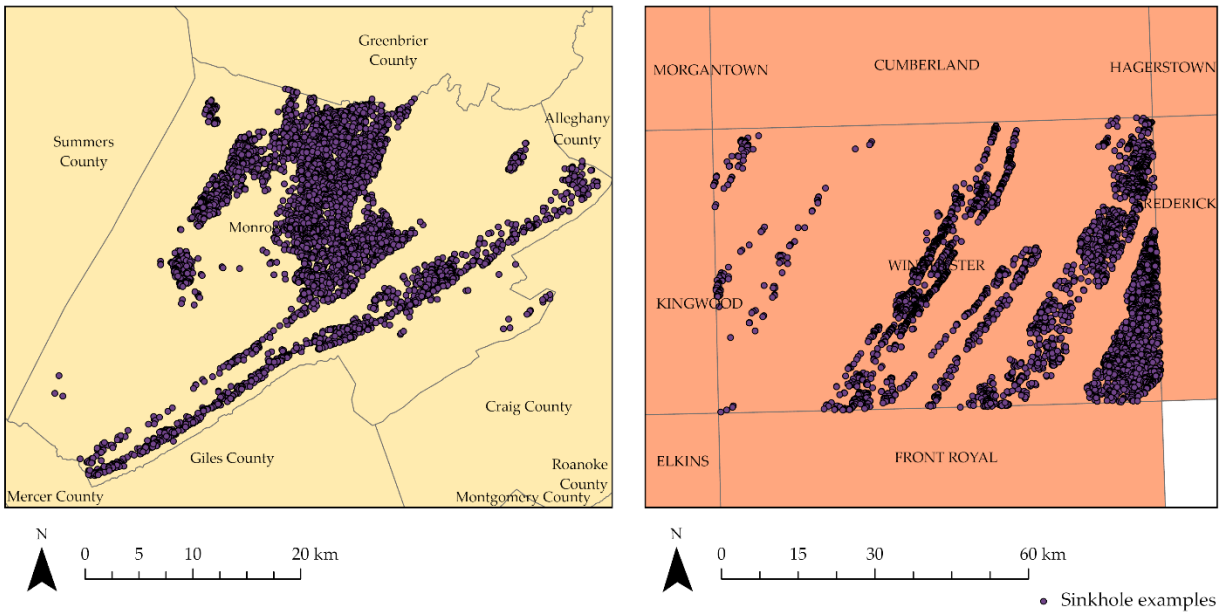

Figure 4. Training data from USGS mapping projects. left = Monroe County, West Virginia; right = Winchester 1:100,000 scale quadrangle.

Predictor variables

For predictor variables, we used the following three land surface parameters (LSPs) derived from a digital terrain model (DTM). These surfaces are visualized in Figure 5 as separate grayscale images and also as a three-layer composite. They were derived from light detection and ranging (lidar)-based DTMs provided by the USGS 3D Elevation Program (3DEP) (Sugarbaker et al., 2014).

1. Hillslope-scale topographic position index (TPI) with a 50 m circular radius
2. Topographic slope in degrees
3. Local TPI with a 2 m inner and 5 m outer radius annulus window

To generate the TPIs, the mean was calculated within the focal windows then subtracted from the center cell value to obtain an index in which larger, positive values indicate topographic high points and negative values indicate topographic low points. The TPI produced using a larger, circular window captured patterns at a hillslope spatial scale while the annulus-based TPI captured more local patterns and surface roughness. For both TPIs, values smaller than −10 were recoded to −10 while values larger than 10 were recoded to 10. The data were then rescaled from 0 to 1 by subtracting −10 then dividing by 20. For topographic slope, we first calculated the square root of slope. Values smaller than 0 were recoded to 0, values larger than 10 were recoded to 10, and all values were then linearly rescaled from 0 to 1 then subtracted from 1 such that flatter areas had higher values and steeper areas had lower values.

