TRANS-REGIONAL PERSPECTIVES ON AGRICULTURAL DEINTENSIFICATION IN THE COLONIAL ANDES THROUGH REMOTE SENSING AND AI-ASSISTED ARCHAEOLOGICAL SURVEY

By

James Robert Zimmer-Dauphinee

Dissertation

Submitted to the Faculty of the

Graduate School of Vanderbilt University

in partial fulfillment of the requirements

for the degree of

DOCTOR OF PHILOSOPHY

in

Anthropology

August 11, 2023

Nashville, Tennessee

Approved:

Steven Wernke, Ph.D.

Tiffiny Tung, Ph.D.

Markus Eberl, Ph.D.

Yuankai Huo, Ph.D.

ACKNOWLEDGMENTS

This research was made possible by an NSF Collaborative Research Grant from the Division of Information and Intelligent Systems: Deep Generative Modeling for Urban and Archaeological Recovery, (Steven A. Wernke, PI, Daniel G. Aliaga, PI, Rajesh Kalyanam, Co-PI, Ian Lindsay, Co-PI, Parker VanValkenburgh, PI.) for which I was a Student Trainee, as well as a Digital Globe Foundation imagery grant and the Vanderbilt Scaling Success Grant (Steven A. Wernke, PI.) Further support was provided by the Tarkington Family Scholarship, and the Vanderbilt Spatial Analysis Research Lab (SARL). In addition to this generous financial and infrastructural support, I would like to thank all of the wonderful people who made this research possible. Thank you:

To Crystal Jackson, Breana Thompson-Guy, Sandy Cherry, and Sarah Nagy, you make the world go 'round and nothing we do at Vanderbilt would be possible without you. Thank you so much for all of the continual support, logistics, reminders, and for making things work when I mess something up. Your kindness and generosity throughout this process have been unending and I am continually amazed by your ability to keep things running smoothly.

To my committee members: Steven Wernke, thank you for your continued support, guidance, and patience, and for providing a space for archaeologists with unusual perspectives and interests to grow and flourish. When I wandered into your office on a casual trip to Nashville, I never realized the adventures it would take me on. To Tiffiny Tung, for always fighting for your students, pushing us to be our best academic selves, and providing us with incredible opportunities. I will never quite understand how I ended up presenting in São Paulo my first semester at Vanderbilt, but it was an unbelievable experience. To Markus Eberl, your support, positivity, and guidance has been something that I could always rely on when I was feeling uncertain. Thank you for making a space that felt safe for me. To Yuankai Huo, your excitement about my project has been contagious, even on those days when I wasn't feeling it myself, and your unerring willingness to contribute time, resources, and even students to the cause has made this work possible.

To the many students who have worked on this project especially: Jiachen "Charlie" Xu, thank you for laying the groundwork to take the machine learning effort to the next level. And thank you to Junlin Guo for running with what I started and helping me to imagine new heights. Thank you also for your constant positivity, support, and reminders that I should be proud of the work I'm doing. To Andrea Gutierrez, Elizabeth "Ella" Wright, and Greta Cullipher, for spending countless hours drawing little boxes around buildings and corrals, tracing terracing, and waiting for images to load. I'm not sure you knew what you were signing up for, but I'm so thankful for your patience, persistence, precision and humor. Not only did you help create the data for my dissertation, you made it fun.

To my family: Mom and Dad (Susan and Matt), thank you for setting me on the path and supporting me all along the way. I have always known that no matter what, you are there to love me, encourage me, and be proud of me. The fingerprints of your curiosity, passion for learning, and creativity are all over this dissertation. Thank you for instilling it in me as well. My sibling Jacob, thank you for sharing with me bravery, persistence, and the best hugs. Thank you to Cynthia, for accepting me as one of your own and always making sure I have what I need to work, even a quiet writing place during Thanksgiving, and to Melissa, for your encouragement and reminders that perfection is for the next life.

To my found family (fellow graduate students, and friends): There are so many things to say to so many of you that to do it all justice would take more pages than the dissertation, so please accept this abbreviated version. Beginning with the graduate students who came before me, Gabriela Oré, thank you for paving the way with your friendship, guidance, and joy, and Carla Hernández Garavito for keeping us grounded with love. Phyllis Johnson for always reminding me that what I can do is more than enough, and Keitlyn Alcántara for always reminding me I am more than my work. Terren Proctor for COVID bubbles, indescribable dinners, and even better company. And Abel Traslaviña for introducing me to Peruvian food, showing me the ropes, and letting me drive that truck. To those who are still fighting the graduate school fight, thank you for trivia, board games, shared food, and community. You are what make the hard times worth having. Samantha Turley, thank you for the weekly family dinners, pizza nights, trading drafts, trading stories, and trading frustrations. Katie McCormack for gardens, nature walks, and biology facts. Sylvia Cheever for great soups, key-lime pie, and shared pets.

To my Spatial Analysis friends: Stacy Curry-Johnson, your guidance in satellite imagery processing was crucial to this whole project. Natalie Robbins, from organizing workers, to organizing files, to digitizing, to brainstorming, to venting. Was there any part of this process you weren't vital for? Thank you for all of it, and for working to help me land on my feet.

Thank you to all of you who have been a part of this process all along the way. Manuel Mamani for showing me the Colca and navigating the snow. Dennis Blanton for introducing me to archaeology and Heidi Altman for helping me figure out how to indulge my loves of Math and Anthropology simultaneously. Ken Kvamme, for pushing the envelope on what archaeology can be and showing me all of the incredible ways I can combine archaeology and technology. Thank you to the many Vanderbilt professors not on my committee who have supported me along the way, particularly Carwil Bjork-James, John Janusek, Monica Keith, Jake Sauer, and Xenofon Koutsoukos. Thank you to my cohort, Maya Krause, Kathryn Peters, Eugenia Zavaleta, and the whole Vanderbilt graduate student community. I love you all.

And finally, thank you to my partner in all things, Cate Faulkner. We did it, my love! None of this would have been possible without your extraordinary patience, and gentle support. We've sure learned a lot of things in the past 7 years, and I'm so proud of us for coming through it stronger than ever. It's been a wild (and not always fun) ride, but we've always known it will be okay because we do it together. I can't wait to start the next phase of our journey together.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

Introduction

1.1 Introducing Andean Agriculture

The western cordillera of the south-central Andes rises to elevations above 5000 meters above sea level less than 100 km from the shores of the Pacific Ocean. These steep and arid mountain slopes are punctuated by narrow river valleys which flow from the glaciated peaks to the Pacific Ocean to the west, and to Lake Titicaca to the East. At first, this rugged landscape does not seem well suited to intensive agricultural production for the Inka empire, the largest prehispanic empire in the Americas. However, Andean people transformed these slopes into productive agricultural landscapes through the construction of monumental-scale agricultural infrastructure. Kilometer-long canals fed (and often continue to feed) the terracing complexes constructed by Andean farmers that cover much of the valley walls and floor. The rapid elevation change also creates rapid changes in ecological zones, enabling farmers to cultivate a wide variety of crops within a few hours or days' walk. Cotton, coca, maize, quinoa, potatoes, and more crops are cultivated as one moves from the coast to the highland planes where herds of camelids are managed (Murra 1972).

By dispersing themselves throughout the valley Andean farmers and pastoralists took advantage of the diverse, vertically-stacked, ecological zones of the Andes (Murra 1972). Murra (1972) described pre-hispanic Andean social and economic systems as a ”Vertical Archipelago” wherein Andean people adapted to the ecological diversity of the Andes and utilized reciprocity amongst kinship groups to ensure access to crucial resources and maintain self-sufficiency. In this model, Andean farming settlements may have been dispersed, but they were also densely connected through complex systems of kinship. These kinship networks bound together communities that spanned the length of the valley, ensuring exchange and access to the diversity of resources they produced. Murra's model dominated

Andean scholarship for decades, however, it has since been problematized by scholars who observed its dependence on being able to recognize ethnicities in the archaeological record (Stanish 2003; Van Buren 1996). Rather, these scholars suggest that rather than supplying food and goods to the population at large, verticality likely functioned to produce food and goods that were useful in the maintenance of elite political power, whether local, Inka, or Spanish (Stanish 2003; Van Buren 1996). In this model, rather than connected by kinship networks, distributed settlements were connected by an extractive political economy. In any case, the resulting small and dispersed farming settlements characteristic of the pre-hispanic Andes lead D'Altroy (2014, p. 33) to describe the Inka state as being mostly a "rural society", though we will see that this form of ruralism is quite different than the isolated rural hinterlands produced by European colonizers (Kagan 2000).

1.2 Colonial Intrusions

In contrast to these small but densely connected settlements, the Spanish crown imagined densely populated, centralized, and independent cities as the epitome of culture. To the Spanish, urban centers were "the preferred arena for social and economic exchange, and the stage for political conflict and accommodation" (Hoberman and Socolow 1986, p. 3). Following centuries of practice during the Reconquest, cities were also a way of legitimizing royal claims to land and formed the lowest level of royal government (Kagan 2000, p. 26). They were seen as a way to bring order, peace, prosperity, and Christianity to new and inhospitable lands, and were therefore crucial to the growth of the empire. Indeed, these ideals became so deeply tied to the city, many Spaniards believed that *policía*,(the complex term used to refer to civic duty and proper life in a republic) could not exist without it (Kagan 2000, p. 27). Spanish colonizers, therefore, aspired to create idealized European settlements, which they believed would foster health, civilized behavior, and obedience in its colonized residents (Málaga Medina 1975, p. 9; Mumford 2012). These new cities were designed in regular gridded patterns with central plazas, churches, and a surrounding hinterland for agricultural production (Bethell 1984, p. 81; Kagan 2000; Matienzo 1967; Mumford 2012), defining the ideal Spanish colonial city. In the Andean region, these aspirations reached their apotheosis in the 1570s CE as Viceroy Francisco de Toledo forcibly resettled indigenous populations into gridded urban settlements known as reducciónes. This resettlement project, known as the Reducción General de Indios, was one of the largest forced resettlements in human history (Kagan 2000; Mumford 2012; Wernke 2007a).

Nevertheless, the Spanish colonists couldn't impose their will for urbanization with impunity. A fundamental paradox of Spanish colonization in the Andes is that the colonizers were dependent on Andean farmers for their own survival and prosperity. The food, goods, and wealth of colonial Peru were produced by Andean people following millennia-old Andean traditions to effectively manage an Andean landscape. Therefore, while the Spanish crown may have sought to transform the people and the landscape to match its own idealized structure of a Spanish colony, it disrupted existing systems at its own peril. Where Spanish policy and the practical realities of Andean life conflicted, the resulting tension demanded resolution.

Such a resolution certainly did not come from the upper levels of Spanish colonial administration. Toledan guidance with respect to the resettlement was often vague, and at times contradictory (Wernke 2013, p. 283). For example, Toledo wanted to distance (physically as well as metaphorically) Andean people from their religious lives, contexts, and practices, which he described as "demonic" (Wernke 2013, p. 283). To do so, he demanded that reducción towns be built far from existing settlements, which were to be destroyed and completely abandoned. However, to do so would have also forced farmers away from their agricultural lands, rendering the entire system of agricultural production and exchange untenable.

Aware of the tenuous position of the colony, Toledo's guidance also recommended minimizing any disruption to the local economy, in clear contradiction to the previous admonition (Mumford 2012, pp. 119–121). In such contexts, where general administrative policy is unclear and conflicts with practical demands, it likely fell to local decision-makers, both Andean and Spanish, to decide how to resolve the resulting paradox.

1.3 Street-Level Bureaucracy in the Colonial Andes

Paradoxes between high-level policy and on-the-ground realities are not unique to 16th-century colonial Peru. Modern governmental policies designed to create uniform ways of generating social order often conflict with the need for flexible responses to particular circumstances. The term street-level bureaucracy was coined to examine the ways in which modern public-service professionals such as social workers, teachers, or police officers shape public policies through their exercise of personal discretion in challenging environments (Lipsky 2010). As on-the-ground workers are forced to contend with the daily realites of the context in which they are working, they often resort to compromises or coping strategies which may result in their work conflicting with the governmental, administrative, or personal ideals they seek to promote.

These street-level bureaucrats, as the individuals who directly interface with the public they are intended to serve, discipline, or control, fundamentally define the relationship between citizens and the state (Lipsky 2010). Street-level bureacracy therefore claims that policy is not created at higher (often political) administrative levels and passively implemented by mindless workers. Rather, practical policy is created on the ground, as a part of a dynamic and contested process between local administrators and the citizens they encounter.

Applying the concepts of street-level-bureacracy to Spanish colonial contexts encourages us to understand the ways in which local and regional administrators may have been guided by the administrative ideals of the resettlement, while simultaneously creating a wide variety of street-level policies. Toledo's vague and conflicting guidance left copious room for discretion by local administrators charged with ensuring the execution of the resettlement. Indeed, it is questionable the degree to which local spanish administrators may have desireds to conform to the ideals of the resettlement espoused by Toledo and the Spanish Crown. Their own lives and prosperity were deeply tied to the degree to the goods produced by Andean peoples. They were therefore highly incentivised to preserve as much as possible the highly effective systems of production perfected by Andean farmers and the Inka empire over the previous hundreds and thousands of years. Likewise, indigenous Andean people, experienced at imperial politics under the Inka empire prior to the Spanish, and the primary agents of reducción construction, had little incentive to abandon their own systems of production and socialization at Spanish behest.

In Huamanga, (modern day Ayacucho), research suggests that full-time residence in reducción settlements was not strictly enforced in spite of the imperial policy. Rather farmers were allowed to maintain their primary residences outside of the town (Stern 1993, p. 90). Relaxed colonial enforcement of residency in reducciones could have been beneficial to both the colonists and indigenous communities, permitting the exploitation of an expanded portion of the rural landscape outside of the daily accessibility of the town. Further south, in the Arequipa Valley, Davies' (2014) historical research suggests that the Spanish effectively separated the colonial and indigenous economies, with the Spanish initially focused on vineyards (and later, wheat) as the city of Arequipa grew and demand for food increased. Meanwhile, the indigenous populations retained usufruct rights to lands both local and distant and tried to maintain traditional methods of exchange and subsistence Davies 1984.

Between Ayacucho and Arequipa is the Colca Valley, a semi-arid river valley with extensive terracing in the middle and upper valleys which has been the subject of some of the most intensive archaeological investigations into agricultural deintensification in southern Peru (Denevan 1986; Donkin 1979; Gelles 2000; Guillet et al. 1987; Treacy 1990b; Wernke and Whitmore 2009; Wernke 2013). Here, the local population did indeed (at least for a time) retain primary residences at reducción settlements, resulting in higher than average rates of agricultural abandonment (as we show in Chapter 4). As local agents made decisions about which fields to maintain or abandon, distance from colonial settlements became a fundamental concern, in addition to field productivity and reliability (Wernke 2013). In the Colca Valley, patterns of abandonment were largely an outcome of risk assessment, as higher elevation terrace complexes with higher frost risk were abandoned at higher rates than lower elevation terraces. However, a sliding scale of elevation versus distance was evident, wherein beyond certain thresholds, sheer distance from a reducción outstripped elevation as the prime driver of terrace abandonmentWernke 2013.

Nevertheless, the effects of street-level policy are evident, even in the Colca. As will be examined in greater depth in Chapter 5, Wernke (2013) has demonstrated that prehispanic community organization played a significant role in patterns of agricultural deintensifiction. The reducción settlement of Corporaque was positioned such that it was closer fields managed by the higher-status (right-side) communities, while fields managed by lower-status (left-side) communities were more distant, shaping differing patterns of field abandonment between local communities. This is powerful evidence that even when forced to move into new ”reduced” towns, indigenous communities retained a significant measure of control over street-level policy.

1.4 Urbanization and Ruralization

In spite of indigenous shaping of street-level policy, and whatever discretion was available to street-level Spanish bureaucrats, the resettlement had an undeniable effect on the existing agricultural infrastructure. The demographic pull of colonial settlements often made the required regular maintenance of the infrastructure untenable (Denevan 1986; Donkin 1979; Gelles 2000; Guillet et al. 1987; Treacy 1990b; Wernke and Whitmore 2009; Wernke 2013) and a mix of environmental risks and distance from reducciones encouraged a settlement pattern characterized by central cores with a proximate agricultural ruralized hinterland with zones of unproductive and mostly unpopulated land between. Therefore, as the Spanish sought to enforce a colonial logic of urbanization, they implicitly constructed along with it a corresponding rural landscape.

If the city was the moral and physical manifestation of order, effective administration, and civil society, the non-urban (*rusticus*) was seen as the opposite of these traits (Kagan 2000, p. 27). The insistence of Spanish authorities on enforcing urbanization therefore forms a twin process, which I term ruralization. Ruralization results from the Spanish attempt at reformulating pre-hispanic dispersed settlements into places that exist outside of society, governmental administration, and morality. Infrastructure, including agricultural fields and canals, that existed outside of the urban sphere were to be deprioritized and, ultimately, abandoned. In broad terms, the consequences of these transformations are known: exacerbated mortality of the indigenous population and agro-pastoral de-intensification; the mass abandonment of large swaths of agro-pastoral infrastructure.

While we understand these consequences at the local level, we lack knowledge of the contours of these processes at regional and trans-regional scales. To what extent did ruralization take hold across the western highlands of the south-central Andes, and what is the extent of the resulting abandonment of agricultural infrastructure? As ruralization generated vulnerabilities as indigenous communities faced new stresses to their longstanding arrangements of agro-pastoral logistics and logics, where and how were street-level policies able to alleviate these burdens? How might we be able to identify these variations in policies in surviving features on the modern landscape? Hypothesizing that there was not a uniform implementation of Spanish policy, or a single solution to conflicting demands, this dissertation examines patterns of agricultural field abandonment and their relationships to reducción settlements at a trans-regional scale. By measuring the landscape at this scale, it is possible to establish general understandings of the phenomena, against which we can compare the local circumstances, thereby providing a more complete picture of both.

