When Computers Dream of Charcoal

Using Deep Learning, Open Tools, and Open Data to Identify Relict Charcoal Hearths in and around State Game Lands in Pennsylvania

Benjamin P. Carter ⦿*, Jeff H. Blackadar* ⦿*, and Weston L. A. Conner* ⦿

ABSTRACT

This research employs machine learning (Mask Region-Based Convolutional Neural Networks [Mask R-CNN]) and cluster analysis (Density-based spatial clustering of applications with noise [DBSCAN]) to identify more than 20,000 relict charcoal hearths (RCHs) organized in large "fields" within and around State Game Lands (SGLs) in Pennsylvania. This research has two important threads that we hope will advance the archaeological study of landscapes. The first is the significant historical impact of charcoal production, a poorly understood industry of the late eighteenth to early twentieth century, on the historic and present landscape of the United States. Although this research focuses on charcoal production in Pennsylvania, it has broad application for both identifying and contextualizing historical charcoal production throughout the world and for better understanding modern charcoal production. The second thread is the use of open data, open source, and open access tools to conduct this analysis, as well as the open publication of the resultant data. Not only does this research demonstrate the significance of open access tools and data but the open publication of our code as well as our data allow others to replicate our work, to tweak our code and protocols for their own work, and reuse our results.

Keywords: relict charcoal hearths, landscapes, Mask R-CNN, deep learning, Pennsylvania, light detection and ranging (lidar), airborne laser scanning (ALS)

Esta investigación emplea el aprendizaje automatizado (Redes Neuronales Convolucionales basadas en Regiones "Máscara" [Mask R-CNN; en sus siglas en inglés]) y el análisis de agrupamientos o clústers (Agrupamiento Espacial Basado en Densidad de Aplicaciones con Ruido [DBSCAN; en sus siglas en inglés]), para identificar más de 20,000 áreas de combustión de hornos de producción de carbón (RCHs; en sus siglas inglés), dispuestos en "campos" amplios dentro y alrededor de Campos de Caza Estatales (SGLs; en sus siglas inglés), en Pensilvania. Esta investigación tiene dos importantes desafíos que esperamos que desarrollará el estudio de los paisajes en arqueología. El primero es el impacto histórico significativo de la producción de carbón, una industria poco entendida de la época temprana del S. XVIII e inicios del S. XIX, del paisaje histórico y actual de Estados Unidos. No obstante, esta investigación se centra alrededor de la producción de carbón en Pensilvania, tiene una aplicación amplia para la identificación y contextualización de la producción de carbón histórica en todo el mundo y para lograr un mejor entendimiento de la producción moderna de carbón. El segundo desafío es el uso de las herramientas de datos libres, fuentes libres y accesos libres para llevar a cabo este análisis, así como la publicación libre del dato resultante. Esta investigación no solamente demuestra el significado de las herramientas y los datos libres, sino que además la publicación libre de nuestro código, así como nuestros datos, permitirá a otros replicar nuestro trabajo, refinar nuestro código y protocolos para su propia investigación, así como reusar nuestros resultados.

Palabras Clave: áreas de combustión de hornos de producción de carbón, paisajes, Mask R-CNN, aprendizaje profundo, Pensilvania, detección de imágenes y distancias mediante láser (lidar), escaneo láser aerotransportado

To some, charcoal may seem a dead or insignificant technology largely limited to backyard barbeques, but charcoal, as the primary fuel for iron production in the eighteenth and nineteenth centuries, powered the Industrial Revolution and the economic rise of the United States (e.g., Bonhage et al. 2020; Johnson and Ouimet 2018, 2021; Kemper 1941; Knowles 2012; Rolando 1992; Straka 2014; Straka and Ramer 2010; Williams 2020). Yet, the significance of charcoal production to history is not geographically limited to the United States (e.g., Bekele and Girmay 2014; Carrari

et al. 2017; Raab et al. 2019; Schneider et al. 2015; Werbrouck et al. 2011). Nor is charcoal production solely an industry of the past—it continues around the world today (e.g., Alfaro and Jones 2018; Kato et al. 2005; Rembold et al. 2013; Rodrigues and Braghini Junior 2019). The Food and Agricultural Organization estimates that approximately 53 million tonnes (58 US tons) of charcoal were produced worldwide in both 2018 and 2019 (FAOSTAT 2021). Historical research on the charcoal industry helps to contextualize modern processes and may provide options that no longer exist but remain viable. Similarly, this research can aid our understanding of how past practices have impacted modern landscapes (e.g., Deforce et al. 2020; Donovan et al. 2021; Mikan and Abrams 1995, 1996; Schneider et al. 2019; Tolksdorf et al. 2015). The following are the questions at the heart of this research:

(1) How was charcoal produced?
(2) What was the impact of charcoal production on the landscape?
(3) How do we recognize that impact?
(4) How do we better understand the social aspects of charcoal (and iron) production?

The answers to Question 1 are relatively well understood, and the Pennsylvania variant is described below. Although Question 2 is the proximate impetus behind this research, at this point, we primarily address Question 3. Our research focuses on identification of relict charcoal hearths (or RCHs) and ignores other components (such as transportation). RCHs are the relatively small (ca. 10–15 m in diameter) features that remain from the process of making charcoal that are distributed (ca. 80–150 m from each other) across the landscape. Similarly, through this research and that of others, we are slowly improving our ability to address the social aspects of charcoal. Because charcoal is produced deep in the woods, well away from society, it is often associated with strangeness, wildness, and difference. Though not addressed here, this is one of the central reasons why charcoal production played an important role in the Underground Railroad in the mid-1800s (LaRoche 2017).

Given these questions, the primary goal of this research is to recognize the impact of charcoal production. First, we discuss ways that the impact of charcoal production has been considered. Second, we discuss the process of charcoal production and how a single instance of charcoal production leaves behind identifiable remains on the landscape. Then, we use lidar (light detection and ranging, also known as airborne laser scanning, or ALS) data openly provided by the state of Pennsylvania. Lidar is particularly useful because this remote sensing technique allows us to model the ground surface, especially in densely forested areas, across broad swaths of the landscape that are difficult to comprehend otherwise. Then, we utilize an array of open source tools that enable us to use a type of artificial intelligence known as "deep learning" to identify RCHs across the state of Pennsylvania.

Deep learning is a type of machine learning that employs user-provided identification (in this case, user-recognized RCHs in digital images) to create a generalized model. It then uses this abstraction to identify arrangements of pixels similar to the training data. Deep learning is particularly valuable in this situation because we are able to use it to identify additional RCHs that are generally like those in the training set but not precisely the same as any individual RCH in the training set. This also means that with the appropriate computer power, we are able to identify RCHs

across large extents of area (in this case, 28% of the state of Pennsylvania).

In order to carry out this research, we employ open data, open tools, and open publishing. The original data are supplied by the state of Pennsylvania. All of the tools used to process the data are open, including QGIS (with plugins noted below), GeoPandas, GDAL, Mask R-CNN, Keras, and TensorFlow. All code constructed to use these tools has been published via Zenodo, as has the resultant data.[1]

The project employed a computer with the following specifications: GPU: TITAN V / 5120 CUDA Cores / 12G; vCPU: 6; RAM: 24GB; Disk Space: 64GB; and O/S: Ubuntu 18.04. Although consumer-grade laptops and desktops have less power, training a deep learning model is possible but would take longer. A cloud-based service that offers GPU processing, such as Google Colaboratory, is another option. Although the free version has some limitations, the paid Pro version ($10/month) does not. With the setup described above, processing time varied by the number of epochs a model was trained for. The processing time for an individual epoch, with 663 training images, was approximately 3 minutes. Most models were trained in the range of 16–40 epochs, which equates to around 1–2 hours. These were often left to run overnight, hence the "dreaming" computers of the title.

ESTIMATING THE IMPACT OF CHARCOAL PRODUCTION

It is easy to realize that an industry that converted large numbers of trees into charcoal had a significant impact on the landscape. Yet, the full impact of charcoal production on the landscape is difficult to determine. It is clear that the amount of charcoal consumed varied by iron furnaces and forges, and their impact on forests and landscapes varied (Straka 2014). Because little historical data was recorded specifically on charcoal production, iron production is often used as a proxy. This is reasonable because the primary—though not sole—purpose of charcoal was as fuel for smelting and forging iron. Care must be exercised because in the 1830s, some iron furnaces began using coal, so we must be able to differentiate between coal- (and coke-)fired furnaces and charcoal furnaces (e.g., Knowles 2012:Figure 10; Williams 1989: Figure 10.5). However, even with the initial use of mineral coal in iron furnaces, charcoal continued to be used to fuel iron production—especially for certain components such as railcar wheels—into the early twentieth century. Although the overall production of charcoal furnaces varied, their contribution to the total iron production of the nation declined (Schallenberg 1975, 1981; Williams 1989:Figure 10.5). The last charcoal furnace "went out of blast" (i.e., was shut down) in 1945 (Schallenberg 1975:341).

We have a fairly good accounting for iron production (e.g., American Iron and Steel Association 1876, 1878, 1880, 1882, 1884, 1886, 1890, 1892, 1894, 1896; Lesley 1866), but how much charcoal was needed per unit (i.e., ton) of iron? Straka (2014) finds that estimates for the amount of charcoal needed per ton of iron range greatly, from approximately 25 bushels to 200 bushels (see also Whitney 1994:Table 9.4). This broad range is due to wide variation in a range of variables including wood type (e.g., pine vs. oak), ore composition, technology used (e.g., cold blast vs. hot blast),

expertise of the iron master and the collier (the "charcoal burner"), the size of a "bushel" (Birkinbine 1879:150; Straka 2017:60), and even the location of the furnace. Additionally, the size and efficiency of charcoal furnaces changed through time. Schallenberg and Ault (1977:445; Williams 1989:339) report that charcoal consumption at "advanced" iron furnaces from the Midwest in the late 1800s ranged from around 73 to 114 bushels of charcoal per ton of iron (Hillstrom and Hillstrom 2006:165–166; Whitney 1994: Table 9.4). These increasingly efficient furnaces were significantly larger and often located in the Midwest (especially Michigan) and the South, replacing the older, smaller, and less efficient furnaces of the East (especially Ohio and Pennsylvania; Schallenberg and Ault 1977; Williams 1989:339). Given this variation, using the amount of iron produced to determine the amount of charcoal required for its production would provide an estimate with a very wide range. Trying to understand the impact of that charcoal production on the landscape is even more difficult.