A wide variety of LSPs can be calculated from DTM data (Franklin, 2020; Maxwell and Shobe, 2022). The choice of LSP in this study was guided by our prior experiences extracting

geomorphic features from digital terrain data. The specific combination used here was suggested by Drs. Dan Doctor and Will Odom of the USGS. In a prior study, we documented that these variables supported a more accurate deep learning-based classification than more commonly used terrain visualization surfaces including slopeshades, hillshades, and multidirectional hillshades (Maxwell et al., 2023).

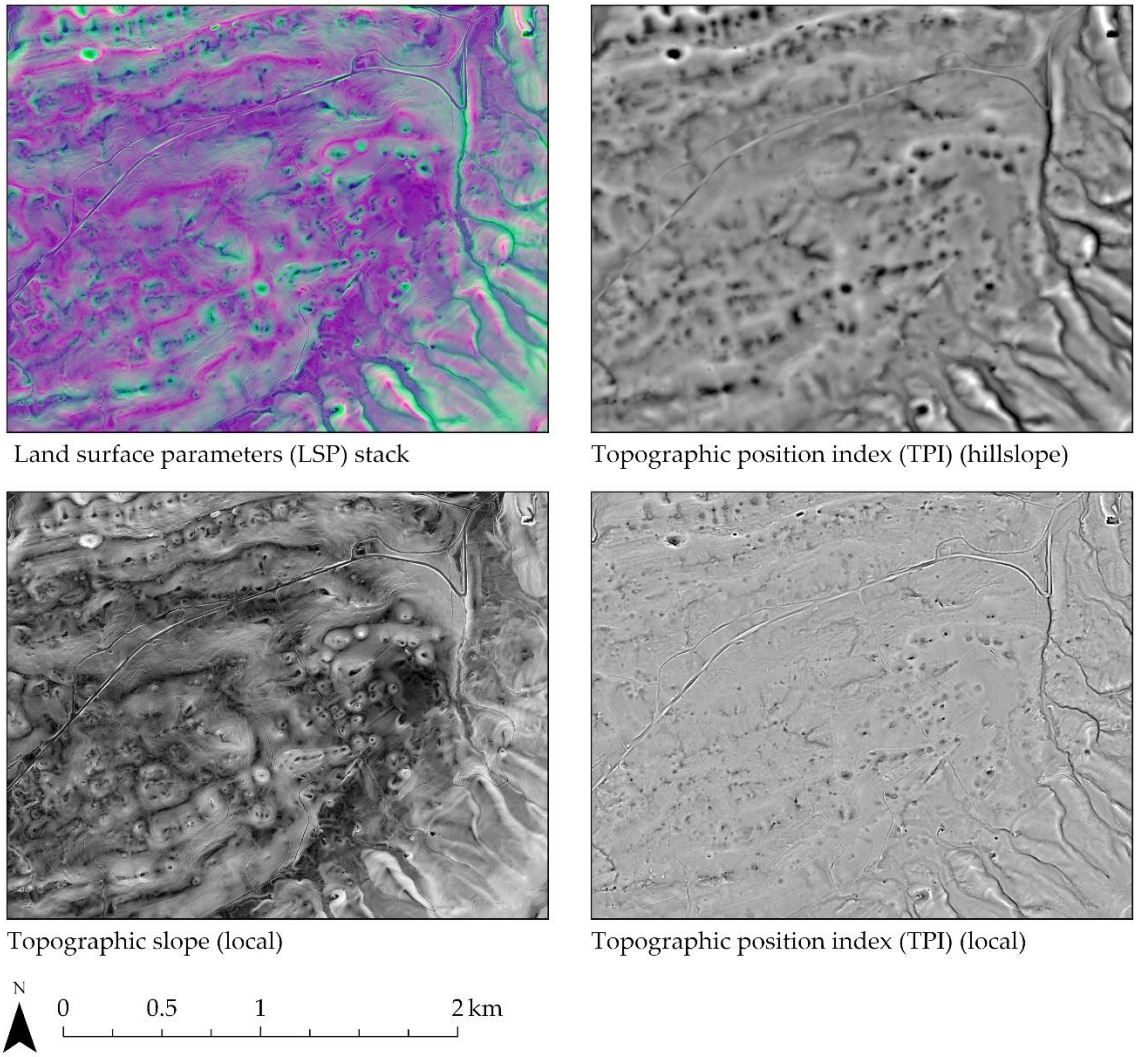

Figure 5. Input predictor topographic/LSP predictor variables.