Archaeological investigation of these processes at a trans-regional scale requires innovations in archaeological methodology. Excavation and pedestrian surveys provide a wealth of local and regional detail, but they are challenging to scale to dimensions covering tens-of-thousands of square kilometers (Alcock and Cherry 2016; Banning 2002; Wernke, VanValkenburgh, Zimmer-Dauphinee, et al. in press). It is therefore important to investigate the extent to which the information offered by these methods is applicable outside of their particular contexts. Archaeological survey through remote-sensing data such as satellite imagery enables us to explore patterns at these larger scales, though at the cost of much of the detail provided by local methods. However with this expanded scope, we can explore whether lessons learned in local contexts appear consistent with broader patterns and, perhaps more importantly, identify where they may be different, thereby driving new directions for fruitful local research. Recent advancements in computer vision and AI (Lecun et al. 1998; Voulodimos et al. 2018) make it possible for computers to aid researchers as they conduct such large scale surveys, making them more robust, consistent, and reliable.

This dissertation therefore seeks to provide two major contributions, one anthropological, the other methodological. Anthropologically, I seek to examine the transformation of the Andean landscape under Spanish colonial rule, and to demonstrate how the interaction between environmental constraints, imperial policy, and local actions shaped, distorted and transformed the twin processes of urbanization and ruralization in the western cordillera of the south-central Andes. Methodologically, I introduce the concept of an AI-assisted imagery survey, where AI tools are used to identify archaeological features that may have been missed during a fully manual survey, or would take too long to manually record. The researcher recursively edits, enhances, and refines the survey results, providing insights that cannot be accessed through automated tools, and editing errors in the data.

1.5 Reimagining Archaeological Imagery Survey: AI-Assisted Survey

This dissertation expands on the work of GeoPACHA, a "browser- based geospatial platform for discovering and documenting archaeological sites in the Andes through systematic visual survey of satellite and historical aerial imagery by a network of trained teams" (Wernke, VanValkenburgh, and Saito 2020). GeoPACHA researchers undertook a federated collection of research-driven projects, enabling teams to work independently to conduct imagery surveys to answer specific research questions. By looking at larger regional patterns, GeoPACHA hopes to be able to provide important comparative data to better understand the local variation that is missing in more localized research (Wernke, VanValkenburgh, Zimmer-Dauphinee, et al. in press, p. 2). Furthermore, Wernke et al. observe that past people also operated across scales, understanding, and local, regional, and trans-regional processes to interpret and control their own lives (Wernke, VanValkenburgh, Zimmer-Dauphinee, et al. in press, p. 2). Archaeologists must join their subjects in adopting multiscalar perspectives if we hope to develop a more complete understanding of past processes. Local, regional, and trans-regional approaches are therefore complementary, each providing critical perspectives that are missing from other scales of research.

In this spirit, this dissertation proposes and employs AI-assisted remote sensing as a new method of archaeological imagery survey for quickly investigating trans-regional scale areas while maintaining a high level of data quality and local detail. This approach complements and expands so-called "Brute Force" manual techniques which rely on teams of highly trained archaeologists to inspect very large areas by hand (Casana 2020:95) and identify archaeological features. In applications such as the detection of archaeological structures, this can help dramatically reduce the amount of time archaeologists must spend examining the majority of the landscape that does not include archaeological features. Not only does this increase efficiency by reducing the amount of work required to record the features in an area, it also helps researchers maintain attention, decreasing the rate of missed features. For applications where the extent of the features are important, such as the inventory and mapping of agricultural infrastructure, this dissertation shows that AI-assisted surveys can precisely map the boundaries of the features, a laborious and time-consuming process beyond merely marking the presence or absence of a feature.

While such digitization is possible by hand, maintaining high levels of detail for large spatial extents is challenging due to fatigue and time constraints. AI-assisted surveys allow trained researchers to generate high-resolution data for a small fraction of the area of interest, and use that data to produce maps for very large areas at a higher level of detail than would be practical by hand. In each case, researchers then inspect the results of the AI model, editing features where necessary, removing false positives, and adding insights that could not have been produced by the AI model alone, thereby gathering higher quality data in less time than would be possible with brute force methods alone.

Such promises to dramatically reduce the work and improve the efficiency of archaeological practice should properly be met with skepticism. Much of the discourse surrounding the use of technology in archaeology is about making archaeology faster, or more efficient. However proponents of ”slow archaeology” argue that important aspects of archaeological knowledge result from the experience of practicing archaeology (Caraher 2016). It is feared that this focus on efficiency often seeks to automate processes that were previously arduous and time-consuming but may do so at the expense of the experiential knowledge that accompanies tedious archaeological work (Caraher 2016). Casana (2020) echoes this concern, observing that automated surveys are incapable of identifying unique features (features of a kind previously unobserved) or of recognizing and recording the complex relationships that exist between archaeological features. For example, the significance of a canal or roadway connecting to settlements of archaeological importance (or even the lack of such a canal or roadway) can only be intuited by an archaeologist who has spent extensive time observing many different combinations of settlements, canals and roadways during many hours of tedious investigation.

I agree with these concerns. From my perspective, attention to the importance of experiential knowledge gained during data collection combined with a dedication to deeply exploring and enhancing the data produced through automated techniques is what distinguishes an AI-assisted survey from an automated one. AI-assisted surveys place human researchers at the center of the survey process, utilizing their unique skills where they are most valuable, to identify subtle relationships between features, and evaluate edge-cases where automated models struggle. In this way, AI-assisted survey methods encourage researchers to seek out and carefully consider edge-cases and places where it is difficult to determine precisely how to characterize the features identified. The greater attention is paid to these features, the better the AI will perform.

This leads naturally to a recursive process, where the researcher carefully considers their features of interest and generates data to train the AI (Davis 2020), the AI directs the researcher towards similar features, and the researcher evaluates the result. This evaluation quickly demonstrates to the researcher shortcomings in their definitions and examples: ”I haven’t included enough examples of round buildings, the model is biased towards square ones.” ”This terracing looks very different to that of the next valley over, that’s why the model is struggling to identify it. Also, I wonder why it’s so different?” The researcher must now identify further examples of features of interest (and features of non-interest) which improves the AI performance and suggests yet more locations for deeper investigation. The AI-assisted survey process, therefore, requires a deep engagement with the data, rather than the ”deskilling” which is a common concern with automation (Caraher 2016). Modern AI models are data-hungry, requiring thousands of training examples to learn effectively. The design of this data demands the researcher develop a clear understanding of the breadth of possible features they wish the AI to identify. It is for this reason I find it vitally important that any researcher conducting an AI-assisted survey spend substantial time generating the training data themselves. It is not enough to simply hire others to create the data. Otherwise it is impossible to grasp the subtleties between a cluster of corrals and a set of walled agricultural fields where alpacas are grazing (or indeed the arbitrariness of the distinction). Without this experience it is impossible to faithfully interpret the data produced by the AI model, and to grasp the ambiguity that may exist therein, in the same sense an archaeologist cannot understand the subtleties of an excavation map if they have not themselves dug an excavation unit (Caraher 2016).

AI-assisted survey therefore should not be seen primarily as a way of reducing effort or tedium. The research for this dissertation contained both of these in multitudes. Rather, the AI-assisted approach is a method for directing attention where it can be most productive, enabling the researcher to provide the deeper insights that are unique to human capacity while highlighting variability in the landscape by where the model fails, and thereby encouraging the researcher to challenge assumptions of uniformity.

1.6 Structure of the Dissertation

This dissertation consists of six chapters, including four articles produced in the course of this study. These articles cover the application of automated survey techniques to identify archaeological features in satellite imagery, AI-assisted satellite survey methods, the distribution of active and abandoned agricultural fields, and the dynamics of agricultural de-intensification in the highlands of the western cordillera of the south-central Andes. Though composed of independent articles, the work is cumulative, building on the findings and lessons learned from the previous articles.

Chapter 2 presents and evaluates the results of an automated survey for archaeological structures (defined as structures without evidence of modern roofs or maintenance less than 30 m in diameter) in satellite imagery. Furthermore, it discusses the importance of sampling in model construction and evaluation, showing the effects of spatial-autocorrelation on evaluations of model performance. Accounting for this effect is vital to properly measure the quality of data created by automated models and build trust in their results, yet is rarely discussed in archaeological efforts at automation. The results of this paper are therefore foundational to the credibility of the results in chapters 3 and 4.

Chapter 3 introduces the framework of AI-assisted satellite surveys while comparing the results of the automated survey for archaeological structures to data collected by GeoPACHA, a ”brute force” satellite survey project undertaken by a team of trained archaeologists. Systematic manual inspection is the gold standard for satellite archaeological survey, however this paper demonstrates that both manual and automated surveys missed a similar percentage of visible features. Importantly, each survey identified features that were missed by the other, though the automated survey was prone to false-positives, which must be removed by human researchers. I therefore propose an AI-assisted approach for future research, where both human researchers and AI models search for features which are then evaluated for accuracy and enriched with important contextual data by the researcher. These edited data can then be iteratively supplied to the AI model to further enhance its abilities, allowing researchers to obtain high quality data while limiting errors caused by fatigue.

Chapter 4 applies an AI-assisted survey approach to the identification and classification of agricultural fields in the western cordillera of the South-Central Andes, measuring and characterizing the distribution of currently in use (active) and abandoned agricultural fields. The data is validated through comparison to independently digitized test data, manual inspection, and comparison to very large scale regional surveys of agricultural terracing undertaken by the Peruvian Ministry of Agrarian and Irrigation development. The final results are a trans-regional inventory of active and abandoned fields at a significantly higher resolution than was reasonably achievable by hand at that scale. This data demonstrates that common estimations of the rate of field abandonment by archaeologists are often much higher than the actual rate at a regional level. I suggest that this is potentially the result of an archaeological sampling bias towards regions where settlements and their associated fields are especially abandoned. Finally, exploratory analyses are completed to understand the relationship between active and abandoned agricultural infrastructure and various environmental covariates in the western valleys of the study region and Titicaca basin.

Chapter 5 brings together data from the Linked Open Gazetteer of the Andean Region (LOGAR) (Wernke and Saito 2019), GeoPACHA (Wernke, VanValkenburgh, Zimmer-Dauphinee, et al. in press), and the agricultural field inventory completed in Chapter 4 to examine the factors influencing Andean decisionmaking in the construction and abandonment of agricultural fields and placement of reducción settlements in the river valleys of the western cordillera of the south-central highlands. Expanding on previous work in the Colca Valley, this work explores whether local patterns identified in this deeply studied region are applicable to the broader region and searches for evidence of regional variation. Key findings show that elevation, slope, and distance from existing fields were important factors in the decision of where to place reducción settlements, suggesting that indigenous populations likely drove the selection of construction sites. Furthermore, elevation, slope, distance from streams and distance from reducción settlement played important roles in agriculturalist decisions to maintain or abandon agricultural fields for the region, demonstrating the additional burden imposed by the resettlement. Finally, preliminary local analyses show evidence of variation in the importance from the distance to reducciones for patterns of agricultural abandonment, suggesting that the political and practical realities of the resettlement were handled differently in different regions.

Chapter 6 synthesizes the results of the previous papers, examining the regional agricultural transformations shaped by the Reducción as well as the promise offered by new AI-assisted approaches to large-scale regional survey. It also proposes promising directions for future research and expansion on these successes to further enhance our understanding of regional and trans-regional variation.

CHAPTER 2

Progress and Prospects in Automated Image-Based Archaeological Survey: An Assessment from the Southern Peruvian Andes

Authors: James Zimmer-Dauphinee, Steven A. Wernke

2.1 Introduction: Evaluating Automated Surveys

Interest in automated survey of remotely sensed data for archaeological features has grown markedly in recent years (Davis 2020). Enthusiasm for AI-based feature detection in satellite and aerial imagery is understandable, as it holds the potential to dramatically expand scales of analysis to interregional- and even continental-scale views of archaeological distributions. Yet assessing AI-based feature detection accuracy is key to establishing reliable AI-based imagery surveys. To do so, the results of automated survey techniques must be compared to independent data sets collected by more traditional means such as systematic visual survey of satellite imagery. The practice of withholding a randomly selected partition of the data for the purposes of model evaluation is standard practice in machine learning workflows (Tan et al. 2021), however, this practice must be further examined with remotely sensed geographic data, where spatial dependency also must be considered in model evaluation. Spatial autocorrelation (systematic variation as a function of distance) suggests that adjacent image tiles cannot be treated as independent samples (Miller 2004). Rather it should be expected that adjacent tiles are correlated to each other, merely due to their proximity, resulting in overconfidence in the model if adjacent or spatially proximate tiles are placed in both the training set and the data sets reserved for model evaluation.

This paper forms the first step towards a credible, reliable, and thoroughly tested automated archaeological survey in the south-central Andean Highlands. To evaluate the effects of spatial autocorrelation on model evaluation, a convolutional neural network is trained and evaluated on data from the western cordillera of the southern Peruvian Andes that has been treated in two different ways. First, the data was split into training and validation data sets using a naive random sampling strategy. As a result, some evaluation tiles are spatially adjacent to the image tiles used for training. The data is then split a second time, with steps taken to ensure spatially proximate tiles are not split between the training and test sets. Evaluations of the model performance are then compared to examine the effects of spatial autocorrelation. This comparison of training data and model validation methods demonstrates that failing to account for spatial autocorrelation in model training data can lead to a substantial overestimation of model performance. Beyond this cautionary tale, this paper contributes a workflow for producing reliable training data and estimates of model performance for AI-assisted archaeological imagery survey.

2.2 Archaeological Context: Applications in the Peruvian Southern Highlands

This project seeks to explore the feasibility of leveraging large-scale imagery-based autonomous archaeological survey to enable regional- to interregional scale analysis of settlement patterning and land use through the transition from Inka imperial rule (ca 1450-1532 CE) through the Spanish invasion and colonization of the Andes. It is situated in the western cordillera of the southern Peruvian Andes (Figure 2.1). Here, the cordillera rises from the hyperarid coast through vast expanses of high-altitude steppe (puna), which is in turn punctuated by glaciated peaks and stratovolcanoes and bounded in the east by the Lake Titicaca Basin (Sandor 1992). Pacific drainages cut through the puna, producing montane valley oases of high potential agricultural productivity (Sandor 1987). By late prehispanic times, these highland valleys were home to large ethnic polities, which through Inka imperial expansion in the 15th century CE were variously incorporated as provincial administrative units (Neira Avendaño 1990). Population and agro-pastoral production in most localities in the region reached their apogee under Inka occupation (Doutriaux 2004; Linares Delgado 1989; Neira Avendaño 1990; Chavez 2019; Wernke 2006).

These processes were truncated abruptly with the Spanish invasion. Following an early

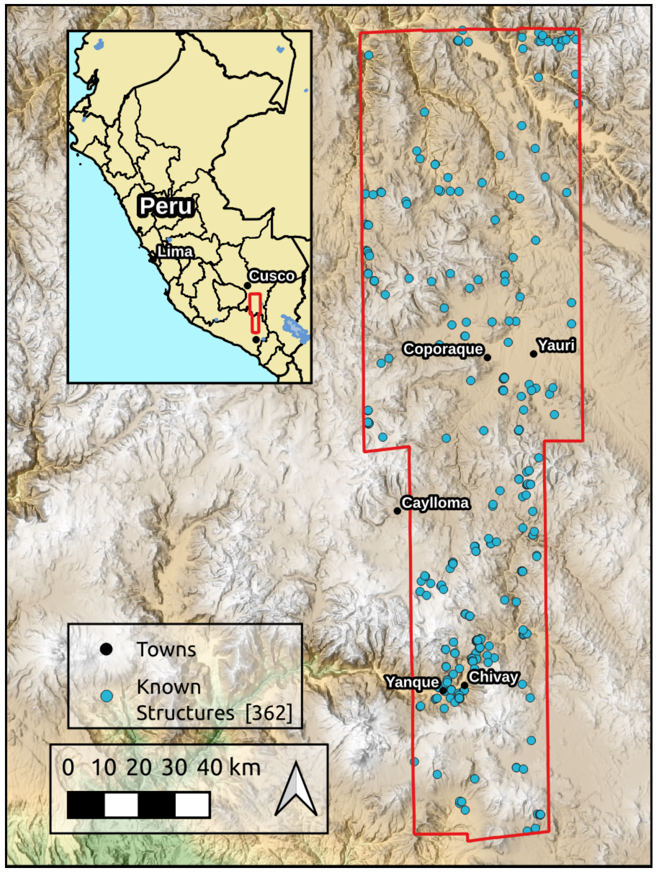

Figure 2.1: Autonomous Archaeological Survey Region with locations of positive image tiles used in model construction.

period of plunder, violence, indirect rule, and internecine conflict among Spanish factions, the Spanish Crown dispatched the Viceroy Francisco de Toledo in 1569 to establish a stable system of governance, evangelization, and economic extraction from the Viceroyalty of Peru (Cook and Cook 2007; Mumford 2012). The centerpiece of Toledo's reforms was a General Inspection and with it the Reducción General de Indios—the General Resettlement of Indians—one of the largest forced resettlement programs by a colonial power in world history. The General Resettlement forcibly displaced over 1.4 million indigenous people from their traditional homes into compact, gridded towns known as reducciones de Indios or literally "Indian reductions." The goals were to increase Spanish control and surveillance of indigenous daily life, improve the efficiency of taxation, and enforce Spanish concepts of urban living that they saw as crucial to civilization and Christianization

(Mumford 2012; Wernke 2007a,b). The broad contours of what resulted are clear enough: the immiseration of large segments of the native Andean populace through the disarticulation of links between community, landscape, and longstanding logistics of production (Mumford 2012; Stern 1993; Wernke and Whitmore 2009). In an arid environment where mobility and dispersion of settlement among vertically distributed ecological zones was a long-standing feature of Andean community and land-use arrangements, the centralization of people into nucleated settlements had the effect of displacing many away from their productive lands, and the close living quarters of reducciones exacerbated the transmission of European-introduced diseases (Cook 1992, 1998).

Nevertheless, the dynamics of this process at large scales remain poorly understood. Historical data and local archaeological research suggest that Spanish influence and power were differentially implemented across the landscape, sometimes with counterintuitive consequences (Davies 1984; Wernke 2007b, 2013). For instance, the city of Arequipa, the southern seat of Spanish colonial power, was dependent on agricultural surpluses both for its own sustenance and the production of wealth through regional exchange (Boza Cuadros 2021; Davies 1984). It seems incongruous, then, that a mere 30 km to the southeast of Arequipa vast complexes of agricultural terracing and their associated irrigation systems, along with the major late prehispanic settlement of Churajon were abandoned shortly following the Spanish invasion (Szykulski 1998, 2008; Szykulski et al. 2000). Simultaneously, agricultural settlements more distant from the city appear more likely to survive or prosper. This project represents a first step toward collecting the data to understand the processes that drove this penumbra of de-population and agricultural deintensification in the outskirts of one of the seats of colonial power.