Assuming that we can estimate the amount of charcoal needed to produce a ton of iron, converting the number of bushels of charcoal used to acres of land is made difficult by a wide range of factors, such as the size/age of the trees, the expertise of the collier, the water content of the wood, and the species of the tree (Straka 2014; Svedelius 1875). Gordon (1996:40) draws out an example of a calculation: a furnace that consumes approximately 356,000 bushels of charcoal annually requires approximately 600 acres of forest per year. Therefore, assuming reuse after 20 years, the furnace requires roughly 12,000 acres of land to be sustainable (see also Whitney 1994:222, Table 9.4). To understand the scale of forest depletion across the United States, Birkinbine (1879:150) calculated that the production and forging of iron in 1877 resulted in the cutting of approximately 147,000 acres of woodland and, in 1879 (Birkinbine 1881:66), approximately 50,000–65,000 acres. Yet, it is unlikely that the production decreased threefold in two years. Williams (1990:163) estimates that between 1855 and 1910, 3,000–4,800 square miles (ca. 2–3 million acres) of forests were converted into charcoal to feed iron furnaces, but he also reminds us that this is only about 0.8%–1.3% of the forest cleared at that time—most clearance was due to agricultural expansion (Whitney 1994:222). But because of the variables described above, these are all rough and divergent estimates that are difficult to apply.

The impact of charcoal production was also not equal across landscapes but focused on specific types. Until 1870, the majority of charcoal used in a charcoal iron furnace was from woodlands within 3–8 km (2–5 miles) of the furnace (Whitney 1994:224). This suggests that the preferred location of iron furnaces would be on landscapes with approximately 3–8 km (2–5 miles) of forest that were less than desirable for agriculture—what Bernhard Fernow (1883:153–155; Rodgers 1991:24–25) called "waste places." The use of these "waste places" would be especially true for landscapes where the majority of arable land was already occupied. Throughout much of Pennsylvania, charcoal production centered on dry, nonarable lands such as rocky mountains (see Johnson and Ouimet 2021). Fernow observed that the Lehigh Furnace lands on the Blue Mountain, which he managed from 1879 to 1887, remained unsuitable for agriculture. This is not necessarily true for some of the earliest furnaces—such as Henry Furnace in Nazareth, Pennsylvania (Jones 1976)—where, when it was in operation, the landscape included limited agriculture along with charcoal lands. Today, most of the charcoal lands around the Henry Furnace have likely been converted into farmland and residential developments.

All in all, researchers have had limited success using iron production to quantify and clarify the impact of charcoal production on landscapes. The research reported in this article takes a different approach. We turn to the archaeological remains of charcoal production. In particular, we focus on "relict charcoal hearths" (RCHs), which represent past charcoal production. But to understand the impact of charcoal production on the landscape, the reader must understand how charcoal is produced.

CHARCOAL PRODUCTION

Charcoal production is a relatively simple technology that requires dirty, physical labor as well as a detailed knowledge of a wide range of variables that can be only partially controlled (Diderot and d'Alembert 2010 [1765]; Kemper 1941; Overman 1854; Straka 2014; Straka and Ramer 2010; Svedelius 1875). Until the late nineteenth century, the majority of charcoal was made by placing hearths close to the wood. Wood was cut during the winter when it contained less sap, allowing it to dry faster and more completely and when woodcutters, who were normally off-season farmers and colliers, were available. The wood was stacked in 1.2 m (4 ft.) lengths near where it was cut—often on the downhill side of the lot.

Then, a collier built a hearth near the wood. Hearths were flat, level circular areas cleared of stones and roots approximately 10–15 m in diameter (Kemper 1941:8; Straka 2014:106; Walker 1966:242). Hearths were flat, smooth, compacted surfaces to minimize moisture, ensure even burning throughout the mound, and avoid collecting soil and other contaminants when raking out the charcoal (Straka 2014:106). On slopes, this often meant digging into the hillside and pitching the dirt downhill, which not only created the round, flat hearth but also resulted in steep banks on the uphill and downhill sides of the hearth (Figure 1a) (e.g., Donovan et al. 2021).

On the hearths, the collier created a mound of wood stacked in a very specific manner known as a "meiler." First, the collier built a triangular chimney composed of overlapping lengths of wood with a central opening of approximately 20 cm (8 in.) in the center of the hearth. Three layers of densely packed wood was stacked vertically, leaning inward against and around the chimney and extending out to the edges of the hearth (Kemper 1941:10). The result was a large, dome-shaped pile of tightly packed wood, which was covered with a layer of leaves, hay, or pine needles. This was subsequently covered with a layer of earth (Figure 1b) and, ideally, charcoal dust from a previous use (Kemper 1941; Straka 2014; Straka and Ramer 2010; Svedelius 1875; Walker 1966).

Once the meiler was constructed, it was fired. The goal, however, was not to burn the wood but to distill it by separating the solid carbon (as charcoal) from the water, sap, and volatile compounds (Overman 1854:102–117). The resultant charcoal is relatively free of contaminants, is solid, and burns hot, making it ideal for iron furnaces. To distill the charcoal, the collier first ignites the pile through the chimney, but once the fire catches, it is also closed off. To keep the pile hot enough (between 270°C [518°F] and 400°C [752°F]; Straka 2017:59), the collier opens and closes air holes near the base of the meiler. Although the meiler must be monitored and managed constantly for the first 24–48 hours, after that, the collier checks on it periodically over the next 10–14 days. Colliers frequently made charcoal eight months out of the year (with the
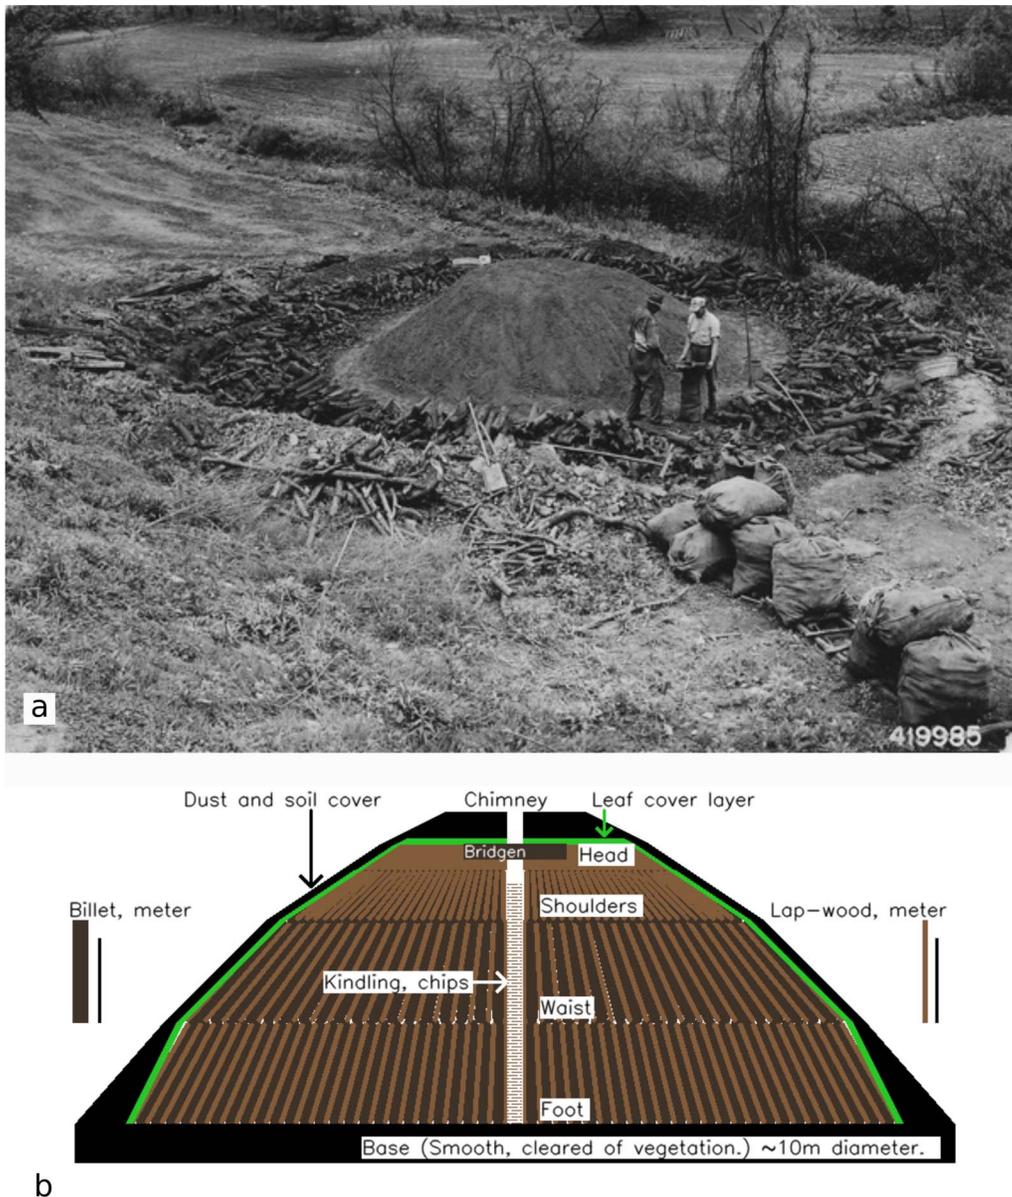
FIGURE 1. (a) Two colliers removing charcoal from the meiler. The underlying flat, circular hearth can be seen excavated into the slope; Wayne National Forest, Ohio, May 1942 (National Archives, Record Group 95: Records of the Forest Service, 1870–2008, National Archives Identifier: 2129419; Local Identifier: 419985). (b) A simplified cross section of a meiler ready to be ignited (image by Jeff Blackadar).

exception of the winter), with up to eight meilers in process at any given time (Kemper 1941; Straka 2014; Straka and Ramer 2010; Svedelius 1875; Walker 1966).

When the wood in the meiler was converted into charcoal, the charcoal was separated from the earth and debris using rakes with large wooden tines (Figure 1a). This is done carefully because the charcoal can easily catch fire, potentially destroying all of the charcoal. The earthen covering often ends up as a ring of earth around the edge of the hearth. Charcoal was then loaded onto specialized carts and transported to the furnace or forge (Kemper 1941; Straka 2014; Straka and Ramer 2010; Svedelius 1875; Walker 1966).

Visible remains from the process described above are present on the modern landscape. Primary among these are RCHs, but others include the remains of roads and colliers' huts. Importantly, RCHs have been recognized using derivatives of high-resolution lidar scanning (e.g., Bonhage et al. 2020, 2021; Carter 2019a; Donovan et al. 2021; Hirsch et al. 2017; Johnson and Ouimet 2018, 2021; Kazimi et al. 2019; Raab et al. 2015; Risbøl et al. 2013; Schmidt

et al. 2016; Schneider et al. 2015). To simplify, lidar is a technology that uses plane-mounted lasers to measure the altitude of the earth's surface and create a point cloud of the surface (for a detailed description, see Fernandez-Diaz et al. 2014; Opitz 2013; see also Opitz and Hermann [2018] for a review of remote sensing in archaeology). In Pennsylvania, the resolution of this scan is approximately one laser shot per 1.5 m^2. The data can then be converted to a digital elevation model (DEM), which is similar to a digital photograph in that each pixel, instead of being associated with numbers that represent colors, is the altitude of that pixel. Two of the most common ways a DEM can be visualized is as a hill-shade (or shaded relief) and as a slope analysis. From these, RCHs can be manually recognized (e.g., Carter 2019a, 2019b; Raab et al. 2015).