Once the example data were processed and the predictor variables were generated, the available data were partitioned into image chips and associated labels. A subset of the chips was used to train the model while the remaining chips were used to validate the model performance during the training process. Due to class imbalance issues, we used a focal Tversky loss (Abraham and Khan, 2019). During the training process, random vertical and horizontal flips were applied in an attempt to combat overfitting. We trained the model for 25 epochs and selected the model state that provided the highest F1-score for the validation set as the final model.

Once the model was trained, we generated the input LSPs for all karst areas in the state then used the model to make predictions for all areas within the extents shown in Figure 1 above. The raster-based classification output was converted to polygons, merged to a single dataset, and post-processing smoothing was applied. This output is provided as the *deepLearningOuput* file.

Manual clean-up

Once the deep learning output was obtained, manual editing was performed to improve the mapping product. This was based on manual interpretation and heads-up digitizing. We made use of the following layers. These layers are shown in Figure 6 along with the final map product. All of these layers are provided as web services by the WVGISTC.

1. 1 m spatial resolution lidar-derived hillshade
2. 2 m spatial resolution three-layer terrain visualization (i.e., the input LSP predictor variables used in the deep learning model)
3. Best available, leaf-off, high spatial resolution orthophotography (WV aerial photo mixed resolution service)

We developed a tiling scheme of 1,458 4-by-4 km tiles, which has been provided as the *checkGrid* file. Each tile was fully scanned and checked by one mapper then subsequently validated by a second mapper so that all areas were interpreted by two individuals. The following general edits were made:

1. Delete false positive samples (i.e., features that were interpreted to be incorrectly predicted as sinkholes by the deep learning algorithm)
2. Add missed or false negative features (i.e., manually observed sinkholes not detected by the deep learning algorithm)
3. Edit features that were generated by the deep learning algorithm (i.e., features that were detected, but the mapper felt that the boundary was incorrect or could be improved via manual digitizing)

We also made use of a dichotomous key that was provided by the team that completed Virginia's sinkhole mapping project. This has been provided as Figure 7.

During the mapping process, we held meetings every week to every other week with a focus on allowing the mappers to share their observations and get input on uncertain features or issues.

The manual editing was the most time-consuming part of the process.

The manually checked and edited product is provided as *sinkholesFromWVGISTC*.

WVDEP data integration

While we were completing our mapping project, the WVDEP completed a separate sinkhole mapping for Greenbrier, Pocahontas, Monroe, Jefferson, and Berkely counties. We integrated features from their dataset that did not overlap with our features in an attempt to reduce false negative errors. This final, merged product has been provided as *sinkholesWVGISTC_WVDEP*.

This is the primary deliverable for the project. In the final dataset, a total of 85,991 sinkhole features were mapped. Of these features, 37,072 were generated by the deep learning algorithm, 27,788 were manually digitized by our mappers, and 21,131 were derived from the WVDEP data.