A first step in understanding these processes is to identify the settlement patterns of both reducciones and prehispanic settlements, to better model colonial dislocation and its effects on agricultural productivity, urbanization, ruralization, and economic exchange systems. These changes in settlement patterns and interactions occurred on a tremendous scale over the expanse of the viceroyalty as a whole, encompassing hundreds of thousands of square kilometers. Even regionally, the scale of resettlement and its effects surpasses areas that can be effectively studied through field methods. To address this problem, we seek to generate machine learning models capable of identifying archaeological settlements (via the identification of relict buildings) in high-resolution satellite imagery, allowing us to rapidly catalog archaeological features at scales that are otherwise difficult to achieve. This approach complements and extends recent efforts using "brute force" (manual inspection) techniques of satellite imagery surveys (Casana, Kantner, et al. 2014; Casana and Cothren 2013; Parcak 2019; Wernke, VanValkenburgh, and Saito 2020). In particular, the area of interest overlaps with one of the manual imagery surveys conducted as part of the GeoPACHA (Geospatial Platform for Andean Culture, History, and Archaeology) project (Wernke, VanValkenburgh, Zimmer-Dauphinee, et al. in press). The project discussed here functions as a pilot study for broader implementation using large datasets of manually-tagged features from GeoPACHA as training data.

2.3 Materials and Methods: Deep Learning and Remote Sensing

Deep Convolutional Neural Networks were used on satellite imagery to test the feasibility and understand the potential challenges of automatically identifying archaeological features in satellite imagery. Imagery from Digital Globe's WorldView I and WorldView 2 satellite constellations (now owned by MAXAR), covering approximately 11,500 square kilometers between the cities of Arequipa and Cuzco, was used to produce and test models to survey for archaeological structures utilizing transfer learning on a CNN with a ResNet-50 backbone. The imagery was radiometrically, atmospherically, and geometrically corrected, and processed using a coarse elevation model as a DigitalGlobe Standard 2A level product by the DigitalGlobe Foundation (Pour et al. 2021). It should be noted however that the coarse elevation model is insufficient to remove all topographic distortions, some of which are apparent in the data. The imagery was mosaiced and pan-sharpened by the researchers, (a standard preprocessing procedure which combines the panchromatic and multispectral images to generate high-resolution color images) and the Red, Green, and Blue (RGB) bands were extracted for analysis. Future research will leverage more spectral bands to improve the results further. Finally, the images were scaled from 32-bit to 8-bit to reduce the computational storage and processing power requirements. ImageNet, the dataset used for initial training of the ResNet-50 model used in this analysis, is composed of 8-bit RGB images (Deng et al. 2009). The above treatment of the data, therefore, brings it into closer alignment with the data used to pre-train the ResNet-50 model.

For analysis, the imagery was chipped into 76.8 x 76.8m tiles (256 pixels at 0.3m resolution, an efficient size for computational processing), for a total of 1,946,106 tiles. Training deep learning models requires large amounts of imagery data and, generally, the more data, the better the results. However, this must be balanced with the time required for collecting training data. Ultimately, a total of 5245 labeled tiles were selected for model training and testing. An initial 5,000 tiles were randomly selected for manually coding the presence or absence of archaeological structures, (defined as round or rectangular structures, no more than 30m in their largest dimension that appear abandoned and/or appear to lack modern roofing and maintenance). Of the 5,000 tiles, approximately 1% (n=47) yielded examples of archaeological structures. This data was then augmented with an additional 308 tiles containing archaeological structures to provide the model with sufficient examples of structure presence for training and testing. Finally, tiles that fell on the edge of the image, and so did not form a complete tile, were removed to produce the final dataset.

The labeled image tiles were then split into training, validation, and testing sets. This is a standard practice in machine learning that has been practiced for decades, however, recently researchers have encouraged a closer examination of the assumptions of random splitting (Tan et al. 2021). In our case, the data split was performed in two ways to test for the effects of spatial autocorrelation on the model's performance, and the model was trained twice, once for each division of the data. The first split (Split A) was via naive random sampling of image tiles. That is, individual tiles were placed in the training, validation, or testing set, and no tile was placed in multiple sets. This method fails to account for spatial autocorrelation (the tendency for proximate phenomena to be more similar than distant ones) because it may split neighboring tiles in the training and validation or testing sets. The presence of spatially-autocorrelated features in both the training and validation data limits the ability of model evaluations to accurately describe the performance of the neural network which may be vulnerable to overfitting on the image background, rather than the features of interest.

The data was then re-split between the training, validation, and testing sets (Split B). To avoid spatially proximate tiles being divided between sets, the second split took a stratified random sample which defined any tiles within 400m of each other as being a part of the same "locus." The resulting loci were split with 80% of loci placed in the training set, 10% in the validation set, and 10% in the testing set. All tiles from each locus were placed in the same data set. This prevents neighboring tiles from being split between sets and reduces the effects of spatial autocorrelation on the model's confidence. The resulting number of training, validation, and testing samples are very similar between the two splits, shown in Table 2.1.

Table 2.1: Number of positive and negative examples in each data split

	Split A		Split B	
Training	Positive: 292	Negative: 3904	Positive: 298	Negative: 3919
Validation	Positive: 27	Negative: 497	Positive: 24	Negative: 486
Testing	Positive: 36	Negative: 489	Positive: 33	Negative: 485

Finally, for each training dataset, the training data was augmented using techniques including vertical and horizontal flips, random rotations, and transformation to black/white images, a standard process for improving the results of computer vision models (Hauberg et al. 2016; Simard, Steinkraus, and Platt 2003). This artificially increases the size of the data available for training which in turn can reduce overfitting and improve the model's ability to respond to a wide variety of positions, rotations, and color combinations of archaeological structures in the imagery.

The open-source Raster Vision (Azavea 2022) framework was used to extract tiles from the imagery, train the deep learning model, and compare the model's results to the validation set. The commonly used ResNet-50 convolutional neural network, pre-trained on more than a million images from the ImageNet database was selected as the backbone architecture for our model. We then transfer-trained the model on our own dataset for 20 Epochs with a Batch size of 20 images using a Learning rate of 0.0001.

2.4 Results: The Importance of Sampling

Initial results of the model using the randomly selected tiled dataset (Split A) appeared promising when compared to the validation set, achieving an approximately 88% recall (percent of true positives identified) and 82% precision (percent of predicted positives that are true positives) for a balanced accuracy of 93%. However, when the model was applied to the survey area, a manual review of 1000 randomly selected tiles that the model coded as containing archaeological structures showed these promising metrics to be inaccurate. With an 82% precision, we would expect 820 of the "positive" tiles to contain archaeological structures. However, only 320 of the extracted predictions actually contained such structures, resulting in a true precision of approximately 32%. This suggests that the model's performance on the validation set in Split A is not representative of how it performs on all satellite tiles in the region and that the model was overfitting to the validation set, despite not having access to the validation set for training. This is likely due (at least in part) to spatial autocorrelation, as will be discussed further below. Furthermore, a manual review of the data suggests that the model frequently confused modern architecture with archaeological structures, suggesting that modern architecture may not be sufficiently represented in the data (Figure 2.2).

Split B was designed to mitigate the effects of spatial autocorrelation. This way, the model was no longer able to use the similarity between spatially proximate tiles to assist

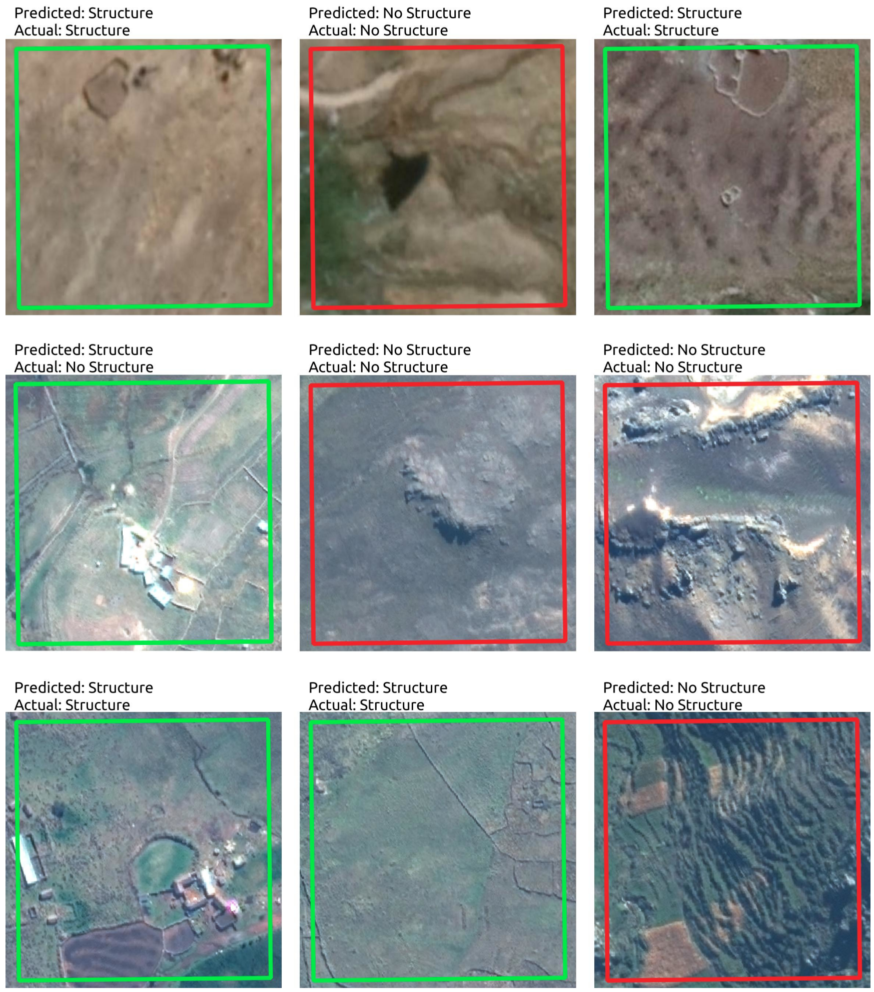

Figure 2.2: Selected Images with Autonomous results and Ground-truth labels. Boxes are 76.8m x 76.8m (256 pixel x 256 pixel at 0.3m resolution)

in its predictions and, the model was much less effective at correctly identifying tiles with structures in the validation dataset. For Split B, the model produced a much lower (and more realistic) recall of 54% and precision of 93% for a balanced accuracy of only 77%. All results are shown in Table 2.2. This suggests that the validation set is now more accurately representing the model's performance in the real world. However, such results are insufficient for application to archaeological work, suggesting that further model improvements are necessary to produce a useful model. We are now pursuing those improvements, the results of which are forthcoming (Zimmer-Dauphinee, VanValkenburgh, and Wernke in press). The insights from our experience here are well worth conveying, since to our knowledge, issues of spatial autocorrelation in the training and validation steps have not been discussed systematically in the literature on machine learning and archaeological satellite remote sensing.

Table 2.2: Evaluation metrics for each data split and precision compared to the manual inspection of Split A

	Split A	Manual Inspection of Split A	Split B
Recall	0.88	Not Evaluated	0.54
Precision	0.82	0.32	0.93
Balanced Accuracy	0.93	Not Evaluated	0.77

2.5 Discussion: Lessons Learned

Accurate evaluation of automation techniques is crucial to their application in archaeological contexts. A model with poor precision fails at sufficiently reducing the manual labor of settlement identification, as it is still necessary for experts to review many false positives to identify the true positives. A low recall rate (one that produces a high rate of Type II errors) such as the result of the second model is perhaps even more concerning as it risks missing important archaeological evidence and gives a false impression that an absence of evidence does indeed imply evidence of absence. Accurately measuring the model's effectiveness on previously unseen, "real world" data is therefore vital to understanding whether and how a particular model is useful.

This research, therefore, highlights the need for careful research design and rigorous evaluation of automated model results and suggests fruitful paths forward for selecting spatially unbiased training data samples. The naive random sampling of image tiles for partitioning into training and validation datasets produced a poor representation of the results the model would have if applied to the "real world." This poor representation may result from a variety of sources. From a machine-learning perspective, archaeological data sets tend to be relatively small. With only 6,000 labeled image tiles, and 1000 validation tiles, it is possible that the validation set did not have sufficient representation of the landscape to accurately depict the results of the model when applied to it. This problem is exacerbated by the highly imbalanced nature of the data, with only approximately 1% of the tiles containing positive examples naturally. Therefore, even with oversampling of tiles with archaeological structures, there were very few (only a few dozen) examples of archaeological structures in the validation dataset. This may not be enough to truly evaluate whether the model is successfully capturing all archaeological structures.

The infrequency of modern buildings in the landscape may also have been a challenge to the accurate evaluation of the model. While more common than relict archaeological structures, the vast majority of the landscape is not covered by modern architecture, resulting in relatively few examples in the validation dataset. However initial models show a strong bias towards false positives when evaluating whether a tile containing modern structures also contained archaeological structures. This is unsurprising. Due to the similarities in size, shape, and material of modern and archaeological structures, it is very likely that the model will have difficulty distinguishing archaeological and modern structures. Furthermore, many modern settlements are located in close proximity to archaeological settlements, particularly in the highlands where pastoralists and modern ranchers make use of the same persistent and limited ecological resources. This means that some tiles contain both archaeological and modern structures, making it further difficult for the model to disambiguate between the two. Therefore, locations with many modern structures such as modern towns or cities are likely to be misidentified. If the number of modern structures in the validation set was low due to the relatively small sample size and random chance, the validation dataset may not have captured sufficient examples of modern structures to adequately test the model's efficacy of rejecting these buildings as archaeological structures.

However, the large discrepancy between the precision calculated using the naive random sampling validation set and the precision evaluated post hoc from applying the model to large-scale satellite imagery suggests that insufficient examples of archaeological and modern structures in the training and testing data are not the only source of concern. Rather, it appears that the model is overfitting to the naive random sampling validation dataset, despite not having trained on this data. This is possible because of spatial autocorrelation. While none of the image tiles in the naive random sampling validation dataset were used to train the model, in some cases neighboring tiles were split between the training and validation sets by random chance. As a result, the model was trained on tiles that were very similar to those in the validation set, and overfitting in the training process was not captured in the validation dataset. Therefore, the model generalized to previously unseen (and less proximate) data more poorly than expected from the validation set's results.

The solution to this problem is to establish a stratified random sampling method for partitioning the dataset into training, validation, and test sets. Previous research in the Andes has used 100m as a minimum distance to designate a boundary between separate archaeological loci (Wernke, VanValkenburgh, and Saito 2020). Following this metric and using geographic information systems, it is possible to cluster any tiles within a 100 m radius from each other as belonging to the same locus. Then, rather than partitioning tiles, one can randomly partition loci between training and evaluation datasets and reduce the effects of spatial autocorrelation on evaluations of the model's performance. The resulting training, validation, and test sets retain an approximately 80,10,10 split while retaining all spatially proximate tiles within a single dataset.

The evaluation of the model using the stratified random partitioning of tiles by locus is a much more accurate depiction of the model's performance, though that performance is somewhat less impressive than hoped. The greatest barrier to deep learning success in the location of archaeological architecture is the extreme imbalance in the number of images that contain an archaeological feature, and those that don't. Clearly, the percentage of the landscape devoid of surface archaeological features discernable in satellite imagery is much larger than the percentage that contains them. In fact, given the random sampling strategy for our initial 5000 image tiles, it appears that for the region in question, approximately 1% of the landscape appears to have some form of archaeological features on it. Working with such fundamentally unbalanced datasets is quite challenging, and an area of active research throughout the machine learning industry.

2.6 Conclusions and Future Research

Convolutional neural networks and satellite imagery are promising tools for identifying archaeological structures in the Andes, but further work is needed to improve model accuracy before they can be deployed at large scale with confidence in their performance. Given the challenges presented by the highly unbalanced nature of the dataset and the limited number of samples available for training, more sophisticated modeling approaches are needed to increase the balanced accuracy to a feasible level for archaeological applications. Ongoing work suggests that a modified SimSiam Representational model (Chen and He 2020; Chen, Kornblith, et al. 2020) can achieve balanced accuracy sufficient for archaeological applications by leveraging large unlabeled datasets in the modeling process.

In any case, achieving a model capable of identifying the location of archaeological sites is only a first step. Even preliminary models suggest that tens or hundreds of thousands of archaeological structures will be identified, scattered across thousands of square kilometers covering the western cordillera of the Andes. To fully utilize data at this scale and begin asking meaningful archaeological and anthropological questions, one needs more than mere location information. Settlement characterization and classification will be crucial. Even simple classifications such as whether a structure or settlement is more likely associated with pastoral or agricultural activity may allow researchers to begin exploring the relationships between these activities in the Andes on a trans-regional scale.

It should also be acknowledged that “structures” (as defined by this paper) are not the only features of the Andean landscape that are of archaeological interest. Landscape features such as the ubiquitous agricultural terraces and canals are vital to understanding the production of food, wealth, and power in the Andes. It has been estimated that as much as 1,000,000 hectares of land in Peru has been terraced (Guillet et al. 1987). In the Colca Valley, it is estimated that 61% of this once-productive agricultural land has been abandoned (Denevan 2003; Donkin 1979; Treacy 1990a) and in northern Chile, that number reaches as high as 80% (Wright 1962). The causes and effects of this abandonment in localities such as the Colca Valley have been thoroughly debated, yet the applicability of these discussions and the potential of local variation across the broader landscape remain understudied. Comprehensive mapping of active and abandoned agricultural terracing from satellite imagery using machine learning as conducted in Chapter 4 will provide crucial information for developing a global understanding of these processes, and direct future research to better understand local variations in communal decision-making. Chapter 5 of this dissertation begins these explorations.