A DEEP-LEARNING APPROACH TO RCH IDENTIFICATION

Our approach here employs these lidar derivatives on a large scale. The rich detail across a large region (ca. 37,000 km^2) makes it difficult, if not impossible, to identify these RCHs manually. Visually scanning images for examples of RCHs is labor intensive (but see Bonhage et al. 2021). Due to advances in deep learning, an automated scan of images is more feasible for large land-scapes. Although our purpose is to assess charcoal production across Pennsylvania, rather than including the entire state of Pennsylvania in our sample, we focused our attention on State Game Lands (SGLs). The SGLs are a proxy for "waste lands" discussed above. It was clear that, like SGL 217 (Carter 2019a, 2019b; Conner 2018), many of these lands were previously used for charcoal (and lumber), but when they became unproductive, they were acquired by the state. Additionally, because much of the area outside SGLs are developed, RCHs are difficult to identify. We included a 1 km buffer around all SGLs to both catch some undeveloped land not owned by the state and demonstrate whether or not we could identify RCHs in residential areas.

Deep learning is a type of artificial intelligence that can recognize similar patterns directly from data, such as images. Generally speaking, deep learning recognizes patterns using multiple processes that abstract and amplify common attributes in a set of images (LeCun et al. 2015). These multiple processes are known as "layers," and they represent the depth of deep learning. An assembly of layers represents a Convolutional Neural Network. These are also known as ConvNets or CNNs, and they are computer representations of patterns.

As each layer of representation amplifies and abstracts aspects of an image, it builds a generalized model to recognize other similar patterns. With the composition of enough layers and repetitions of training, this computer model can learn to recognize complex patterns and objects. This is important, given that we want to generally recognize the features of RCH that vary due to topography or size, not just a particular RCH.

We structured our sample around available lidar data provided by Pennsylvania and the organization of that data. We used the vector file of SGLs provided by the state and added a 1 km buffer. We identified overlapping lidar data using the tile index vector. Because the tile index contained the download link for the lidar

data within its attributes, this allowed us to automate downloading and organizing the data by SGL. The result was one file folder per SGL that contained all lidar tiles. Over 4,000 LAS files (the native format for lidar)—a total of about 350 gigabytes of data in zipped files—was downloaded and unzipped. Because there was some overlap in the tiles included between SGLs, 3,925 of these tiles were unique, representing 28.4% of the entire state (Figure 2).

This data was organized and converted into the desired format. The goal was to produce a Digital Elevation Model (DEM), a hill-shade, and a slope analysis for each SGL (all in TIFF format; see Carter [2019a] for a description on how to do this manually). This was automated with the assistance of Moritz "Moe" Schiesser who was able to process this data programmatically so that it was organized in folders by SGL and processed using LAStools (specifically las2las.exe, lastile.exe, BLAST2DEM.exe; and QGIS. Resolution of the resultant DEMs is 1.0 m^2. It should be noted, because Pennsylvania is broken into North and South Tiles (based on different coordinate reference systems; State Plane 1983 PA South and North), some of the SGLs have two DEMs, two hill-shades, and two slope analyses. In order to ensure that we had a complete set of images and that they loaded correctly, we wrote utility programs in QGIS to load each TIFF (Tag Image File Format) that corresponded to our list of SGLs.[2]

The deep-learning process described below needs to be trained on known RCHs. Therefore, we reviewed and marked RCHs in 62 SGLs. To identify these RCHs, we primarily utilized the slope analysis over the Google Maps satellite photos. The latter helps identify non-RCHs, such as small buildings. In situations where identification was unclear, we also examined the hillshade. If we were still unsure of our identification, we also employed a QGIS plugin called Profile Tool, which allowed us to see the profile of the suspected RCHs in the DEM (Figure 3). In our review, we looked for features that represented two distinct but overlapping signatures (see also Bonhage et al. 2021). RCHs on slopes were dug into the hillside and therefore have steep slopes above and below a flat, level area between 10 and 15 m in diameter. The flat area tends to be ovoid rather than circular likely because the collier used the portion level with the rest of the slope as a path and because the steep upper slope eroded over time, partially covering the circular RCH (Figures 3b, 3d, and 3f). RCHs on flat areas are more difficult to identify. These tend to be primarily identifiable through a very flat, round area around 10–15 m in diameter that normally has a circular berm around the outer edge that represents the earth from the top of the meiler that had been raked to the edges. They are also often associated with borrow pits (Figures 3a, 3c, and 3f). This process resulted in the manual identification of 4,376 RCHs to serve as our training data. We attempted to ground truth 174 RCHs identified in this manner. Of these, 170 (97.7%) were confirmed in the field to be true positives, and only four (2.3%) were identified as false positives. These features are quite distinctive on the landscape and easy to confirm by their shape (as described above), but all of these were also confirmed by digging an informal hole approximately 10 cm deep on the lower slope that revealed large quantities of charcoal. Most of those confirmed through fieldwork are located on slopes, so they are easier to identify in the slope analysis than those on flat areas. Conner (2018) utilized intensive survey on flat areas to identify charcoal RCHs. He identified potential hearths in the slope analysis using a confidence rating system (from 1 = highly confident to 3 = possible but unlikely). Approximately 50% (6/13) of the
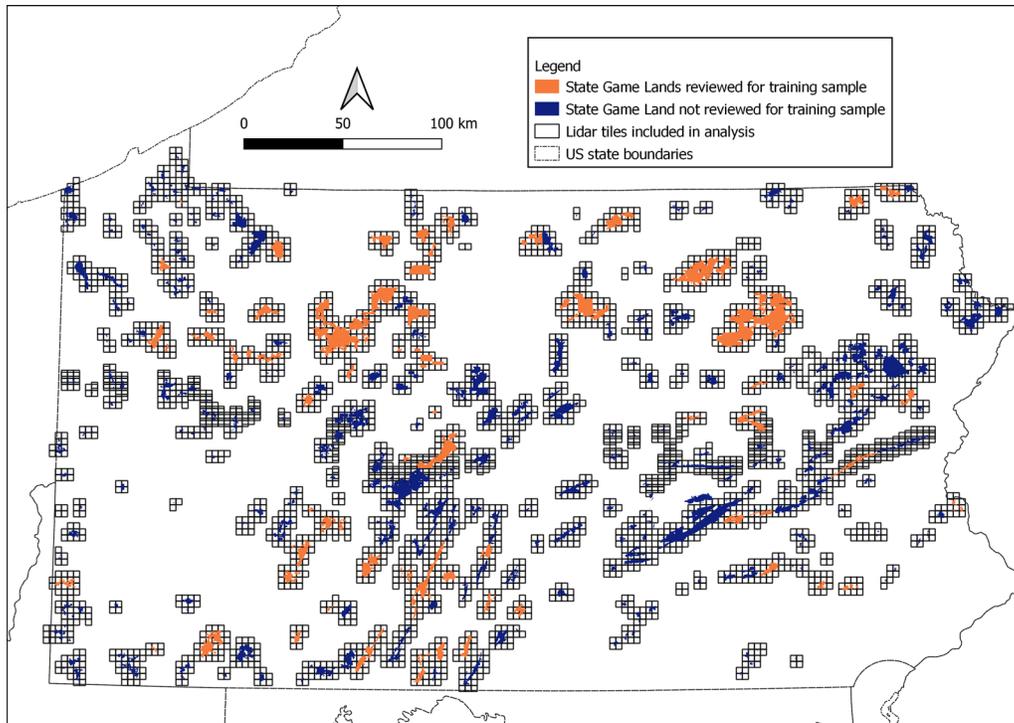
FIGURE 2. Map indicating all Pennsylvania State Game Lands showing those included in training the model and the lidar tiles used for this study (image by Benjamin Carter and Weston Conner).

locations that he rated highly (1 or 2 in his system) were confirmed in the field, and none (out of 8) of the low-confidence-rated potential RCHs were confirmed. Using systematic field survey of 0.2445 km², he was also able to identify three RCHs that had not been identified in the slope analysis (i.e., false negatives). Therefore, it is likely that our methods for identifying RCHs in the slope analysis are significantly more accurate for those on slopes compared to those on flat areas.

DEEP-LEARNING OBJECT RECOGNITION USING MASK R-CNN

We selected Waleed Abdulla's Mask R-CNN as a deep-learning platform for object detection (see also Bonhage et al. 2021). Mask R-CNN is suitable for identifying objects that appear only in a small part of an image (Brownlee 2019a, 2019b; He et al. 2020; Rosebrock 2019), such as RCHs in images of much larger landscapes. Additionally, a guiding principle for this project is to use accessible open source technology, such as Mask R-CNN, where possible to lower the barriers for other researchers so that they can build on our work.

Abdulla's Mask R-CNN uses Keras on TensorFlow, programmed in Python. We implemented Mask R-CNN on a virtual Ubuntu Linux computer GPU to allow for efficient processing. Anaconda Navigator was used to manage the required packages for Keras and TensorFlow and the compatible version of Python (Dombrowski et al. 2019). The use of Anaconda to manage the Python environment reduced complexity in the project.

Data Preparation

Each slope file (in TIFF format) for a SGL is large. For example, the file for SGL 33 is 1.5 GB in size. To make these images easier to process by Mask R-CNN, we wrote a program[3] to divide each TIFF into smaller 1024 × 768 pixel tiles that Mask R-CNN can consume efficiently. The program divided the TIFF into a grid of tiles and saved each tile as a JPEG image. Some tiles were empty due to the irregular shape of SGLs, and the program discarded these voids.

The shape file containing the points of the 4,376 manually recognized RCHs was used to identify tiles to be set aside for training, given that only a portion of tiles contained known RCHs. If a tile contained the point of a known RCH, the program copied the tile JPEG to a separate folder for training.