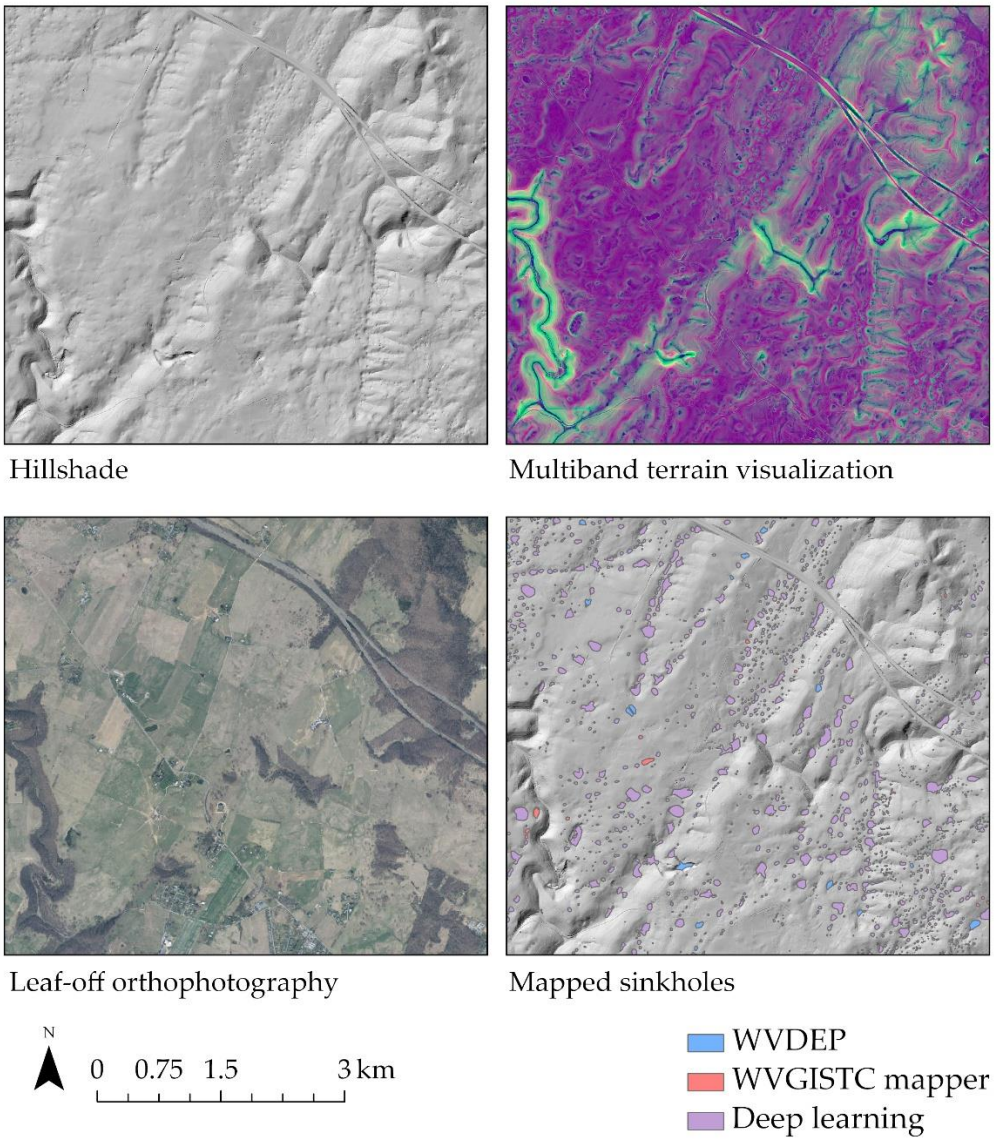

Figure 6. Reference data used during manual clean-up and checking along with example mapped sinkholes.