Finally, it bears reiteration that satellite surveys, whether by hand or by machine, are only capable of identifying features that are visible within satellite imagery. Such methods will not capture archaeological materials that are obscured by vegetation such as may be common in certain regions, in the North and East of Peru. Nevertheless, similar approaches utilizing regional scale LiDAR data have proven effective at transforming archaeological understandings of settlement and economic patterns in Cambodia and the Maya lowlands (Canuto et al. 2018; Chase et al. 2012; Guyot et al. 2021; Verschoof-van der Vaart and Lambers 2019). Neither of these data sources will be able to capture more transient structures which lack surficial remains to be observed from the air, nor archaeological features which may never have had structures to begin with. For these reasons, further comparison to pedestrian survey samples is vital to evaluate and account for the bias that is introduced into our understanding of the archaeological record through primarily remote sensing and computational approaches.

Trans-regional approaches promise to change our understanding of interactions at scale, and to better highlight regional and local variation across the broader landscape, however, such approaches should be seen as effective tools to broaden our understanding of the scale and scope of the archaeological record, not as replacement methodologies for more traditional archaeological approaches.

CHAPTER 3

Eyes of the Machine: Toward AI-assisted Satellite Archeological Survey in the Andes

Authors: James Zimmer-Dauphinee, Parker VanValkenburgh, Steven A. Wernke

3.1 Introduction: The Potential of AI-Assisted Surveys

Archaeological surveys conducted through the manual inspection of high-resolution satellite imagery hold transformational promise for regional and supraregional research. Traditional pedestrian surveys covering many thousands of square kilometers can be prohibitively expensive and time-consuming, especially when conducted in physically challenging landscapes such as the Central Andean Cordillera. In contrast, manual satellite prospection allows archaeologists to efficiently survey hundreds of thousands of square kilometers for visible features and to generate reliable data (Casana and Cothren 2013; Parcak 2017; Ur 2013; Wernke, VanValkenburgh, and Saito 2020). Furthermore, it allows regions that were once geographically or politically inaccessible to be investigated through the inspection of freely available satellite imagery (Casana and Cothren 2013).

Nevertheless, manual satellite survey methods are labor-and time-intensive, and the low frequency of archaeological features in many landscapes can make for monotonous work. Far from being a trivial issue, observer fatigue dramatically increases the likelihood of false negatives (Körber et al. 2015), meaning that important archaeological features may go unrecorded. Reliable satellite survey also requires surveyors to have extensive training in both archaeology and satellite image interpretation (Casana 2020, p. 95), increasing the cost of inspecting large areas without archaeological features. Facing these challenges, researchers have sought to automate satellite prospection to cover large areas more quickly, completely, and efficiently (Davis et al. 2021; Lambers, Verschoof-van der Vaart, and Bourgeois 2019; Lasaponara and Masini 2007; Somrak, Džeroski, and Kokalj 2020; Trier, Cowley, and Waldeland 2018; Zingman et al. 2016).

Despite their potential benefits, there are reasons to be skeptical about the potential of fully automated approaches to replace manual archaeological imagery survey. Archaeological features missed by automated detection (false negatives) may contain vital information, while locations incorrectly identified as archaeological loci (false positives) may mislead the researcher as they examine large-scale patterns. Automated surveys are also incapable of capturing unique features—or of recognizing the complex relationships that exist between archaeological features (Casana 2020).

Here, we propose an approach that seeks to avoid some of these downsides. Rather than fully replacing manual with automated methods, we promote an "AI-assisted" approach that employs deep learning to augment the results of manual surveys by directing the surveyor's attention toward locations that are most likely to contain archaeological features. To evaluate its potential, we compare the results of a deep learning model for identifying archaeological structures in the southern Peruvian highlands to data collected and edited by a team of experts through the Geospatial Platform for Andean Culture, History, and Archaeology (GeoPACHA). Our results suggest AI-assisted survey provides an additive check on features of interest and reduces the burden of examining survey grids devoid of visible archaeological features, which in the case of Geopacha approached 95% of grid cells. We envision a near-term future in which human-machine teaming approaches will surpass the sensitivity and specificity rates of either approach when deployed exclusively. In turn, these results can be further augmented and verified, through on-the-ground pedestrian surveys.

3.2 Human vs. Machine: Comparing Brute-Force and Automated Methods

3.2.1 Manual (Brute-Force) Methods

Manual satellite surveys have been referred to as "brute force" techniques, in contrast to automated surveys (Casana and Cothren 2013; Wernke, VanValkenburgh, and Saito 2020). Some brute force methods, such as Albert Lin's search for Genghis Khan and Sarah Parcak's GlobalXplorer utilize "citizen scientists" to quickly survey large areas and obtain measures of confidence in their findings through repeated observations by non-specialists (Lin et al. 2014; Parcak 2019). Such approaches are suited for generating presence/absence information, but non-specialists cannot leverage contextual knowledge to record detailed metadata (Casana 2020). Presence/absence data may flag areas of interest for further investigation by specialists, but without specialist intervention, applications are limited to broad and atemporal settlement pattern analyses. Other research programs, such as the CORONA Atlas project and GeoPACHA (Casana and Cothren 2013; Wernke, VanValkenburgh, and Saito 2020), rely on trained specialists to ensure the quality of the collected data and provide deeper and more rich metadata about the observations made. To date, these methods have proven highly effective in identifying archaeological features over vast areas but remain time-consuming and require team members to have significant training and domain knowledge. They are also subject to human vigilance constraints (attention span)—a factor that may be especially significant in large-scale, extended imagery survey campaigns.

Research on human vigilance has repeatedly demonstrated, across diverse domains, that even trained specialists can be poor at detecting "rarely occurring, low-signal-to-noise-ratio signals embedded in the context of varying background configurations" (Harris 2002; Körber et al. 2015; Shingledecker et al. 2017). The distribution of archaeological features in the southern Peruvian highlands follows such a pattern: GeoPACHA surveyors identified archaeological features in only 5% of survey grids (Wernke, VanValkenburgh, Zimmer-Dauphinee, et al. in press). In such circumstances, even well-trained experts may miss visible archaeological features due to lower vigilance from distraction, boredom, and fatigue. Automated survey methods may help to alleviate these problems by reducing time spent examining locations that lack features and refocusing surveyors' attention on possible locations of missed features.

3.2.2 Automated Methods

Traditional remote sensing approaches to automated object detection/classification have relied on pixel-wise spectral comparisons to identify features of interest (Comer and Harrower 2013; Parcak 2009). However, archaeological features vary widely in their construction materials or (as in the case of the Andes) are constructed of similar materials to the background landscape, resulting in little to no spectral difference between features of interest and their surroundings (Alexakis et al. 2009; Garrison et al. 2008). This has led to the adage that there is ”no generalized spectral signature for archaeology,” though there may be spectral signatures for particular cases or archaeological correlates (Agapiou et al. 2013; Lasaponara and Masini 2007; Parcak 2009, 2017; Saturno et al. 2007). In contrast, state-of-the-art computer vision models such as Convolutional Neural Networks (CNNs) and Vision Transformer (ViT) Models examine the correlations between proximate pixels, enabling the detection of morphological variation at the scale of the objects of interest, in addition to spectral variation at the pixel level (Sevara et al. 2016). These technological improvements open up new opportunities for detecting archaeological features using computational methods. Furthermore, rapid improvement in computer vision in the past decade has allowed models in select fields to meet (and in a rare, but growing number of cases, to surpass) specialist capabilities (Bewes et al. 2019; Buetti-Dinh et al. 2019; Byeon et al. 2019; Dodge and Karam 2017; He et al. 2015). Archaeology may similarly benefit from these methods.

Even with such improvements, there remain concerns with automated surveys. Current, state-of-the-art models in archaeological site detection are more akin to citizen scientists than to trained specialists. False positives/negatives are persistent problems, and the best models only provide location and morphological information with little to no metadata concerning context, or relationships between features. Such results may be useful for guiding broad analyses or to assist human specialists, but may miss important dimensions necessary for detailed analyses.

In addition to practical concerns with automated techniques, conceptual problems must be addressed. Due to their complexity, deep learning algorithms are ”black boxes” (Latour and Woolgar 1979) that are impossible for researchers to completely understand. If one cannot concretely explain why a given model makes any particular evaluation about the presence or absence of archaeological features, one may doubt the trustworthiness of the information the model produces. Fortunately, computer science researchers in the field of Explanatory AI (XAI) have provided researchers with a series of tools such as Grad-CAM (Selvaraju et al. 2020), Layer-Wise Relevance Propagation (Montavon et al. 2019) and others (Rai 2020; Tjoa and Guan 2021) to evaluate model behavior at the level of individual predictions as well as the global results of the model. Furthermore, this concern is not unique to data produced through automated means (Davis 2020, p. 3), humans are also black boxes. A researcher working with data collected via brute force survey by volunteer citizen scientists or even trained professionals may struggle to delineate an explicit hermeneutics of archaeological imagery survey, which is a complex interplay of visual cues, prior knowledge, and decision-making resulting in the evaluation of the presence or absence of archaeological features.

Nevertheless, some aspects of archaeological survey remain inaccessible to computational methods. As recognised by Casana (2014, p. 228), a human researcher is “engaging in a discursive, analytic process, thinking creatively about features we see” and is capable of “identifying unique or unusual features.” Neither of these vital characteristics is possible using automatic techniques. In contrast to automated identification of the presence or absence of a feature, archaeologists are searching for what such presences and absences mean concerning their relationships to each other and the archaeological record. Given this fact, researchers cannot relinquish our ability to think analytically and creatively about the features we see. Humans must remain active participants in the acquisition and analysis of archaeological data and take the time to examine the features we identify in their larger context.

Despite these caveats, we argue that automated surveys hold great potential. The automated identification of archaeological features at regional and supraregional scales provides contexts for deciding where best to focus human creativity and analytical thinking. Much like a trowel, shovel, or backhoe, automation tools are not appropriate in all circumstances, but AI-Assisted surveys may be invaluable to address questions at a supraregional scale.

3.3 Data Sources

3.3.1 Satellite Imagery

For our preliminary automated survey, satellite imagery from DigitalGlobe's WorldView 2 and WorldView 3 satellite constellations (since purchased by Maxar) were used to produce and test models to detect archaeological structures. The imagery, covering approximately 11,500 square kilometers between the cities of Arequipa and Cuzco (Figure 3.1) is a Standard 2A level product, radiometrically and atmospherically corrected, and geometrically pre-processed using a coarse elevation model by DigitalGlobe (Pour et al. 2021). The images were then mosaiced, pan-sharpened using the OrfeoToolbox Bayesian Fusion algorithm, (Grizonnet et al. 2017) and scaled from 32-bit to 8-bit to reduce computational storage and processing requirements. While it is possible that the increased bit-depth could provide additional information for model training, traditional machine learning datasets such as ImageNet are composed of 8-bit Red, Green, and Blue (RGB) images. Therefore, to facilitate transfer learning on pre-trained networks the RGB spectral bands were extracted to bring the images more in line with these data (Deng et al. 2009).

3.3.2 Archaeological Data for Automated Survey

Archaeological features of interest in the automated survey were relict buildings (Figure 3.2), defined as round or rectilinear structures, no more than 30m in their largest dimension, lacking modern roofing and maintenance. Using the QGIS Geographic Information System (QGIS Development Team 2009), a grid of 76.8m x 76.8m squares (256 x 256 pixels at 0.3m resolution) was generated covering the imagery. Ultimately, a total of 6,428

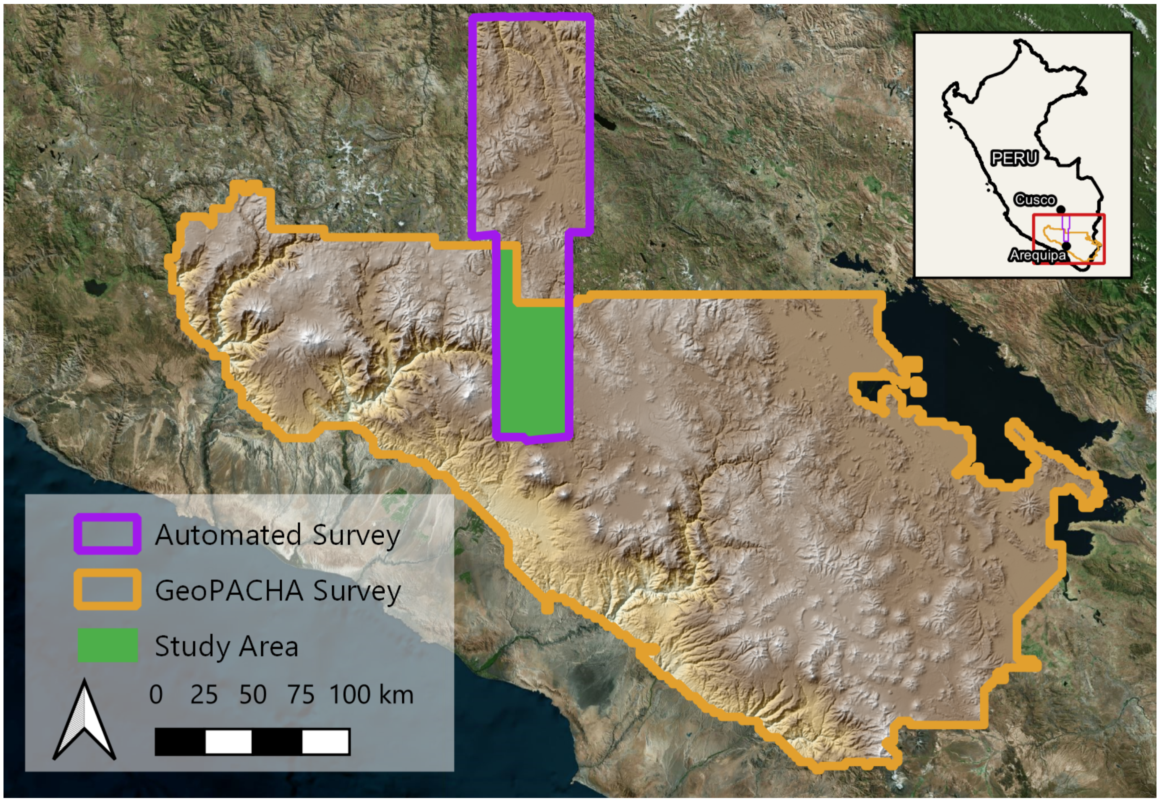

Figure 3.1: GeoPACHA and Automated Survey Study regions

labeled tiles were selected for model training and testing. Initially, 5000 squares were randomly selected for manual inspection. The presence or absence of archaeological structures was recorded for each square. Of these, approximately 1% (n=47) yielded examples of archaeological structures. This data was augmented with an additional 308 squares containing archaeological structures previously known to the authors to provide sufficient positive examples for model training and validation. Initial modeling efforts revealed that modern structures were often confused with archaeological structures, so an additional 1,183 squares containing modern structures were added as negative examples to the dataset. Removing squares that fell on the edge of an image, or where the imagery was obscured by clouds, resulted in a total of 6,428 labeled squares for model training, validation, and testing. Finally, image chips corresponding to the label squares were extracted using Raster Vision (Azavea 2022).

Figure 3.2: Examples of archaeological structures used for training the automated survey model. Note that archaeological structures may be isolated or associated with other archaeological or modern features.

The labeled image chips were split into training (n structures=298, n negatives=3919), validation (n structures=24, n negatives=486), and testing (n structures=33, n negatives=485) sets. Any chips within 400m of each other were grouped and the resulting groups were split with 80% of groups placed in the training set, 10% in the validation set, and 10% in the testing set. This ensured that all chips from within a single locus, or located in closely neighboring loci, were not split between training and validation sets, thereby minimizing the effects of spatial autocorrelation on model evaluation. Finally, the training data were augmented using vertical and horizontal flips and random rotations. This increases the number of images available for training to reduce overfitting and improve the model's ability to identify archaeological structures in any orientation. The validation set was used to monitor and evaluate the training progress. However, as one modifies model hyperparameters to optimize performance the model becomes fitted to the validation set. A separate test set was therefore reserved for independent evaluation of future versions of the model.

3.3.3 GeoPACHA Data

The automated survey results were compared to those acquired through brute force methods via GeoPACHA. As described by Wernke and colleagues, GeoPACHA is a "browser-based geospatial platform for discovering and documenting archaeological sites in the Andes through systematic visual survey of satellite and historical aerial imagery by a network of trained teams" (Wernke, VanValkenburgh, and Saito 2020). In this case, GeoPACHA data from a survey area overlapping with our automated survey area offer an independent means of comparison of the two methods. The GeoPACHA database currently contains 36,248 recorded loci. The south-western survey zone, edited by Wernke, covered 78,372 km² and recorded 14,685 loci with attribute data. This dataset was collected entirely independently from the automated survey data described above and provides an independent test of the model for the approximately 3,000 km² where the two surveys overlap, in which GeoPACHA identified 844 loci. (Figure 3.3) This data should therefore be both sufficiently high-quality and high-quantity to evaluate the capabilities of the automated survey.

3.4 Methods: Building a Model

3.4.1 Deep Learning (DL)

Innovations in DL in recent decades have dramatically improved the capabilities of computer vision models (Voulodimos et al. 2018). However, training such models from scratch requires datasets of thousands or millions of labeled example images. Such datasets are often beyond the scale available for archaeological data. Fortunately, a method known as transfer-learning makes it possible to repurpose models that have been pre-trained on standard datasets such as ImageNet (Deng et al. 2009) to work on new problems with much smaller data sets. This research uses transfer-learning on the well-known ResNet-50 (He et al. 2015) computer vision model, trained on ImageNet, to classify satellite imagery chips for the presence or absence of archaeological features.