Next, we told Mask R-CNN where to look on each training image for examples of RCHs. This is done through annotation, which documents the *x,y* pixel boundaries (in xml format) that contain the region of interest (ROI)—in this case, an RCH. The program used the shapefile of known RCHs and expanded each point to a rectangular ROI bounding the RCH.[4] The size of the ROI was set to include all of a typically sized RCH (ca. 10–15 m in diameter) as well as several additional meters of surrounding land for a total size of 30 × 30 m (Figure 4). Experimentation showed that including the area surrounding the RCH was part of what distinguished it from other similar looking objects in the slope analysis. When the ROI was made smaller, and the boundary was close to the edge of a typical RCH, training results were less accurate.
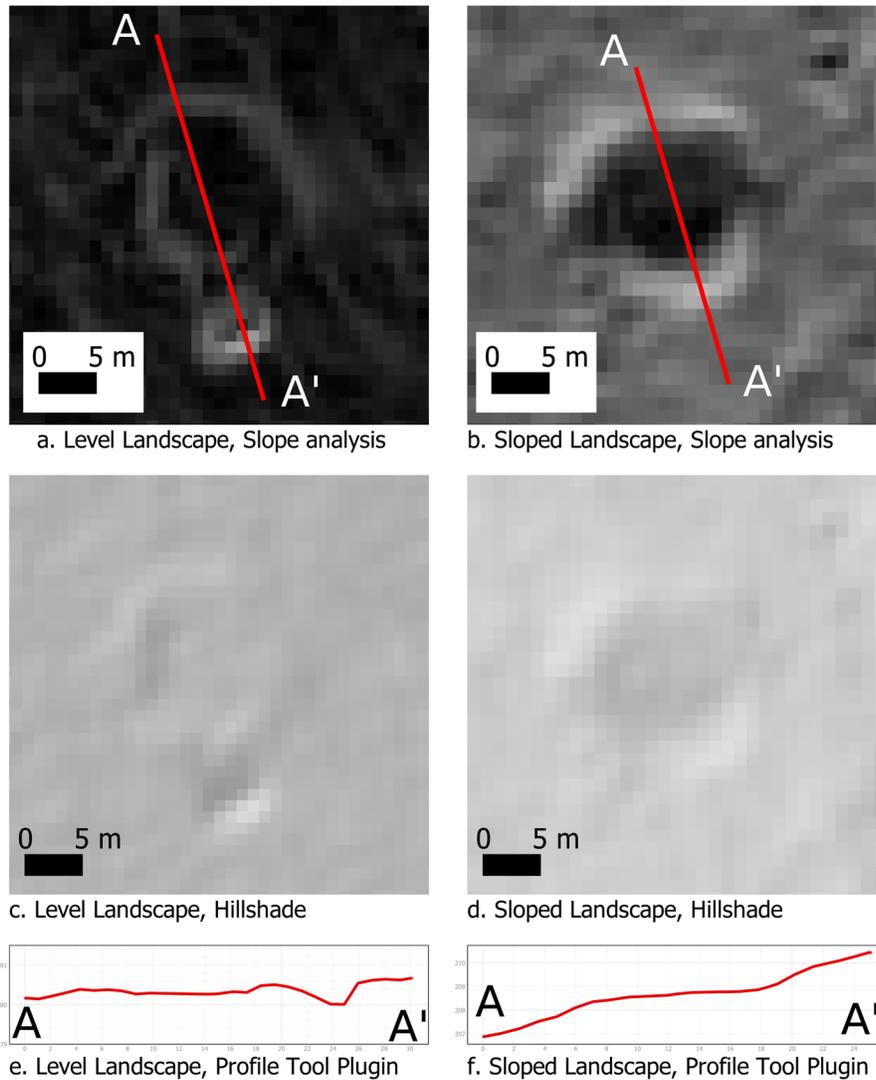
FIGURE 3. Comparison of relict charcoal hearths on level and sloped landscape in slope analysis, hillshade, and using the Profile Tool plugin for QGIS (image by Benjamin Carter).

Training the Model

To train the model, the JPEG images mentioned above were split randomly into training (80%) and testing sets (20%). Inside each training tile is at least one 30 × 30 m ROI (Figures 4 and 5). The ROIs in these images allow Mask R-CNN to learn to find other objects like those in the ROI. Once the model has learned to identify RCHs based on known ROIs, it then runs the model to see how successfully it can also identify known RCHs in the smaller "testing" set of known ROIs. In this way, the model recognizes how well it is working (or not) and makes adjustments.

Models are trained for a number of training/testing cycles called "epochs" as a way to improve accuracy with repetition. Models start with a random set of values known as weights that represent the model's recognition of a general pattern. The weights are adjusted as recognition improves with subsequent epochs. This improvement in learning reaches an optimal point. If a model is

trained for too many epochs, it becomes "overfit" and is biased to recognize only data it was trained on rather than generally recognizing similar objects.

The process of training and deciding when an object detection model is at peak learning performance is iterative. Because initial models start out with random weights, outcomes vary. We adjusted parameters to tune how the model learns in order to get the best outcome. The learning rate parameter represents how much weights are adjusted between each epoch of training. The greater the learning rate, in theory, the faster a model can be trained. However, high learning rates can quickly lead to a model being overfit and consequently useless. We used a learning rate of 0.001 for 8 epochs and then reduced it to 0.0001 for epochs 9–18 (Khandelwal 2019; Rosebrock 2019).

We experimented with "detection min confidence," a threshold value Mask R-CNN uses while it is training to decide if it has
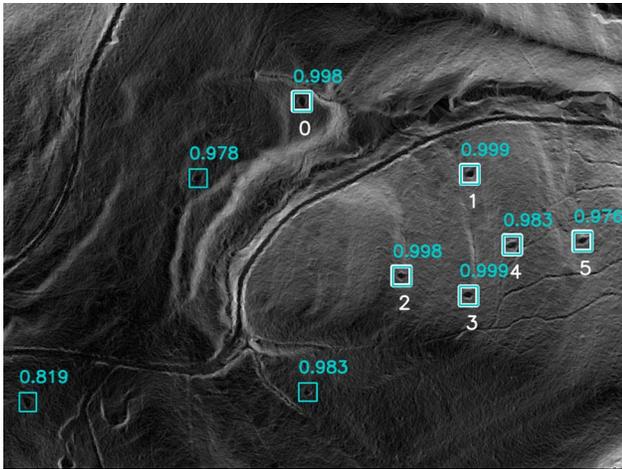
FIGURE 4. Slope analysis of a portion of SGL 33 with manually recognized relict charcoal hearths (RCHs) in white numbered squares and Mask R-CNN–identified RCHs in light blue squares. RCHs with both blue and white were recognized by both methods. The number above the square is the confidence score (image by Benjamin Carter).

correctly detected an object. We tried values of 0.7 and 0.9, and we found that 0.9 produced a more accurate model.

We set a maximum limit of 46 on the number of RCHs detected in a 1024 × 768 jpg. Sampling showed that this was a reasonable upper limit of RCHs visible in an image that did not risk loss of true positives, but it also put a ceiling on possible false positive results. After each training run, the results were validated with a calculation known as "average precision" (AP)—a performance metric based on the model's ability to detect true positives and minimize false negatives (Bonhage et al. 2021; Hui 2018). Mask R-CNN uses mean average precision (mAP), which averages across multiple classes of objects as a quick way to judge a model. However, because this model employs only one class of object, mAP = AP. Models with higher mAP, which is expressed as a proportion (0–1; see Hui 2018), suggest increased recognition of true positives and minimized inclusion of false positives.

We trained numerous models after adjusting learning rate and "detection min confidence." We also adjusted the number of epochs used with different learning rates. For example, on the last training run, for epochs 19–28, we used a value that was 1% of the learning rate for epochs 1–8 noted above. In the end, this variant was not relevant because our model's accuracy peaked at epoch 16.

Once we trained a set of models, we graded it using a four-stage process in order to select the best. Models that had a peak mAP of <0.5 were discarded immediately. Models with a mAP of >0.5 were used to detect RCHs in 20 images. After visual inspection of the predictions on 20 images, models that failed to find RCHs or had excessive false positives were discarded. If a model passed the visual inspection of 20 images, it looked promising, and we formally scored it using a set of 100 randomly selected images. The score was analyzed for the model's predictive value (AP).[5] Finalist models were used for predictions on all images, and large samples of these

predictions were manually examined to determine the model's value. For the final model, the mAP peaked at epoch 16 (Train mAP = 0.622 and Test mAP = 0.554) with the settings LEARNING_RATE = 0.001 and DETECTION_MIN_CONFIDENCE = 0.9.

Running Predictions

Once we determined appropriate settings, we ran the model on all JPEG tiles.[6] The resultant data comprised rectangles using the *x,y* pixel coordinates of the image stored in an XML file (Figures 4 and 5). The program converts the resulting predictions from pixel coordinates into geolocated 30 × 30 m rectangles stored in a shape file.

Due to the overlap of images from the two different coordinate reference systems used for the Pennsylvania SGLs, some RCHs appeared on two different images and showed up as duplicate predictions. These duplicates were programmatically eliminated.[7] The results also provided a predicted confidence for each RCH— that is, based on the model, how likely was the predicted RCH a true RCH (from 0–1). We set our lower limit at 0.7. The final result is a list of unique geographic points of 52,884 predicted RCHs.

ASSESSING THE MASK R-CNN PREDICTIONS

Once Mask R-CNN completed its predictions, a quick review helped us realize that, although lots of the predictions were true positives, many were not. We used the following methods to separate these.

For our first attempt, we chose eight SGLs (13, 33, 37, 39, 51, 76, 83, 217) distributed across the state and physiographic provinces. Rather than examining each predicted RCH, we first divided them into "confidence" bins (Table 1). We randomly selected four predicted RCHs from each bin for each SGL (except SGL 83, which only had two RCHs in the 0.75–0.80 bin; Table 1), for a total of 190 predicted RCHs. We used the methods for manually identifying RCHs in slope analyses described above to determine if the predicted RCH was a true positive or a false positive. This resulted in the recognition of 150 false positives (78.9%) and 40 true positives (21.1%, see Table 1). Although this initially seems to be an indication of poor results, this helped us realize that we needed alternative methods for assessing the predicted sample. For example, of this sample, two SGLs (13, 37) had no true positives. There was likely little or no historic charcoal production in these SGLs. On the other hand, the majority of the true positives were in the two upper "confidence" bins (27/40, or 67.5%; Table 1).