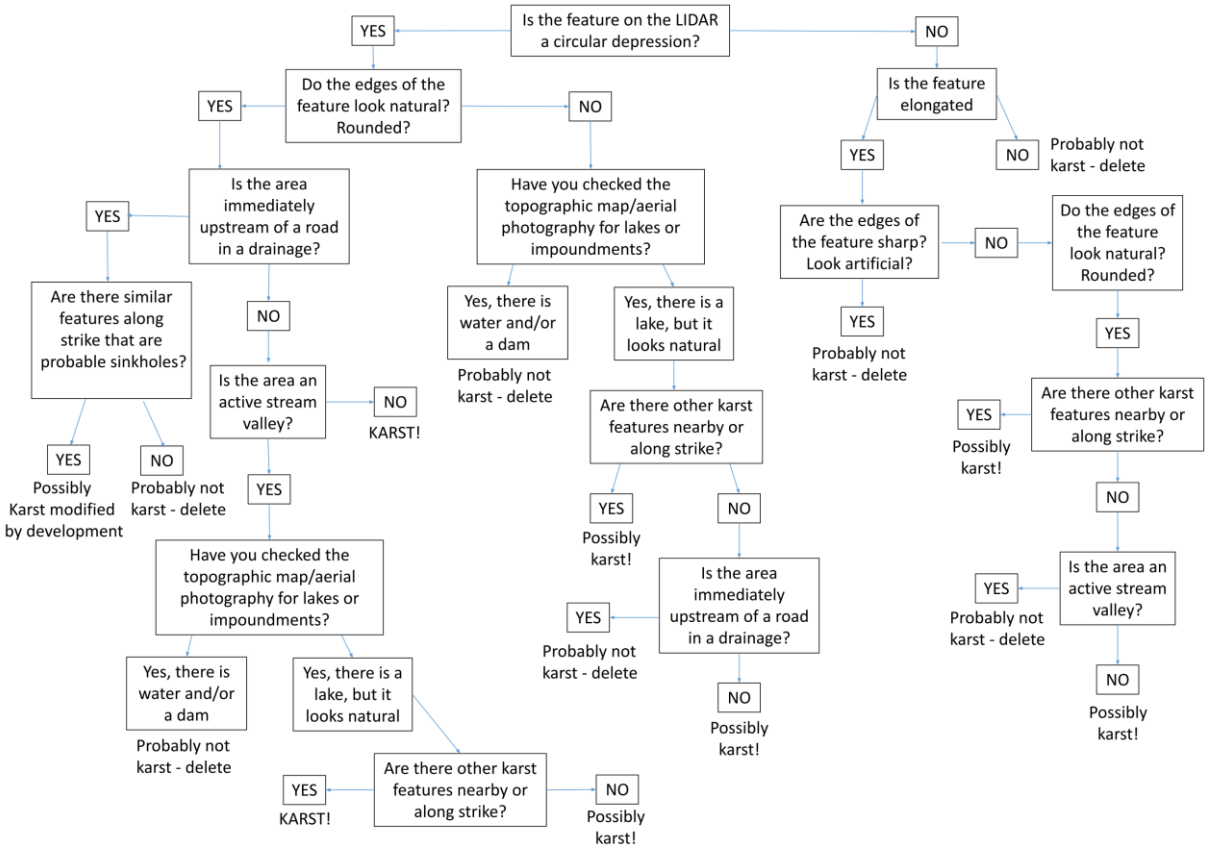

Figure 7. Dichotomous key for manual interpretation provided by the Virginia mapping group.

Post-processing

Once the final set of features were generated, we performed post-processing to generate attributes and summary information for each feature.

Using the Calculate Geometry tool in ArcGIS Pro, we calculated the total area in square meters and total perimeter length in meters for each feature. We used the Zonal Statistics as Table tool to obtain the mean elevation of the sinkhole. The following process was used to estimate the maximum depth, mean depth, and total volume of each feature.

1. Areas within sinkholes were erased from the DTM data so that only cells outside of sinkholes contained elevation measurements while cells in sinkholes were coded to NODATA.
2. The gaps in the DTM were then filled using a nearest neighbor method, effectively replacing missing cells with the elevation of the closest, non-sinkhole cell.
3. The original DTM was then subtracted from the modified DTM.
4. The Zonal Statistics as Table tool was used to extract summary information from the differenced grid within each sinkhole feature. The maximum difference was recorded as the maximum depth while the total volume was calculated as the sum of the depths of all the cells within the sinkhole multiplied by the cell size.

We also noted the geologic formation or group in which the feature occurred using the 1:250,000 scale geology for the state (https://wvgis.wvu.edu/data/dataset.php?ID=43). Due to potential errors in the boundaries, we also extracted the nearest karst formation. This was accomplished using the Nearest tool.

The WVGES is currently working on generating 1:24,000 scale bedrock geology data via the STATEMAP project. This project is not yet complete. We extracted the unit name from this new dataset where data were available. Where data were not yet available, "Missing" was recorded for those sinkholes. We were not permitted to share these geologic boundaries as part of the project deliverable.

We also noted the county; major land resource area (MLRA); EPA ecoregion; and hydrologic unit code (HUC) 8, 10, and 12 watershed in which each feature occurred using a simple spatial join.