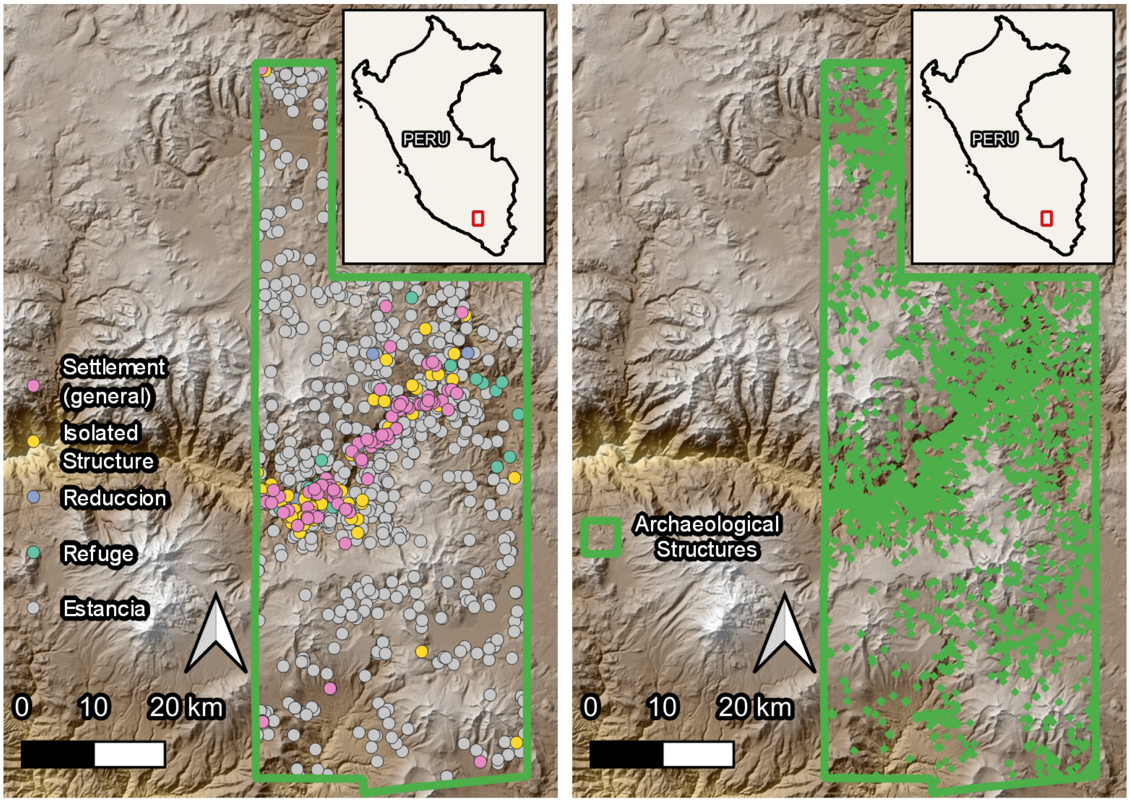

Figure 3.3: Data produced by the GeoPACHA and Automated surveys for the region of overlap overlaps with the automated survey. Note the similarity in distribution between the surveys.

We used an open-source framework for deep learning on remote sensing data, known as Raster Vision (Azavea 2022), to train and deploy our automated survey. The training data was used to transfer-train a ResNet-50 model using the standard Adam optimizer for 20 epochs (full passes through the training data), with a learning rate of 0.0001 (for rapid and stable optimization), and a batch size of 20 images (to speed up computation). The model was checked against the validation set to evaluate overfitting. Methodological details, code, and model parameters are included at github.com/geopacha.

3.4.2 Making Data Comparable

The GeoPACHA dataset enables comparison of the results of the automated survey directly with high-quality, independently produced data. However, some translation is necessary to make the datasets comparable. The automated survey algorithm marks 76.8 x 76.8m chips that contain archaeological features, while the GeoPACHA data records a single point for each locus, located near the surveyor's estimation of the center of the locus. In the GeoPACHA dataset, all features within 100 m of each other are designated as a part of the same locus. Therefore, a locus with two components separated by a space of 80m, would be marked in the space between the components, where there are no archaeological features for the automated survey to find. Furthermore, differences in preprocessing between the Worldview imagery and the imagery sources (Bing, Google) used for GeoPACHA produced differences in loci locations. To account for these differences, any chip identified by the automated survey as containing an archaeological feature which was located within 100m of a GeoPACHA locus was designated as having identified that locus.

Of the 844 loci identified in GeoPACHA for the survey area, 391 locations marked in GeoPACHA were not identified by the automated survey. A manual review of these 391 loci found 102 which were not visible in the satellite imagery used for the automated survey. These loci were either obscured in the automated imagery dataset, (due to clouds, shadows, or destruction) or had been mapped from their record in published works (and labeled as such in GeoPACHA). These loci were not visible in the Worldview imagery, and so could not be identified by the automated survey. Excluding these loci from the analysis left 742 loci in the GeoPACHA dataset that are also visible in the Worldview imagery for the automated survey region. There are several ($/sim20$) features identified in the GeoPACHA dataset that do not meet the working definition of archaeological structures used in the construction of the automated survey. These are large-scale constructions such as fortifications or corrals that were clearly of archaeological relevance but did not contain structures smaller than 30m in dimension. While it is expected that the model would not identify these features because the training data was not designed to do so, these features were not excluded from the following analysis because there are indeed visible archaeological features at these locations which were not identified by the automated survey. As discussed above, one of the advantages of human surveyors over automated methods is their ability to engage in the creative process of identifying features of archaeological interest that fall outside narrow definitions. The choice to penalize the model for not identifying visible archaeological features for which it was not designed reflects this shortcoming and offers a more realistic picture of working with automated survey data. While the model can (and will) be re-trained to identify such features, the following analysis reflects its current capabilities and is therefore a conservative estimate of the automated survey's potential performance.

3.5 Results: Comparing Deep Learning to Brute-Force Surveys

The automated survey successfully identified 453 visible loci that were also identified in the GeoPACHA dataset. In these locations, GeoPACHA and the automated survey were in agreement and we have high confidence in the presence of visible archaeological features. This leaves 289 loci visible in the GeoPACHA data that the automated survey failed to identify. However, in addition to the loci identified by GeoPACHA, the automated survey also identified 1031 other locations as having archaeological structures that were not identified in the GeoPACHA dataset. A manual review of these locations found that 315 locations did indeed contain structures that met the definition of an archaeological structure used in this analysis. This brings the total of known loci in the area identified through the combined surveys to 1,057. Of these, GeoPACHA surveyors identified 742 ($recall \approx 70\%$), while the automated survey identified a comparable 766 ($recall \approx 73\%$). The results are summarized in Figure 3.4. Of course, there may be more archaeological features visible in the imagery which were not identified by either the GeoPACHA or the automated surveys, however, we expect that the number of unidentified visible features in the region is very small.

While the two surveys had a similar recall rate, GeoPACHA surveyors far outperformed the automated survey in terms of precision. It is difficult to evaluate how many of the 742 loci identified by GeoPACHA surveyors were incorrectly marked as archaeological features without visiting these features on the ground. However, following our inspection

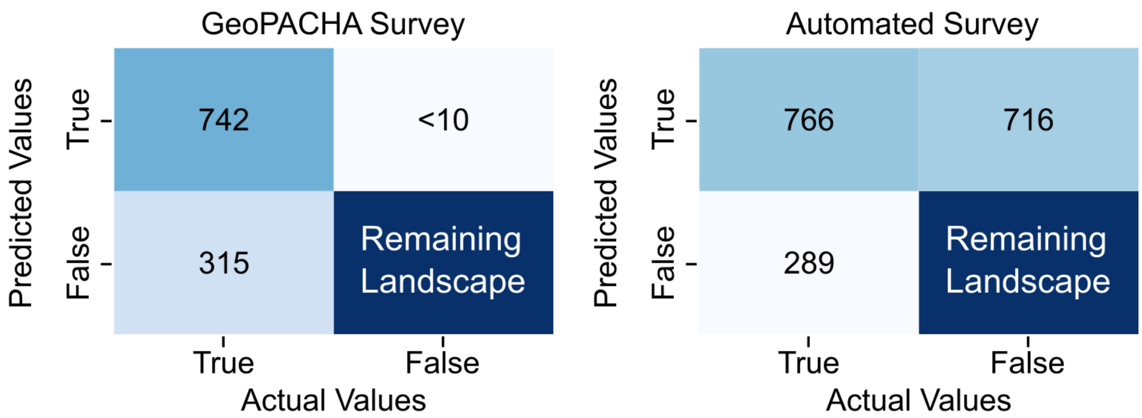

Figure 3.4: Confusion matrix for the GeoPACHA survey shows decent recall and very few false positives, while the automated survey had a decent recall but many false positives.

of the data, we expect that the number of falsely identified loci is negligible, likely less than 1%. In the automated survey, there were 716 false positives wherein the model mistakenly identified modern structures, rocks, or other natural formations as archaeological structures. Therefore, 48% of all the loci identified by the automated survey were false positives. As a result, the F1 score (used for evaluating model performance) for the automated survey is $\sim$61% while GeoPACHA surveyors achieved an F1 score of $\sim$82%.

3.6 Blending Automation and Brute-Force: Introducing the AI-assisted Survey

3.6.1 Evaluating False Negatives

We found it surprising that the team of expert GeoPACHA surveyors and the automated algorithm had similar recall rates, with each missing approximately 30% of the visible archaeological features identified between the two surveys. Human vigilance constraints may have played a role in producing this pattern. Most of the features identified in the study area are estancias (pastoralist corral and residential complexes), which often have both modern and archaeologically relevant components. GeoPACHA surveyors for this region were asked to record an estancia if it appeared to contain components that lacked roofs or evidence of maintenance for part or all of the estancia complex. After viewing many modern estancias that did not contain unmaintained structures, researchers may have been less likely to observe archaeological features when they were present due to fatigue or uncertainty about whether a structure met the standard of "unmaintained."

Regardless of the underlying factors, the relatively low recall rate for both surveys suggests that future analyses must account for biases caused by missing data, regardless of whether the data comes from an automated survey or expert surveyors. For more rigorous analyses, it may be necessary for teams to re-survey areas that have already been inspected to ensure that as complete a dataset as possible has been collected, establish quantifiable metrics for how much data is likely missing, and evaluate biases. Comparisons to terrestrial surveys are also necessary to account for the archaeological features that were missed because the features were not visible from the satellite imagery. Fortunately, archaeologists are accustomed to these sorts of considerations as archaeological data is always partial, contingent, and uncertain due to taphonomic and sampling biases.

3.6.2 Evaluating False Positives

Our results suggest that while the automated survey was not subject to fatigue or uncertainty, it was also less discerning. That is, the automated survey identified some archaeological features missed by GeoPACHA surveyors, but it often misidentified modern structures and complexes as archaeologically relevant loci, and it occasionally misidentified rocks, trees, and other "natural" landscape features as archaeological features. This issue might be addressed by increasing the size of the training dataset; enhancing data augmentation; and using all eight available spectral bands of the satellite imagery to provide the algorithm with more information about the materials of the identified features.

In any case, the problem with false positives will likely persist as the unbalanced nature of the data is a fundamental problem. CNNs tend to perform best when there are similar numbers of positive and negative examples, while archaeological features are uncommon on the landscape. Therefore, it remains necessary for human researchers to manually check the results of even highly reliable automated models to eliminate false positives. Rather than surrendering control to an algorithm, using AI to assist in the survey process appears optimal, with humans acting as the ultimate arbiters of presence or absence, and identifying unique or unusual features the model may have missed.

3.7 Conclusions

While manual and automated methods have been counterpoised as mutually exclusive approaches to archaeological imagery survey, our experiences suggest they can be productively combined along a spectrum of human-machine teaming approaches. In this paper, our AI-assisted approach is shown to have identified several hundred more archaeological loci in the study area than either manual or automated survey alone. These results suggest that when automated methods are used to focus and expand human interpretation, rather than replace it, they can make satellite surveys more robust, consistent, and reliable. Post hoc comparison between different survey methods, such as the one presented above, allows for independent evaluation of results and can help to highlight shortcomings in data produced by different approaches. However, because it requires multiple surveys to be completed independently of each other, it also demands excess work and provides little immediate benefit to the surveyors themselves. A more productive alternative is to use automated and manual survey methods in tandem, to improve results and relieve some of the time and fatigue burdens on expert surveyors. We envision three particularly promising approaches for AI-human teaming in archaeological imagery surveys. First, automated methods can be used for "low-probability filtering" – identifying regions with no clear archaeological features and allowing researchers to exclude them to focus on areas that have a higher probability of containing features of interest. This approach would not only reduce the overall time it takes to survey, but also help the researcher to maintain vigilance as they will discover features of interest more frequently, and suffer less fatigue. Alternatively, automated survey can be employed as "quality control", that is, as a secondary check during the survey process, identifying areas that the researcher should review again before marking a location as not containing features of archaeological interest. One of the great advantages of this second technique is that human researchers maintain control over the process throughout, enabling them to record the kinds of complex metadata that are still beyond automated approaches. Third, and perhaps most excitingly, the automation/manual survey process can take a “recursive teaming” approach, with manual survey data feeding back into the automated survey algorithm to improve its results, which can then guide further manual survey. The model analyzed above was trained using a mere 298 positive examples of archaeological structures, and the simplest and most powerful way of improving deep learning algorithms is to provide them with more training data. GeoPACHA contains over 36,000 loci with archaeological features. Using this data as training data, could dramatically improve the capabilities of the automated system, allowing it to better guide manual surveyors to identify features missed in their initial efforts, and to expand to new areas. This loop can be condensed further: human-guided incremental machine learning enables models to improve dynamically as new data is added, rather than having to repeatedly train from scratch (Gil et al. 2019; He et al. 2015; Yang and Tang 2020)(Gil et al. 2019; He et al. 2015; Yang and Tang 2020). Such approaches allow the researcher and the automated system to work together, simultaneously improving the results of both. Indeed, this is the path we are pursuing going forward. Advances in computer vision techniques in the last decade have made it important to revisit their usefulness in conducting archaeological surveys using satellite imagery. In some domains such as medical imaging, automated search/classification algorithms are approaching or surpassing human capacity for identifying features of interest. However, their application in archaeology is just beginning. Our research demonstrates that DL methods, when used in tandem with “brute-force” manual surveys, show great promise for large-scale regional and supraregional surveys.

CHAPTER 4

Chapter 4: AI-Assisted Satellite Imagery Survey and Trans-Regional Evaluation of Field Abandonment in the South-Central Andes

Authors: James Zimmer-Dauphinee, Andrea Gutierrez, Mary Wright, Greta Cullipher, Grecia Roque, Parker VanValkenburgh, Steven A. Wernke

4.1 Introduction: Trans-regional Agriculture

Terracing is one of the defining human constructions in the Andes. Predominantly constructed for agricultural purposes, terracing's primary benefit is its facilitation of irrigation (Treacy 1990a), though it offers a wide array of other physical benefits including the creation of level ground, slowing erosion, and even moderating temperature at high altitudes (Donkin 1979; Morlon 1996; Treacy and Denevan 1994). In addition to shaping the physical landscape, the prehistoric construction of large-scale terrace complexes shaped the social landscape through the communal labor required for their construction and continual maintenance and management (D'Altroy 2014; Janusek and Kolata 2004; Kolata 1991; Langlie 2018; Van Buren 1996). Spanish colonists sought to commandeer this physical and social infrastructure as a foundation for wealth and power. Following the Spanish invasion, much of the monumental-scale agricultural infrastructure fell into disrepair and was abandoned. Substantial research has been done on this process at the local level, particularly in the Colca Valley (Denevan 1988; Guillet et al. 1987; Treacy 1987, 1990b,a; Wernke 2007b, 2010, 2013). However, post-invasion regional and trans-regional dynamics of Andean agricultural infrastructure remain much less documented and understood, in significant measure because of the sheer scale of the systems and the difficulties involved in mapping them at scale. (Erickson 2000; Stern 1993). This has led to a lack of consensus amongst archaeologists with widely varying estimates of the proportion of agricultural infrastructural abandonment between 40% (Denevan 1988) and 80% (Wright

1962). Government-sponsored studies have tended to place the rate of infrastructural abandonment much lower, around 15% to 25% AgroRural and CERA 2021.

This research addresses the scalar challenges for registering a trans-regional inventory of terraced field systems through an AI-assisted satellite imagery survey of currently in-use and abandoned agricultural fields in the southwestern Peruvian highlands. Recording agricultural data at this scale, we offer a trans-regional evaluation of the rates and distribution of terrace abandonment, expanding on local research to evaluate its generalizability and identify potential locations with alternative historical pathways. To do so, we use a sample of high-accuracy manual survey data to train a Convolutional Neural Network (CNN) to segment high-resolution (Worldview 2 and Worldview 3) satellite imagery into three classes: active agricultural fields, abandoned agricultural fields, and no agricultural fields, then deploy the trained model over 81,000 km^2 of the south-central Andes.

This work complements and builds on "brute force" (manual) surveys of satellite imagery such as those conducted by GeoPACHA (Wernke, VanValkenburgh, and Saito 2020; Wernke, VanValkenburgh, Zimmer-Dauphinee, et al. in press) and other satellite archaeological surveys (Langlie 2022), or the Peruvian Ministry of Agriculture (AgroRural and CERA 2021). This AI-assisted approach affords the expansion of analyses to new regions and imagery across the western cordillera of the Andes and deployment on new imagery as it becomes available. This has implications not only for archaeologically understanding processes of the past at very large scales, but will also contribute to understanding the effects of climate change on sustainable agricultural development in the present and future. By automating the bulk of the labor burden in generating inventories of agricultural fields, the inventory can quickly be repeated on a periodic basis to track changes in agricultural infrastructure over time. The region's aridity and dependence on irrigation for agricultural production implies that the extent of agricultural fields will be strongly correlated to water availability from glacial melt (Guillet 1992; Guillet et al. 1987). Tracking the expansion and contraction of agricultural fields over time may therefore provide important insights into glacial dynamics as they are affected by global warming.

Following manual data collection and model training, we deployed the model on imagery covering approximately 70,000 km² of the western, south-central highlands of Peru. Comparison to existing datasets from GeoPACHA show that the model was able to map agricultural features with substantially more detail than was achieved by hand at that scale, and even identified terracing systems that were missed in a large-scale inventory of terracing by the Peruvian Ministry of Agrarian and Irrigation development, making this the most detailed, accurate, and complete inventory of agricultural fields in the western south-central Peruvian highlands to date. Finally, we present exploratory spatial analyses to characterize the ecological contexts in which agricultural terracing was built and where it continues to be used or has been abandoned.