To better assess this, we used an alternative approach. In the field, RCHs tend to be clustered in groups. This makes sense because a collier frequently tends multiple charcoal hearths at one time. To limit travel time and facilitate working many hearths, hearths are often located near each other (ca. 80–150 m, but with great variation). Additionally, in our manual review, we noticed that many of the false positives were not near other predicted RCHs. Therefore, we employed cluster analysis as a way to separate true positives from false positives. We employed the DBSCAN plugin in QGIS (for the use of DBSCAN in archaeology, see Argote-Espino et al. 2012; Caspari and Jendryke 2017; Maddison 2020; Maddison and
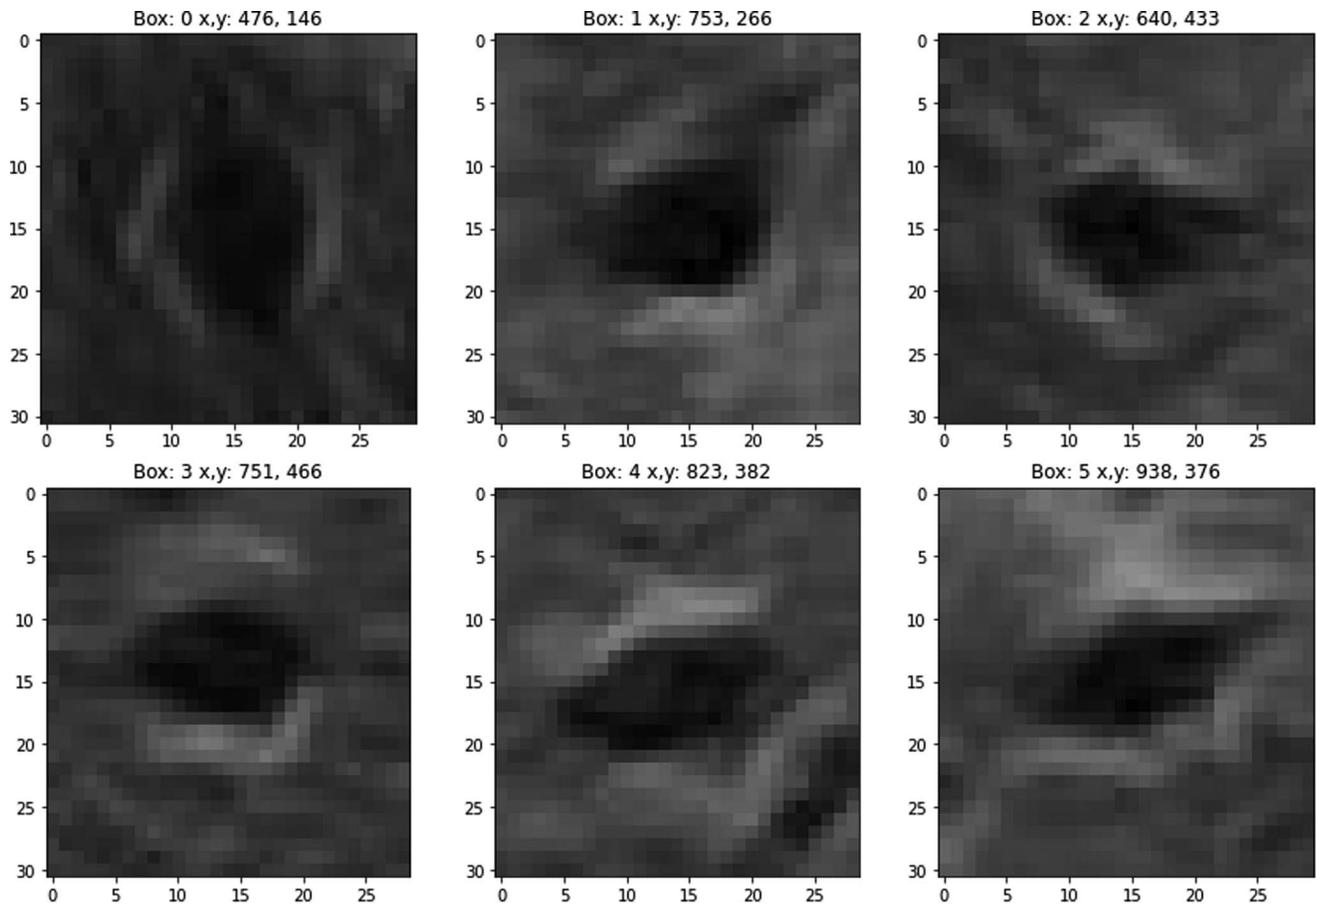
FIGURE 5. A close-up of six manually recognized RCHs shown in Figure 4 demonstrating similarities and variation (image courtesy of Jeff Blackadar).

TABLE 1. Manually Recognized True Positives at a Selected Set of SGLs by Confidence Bin.

Confidence Bin	State Game Land								Total
	SGL 13	SGL 33	SGL 37	SGL 39	SGL 51	SGL 76	SGL 83	SGL 217	
0.70–0.75	0	0	0	1	1	0	0	1	3/32 (9%)
0.75–0.80	0	0	0	1	0	0	0	1	2/30 (7%)
0.80–0.85	0	1	0	0	0	1	2*	1	5/32 (16%)
0.85–0.90	0	0	0	1	0	1	0	1	3/32 (9%)
0.90–0.95	0	3	0	1	0	1	2	2	9/32 (28%)
0.95–1.0	0	1	0	3	2	4	4	4	18/32 (56%)
Total	0/24 (0%)	5/24 (21%)	0/24 (0%)	7/24 (29%)	3/24 (13%)	7/24 (29%)	8/22 (36%)	10/24 (42%)	40/190 (21%)

Note: For example, in the 0.70–0.75 confidence bin, 0 out of 4 of the RCHs recognized by Mask R-CNN were visually confirmed to be true positives.

Schmidt 2020). Density-based spatial clustering of applications with noise, or DBSCAN, creates clusters based on two variables: the maximum distance and the minimum number of individuals in a cluster. Starting with a random RCH, the plugin determines if there are other points within the maximum distance. If so, that point is included in a cluster, as long as that cluster reaches or exceeds the minimum number of RCHs. Anything not within a cluster is described as "noise." To start, we experimented with a variety of variables and observed the results by visualizing the clusters in QGIS. We learned that there is no combination of DBSCAN settings that reliably distinguishes individual true positives from false positives. For example, an RCH on the edge of a cluster may be identified as noise if it is not close enough to another RCH. Alternatively, false positives will be included in a cluster if they are within the parameters. Therefore, we could not use clustering to determine if an individual predicted RCH was a

TABLE 2. DBSCAN Clusters of Mask R-CNN–Identified RCHs Compared to Manual Recognition at SGL 43.

Number of Clusters	RCHs Predicted by Mask R-CNN		
	False Positive	True Positive	Total
0	33 (75%)	11 (25%)	44 (100%)
1	37 (63.8%)	21 (36.2%)	58 (100%)
2	27 (45.0%)	33 (55%)	60 (100%)
3	123 (13.7%)	778 (86.3%)	901 (100%)
Total	220	843	1,063

Note: The total number of predicted hearths that fell into at least one cluster is 1,019 (901 + 60 + 58), of which 832 (778 + 33 + 21) were identified as true positives and 187 (123 + 27 + 37) were false positives.

true or false positive, but we could use it to get a much better idea of where clusters were located. Consequently, we used three different combinations of variables (minimum number / maximum distance = 5 / 300 m, 10 / 500 m, 20 / 1,000 m) to visualize these clusters.

The results of the DBSCAN analyses allow us to distinguish likely true clusters from "noise" (i.e., false positives) even if a determination cannot be made for each RCH. A total of 26,873 RCHs (50.8% of all predicted RCHs) were placed in a cluster based on at least one of the three variable combinations. A total of 18,202 predicted RCHs (34.4%) were organized into all three cluster variants. It is likely that most of these predicted RCHs are true positives. Predicted RCHs not included in these clusters are more likely to be false positives. For example, none of the predicted RCHs in SGL 13 and 37 (Table 1), where we failed to manually recognize any RCHs in the slope analysis and therefore know that they are all definitively false positives (see above), are included in the three clusters.

In order to address how well DBSCAN effectively distinguished true and false positives, we reviewed the predicted RCHs in and around SGL 43 using the methods discussed above. We chose this area because it is located near Hopewell Furnace, a National Historic Site as well as one of the best-known and well-researched iron furnaces in the country (e.g., Kemper 1941; Straka and Ramer 2010; Walker 1966), and we have begun to conduct fieldwork there. Table 2 demonstrates that 832 (out of 1,019, or 82.6%) of the predicted RCHs that fall within one, two,
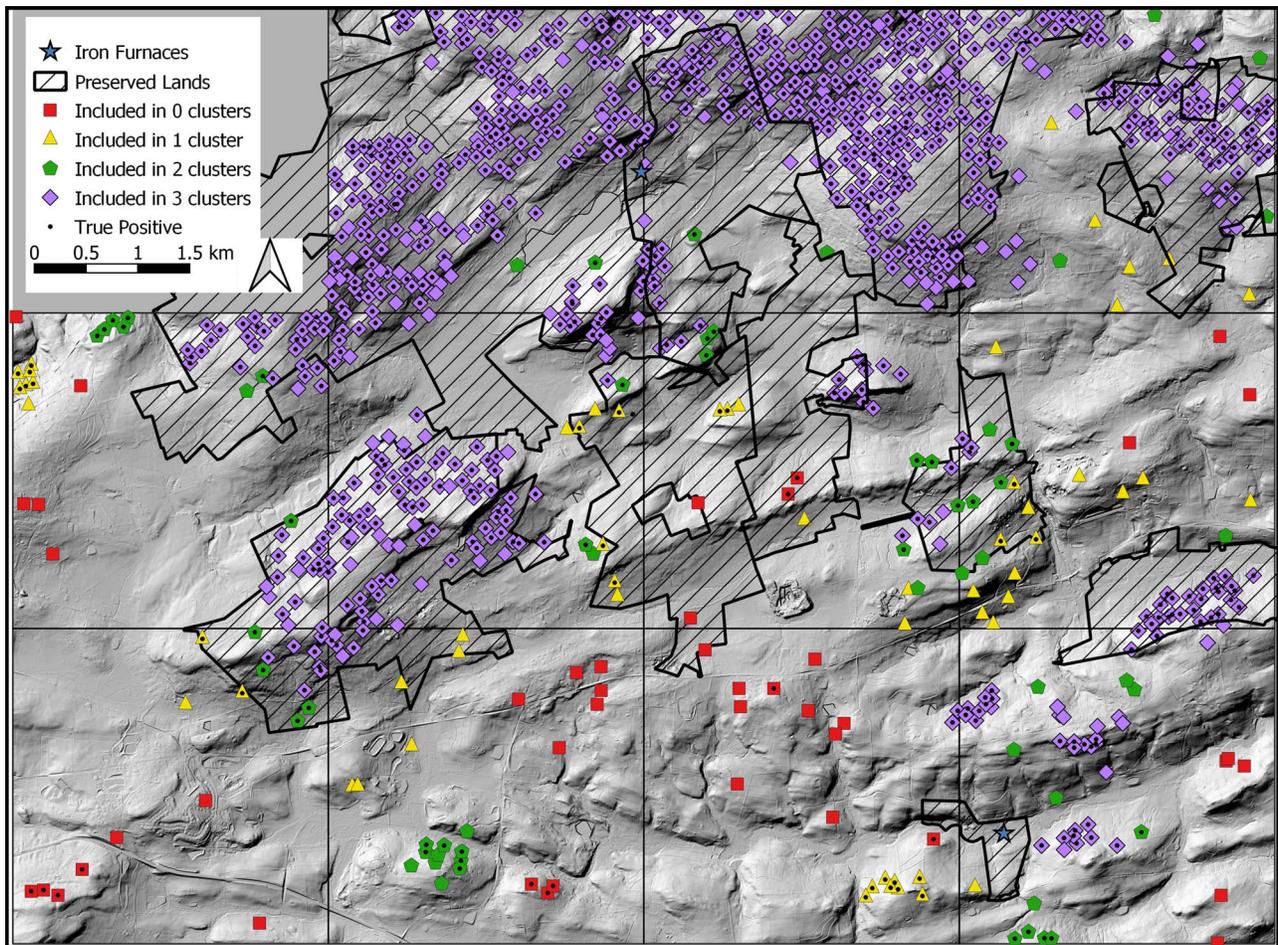