Lastly, polygon shape metrics were generated using the Patch Shape Tools from WhiteboxTools (https://www.whiteboxgeo.com/).

Attributes

The final set of feature attributes are summarized below.

Unique ID and source

1. *sinkID*: unique ID for sink feature
2. *source*:
    - "deep learning": output from deep learning algorithm maintained after manual checks
    - "WVGISTC mapper": data manually created by mapping team
    - "WVDEP": data generated from the WVDEP project (not created by this project)

Geometry, depth, and volume

3. *area*: total area estimated using Calculate Geometry in ArcGIS Pro (square meters)
4. *perimeter*: total perimeter length estimated using Calculate Geometry in ArcGIS Pro (meters)
5. *meanElev*: mean elevation of the cells occurring within the sinkhole (meters)
6. *maxDepth*: maximum depth of the sinkhole estimated using the methods described above (meters)
7. *meanDepth*: mean depth of the sinkhole estimated using the methods described above (meters)
8. *totVol*: total volume of sinkhole estimated using the methods described above (cubic meters)

Path shape

*The following attributes were calculated using Patch Shape tools from WhiteboxTools. Documentation is available here: (https://www.whiteboxgeo.com/manual/wbt_book/available_tools/gis_analysis_patch_shape_tools.html#gis-analysis--patch-shape-tools)

9. *COMPACT*: compactness ratio = ratio of polygon area to its perimeter
10. *ELONGATION*: elongation ratio ($E = 1 - S/L$), where S is the short-axis length and L is the long-axis length; axes lengths are determined by estimating the minimum bounding box
11. *ORIENT*: patch orientation of polygon feature based on the slope of a reduced major axis (RMA) regression line; the regression analysis uses the vertices of the exterior hull nodes of a vector polygon; orientation values are measured in degrees from north
12. *P_A*_RATIO: perimeter area ratio = ratio of polygon perimeter length to its area
13. *LINEARITY*: linearity index = the coefficient of determination (R-squared) calculated from a regression analysis of the x and y coordinates of the exterior hull nodes of a vector polygon; measure of how well a polygon can be described by a straight line
14. *COMPLEXITY*: shape complexity index ($1 - A/Ah$), where A is the polygon's area and Ah is the area of the convex hull containing the polygon
15. *RC_CIRCLE*: related circumscribing circle: ($RCC = 1 - A/Ac$), where A is the polygon's area and Ac is the area of the smallest circumscribing circle for the polygon

Ponds and compound sinks

16. *pond*:
    - "yes": object contained standing water in reference aerial orthophotography
    - "no": object did not contain standing water in reference aerial orthophotography
17. *inComp*:
    - "Yes": object occurred within one of the composite features mapped by the WVDEP; used to differentiate compound sinks in this mapping project
    - "No": object did not occur within one of the composite features mapped by the WVDEP

Location information

18. *huc12*: HUC 12 watershed in which feature occurs
19. *huc10*: HUC 10 watershed in which feature occurs
20. *huc8*: HUC 8 watershed in which feature occurs
21. *county*: county in which feature occurs
22. *state*: state in which feature occurs
23. *MLRA_NAME*: major land resource area (MLRA) in which feature occurs (https://www.nrcs.usda.gov/resources/data-and-reports/major-land-resource-area-mlra)
24. *LRR_NAME*: EPA land resource region in which feature occurs (https://www.epa.gov/eco-research/ecoregions)
25. *NA_L3NAME:* EPA level 3 ecoregion in which feature occurs (https://www.epa.gov/eco-research/ecoregions)

26. *NA_L2NAME:* EPA level 2 ecoregion in which feature occurs (https://www.epa.gov/eco-research/ecoregions)
27. *NA_L1NAME:* EPA level 1 ecoregion in which feature occurs (https://www.epa.gov/eco-research/ecoregions)

Geology (STATEMAP) (https://www.wvgs.wvnet.edu/www/statemap/statemap.htm)

28. *wvgesMU*: map unit code from WVGES STATEMAP project
29. *wvgesNAME*: unit or formation name from WVGES STATEMAP project