4.2 Archaeological Context

The research conducted in this study focuses on Southwestern Highlands study region defined by GeoPACHA, a "browser-based geospatial platform for discovering and documenting archaeological sites in the Andes through systematic visual survey of satellite and historical aerial imagery by a network of trained teams" (Wernke, VanValkenburgh, and Saito 2020). This region encompasses diverse regions including the western Titicaca Basin and several Pacific drainages including the Colca, Arequipa, and Moquegua Valleys (Figure 4.1). The high and arid valleys of the southern Peruvian highlands stand out as some of the most intensively terraced landscapes in the Andes, and are dependent on irrigation from streams and rivers to ensure consistent yields of a diverse variety of crops, with a focus on maize, quinoa, and potatoes (Guillet et al. 1987; Treacy 1990b). In contrast, the terraced fields in the Titicaca Basin predominantly rely on rainwater for irrigation, while simultaneously benefiting from temperatures moderated by the large body of water(Smith, Denevan, and Hamilton 1968). The Titicaca floodplain also provides extensive areas with minimal slope, where agriculture thrives, often working to limit excessive water, rather than bringing distant water in (Kolata 1991; Stanish 1997). These two regions therefore have very different dynamics, which we begin to explore in this chapter.

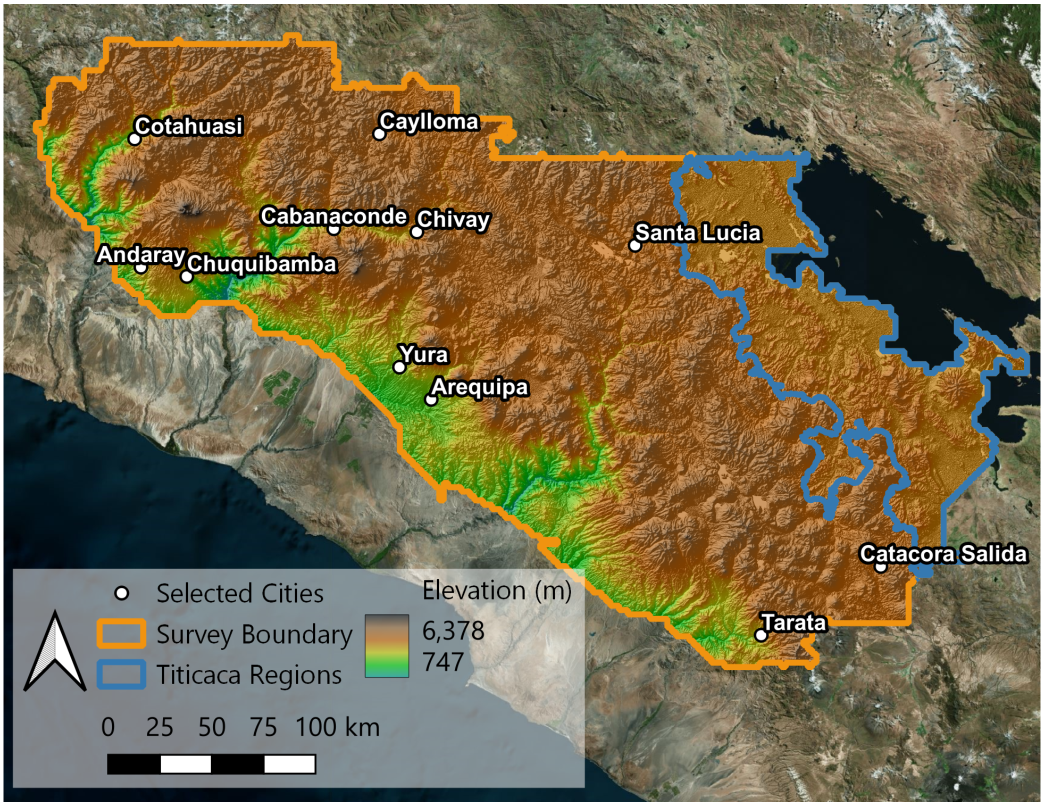

Figure 4.1: Study region, including highland river valleys and western Titicaca agricultural zones. The blue region represents the contiguous region with a 0.75°C or higher increase in temperature due to the lake effect of Lake Titicaca.

4.2.1 Terraces: Unlocking fertility in Andean landscapes

> "Considering the small extent of land around Cuzco suited for cultivation, [Pachacuti Inka Yupanqui] supplied by art what was wanting in nature" - Pedro Sarmiento de Gamboa

The steep and rocky slopes of the arid-to-semi-arid western cordillera of the Andes in southern Peru do not upon first impression appear to lend themselves to cultivation. However, *sucre* (Quechua)/ *taka* (Aymara)/ *andenes* (Spanish)/ terraces (English) were constructed by prehispanic Andean agriculturalists to enable the cultivation of a wide variety of crops in the region. Among the most obvious of the physical benefits of terracing is their creation of nearly level surfaces which helps to prevent soil erosion, improves soil depth and increases water retention (Donkin 1979, p. 131; Treacy and Denevan 1994). It has also been proposed that terracing complexes can create their own microclimates, shaping the environment to obstruct the flow of cold air and thereby protecting crops from the risk of frost (Donkin 1979, p. 131; Erickson 2000; Morlon 1996). Through these transformations, terracing allows cultivation to occur throughout the river valleys and take advantage of the rich diversity in microclimates offered by rapid elevation changes. At high altitudes, unirrigated terraces in the altiplano are still used to grow crops such as quinoa, potatoes, cañahua, oca, and ulluca (Donkin 1979, p. 34). In the lower valleys, where irrigated bench terracing is more common, maize, coca, and cotton are grown, as well as (following the Spanish introduction of new cultigens) wheat, grapes, and olive trees (Fisher 1998, p. 32). Bench terraces and contour terraces are also ubiquitous in the Titicaca Basin, and cover nearly every slope surrounding the Lake's shores (Erickson 2000). Erickson estimates that precolumbian terracing spans 500,000 hectares in the hills immediately surrounding the lake and the Basin valleys alone (Erickson 2000, p. 329).

Apart from their physical benefits, terraces have also shaped social relations, requiring immense communal labor for their construction and maintenance, and requiring supra-household coordination of irrigation water apportionment (Guillet 1992). Traditional models for understanding monumental-scale agricultural infrastructure propose that highly centralized authority is required to organize the labor networks required for such major landscape transformation projects (Wittfogel 1957; Stanish 2003). Under Inka rule, highly formalized and irrigated bench terraces were part of "a systematic and nationwide policy of land improvement and colonization" with the goal of expanding maize cultivation (Donkin 1979, pp. 33–34) Erickson notes that many terraces constructed under Inka auspices "appear over constructed or over-engineered" (Erickson 2000, p. 332)(to symbolize and enforce Inka power. The imposing complexes of meticulously designed and strictly managed terracing were constructed by local labor that was mobilized by Inka elites. This formed a powerful demonstration of Inka control over the land and the people, shaping them both to conform to Inka rule.

However, recent research has challenged the idea that monumental agriculture projects could only be achieved through centralized, hierarchical rule, showing that large-scale agricultural systems in the Titicaca Basin were constructed prior to the establishment of any state-level societies (Erickson 1999). Indeed, Langlie (2018, p. 3) argues that long-term collective action by farmers over generations created monumental-scale agricultural landscapes that encouraged food independence and local management, thereby forming a powerful tool of resistance to hegemonic authorities. She emphasizes that understanding local practices are fundamental to understanding broader regional and political patterns. The following research bridges local and regional dimensions, by generating accurate local data about agricultural feature construction and management at a broad regional and supraregional scale.

4.2.2 What is a Terrace?

Agricultural terraces are anthropogenic platforms of earth formed behind walls constructed to retain the soil and provide a space for planting crops. For certain features, such as the highly formalized and meticulously constructed Inka bench terraces described above, their designation as "terracing" is clear (Figure 4.2). However, there are more marginal cases which can be difficult to distinguish from other parcels of land. In his typology of terracing in the Colca Valley, Denevan designates "Valley bottom walled terraces" as anthropogenic terraces "at lower elevations on the nearly flat river terraces. They are enclosed by stone or adobe walls and are interspersed with unterraced walled fields. These are very broad terraces with low retaining walls, divided by wall field boundaries."(Denevan 1986, p. 39) It can therefore be difficult to distinguish between "valley bottom walled terraces" and "unterraced walled fields," especially when identifying them from aerial or satellite imagery.

Similarly, Denevan describes upland walled terraces as sloping-field, self-filling terraces which “do not appear to have been corrals for livestock” and require “careful examination to determine function.” (Denevan 1986, p. 39). In this case, not only is it difficult to determine whether a field is terraced or unterraced, it can be difficult to determine whether it is a field or a corral. For his part, Denevan’s survey of the valley using aerial imagery acknowledges that the visual differentiation between categories is often arbitrary and does not attempt to distinguish between valley bottom walled terraces and unterraced walled fields, simply designating all fields on the valley floor as “Bottomland fields” and fields on the valley slopes as “terraces” (Denevan 1986, 1988).

Figure 4.2: Example of Agricultural Terracing near Yura, Peru

For the purposes of this research, the semantic differences between a terrace and a walled field are outside of our scope. Rather, we are interested in where there is visual evidence that agricultural production is happening or has happened in the past. That is, where agricultural fields (regardless of whether they are on an anthropogenic terrace or not) are being maintained, or have been abandoned. Therefore, while the literature on agricultural fields in the Andes is generally framed in terms of terracing, we have taken a more expansive approach to inventory all areas of agricultural production visible in the satellite imagery.

4.2.3 Agricultural Abandonment

A casual observer in the Andes today will notice that much of the once-productive agricultural landscape now appears unused, though precise measurements of how much has been abandoned have been elusive. The reasons for this large-scale abandonment have been thoroughly debated, with proposed causes including environmental factors such as climatic shifts (Guillet et al. 1987), tectonic uplift (Clement and Moseley 1991; Moseley 1983), or environmental catastrophe, as well as factors such as depopulation, new agricultural methods, and Spanish colonial policies (Cook 2004; Cook and Cook 2007; Guillet et al. 1987; Wernke 2010). Guillet (1987) proposed that periodic droughts were major causes of terrace contraction and abandonment, while periods of relative water abundance caused expansion. To Guillet and Donkin, one of the primary benefits of terracing in the semi-arid western Andean cordillera is its facilitation of canal irrigation (Donkin 1979, p. 34; Guillet et al. 1987), so a lack of water renders terracing useless. In contrast, Treacy argues that drought may have the opposite effect, stimulating irrigation system expansion to account for the lack of water, thereby creating terrace system expansion (Treacy 1990b). Denevan sought to better understand terrace expansion and contraction in the Colca Valley by comparing active and abandoned terracing around the town of Corporaque in aerial imagery collected during the Shippee-Johnson aerial photography expedition in 1931 to that in imagery collected in 1974 (Denevan 1988, 2003). The results showed that there had been a clear increase in cultivation since 1931, though Denevan does not discuss the potential reasons for this increase.

In addition to water management, D'Altroy points out in his comments on Guillets's work that social and political decisions shape the agricultural landscape (Guillet et al. 1987). In the Upper Mantaro and Yanamarca Valleys, the intensity of high-altitude land use decreased dramatically under Inka control as people were moved to lower-elevation areas and preference was given to maize over tubers for cultivation. Lynch agrees, contending that the availability of water itself is more dependent on the social and political decisions of water management than the presence of the water itself, stating clearly that "terrace abandonment in the south-central Andes result in large part from social and political, not natural, causes (Guillet et al. 1987)." Wernke (2013) demonstrates that these physical limitations were also shaped by cultural and social structures in a complex network of decision-making by Andean farmers. Using detailed inventories of agricultural fields, mapped to visita documents, Wernke shows different strategies for different ayllus. Lower-ranked communities in reducción settlements were further displaced from their fields during the resettlement, forcing them to walk further and accept higher and more ecologically marginal fields, while higher-ranked communities were able to prioritize elevation, abandoning more distant fields (Wernke 2013). In this way, the complex interplay between social rank, community membership, and ecological affordances worked together to shape agricultural abandonment in the Colca Valley. In all likelihood, a complex mosaic of each of the factors listed here, in combination with others, continue to shape the agricultural infrastructural dynamics. Though this paper does not claim to resolve these questions, we believe that any attempt to tease apart the palimpsest of causes and decisions that shape the agricultural landscape today must include an analysis of where agricultural fields were abandoned, and where the decision was made to continue maintaining them.

Just as it is important to analytically parse the concept of "terrace" above, it is also necessary to examine the concept of abandonment in greater detail. Researchers who have sought to understand the continued maintenance or abandonment of agricultural features in the Andes have typically identified fields with eroded walls, native vegetation, and particular patterns of color or texture designating a lack of maintenance as "abandoned." However, this may be a poor description of the actual use of the land by those who interact with it on a daily basis (Guillet et al. 1987). In some instances, crops that do not need the thicker soils or increased moisture provided by irrigation canals may be cultivated on fields that are not maintained as formal terraces. In the Titicaca Basin, Erickson claims that the walls and platforms of (primarily unirrigated) Prehispanic terraces are often not maintained by contemporary farmers, resulting in poor preservation and destruction. Nevertheless, fields are still cultivated on the dissolving terraces. In fact, walls are sometimes removed intentionally to increase the area of individual field plots (Erickson 2000, p. 331). Even fields that are no longer cultivated may not be "abandoned," as they may provide excellent spaces for grazing or corralling camelids, due to their existing walls and fertile soils. Therefore, though "abandonment" is used as a shorthand in this paper to match the existing literature, it is important to remember that the phenomenon being studied may better be thought of as an agricultural "deintensification," or simply a land use transformation.

4.2.4 Previous Efforts to Characterize the Extent and Distribution of Agricultural Abandonment

There has long been great interest in characterizing the extent and distribution of agricultural use and abandonment in Peru. Archaeologists in the academic literature sought to understand the social and ecological contexts of what some described as an "agrarian collapse" (Denevan 1986, p. 9). Simultaneously, policymakers in the fields of agriculture and sustainable development wanted to understand the extent and context of agricultural abandonment in the hopes of revitalizing abandoned fields to increase economic opportunity and production (Denevan 1988; Kendall and Drew 2019; AgroRural and CERA 2021). However, the sheer scale of agricultural terracing has been a challenge to their systematic study, and large-scale inventory projects have often been left incomplete. Of even greater concern, estimations of the rate of terrace abandonment vary dramatically. In 1963, A.C.S Wright estimated that as much as 80% of terracing in Northern Chile was abandoned (Wright 1963, p. 73). In 1986, Agricultural Engineer and ecologist Luís Masson estimated that as much of 1,000,000 hectares of Peru has been terraced and that 75% of that terraced landscape is abandoned, though two years later he'd reduced his estimation of the extent by nearly half and estimated that approximately 50% of those were abandoned (Denevan 1988, p. 20). In 1987 the *Oficina Nacional de Evaluación de Recursos Naturales* (ONEM) undertook a survey project to inventory the terraced landscape in Peru, with the goal of revitalizing the abandoned agricultural landscape and developing Peru's agricultural economy (AgroRural and CERA 2021, p. 40). The inventory was never completed, but preliminary results found more than 150,000 ha of terraces across four departments, more than 60% of which was still in use (AgroRural and CERA 2021). Peruvian governmental research into terrace abandonment continued with the *Instituto Nacional de Recursos Naturales* (INRENA) of Peru in 1995 expanding on the efforts of ONEM, and ultimately publishing a survey of agricultural fields in 1999 that included 8 departments, with 256,950 ha of terracing (AgroRural and CERA 2021, p. 41). According to INRENA, only 16% of inventoried terraces were abandoned (AgroRural and CERA 2021, p. 41).

Simultaneously, archaeologists were becoming increasingly interested in understanding the extent of terrace abandonment but also its causes. The bulk of archaeological research on terrace abandonment at this time focused on the Colca Valley in the southern Peruvian highlands. In 1988 William Denevan produced an inventory of active and abandoned terracing in the Colca Valley, created through the manual inspection of aerial imagery collected in 1974 (Denevan 1988). At a scale of 1:17,000, Denevan examined ten districts in the Colca Valley and concluded that overall approximately 42 percent of all the fields are abandoned, and 61 percent of bench terraces are abandoned. Bottomland fields in contrast remained cultivated at a much higher rate with only 7 percent abandoned (Denevan 1988, p. 28).

Finally, in 2008, The Peruvian Ministerio de Desarrollo Agrario y Riego proposed a study called "Programa Andenes" to inventory and characterize the state of terraces in Peru through the visual inspection of satellite and remote sensing data across 11 regions and 720 districts (AgroRural and CERA 2021, p. 45). In contrast to the GeoPACHA manual survey and AI-Assisted survey (this chapter), the Programa Andenes took a more constrained approach to the agricultural landscape, choosing to identify only agricultural terracing, rather than agricultural fields more broadly. To this end, the project generated clear definitions of terraces, differentiating them from other land parcels by the structure of the wall and the slope of the platform, and differentiating in-use terraces from abandoned ones using indicators such as color, vegetation, and wall erosion (AgroRural and CERA 2021, pp. 46–48). The Programa Andenes project estimated that for all 11 regions, approximately 24% (81,400 ha) of terracing was abandoned (AgroRural and CERA 2021, p. 52). Nevertheless, as the Programa Andenes researchers observed in their report, it can be challenging to distinguish between terracing from other agricultural parcels, as the differences can be quite subtle and difficult to see, especially in satellite imagery. This makes the distinction somewhat arbitrary, and may more easily allow variation in observation procedures between research teams or over the course of the study. Among other challenges the researchers cited the very large dimensions of the study area in contrast to the very small size of the units of observation (terraces), and their dispersion across the territory (AgroRural and CERA 2021, pp. 47–48). Each of these challenges is directly addressed by the AI-Assisted survey approach.

4.3 Data Construction: Building data for model construction

Two data sets are required to conduct an AI-Assisted survey of satellite imagery, the first is the satellite imagery, which must be pre-processed using standard procedures to correct for distortions due to elevation and to merge multispectral and panchromatic bands to generate multispectral imagery at high spatial resolution. The second is a collection of "labels" which provide the deep learning model with examples of the features we wish the model to identify, in this case, the boundaries of active and abandoned agricultural fields. Finally, environmental variables of interest were selected for exploratory analysis of the ecological context of agricultural fields in the study region. Details of these data collection procedures are as follows.