FIGURE 6. Map showing predicted RCHs and their inclusion in the three DBSCAN cluster analyses compared to a manual review that identified true positives in and around SGL 43. Base map is the hillshade created using methods described herein. See Figure 7 for preserved land labels (image by Benjamin Carter).
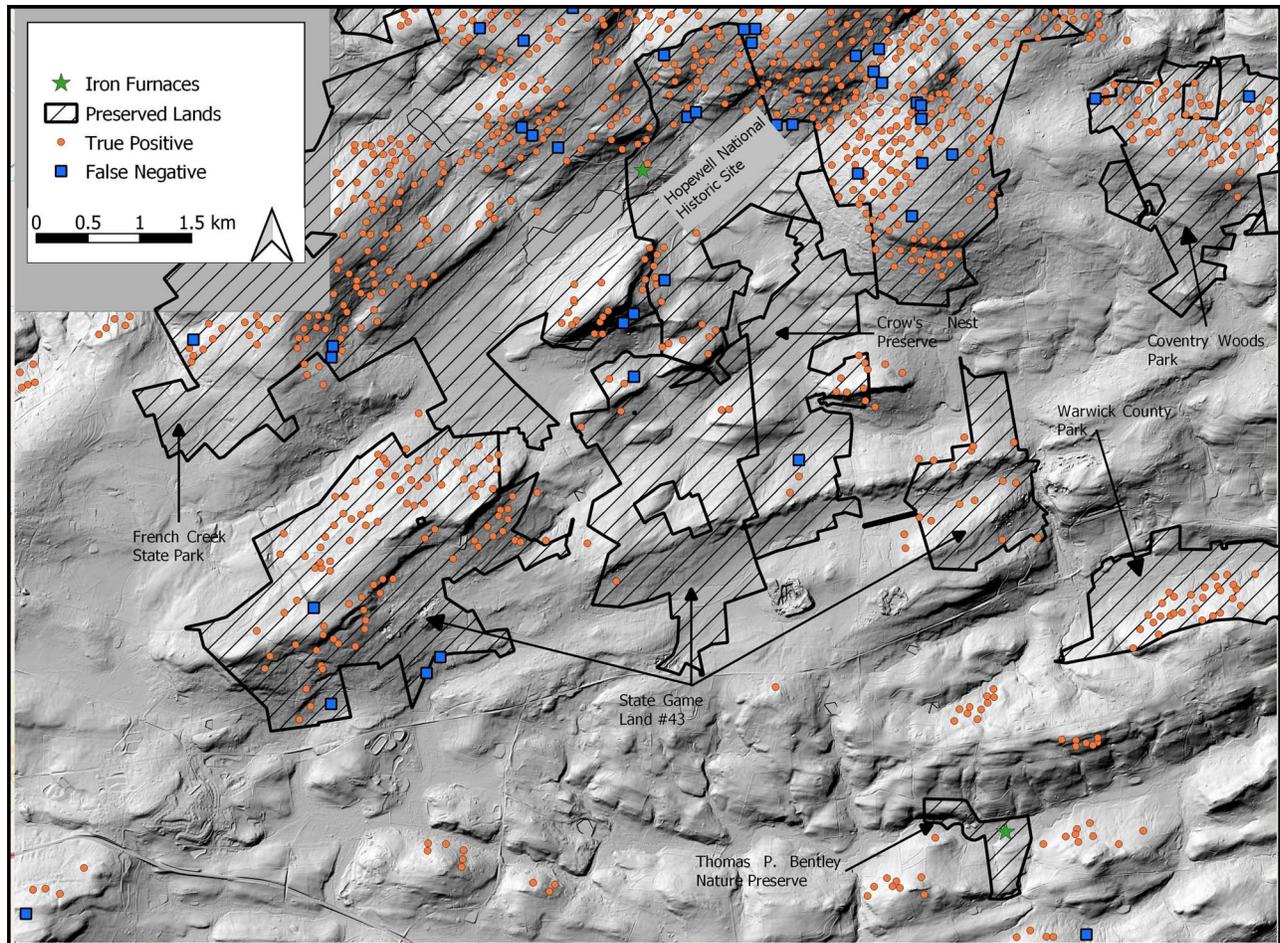
FIGURE 7. Map showing true positives and false negatives for RCHs in and around SGL 43. Note that both are concentrated in preserved areas. Base map is the hillshade create using methods described herein (image by Benjamin Carter).

or all three of the clusters were identified as true positives—187 (18.4%) were false positives. Of those RCHs that were not placed into any clusters, the majority (33/44, or 75.0%) were manually recognized as false positives, and 25.0% (11/44) were true positives. Figure 6 shows the contexts of these RCHs. Most of the RCHs not included in a DBSCAN cluster (Figure 6; red squares) were not confirmed by manual evaluation and are in residential areas outside of preserved areas. Most of the RCHs included in all three clusters were confirmed and were located within local preserved space, including SGL 43, French Creek State Park, Hopewell National Historic Site, Thomas P. Bentley Nature Preserve (the location of Warwick Furnace), Coventry Woods Park, Crow's Nest Preserve, and Warwick County Park. Therefore, cluster analysis can, with reasonably high probability, indicate that a single predicted RCH is a true or false positive.

Consequently, clustering provides us with a broader-scale understanding of charcoal "fields"—areas of the landscape densely dotted with RCHs. Figure 8 shows a comparison between predicted RCHs that fall into all three DBSCAN clusters and the lidar tiles employed in this analysis. If one drew a line from the southwest to the northeast of the state, the southeastern half would

have the vast majority of charcoal fields. This includes a near continuous line of "fields" along the 250-mile stretch of Blue Mountain, along with numerous fields to the north and south of it. Fields are notably absent across much of the northern portion of the state where there were fewer charcoal furnaces and forges. Small fields are located in the northwest, where the lifespan of charcoal furnaces was relatively short (ca. 10 years; Knowles 2012:49; see also Williams 2020).

DBSCAN clustering also allows us to examine how well the Mask R-CNN was able to predict the same RCHs that were used to construct the model. Of the 4,376 RCHs used to train the model fully, 4,174 (95.4%) were predicted by the model, demonstrating both the effectiveness of Mask R-CNN and how the abstraction created by Mask R-CNN does not automatically detect each of the individual training RCHs. Of the 4,174 training RCHs included in the predicted RCHs, 4,004 (96.0%) were placed in at least one of the DBSCAN clusters, and 3,523 (84.4%) were placed in all three clusters. Although most individual SGLs mirror these proportions, localized patterns also exist. For example, the majority of the training RCHs that were also recognized by Mask R-CNN in SGL 45 and SGL 51 were not placed into any of the clusters (28/36, or 78%; 14/16 or 87.5%, respectively).
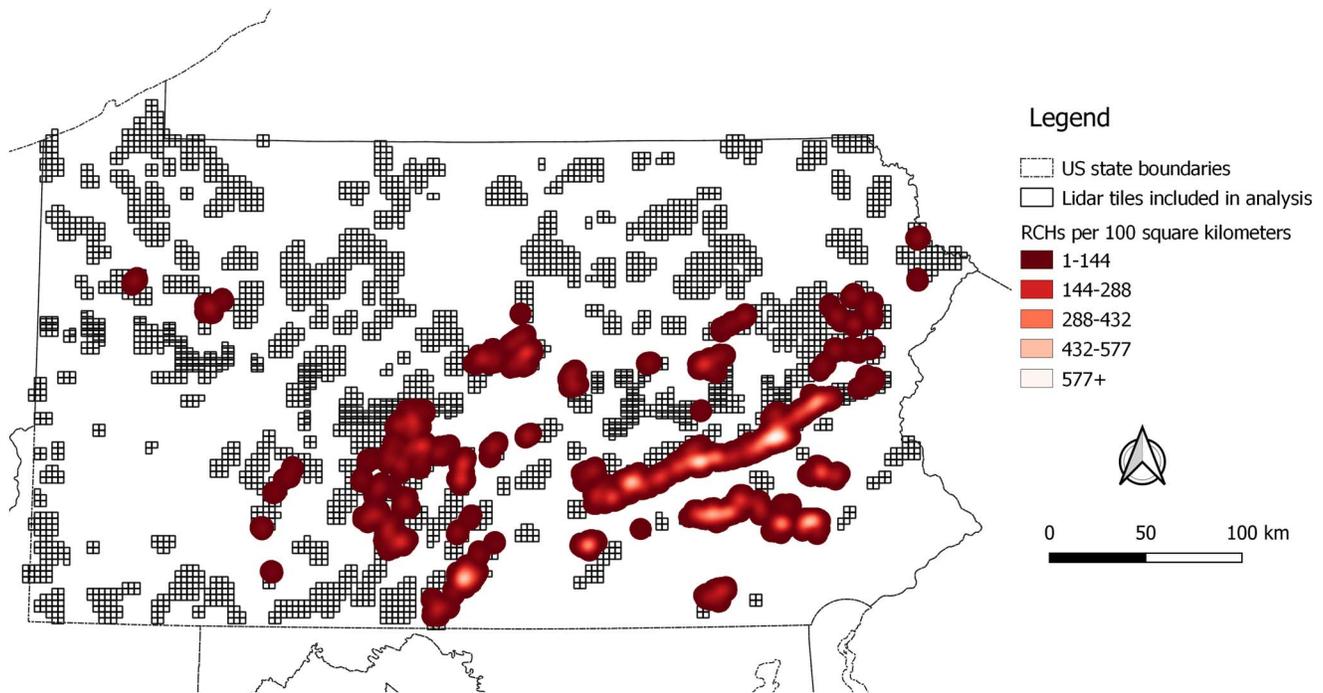
FIGURE 8. Map showing a heat map of predicted RCHs included in all three DBSCAN analyses compared to tiles included in the sample. Darker red indicates lower density, and lighter red indicates higher density (image by Benjamin Carter).

The DBSCAN of the results of the Mask R-CNN provides a much broader view of the immense impact of charcoal production on the landscape of Pennsylvania. It also allows us to differentiate where in the state the majority (if not all) of the impact was located and where it was unlikely. In particular, charcoal production focused on the 400 km (250-mile) long Blue Mountain and the hilly areas within approximately 100 km (60 miles) to the north and west, as well as on the hills to south. This more fully addresses the impact of charcoal production on the landscape of Pennsylvania. There are some limitations, however. First, this is only a partial sample of the entire state—other undeveloped lands (e.g., state parks, county parks, private preserves) could also be examined. Second, the combined use of Mask R-CNN and DBSCAN is more appropriate for identifying individual charcoal hearths that collectively form "fields." This means that individual (or small clusters of) RCHs will be missed. Third, additional work may be able to demonstrate additional variables—such as proximity to furnaces, water, et cetera—that impact both the construction of hearths and the recognition of those hearths in lidar derivatives. These variables may be employed to increase our ability to separate true positives from false. Fourth, we have limited our conversation about false negatives. Though the comparison between the training data and the examination of SGL 43 suggests that false negatives are rare (ca. 5% or less), more work needs to be done. Last, in order to better understand charcoal production and consumption, these fields of RCHs should be reconnected with the forges and furnaces that consumed the charcoal to make iron.