Geology (WVGES) (https://wvgis.wvu.edu/data/dataset.php?ID=43)

30. *ERA*: geologic era
31. *PERIOD*: primary geologic period
32. *PERIOD2*: secondary geologic period (if applicable)
33. *CLASS*: rock class
34. *TYPE*: rock type
35. *group*: geologic group name
36. *FORMATION*: geologic formation name
37. *Mapunit*: map unit code
38. *karstDist*: distance to nearest karst formation
39. *kClass*: class of nearest karst formation
40. *kType*: rock type of nearest karst formation
41. *kFormation*: name of nearest karst formation
42. *kGroup*: name of nearest karst group
43. *kPeriod*: period of nearest karst formation
44. *kmapunit*: karst unit code

Mapping Metadata

45. *mapper*: first name of person that mapped/verified the tile in which the object occurred
46. *checker*: first name of person that verified the mappers work
47. *generalC*: general comments from WVGISTC team
48. *mapperC*: comments from WVGISTC mapper
49. *checkerC*: comments from WVGISTC checker
50. *wvdepC*: comments from WVDEP for objects they generated

Lowest Point

51. *min_elev*: elevation in meters for the cell that intersects the sinkhole extent with the loest elevation
52. *x*: x-coordinate of cell that intersects the sinkhole extent with the lowest elevation relative to NAD83 UTM Zone 17N
53. *y*: y-coordinate of cell that intersects the sinkhole extent with the lowest elevation relative to NAD83 UTM Zone 17N

SSURGO Soil

The following attributes were extracted from the SSURGO unit boundaries using the point location of the lowest elevation in the sinkhole extent

54. *AREASYMBOL*
55. *SPATIALVER*
56. *MUSYM*
57. *MUKEY*

Sources of error

It is important to note that no map product is perfect. Our goal was to generate the best possible product given the complexity of the task and practical constraints. Field validation has proven challenging due to the difficulty of interpreting features in the field and access/time constraints. We are currently discussing field validation options with the WVGES. Our goal is to have their field mappers assist in validating the product.

Here, we want to highlight some of the sources of error that were noted by our mappers during the manual interpretation process. We have conceptualized sources of false positive (commission) and false negative (omission) errors in Table 1 while Figure 8 provides visual examples of some sources of error.

Our mappers noted difficulty in distinguishing sinkholes from tree falls in forested areas where falls created depressions in the DTM. Detections were generally removed if they were well outside of the karst formation extents or were not near other sinkholes. Artificial depressions could often be detected using the available aerial orthophotography. These depressions were often the result of road construction or infrastructure development. We also noted some false positives in areas with rough or bouldery topography since the surface textures in these areas, as represented in the LSPs, could be similar to sinkhole-prone areas. The geologic unit boundaries and aerial orthophotography were helpful in detecting these incorrect detections.

In regards to false negatives, small sinkholes or those that were wide and/or had shallow margins were difficult to detect using both automated methods and manual interpretation. We suspect that some sinkholes were partially covered, such as by colluvium or alluvium, and were not detectable based on surface morphology. Some features were somewhat abnormal in shape, such as those that were modified by fluvial processes or humans.

In the Greenbrier Valley specifically, it was common for sinkholes to be nested or compound. In this study, we focused on mapping the smallest features. The WVDEP project mapped larger, composite features, and we noted which features in our merged dataset occurred within these composite features in the “inComp” field. Composite features were generally less pronounced in other areas of the state.

The ”pond” attribute notes whether standing water was observed within the full extent or a subset of the extent of a sinkhole within the aerial orthophotography. This labelling was conducted for all features, including those manually digitized by our mapping group, generated by the deep learning algorithm, or provided by the WVDEP. It was common for ponds to be detected as sinkholes by the deep learning algorithm. These features were deleted if

there was strong evidence that the feature was not a sinkhole. For example, features were removed if there was a visible impoundment or dam. Ponds that were not near sinkholes were also generally removed. In short, ponds were maintained if there was no clear sign of an impoundment and they occurred within areas where sinkholes were present.