4.3.1 Imagery Description and Processing

The Worldview-2 and Worldview-3 satellite imagery for this research are provided gratis via USGS as an NSF federally-funded project and were provided as OrthoReady Standard 2A products. This means that the data is atmospherically corrected, geographically projected, and provides an RPC sensor model for orthorectification. We acquired 126 images with less than 10% cloud cover per image. Thirty images were from the WorldView 3 (WV3) satellite constellation with 8 spectral bands and a 0.31-meter spatial resolution, and 96 from the WorldView-2 (WV2) constellation with the same 8 spectral bands and a 0.46 meter spatial resolution. Together, these images cover the entirety of the study region. Each image was supplied as a collection of tiles with both multispectral and panchromatic elements and was processed using Orfeo Toolbox (OTB) through the Python API (Grizonnet et al. 2017). The tiles were first orthorectified using their RPC sensor models and a 30m SRTM DEM. The tiles were then mosaiced together to form full Multispectral and Panchromatic images. The images were Pansharpened using OTB's Bayesian fusion algorithm and then resampled to 8-bit to save on storage and computational requirements. The relevant script and details on its execution are provided at https://github.com/JamesZDonline/GeoPACHA_machine_learning. It should be noted that due to variations in processing workflows, the resulting imagery does not always align perfectly with online imagery such as that provided by Google Earth or Bing. Geometric transpositions of approximately 30m are evident in some regions. Furthermore, in areas where the images overlap, there is occasionally a slight shift evident between images in the dataset.

4.3.2 Agricultural Labels

High-precision maps of the presence of active and abandoned agricultural fields are needed to train the AI to detect terracing and distinguish between its status. As an adjunct to the GeoPACHA imagery survey, Grecia Roque and Steven Wernke digitized polygons of active agricultural areas in the southern highlands survey area. This data is still preliminary and under editorial review, however it provides a powerful comparison for the AI-assisted survey. This preliminary GeoPACHA data may be described as "regionally accurate," which is to say that it does an excellent job of recording the broad boundaries of agricultural features, but it also includes other features of the landscape such as settlements, rivers, and hills which were not terraced but are surrounded by terracing. The result is an overestimation of the cultivable area. This is necessary for such large-scale manual surveys as a finer-grained approach that excludes these features takes far too long to practically complete over many thousands of square kilometers. Furthermore, due to the slight variation in processing between online imagery and our imagery dataset, the digitized fields do not always perfectly align with the WorldView II/WorldView III imagery used for this analysis. The lower precision and spatial shift make it inadvisable to use this data for AI training directly.

Data on abandoned agricultural infrastructure was opportunistically tagged by minimal grid cell provenience, as a part of the GeoPACHA survey project. This data is also preliminary and currently under editorial review. It is also much lower resolution than the active data, comprised of marked .005 degree (∼500m) survey grid squares that were identified to contain any amount of abandoned agricultural features. Therefore, a single survey grid can contain both active and abandoned agricultural features, resulting in an overlap of the datasets and a lack of clear labels for each feature. The opportunistic sampling of this data also results in incomplete data for the region, with some areas lacking records.

To generate a dataset that is more precise in its labeling of agricultural features, and is aligned with the processed WVII and WVIII imagery, (and therefore useable for training the AI model) we sampled 50 areas of approximately 30 km² (squares with 0.05 degrees of latitude and longitude to a side) that contained active agricultural fields for a total of approximately 1500 km². A team of researchers then edited the regional survey of active agricultural terracing within these sample regions to refine the field boundaries, exclude unwanted regions, and align it with the processed imagery (Figure 4.4). This resulted in a 14% reduction on average of the measured cultivable area from 29,002.00 to 24,937.60 hectares for the regions surveyed. This was a laborious process. A single 30 km² region, surveyed at a finer resolution, took between a few minutes up to 8 hours to digitize thoroughly, depending on the scale and complexity of the terracing in the region, with the average region taking several hours to complete. A second round of digitizing was then required to generate a detailed map of the abandoned terracing in our 50 survey regions. One of the benefits of a machine-learning approach is that if we can train the model at this higher resolution, it will reproduce this effort over thousands of square kilometers in a matter of hours, rather than the years required to do it by hand. Nevertheless, the field boundaries appear to have a nearly fractal complexity and while we excluded the bulk of the towns, hills, and valleys, we did not manage to exclude all small local structures, roads, and a vast array of other non-terraced features.

The digitized regions were then split into training (n=40) and validation (n=10) datasets, with care taken to ensure that adjoining regions were not split between training and validation to limit the effects of spatial autocorrelation on the model evaluation. Examining the training data it was clear there was significantly more active terracing than abandoned terracing in our data. Therefore, a 51st area was added over the Churajón area, which contains extensive abandoned terracing, to augment the abandoned training data and provide more examples of abandoned fields on which to train the model. This data was then passed to the Free and Open Source (FOSS) RasterVision (Azavea 2022) framework for further processing and model training and deployment.

4.3.3 Ecological data

Several ecological variables were selected for comparison to the model results to better understand the distribution of agricultural fields and their abandonment. Elevation, slope, aspect (the cardinal direction in which the slope faces), and Geomorphon (a geomorphic index; discussed below) were all selected as preliminary characteristics of the environment that may have had an effect on the construction, and later the maintenance or abandonment, of agricultural fields and terracing. The elevation data was a 30m resolution Digital Elevation Model (DEM) from the Shuttle Radar Topography Mission (SRTM) (Center 2017). The Slope (in degrees) and Aspect (in degrees clockwise from North) variables were calculated from the DEM using the r.slope.aspect function in GRASS GIS (GRASS Development Team 2012). The geomorphon dataset was calculated using the r.geomorphon tool in GRASS. While the other variables are commonly used in geospatial analyses and are well known, the geomorphon data may bear further explanation due to its relative lack of use. The geomorphon tool categorizes the landscape into 8 different morphological classes based on sight lines and the local neighborhood of elevations. These classes include: flat, peak, ridge, shoulder, spur, slope, hollow, footslope, valley, and pit. This provides an intuitive way of understanding the shape of the landscape that aligns well with commonly used terms in geographical and archaeological descriptions. For instance, this data provides a clear way to examine the proposal that most active agricultural fields occur in the valley bottom (flat, valley, footslope), while terraces on the valley wall are more likely to be abandoned (slope, shoulder, spur). Finally, relevant climatic variables related to precipitation and temperature were downloaded from the WorldClim 2 Bioclimatic variables dataset (Fick and Hijmans 2017).

4.4 Methods: Training and assisting the AI

4.4.1 Automated Survey using Raster Vision

The processed imagery, along with the boundaries of the 51 squares designated as Areas of Interest (AOI) and the high precision terrace label data form data objects known in Raster Vision as “scenes.” Using Raster Vision, up to 1000, 300 pixel x 300 pixel image and corresponding label chips were randomly extracted from within each area of interest for a total of approximately 51,000 chips. Labels were composed of three categories, “Abandoned

Field" (class id =2), "Active Field" (class id=1) and "Background" (class id=0) as shown in Figure 4.3.

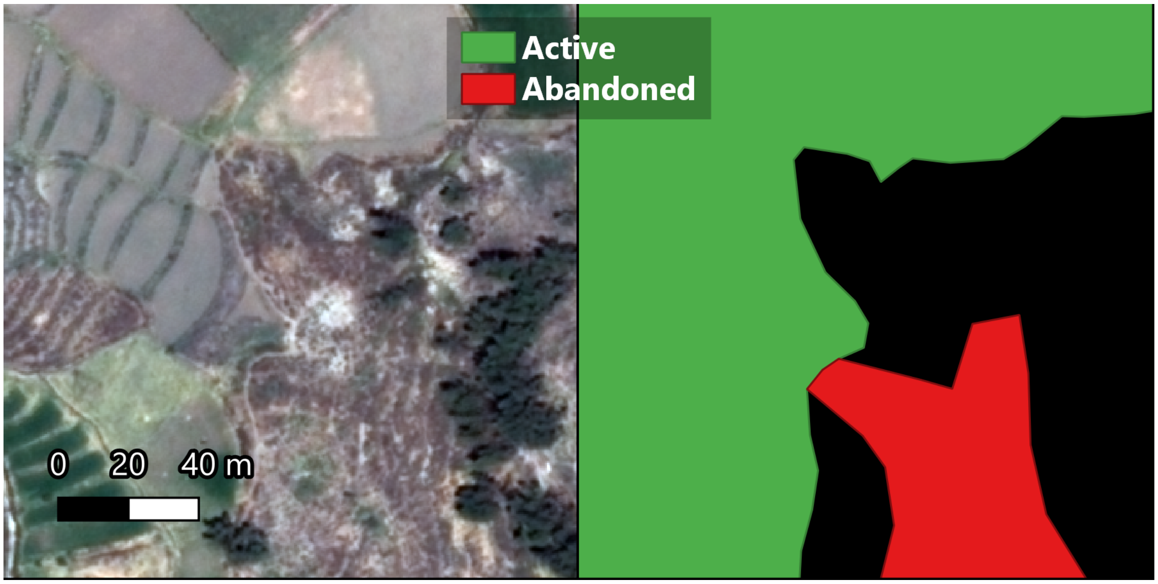

Figure 4.3: An example of a training data "chip." Imagery is shown on the left and manually digitized labels are shown on the right.

Training data was augmented using random 90° rotations, and horizontal and vertical flips, and then used to transfer-train DeepLabv3, a Deep Learning semantic segmentation architecture. A pretrained ResNet-50 was used as the backbone architecture, which does feature extraction for the DeepLabv3 model. The model was trained using the Adam optimization algorithm for 30 epochs with a learning rate of 0.0001 with a batch size of 10. Configuration and log files are included in the Github Repository. After each epoch, the model predictions were compared to the data in the validation set to track training progress and watch for evidence of overfitting. Once the model was trained, Raster Vision was used to create a "model bundle" which contains all of the model parameters and protocols for deploying the model on new imagery. The model bundle was then used to deploy the model across all of the images in the Southwestern Highlands study region, thereby producing pixel-segmented raster maps of the presence of active and abandoned fields for each image.

4.4.2 Data Compilation

Semantic segmentation predictions were created for each of the 126 images covering the study region. However, many of these images overlap with each other. This can be advantageous, because it provides multiple opportunities for the deep learning model to identify features where they are present, thereby avoiding missing features as a result of clouds or shadows which may obscure features in the imagery. Simultaneously, it increases the potential for false positives and creates the potential for conflicting predictions between images. The r.series function in GRASS GIS was used to merge the predictions into a single and cohesive dataset, generating a raster dataset with a 1m pixel resolution for the region. For each pixel r.series evaluated if the model results for any image predicted active or abandoned agricultural fields at that location. If there were multiple predictions, the modal result was selected. In cases of conflict where the mode was equally split between active and abandoned terracing, the result defaulted to active terracing, as these predictions were shown to be much more accurate in the validation data, and therefore were more likely to be correct. This suggests that the model may overestimate the amount of active terracing for locations of conflicting information. However, the accuracy of the active terracing results and visual inspection suggests that the effects of this result are marginal. These biases are also likely countered by the prevalence of false positives in the abandoned fields dataset. Finally, the predictions for the study region were converted from raster to vector format using the GRASS GIS r.to.vect function for data cleaning and analysis.

4.4.3 Data Cleaning

As the first step in the AI-assisted research design, we then undertook data cleaning to reduce the prevalence of false positives and gain a better understanding of the quality and potential shortcomings of the results. The model produced a great many very small features, often due to noise in the imagery. Following visual inspection, a threshold of .01 hectares was selected as the smallest features of archaeological interest, and any features smaller than this were excluded from the analysis. Furthermore, the manual inspection of the results revealed no evidence of agricultural fields above 4,500 meters in the study region. However, the dataset did show terracing above this elevation, primarily recorded as abandoned. These very high-elevation features were primarily modern erosion control measures, rather than active or abandoned agricultural features. Though they do strongly resemble agricultural features, they are not the features of interest for this research, and future versions of the model will attempt to distinguish between abandoned terracing and erosion control to exclude such features. For the purposes of this analysis, any features above 4,500 meters in the dataset were excluded as either noise or modern features. Modern erosion control at lower elevations was manually removed wherever they were observed. Finally, a comprehensive inspection of the data was undertaken, recording which model predictions had been visually inspected and approved, and removing or editing features where necessary. All active and abandoned terracing above 50 hectares in size (n active:541, n abandoned:292) were manually inspected in addition to nearly 70 thousand smaller features distributed across the landscape which were opportunistically evaluated. Abandoned features tended to be smaller in size and more fragmented, making them more difficult to inspect overall. Nevertheless, in total 337,507 hectares ~86% of active agricultural features were inspected, and 85,576 hectares ~60% of abandoned agricultural features were inspected.

4.4.4 Defining regions: Titicaca Basin and Western Andean Valleys

The use, maintenance, and abandonment of agricultural infrastructure in the slopes adjacent to Lake Titicaca is a substantially different phenomenon than in other areas of the Southwestern Highland Study Region. In part, this owes to the distinctive climatological and environmental conditions driven by the lake itself, as the mass of water acts as a heat sink, moderating temperatures relative to altiplano settings beyond the circum-lacustrine area.

Given the biogeophysical diversity and the great depth and diversity of human interventions in the area, field systems in the Titicaca Basin are diverse, including several variants of raised field systems along the lake edge, valley bottom field systems, and unirrigated (rainfed) terrace complexes on upland slopes. In contrast to the narrow vertical ecological zones of the western cordillera's river valleys, the immediate vicinity of the lake provides massive areas of relatively flat, and occasionally waterlogged, land for cultivation. In spite of the high altitude, the large body of water also helps to moderate the temperature of the region, protecting plants from frost risk. Therefore the area of current and abandoned agricultural activity around Titicaca is much greater than in other regions, which can obscure regional and subregional variation in the distribution fields outside of the vicinity of Lake Titicaca. We have therefore divided the data into two regional sets, one which includes the Western Cordillera river valleys, and the other which covers the extensive fields that extend several kilometers from the shores of Lake Titicaca.

It is challenging to rigorously define the boundary between the Titicaca regime and agriculture more typical to the river valleys. An approximate boundary can be determined visually, as there is a sharp decline in the areal extent of agriculture approximately 40 km from the lake shore, however, this boundary becomes less clear in areas along the shore where the density of agricultural production is lower. If we take the temperature-moderating effect of the lake as a primary driver of the agricultural landscape in the region, we can model its boundary. A statistical comparison between the elevation and the mean annual temperature for the region shows the two variables are very strongly negatively correlated, with a covariance of $\sim$-0.96. However, as one approaches the lake, this relationship changes. A linear regression can be used to predict expected temperature based on elevation allows us to map the deviation from this expected value by mapping out the model residuals and generating contour lines to show the magnitude of this effect.This was implemented in QGIS with the GRASS 7 plugin, using the r.regression.line tool to estimate the linear model coefficients. Estimated coefficients were then used in the Raster Calculator to predict temperature as a function of elevation. These predictions were subtracted from the true temperature values as provided by WorldClim to map the residuals. As can be seen in Figure 4.5 the contour line representing 0.75 degrees C above expected temperature appears to map closely to the boundary of the increased density of agricultural activity near Lake Titicaca. Therefore, any agricultural fields within that 0.75 degree contour were marked as a part of the Titicaca phenomena and modeled separately from those outside of it. The one deviation from this pattern is at the southern edge of Huiñaymarca, where the temperature is slightly cooler, possibly due to air currents from the glaciated peaks to the east in Bolivia. These fields, located immediately adjacent to the lake are clearly a part of the same agricultural complex as the rest of the Titicaca data, and were added manually to the Titicaca dataset. Having separated the data, it is now possible to examine each phenomenon independently.

4.5 Results: Evaluating the model and characterizing field ecology

4.5.1 Model Quality

Evaluating the quality of the data produced by the AI-assisted survey is vital to have confidence in the results of any analysis using the data. Of the 51 areas that were manually surveyed in detail, 10 were reserved for the purpose of evaluating data quality. These include a mixture of active and abandoned agricultural features, and background non-agricultural land to provide a thorough test. For semantic segmentation, a common and easy-to-interpret metric for measuring data quality is the Intersection over Union (IOU), also known as the Jaccard Index. This calculates the area of the intersection between the predicted label and true label, and divides it by the total area covered by both the predicted and the true labels. This results in a value between 0 and 1 where 0 represents no overlap between the predicted and true labels, and 1 represents a perfect alignment between the two. IOU is preferable to measures such as accuracy (the ratio between correct and incorrect labels) for semantic segmentation, because accuracy can be skewed by the imbalance between label classes. For example, the data suggests that approximately 6% of the land area in the survey area contains an active or abandoned agricultural field. The model could therefore achieve a 94% accuracy simply by marking the entire region as not containing any agricultural fields. Therefore, I used IOU to evaluate the model's performance, comparing the model output prior to data cleaning to the validation dataset. This provides a conservative description of model performance, with post-hoc manual data cleaning ensuring the results of the AI-Assisted survey are better than those reported for the automated survey alone.

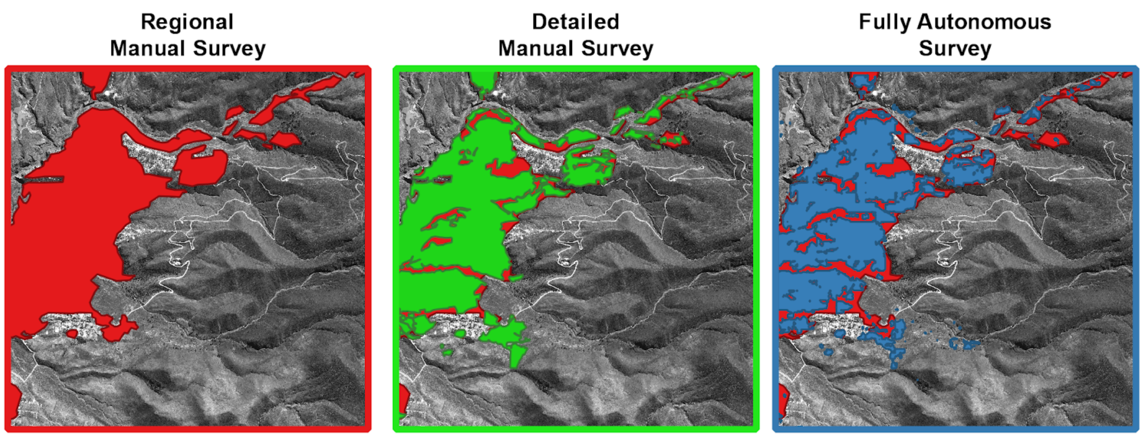

Figure 4.4: Comparing the results of the "Active" agricultural survey. Regional-scale manual survey data collected for GeoPACHA is shown on the left. Detailed "local-scale" survey data collected for model validation is shown in the center, and unedited fully automated survey data is shown on the right.