We also reviewed SGL 43 to identify false negatives—those RCHs missed by the Mask R-CNN model but identifiable in the slope analysis. Two of the authors (WC, BC) scanned the slope analysis of SGL 43 at a scale of 1:1000 looking for the features discussed

above—flat, round areas approximately 10–15 m in diameter that either had clear steep uphill and downhill slopes on sloped terrain or circular berms on flat terrain. Any identifications on which the reviewers disagreed were reassessed, and a final collaborative determination was made. A total of 46 potential RCHs were identified (Figure 7)—in comparison to the 843 RCHs detected by Mask R-CNN and confirmed as true positives. Although we have only confirmed a few of the RCHs in the field at SGL 43, the ground truthing discussed above suggests that these methods are likely quite accurate on slopes, although we are less confident in our identification of RCHs on flat terrain. This is particularly important because Williams (2020) has identified 279 colliers in the 1850 census for Clarion County (in the northwest of Pennsylvania). In our own preliminary analysis, we believe this to be an incredibly high number of colliers. Yet, although the heat map does indicate RCH "fields" in Clarion county, this would seem to underrepresent the number of colliers. We are not arguing, therefore, that the techniques described herein are able to identify all RCH fields—only that, compared to previous attempts, they move us significantly in the right direction.

CONCLUSION

This research demonstrates how the impact of charcoal production in Pennsylvania can be understood by using deep learning (Mask R-CNN) and cluster analysis (DBSCAN in QGIS) to analyze lidar derivatives. Charcoal production is a relatively poorly known industry that has left extensive modifications of the landscape across hundreds of square kilometers of terrain. Whereas a recent spike in research on RCHs (largely due to the availability of lidar) has demonstrated the prevalence and impact of charcoal production, this research is able to show the presence of hundreds of

kilometers of "fields'' of RCHs, as well as large swaths of the state where RCHs are absent, although we may underrepresent the size of these fields. As Bernhard Fernow suggested, based on our research, charcoal production appears to be associated with "waste places," particularly the arid tops of long ridges, such as the Blue Mountain.

Acknowledgments

Data Availability Statement

One of the goals of this research is to provide all of the code/data employed herein in order to both ensure replicability and encourage additional use of our work. The data used for this research originated in open repositories (primarily through the state of Pennsylvania). We plan to further publish both code and data constructed in this research. Until then, links for these are provided through an archived document located at https://zenodo.org/record/4766351.

NOTES

1. For detailed information, please see the following document, which contains links to the data and code used in this work: https://zenodo.org/record/4766351.
2. See programs qgis_1_check_predictions.py, qgis_2_add_poly_layers_for_north_east.py, qgis_3_load_rasters_north_east.py, and qgis_4_check_we_have_all_layers.py in https://github.com/jeffblackadar/charcoalhearths/tree/master.
3. https://github.com/jeffblackadar/charcoalhearths/blob/master/0_split_tifs_refactored.ipynb.
4. See def annotate_tif_if_it_has_hearths in 0_split_tifs_refactored.ipynb.
5. An analysis for model cfg20200826T2315 epoch 16 is here: https://github.com/jeffblackadar/charcoalhearths/blob/master/prediction_results.pdf.
6. Predictions on all images are performed in the later cells of notebook data_5000_3_rcnn_charcoal_hearths.ipynb.
7. The last cell of 2_read_predictions_from_xml_put_into_shp.ipynb eliminates duplicates in the results.

REFERENCES CITED

Alfaro, Jose F., and Brieland Jones
  2018 Social and Environmental Impacts of Charcoal Production in Liberia: Evidence from the Field. *Energy for Sustainable Development* 47:124–132. DOI:10.1016/j.esd.2018.09.004.

American Iron and Steel Association
  1876 *The Ironworks of the United States: A Directory of the Furnaces, Rolling Mills, Steel Works, Forges and Bloomaries in Every State.* Centennial edition. Philadelphia.
  1878 *Directory of the Iron and Steel Works of the United States.* 4th ed. James B. Chandler's Steam Printing Establishment, Philadelphia.
  1880 *Directory of the Iron and Steel Works of the United States.* 5th ed. James B. Chandler's Steam Printing Establishment, Philadelphia.
  1882 *Directory of the Iron and Steel Works of the United States.* 6th ed. James B. Chandler's Steam Printing Establishment, Philadelphia.
  1884 *The Iron and Steel Works of the United States.* 7th ed. Allen, Lane & Scott, Philadelphia.
  1886 *The Iron and Steel Works of the United States.* 8th ed. Allen, Lane & Scott, Philadelphia.
  1890 *The Iron and Steel Works of the United States.* 10th ed. Allen, Lane & Scott, Philadelphia.
  1892 *The Iron and Steel Works of the United States.* 11th ed. Allen, Lane & Scott, Philadelphia.
  1894 *The Iron and Steel Works of the United States.* 12th ed. Allen, Lane & Scott, Philadelphia.
  1896 *The Iron and Steel Works of the United States.* 13th ed. Allen, Lane & Scott, Philadelphia.

Argote-Espino, Denisse, Jesús Solé, Pedro López-García, and Osvaldo Sterpone
  2012 Obsidian Subsource Identification in the Sierra de Pachuca and Otumba Volcanic Regions, Central Mexico, by ICP-MS and DBSCAN Statistical Analysis. *Geoarchaeology* 27:48–62. DOI:10.1002/gea.213891.

Bekele, Melaku, and Zenebe Girmay
  2014 *Reading through the Charcoal Industry in Ethiopia: Production, Marketing, Consumption and Impact.* Forum for Social Studies, Addis Ababa, Ethiopia.

Birkinbine, John
  1879 The Production of Charcoal for Iron Works. *Transactions of the American Institute of Mining Engineers* 7:149–158.
  1881 Our Fuel. *Journal of the United States Association of Charcoal Iron Workers* 2(2):66–79.

Bonhage, Alexander, Mahmoud Eltaher, Thomas Raab, Michael Breuß, Alexandra Raab, and Anna Schneider
  2021 A Modified Mask Region-Based Convolutional Neural Network Approach for the Automated Detection of Archaeological Sites on High-Resolution Light Detection and Ranging-Derived Digital Elevation Models in the North German Lowland. *Archaeological Prospection* 28:177–186. DOI:10.1002/arp.1806.

Bonhage, Alexander, Florian Hirsch, Thomas Raab, Anna Schneider, Alexandra Raab, and Will Ouimet
  2020 Characteristics of Small Anthropogenic Landforms Resulting from Historical Charcoal Production in Western Connecticut, USA. *CATENA* 195:104896. DOI:10.1016/j.catena.2020.104896.

Brownlee, Jason
  2019a How to Use Mask R-CNN in Keras for Object Detection in Photographs. *Machine Learning Mastery* (blog), May 24, 2019, updated September 2, 2020. https://web.archive.org/web/20210202022318/https://machinelearningmastery.com/how-to-perform-object-detection-in-photographs-with-mask-r-cnn-in-keras/, accessed December 28, 2020.
  2019b How to Train an Object Detection Model with Keras. *Machine Learning Mastery* (blog). May 29, 2019, updated September 2, 2020. https://web.archive.org/web/20210130013601/https://machinelearningmastery.com/how-to-train-an-object-detection-model-with-keras/, accessed December 28, 2020.

Carrari, Elisa, Evy Ampoorter, Francesca Bottalico, Gherardo Chirici, Andrea Coppi, Davide Travaglini, Kris Verheyen, and Federico Selvi
  2017 The Old Charcoal Kiln Sites in Central Italian Forest Landscapes. *Quaternary International* 458:214–223. DOI:10.1016/j.quaint.2016.10.027.

Carter, Benjamin
  2019a Identifying Landscape Modification Using Open Data and Tools: The Charcoal Hearths of the Blue Mountain, Pennsylvania. *Historical Archaeology* 53:432–443. DOI:10.1007/s41636-019-00171-1.
  2019b Data for "Identifying Landscape Modification Using Open Data and Tools: The Charcoal Hearths of the Blue Mountain, Pennsylvania." *Journal of Open Archaeology Data* 7:3. DOI:10.5334/joad.53.

Caspari, Gino, and Michael Jendryke
  2017 Archsphere: A Cluster Algorithm for Archaeological Applications. *Journal of Archaeological Science: Reports* 14:181–188. DOI:10.1016/j.jasrep.2017.05.052.

Conner, Weston
  2018 *Industrial Landscapes Lost in the Wild: Creating a Methodology for the*

Remote Identification of Charcoal Hearths. Honors thesis, Department of Sociology and Anthropology, Muhlenberg College, Allentown, Pennsylvania. http://hdl.handle.net/10718/3428, accessed December 7, 2020.

Deforce, Koen, Bert Groenewoudt, and Kristof Haneca
2020 2500 Years of Charcoal Production in the Low Countries: The Chronology and Typology of Charcoal Kilns and Their Relation with Early Iron Production. *Quaternary International* 593–594:295–305. DOI:10.1016/j.quaint.2020.10.020.

Diderot, Denis, and Jean Le Rond d'Alembert
2010 [1765] Agriculture and Rural Economy – Charcoal. *The Encyclopedia of Diderot & d'Alembert Collaborative Translation Project*. Michigan Publishing, University of Michigan Library, Ann Arbor. http://hdl.handle.net/2027/spo.did2222.0001.346, accessed December 13, 2014.

Dombrowski, Quinn, Tassie Gniady, and David Kloster
2019 Introduction to Jupyter Notebooks. *Programming Historian* 8. DOI:10.46430/phen0087.

Donovan, Sally, Mary Ignatiadis, William Ouimet, David Dethier, and Michael Hren
2021 Gradients of Geochemical Change in Relic Charcoal Hearth Soils, Northwestern Connecticut, USA. *CATENA* 197:104991. DOI:10.1016/j.catena.2020.104991.

FAOSTAT
2021 Forestry Production and Trade Database (updated March 1, 2021). Food and Agriculture Organization of the United Nations. Electronic document, https://web.archive.org/web/20210202203523/http://www.fao.org/faostat/en/, accessed February 2, 2021.

Fernandez-Diaz, Juan Carlos, William E. Carter, Ramesh L. Shrestha, and Craig L. Glennie
2014 Now You See It . . . Now You Don't: Understanding Airborne Mapping LiDAR Collection and Data Product Generation for Archaeological Research in Mesoamerica. *Remote Sensing* 6:9951–10001. DOI:10.3390/rs6109951.

Fernow, Bernhard E.
1883 Planting in Waste Places. *American Journal of Forestry* 8:153–155.

Gordon, Robert B.
1996 *American Iron, 1607–1900*. Johns Hopkins University Press, Baltimore, Maryland.

He, Kaiming, Georgia Gkioxari, Piotr Dollar, and Ross Girshick
2020 Mask R-CNN. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 42:386–397.

Hillstrom, Kevin, and Laurie Collier Hillstrom
2006 *Industrial Revolution in America: Iron and Steel*, Vol. 1. ABC CLIO, Santa Barbara, California.

Hirsch, Florian, Thomas Raab, William Ouimet, David Dethier, Anna Schneider, and Alexandra Raab
2017 Soils on Historic Charcoal Hearths: Terminology and Chemical Properties. *Soil Science Society of America Journal* 81:1427–1435. DOI:10.2136/sssaj2017.02.0067.

Hui, Jonathan
2018 mAP (mean Average Precision) for Object Detection. *Deep Learning* (blog), March 6, 2018. https://web.archive.org/web/20210501031623/https://jonathan-hui.medium.com/map-mean-average-precision-for-object-detection-45c121a31173, accessed May 2, 2021.