Table 1. Examples of false positive (FP) and false negative (FN) errors.

		Reference	
		Sinkhole	**Not**
Prediction	**Sinkhole**	True positive	False positive: tree falls, ponds, artificial depressions rough topography, boulders
	Not	False negative: small features, ponds, abnormal shaped sinkholes, fluvial transformed, human modified	True negative

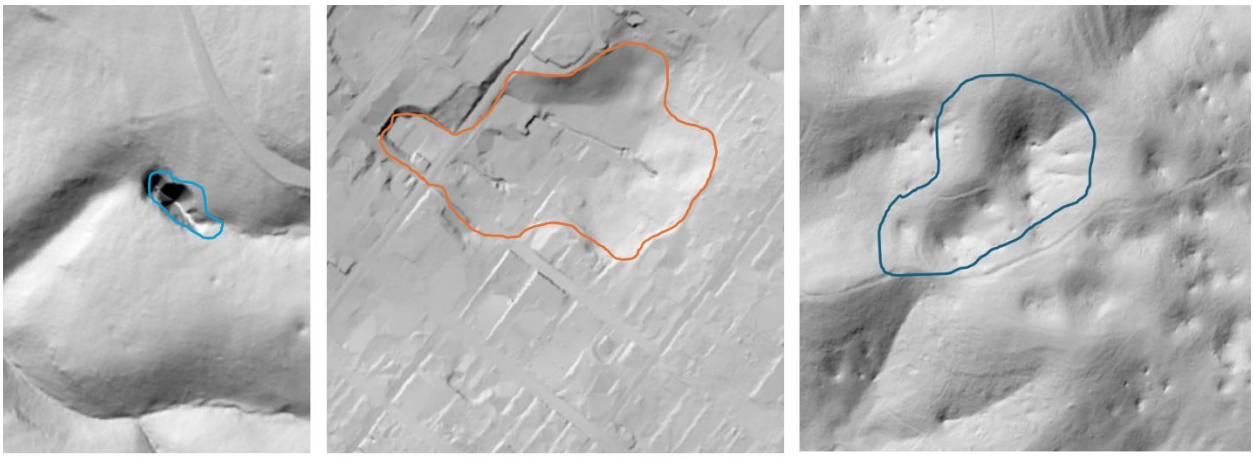

Figure 8. Example sources of error.

References

Abraham, N., Khan, N.M., 2019. A novel focal tversky loss function with improved attention u-net for lesion segmentation, in: 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019). IEEE, pp. 683–687.

Falbel, D., 2023. luz: Higher Level "API" for "torch."

Falbel, D., Luraschi, J., 2023. torch: Tensors and Neural Networks with "GPU" Acceleration.

Franklin, S.E., 2020. Interpretation and use of geomorphometry in remote sensing: a guide and review of integrated applications. International Journal of Remote Sensing 41, 7700–7733.

Hijmans, R.J., 2023. terra: Spatial Data Analysis.

Maxwell, A.E., Farhadpour, S., Das, S., Yang, Y., 2024. geodl: An R package for geospatial deep learning semantic segmentation using torch and terra. PloS one 19, e0315127.

Maxwell, A.E., Odom, W.E., Shobe, C.M., Doctor, D.H., Bester, M.S., Ore, T., 2023. Exploring the Influence of Input Feature Space on CNN-Based Geomorphic Feature Extraction From Digital Terrain Data. Earth and Space Science 10, e2023EA002845.

Maxwell, A.E., Shobe, C.M., 2022. Land-surface parameters for spatial predictive mapping and modeling. Earth-Science Reviews 226, 103944.

R Core Team, 2022. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria.

Ronneberger, O., Fischer, P., Brox, T., 2015. U-net: Convolutional networks for biomedical image segmentation, in: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18. Springer, pp. 234–241.

Sugarbaker, L.J., Constance, E.W., Heidemann, H.K., Jason, A.L., Lukas, V., Saghy, D.L., Stoker, J.M., 2014. The 3D Elevation Program initiative: a call for action. US Geological Survey.