Johnson, Katharine M., and William B. Ouimet
2018 An Observational and Theoretical Framework for Interpreting the Landscape Palimpsest through Airborne LiDAR. *Applied Geography* 91:32–44. DOI:10.1016/j.apgeog.2017.12.018.
2021 Reconstructing Historical Forest Cover and Land Use Dynamics in the Northeastern United States Using Geospatial Analysis and Airborne LiDAR. *Annals of the American Association of Geographers*, in press. DOI:10.1080/24694452.2020.1856640.

Jones, Claude A.
1976 Making Charcoal at Jacobsburg. In *Slate Belt Bicentennial Heritage*, edited by Marjorie May, Maryann Repsher, Robert R. Saltern, Gordon E. Wilson, and Albert M. Toth, pp. 32–33. Heritage Committee of the Slate Belt Bicentennial Commission, Pen Argyl, Pennsylvania.

Kato, M., D. M. DeMarini, A. B. Carvalho, M. A. V. Rego, A. V. Andrade, A. S. V. Bonfim, and D. Loomis
2005 World at Work: Charcoal Producing Industries in Northeastern Brazil. *Occupational and Environmental Medicine* 62(2):128–132.

Kazimi, Bashir, Frank Thiemann, and Monika Sester
2019 Object Instance Segmentation in Digital Terrain Models. In *Computer Analysis of Images and Patterns – CAIP 2019*, edited by Mario Vento and Gennaro Percannella, pp. 488–495. Lecture Notes in Computer Science. Springer, Cham, Switzerland.

Kemper, Jackson
1941 *American Charcoal Making in the Era of the Cold-Blast Furnace*. U.S. Department of the Interior, National Park Service, Washington, DC. Electronic document, http://npshistory.com/series/popular/14.pdf, accessed July 20, 2017.

Khandelwal, Renu
2019 Object Detection Using Mask R-CNN on a Custom Dataset. *Towards Data Science*. Electronic document, https://web.archive.org/web/20201101051548/https://towardsdatascience.com/object-detection-using-mask-r-cnn-on-a-custom-dataset-4f79ab692f6d?gi=8a83931da0e9, accessed March 3, 2021.

Knowles, Anne Kelly
2012 *Mastering Iron*. University of Chicago, Chicago.

LaRoche, Cheryl Janifer
2017 *Free Black Communities and the Underground Railroad: The Geography of Resistance*. University of Illinois Press, Urbana.

LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton
2015 Deep Learning. *Nature* 521:436–444. DOI:10.1038/nature14539.

Lesley, Peter
1866 *The Iron Manufacturer's Guide to the Furnaces, Forges and Rolling Mills of the United States; with Discussions of Iron, etc.* John Wiley, New York.

Maddison, M. Simon
2020 Percolation Analysis. In *Archaeological Spatial Analysis: A Methodological Guide*, edited by Mark Gillings, Piraye Hacıgüzeller, and Gary Lock, pp. 77–92. Routledge, New York.

Maddison, M. Simon, and Sophie C. Schmidt
2020 Percolation Analysis: Archaeological Applications at Widely Different Spatial Scales. *Journal of Computer Applications in Archaeology* 3:269–287. DOI:10.5334/jcaa.54.

Mikan, Carl J., and Marc D. Abrams
1995 Altered Forest Composition and Soil Properties of Historic Charcoal Hearths in Southeastern Pennsylvania. *Canadian Journal of Forest Research* 25:687–696. DOI:10.1139/x95-076.
1996 Mechanisms Inhibiting the Forest Development of Historic Charcoal Hearths in Southeastern Pennsylvania. *Canadian Journal of Forest Research* 26:1893–1898. DOI:10.1139/x26-213.

Opitz, Rachel S.
2013 *Interpreting Archaeological Topography: Airborne Laser Scanning, 3D Data and Ground Observation*. Oxbow Books, Oxford.

Opitz, Rachel, and Jason Herrmann
2018 Recent Trends and Long-Standing Problems in Archaeological Remote Sensing. *Journal of Computer Applications in Archaeology* 1:19–41. DOI:10.5334/jcaa.11.

Overman, Frederick
1854 *The Manufacture of Iron, in All Its Various Branches*. Henry C. Baird, Philadelphia.

Raab, A., A. Bonhage, A. Schneider, T. Raab, H. Rösler, K. -U. Heußner, and F. Hirsch
2019 Spatial Distribution of Relict Charcoal Hearths in the Former Royal Forest District Tauer (SE Brandenburg, Germany). *Quaternary International* 511:153–165. DOI:10.1016/j.quaint.2017.07.022.

Raab, A., M. Takla, T. Raab, A. Nicolay, A. Schneider, H. Rösler, K. -U. Heußner, and E. Bönisch
2015 Pre-Industrial Charcoal Production in Lower Lusatia (Brandenburg, Germany): Detection and Evaluation of a Large Charcoal-Burning Field by Combining Archaeological Studies, GIS-Based Analyses of Shaded-Relief Maps and Dendrochronological Age Determination. *Quaternary International* 367:111–122. DOI:10.1016/j.quaint.2014.09.041.

Rembold, F., S. M. Oduori, H. Gadain, and P. Toselli
2013 Mapping Charcoal Driven Forest Degradation during the Main Period of Al Shabaab Control in Southern Somalia. *Energy for Sustainable Development* 17:510–514. DOI:10.1016/j.esd.2013.07.001.

Risbøl, Ole, Ole Martin Bollandsås, Anneli Nesbakken, Hans Ole Ørka, Erik Næsset, and Terje Gobakken
2013 Interpreting Cultural Remains in Airborne Laser Scanning Generated

Digital Terrain Models: Effects of Size and Shape on Detection Success Rates. *Journal of Archaeological Science* 40:4688–4700. DOI:10.1016/j.jas.2013.07.002.

Rodgers, Andrew Denny
1991 *Bernhard Eduard Fernow: A Story of North American Forestry*. Forest History Society, Durham, North Carolina.

Rodrigues, Thaisa, and Aldo Braghini Junior
2019 Technological Prospecting in the Production of Charcoal: A Patent Study. *Renewable and Sustainable Energy Reviews* 111:170–183. DOI:10.1016/j.rser.2019.04.080.

Rolando, Victor
1992 *200 Years of Soot and Sweat: The History and Archeology of Vermont's Iron, Charcoal, and Lime industries*. Vermont Archaeological Society, Burlington.

Rosebrock, Adrian
2019 *Deep Learning for Computer Vision with Python*. 3rd ed. Vol. 3. Electronic document, https://web.archive.org/web/20210304041851/https://www.pyimagesearch.com/deep-learning-computer-vision-python-book/, accessed March 3, 2021.

Schallenberg, Richard H.
1975 Evolution, Adaptation and Survival: The Very Slow Death of the American Charcoal Iron Industry. *Annals of Science* 32:341–358.
1981 Charcoal Iron: The Coal Mines of the Forest. In *Material Culture of the Wooden Age*, edited by Brooke Hindle, pp. 271–299. Sleepy Hollow Press, Tarrytown, New York.

Schallenberg, Richard H., and David A. Ault
1977 Raw Materials Supply and Technological Change in the American Charcoal Iron Industry. *Technology and Culture* 18:436–466. DOI:10.2307/3103901.

Schmidt, Marcus, Andreas Mölder, Egbert Schönfelder, Falko Engel, and Werner Fortmann-Valtink
2016 Charcoal Kiln Sites, Associated Landscape Attributes and Historic Forest Conditions: DTM-Based Investigations in Hesse (Germany). *Forest Ecosystems* 3:Article 8. DOI:10.1186/s40663-016-0067-6.

Schneider, Anna, Florian Hirsch, Alexandra Raab, and Thomas Raab
2019 Temperature Regime of a Charcoal-Enriched Land Use Legacy Soil. *Soil Science Society of America Journal* 83:565–574. DOI:10.2136/sssaj2018.12.0483.

Schneider, Anna, Melanie Takla, Alexander Nicolay, Alexandra Raab, and Thomas Raab
2015 A Template-Matching Approach Combining Morphometric Variables for Automated Mapping of Charcoal Kiln Sites. *Archaeological Prospection* 22:45–62. DOI:10.1002/arp.1497.

Straka, Thomas J.
2014 Historic Charcoal Production in the US and Forest Depletion: Development of Production Parameters. *Advances in Historical Studies* 3 (2):104–114. DOI:10.4236/ahs.2014.32010.
2017 Charcoal as a Fuel in the Ironmaking and Smelting Industries. *Advances in Historical Studies* 6(1):56–64. DOI:10.4236/ahs.2017.61004.

Straka, Thomas J., and Wayne C. Ramer
2010 History on the Road: Hopewell Furnace National Historic Site. *Forest History Today* Spring/Fall:58–62.

Svedelius, Gustaf
1875 *Hand-Book for Charcoal Burners*. John Wiley and Son, New York.

Tolksdorf, Johann Friedrich, Rengert Elburg, Frank Schröder, Hannes Knapp, Christoph Herbig, Thorsten Westphal, Birgit Schneider, Alexander Fülling, and Christiane Hemker
2015 Forest Exploitation for Charcoal Production and Timber since the 12th century in an Intact Medieval Mining Site in the Niederpöbel Valley (Erzgebirge, Eastern Germany). *Journal of Archaeological Science: Reports* 4(Supplement C):487–500. DOI:10.1016/j.jasrep.2015.10.018.

Walker, Joseph E.
1966 *Hopewell Village; A Social and Economic History of an Iron-Making Community*. University of Pennsylvania Press, Philadelphia.

Werbrouck, I., M. Antrop, V. Van Eetvelde, C. Stal, Ph. De Maeyer, M. Bats, J. Bourgeois, et al.
2011 Digital Elevation Model Generation for Historical Landscape Analysis Based on LiDAR Data, a Case Study in Flanders (Belgium). *Expert Systems with Applications* 38:8178–8185. DOI:10.1016/j.eswa.2010.12.162.

Whitney, Gordon Graham
1994 *From Coastal Wilderness to Fruited Plain: A History of Environmental Change in Temperate North America from 1500 to the Present*. Cambridge University Press, Cambridge.

Williams, Charles E.
2020 Counting Colliers: A Demographic Profile of the Forgotten Charcoal Makers of Clarion County, Pennsylvania, 1850–1880. *International Journal of Regional and Local History* 15:1–18. DOI:10.1080/20514530.2020.1745362.

Williams, Michael
1989 *Americans and Their Forests: A Historical Geography*. Cambridge University Press, Cambridge.
1990 The Clearing of the Forests. In *The Making of the American Landscape*, edited by Michael P. Conzen, pp. 146–168. Routledge, New York.

AUTHOR INFORMATION

Benjamin P. Carter ■ Department of Sociology and Anthropology, Muhlenberg College, 2400 Chew Street, Allentown, PA 18045, USA (bcarter@muhlenberg.edu, corresponding author)

Jeff H. Blackadar ■ Department of History, Carleton University, 400 Patterson Hall, 1125 Colonel By Drive, Ottowa, Ontario, K1S 5B6, Canada

Weston L. A. Conner ■ Office of Admissions, Lehigh University, 27 Memorial Drive West, Bethlehem, PA 18015, USA