Due to the differences in difficulty between identifying active and abandoned agricultural fields, the IOU for active and abandoned fields were calculated separately, and an overall IOU score was obtained by averaging the two metrics together. For active agricultural fields, the model performed remarkably well with an IOU of 0.64. However, the model's performance on abandoned agricultural fields was not as good, with an IOU of 0.43 for a mean IOU of 0.54. This is largely due to the model overestimating the extent of abandoned agriculture, or identifying features such as water drainages, power lines, or other linear features with limited amounts of vegetation as abandoned terracing landscapes. While the comparison of the AI predictions to the validation dataset provides a good measure of the data quality, it can be difficult to judge from the metrics alone the implications of these IOU values. We therefore compared the coarse manual surveys conducted for the

Table 4.1: IOU scores for the Automated Survey and GeoPACHA Survey compared to the high resolution manually coded training data

	Automated Survey	Geopacha Survey
Active	0.64	0.51
Abandoned	0.43	0.06
Mean	0.54	0.30

GeoPACHA project to the refined manual surveys conducted in preparing the data for AI training and validation. This creates comparative metrics to understand the DL model's performance in relation to that of humans working at very large scales. For active agricultural fields, GeoPACHA surveyors achieved an IOU of 0.51 while the low resolution of abandoned fields led to a much lower performance of ∼0.06 for a mean IOU of 0.30. We can therefore confidently claim that the raw data produced by the AI model is of as high or higher quality than that produced by the GeoPACHA survey alone for both active and abandoned fields (Figure 4.4). Quality metrics are summarized in Table 4.1.

4.5.2 Measuring Deintensification

The high quality of the output data allows us to evaluate prevailing understandings of the extent, prevalence, and distribution of agricultural deintensification and field abandonment. In total, the AI-assisted survey identified approximately 512,098 hectares of agricultural fields in the study region (Figure 4.5). Of these, 76.4% (391,158 ha) were under cultivation at the time the imagery was collected, while the remaining 23.6% (120,940 ha) were identified as abandoned. However, this distribution is not consistent across the region. As discussed above, the region around Lake Titicaca dominates the dataset including 399,606 hectares of agricultural land or 78% of the agricultural fields in the study region. Of the Lake Titicaca region fields, 21% (85,488 ha) are abandoned, while 31.5% (35,451 ha) are abandoned in the western valleys. Overlaying a 0.1-degree grid over the survey region, it is possible to map the changes in abandonment rate across the entire survey region (Figure 4.6).

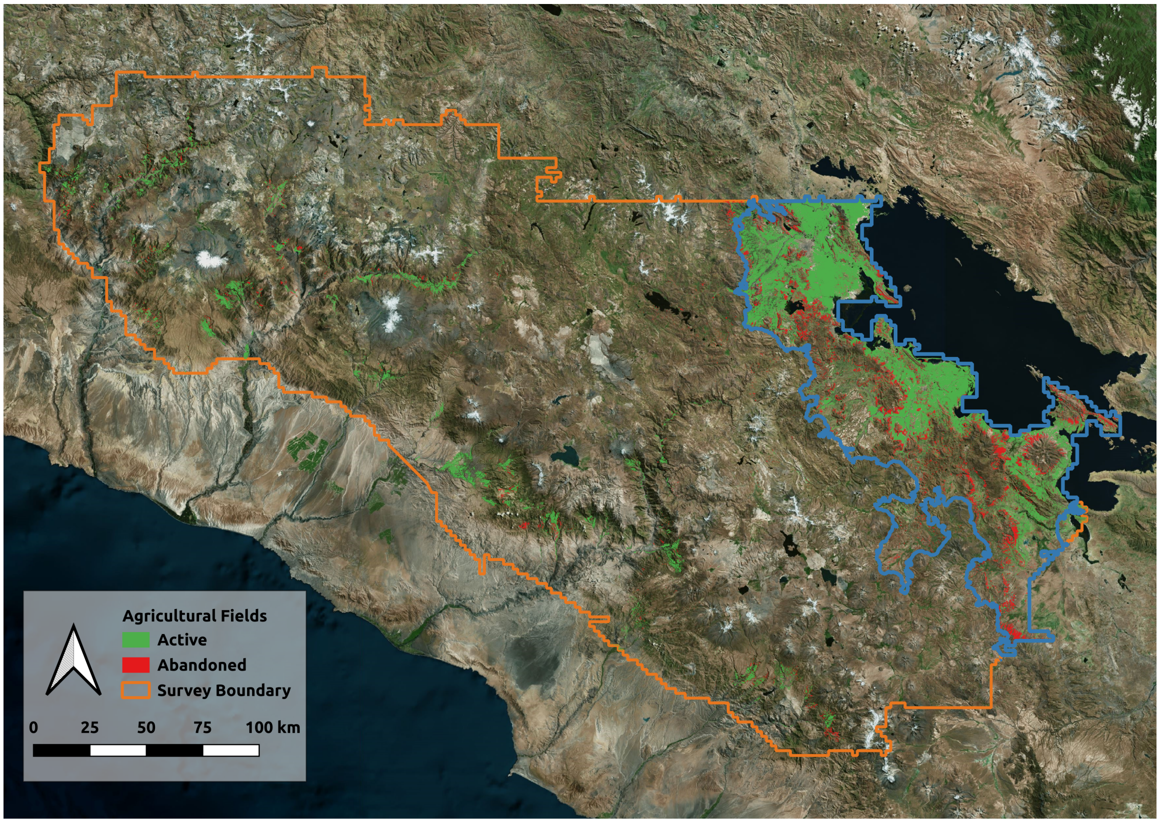

Figure 4.5: Results of the AI-assisted agricultural survey following data cleaning.

4.5.3 Exploratory Analysis Ecologically Contextualizing Agricultural Infrastructure

As an initial exploratory analysis, the distributions of active and abandoned agricultural features were compared to a variety of environmental features expected to be correlated to the construction and maintenance of agricultural fields. Certainly, these are not the only factors of relevance to agriculture, as agriculturalists in the Andes have transformed their environments by working around, with, and through the ecological variety afforded to them. Nevertheless, as the below analysis will show, environmental factors do impact the process of deciding where to invest the time and energy to make these environmental transformations, and so this analysis establishes a null hypothesis, against which such transformations can be understood.

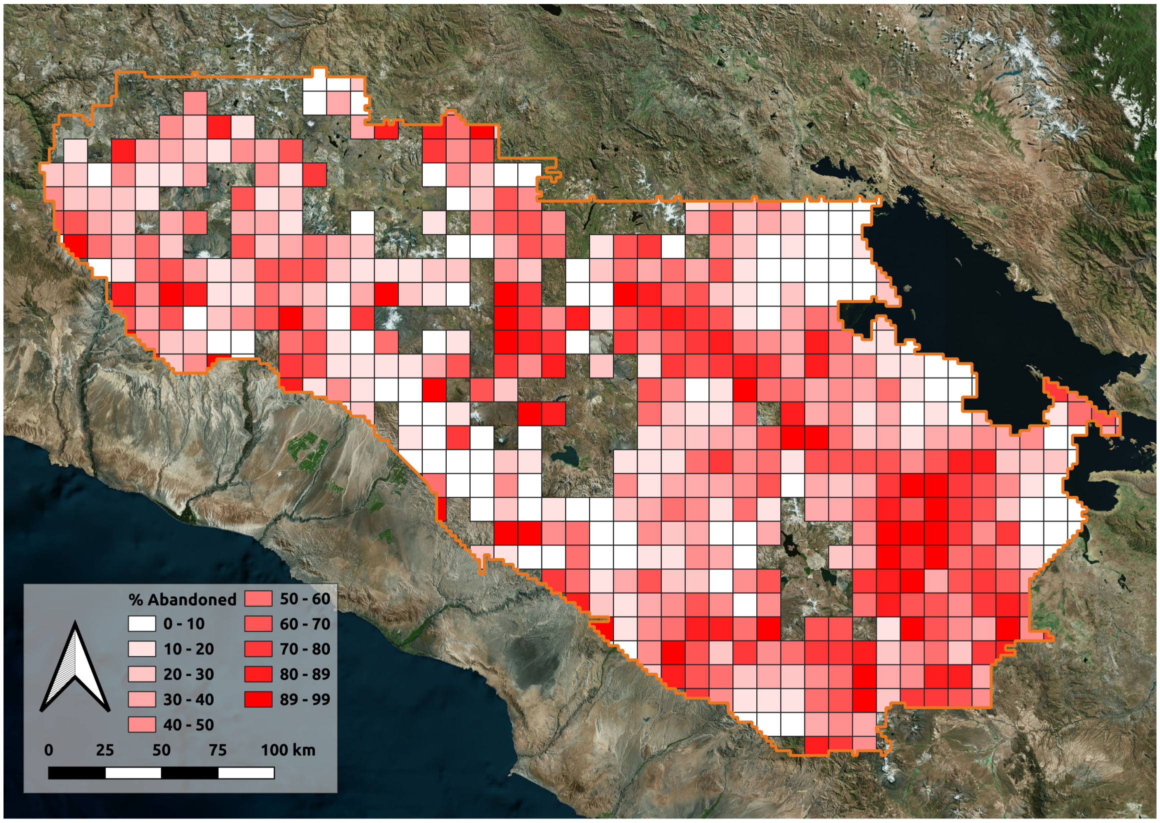

Figure 4.6: Local variations in % abandoned are mapped across the region. Comparing to Figure 4.5, note that most of the regions showing >50% abandonment are in locations with very little agriculture, suggesting vulnerability to noise in the data. An exception to this is Churajon, where abandonment rates do indeed top 70%. Missing squares had no agricultural features recorded.

4.5.3.1 Elevation

One of the commonly cited environmental factors to limit agricultural production in the Andes is elevation. This is primarily due to elevation's strong correlation with temperature, and therefore the risk of frost. Sampling the data, the average elevation of currently in-use agricultural fields is 3707 m, and the average elevation of abandoned fields is 3876 m while the average elevation of nonagricultural land is 4101 m. From an agricultural perspective, it is rational to build fields at lower elevations where possible to limit the risk of frost, thus leading to the overall lower elevation of agricultural fields. Terraced slopes also shed cool air better compared to natural slopes and valley bottoms (where cold air pools at night). Furthermore, when determining whether to maintain or abandon a field, the higher and therefore more marginal fields are likely to be abandoned first, resulting in the overall pattern of abandoned fields at higher elevations than in-use ones. Splitting the data into River Valley and Titicaca datasets, the pattern is even stronger in the River Valleys, with an average elevation of 3,100 meters for in-use fields, while abandoned fields have an elevation of 3,520 (Figure 4.7). Results are summarized in Table 4.2

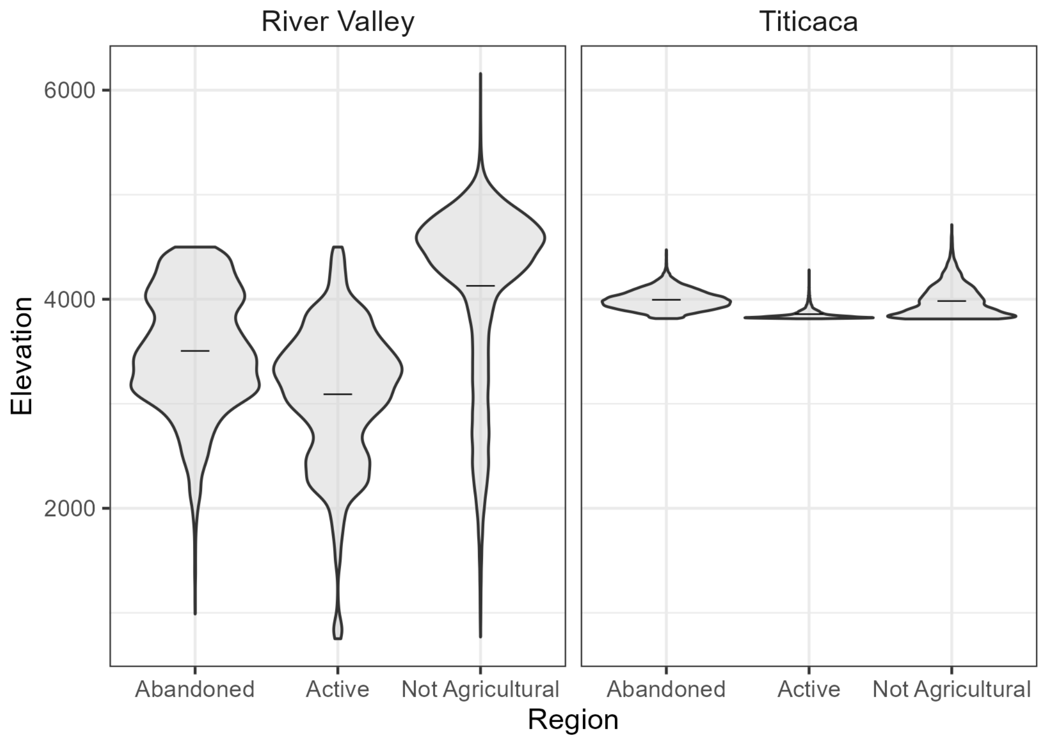

Figure 4.7: Violin plot showing the distributions of agricultural features across Elevation for the different survey regions. Horizontal lines designate means.

A Kruskal-Wallis test shows that statistically significant differences exist between the three, with a p-value of <2.2e-16, while a pairwise Wilcox test with Bonferroni correction shows that all three values are statistically different with a p-value of <2.2e-16. Comparing these numbers to the Titicaca data demonstrates the power of the effect of the lake on frost risk, as the average in-use field at 3,860 meters is higher than the average abandoned ones in the river valleys. The Titicaca abandoned fields are even higher than the average elevation for the region. Kruskal-Wallis and Wilcox tests on these data also show all of these differences to be statistically significant with p-values of <2.2e-16.

Table 4.2: Mean Elevations (m) for each land-use classification and region

	River Valleys	Titicaca	Overall
Active	3091	3857	3709
Abandoned	3505	3994	3874
Not Agricultural	4128	3983	4108

4.5.3.2 Slope

Another environmental factor expected to influence the construction and maintenance of agricultural fields is slope. Throughout the study region, active fields are located on terrain with lower-than-average slopes. This is especially true around Lake Titicaca, where the average slope is only around 4.24° because most of the active fields are located on the lake's floodplain. In the narrow River Valleys, active fields are primarily in the valley bottoms, with a smaller but significant portion on the valley walls, resulting in an average slope of 12.2°. The abandoned fields, in contrast, tend to occur on steeper slopes, with an average of 18.2° in the river valleys, and 13.3° around Lake Titicaca (Table 4.3). These are both higher than the average slope for the regions, which are likely somewhat lowered due to the high-altitude puna (flat grasslands), which are too high for agricultural production (Figure 4.8).

Again, each of these differences is statistically significant with p-values of <2.2e-16 according to Kruskal-Wallis and pairwise Wilcox tests. While terracing allows farmers to cultivate fields on steep slopes, the steeper the slope the more labor is required to construct and maintain terrace walls and fill them with soil to create platforms. Simultaneously, steep slopes require terraces to be narrower, resulting in less growing area than land that is naturally flatter. Therefore, it is unsurprising to see such a strong preference in field construction for flat lands and abandonment rates increasing with slope. This also aligns with Denevan's work in the Colca Valley, which found that bottomland fields were nearly all maintained, while terraces were much more likely to be abandoned (Denevan 1988).

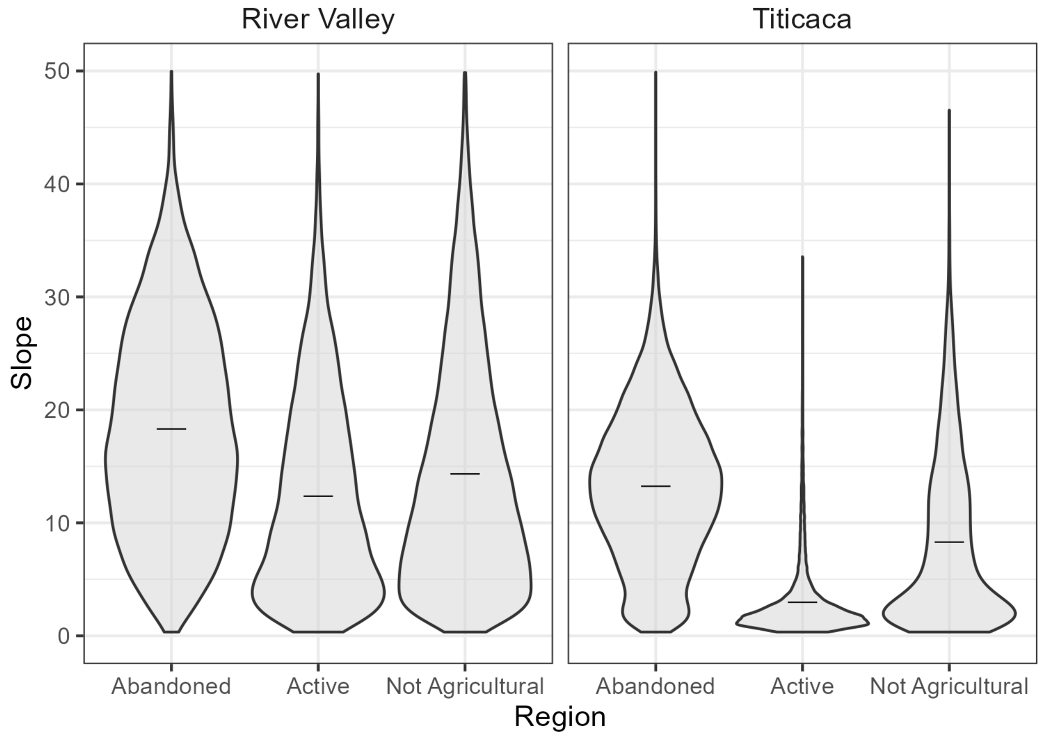

Figure 4.8: Violin plot showing the distributions of agricultural features across Slope for the different survey regions. Horizontal lines designate means.

Table 4.3: Mean Slope for each land-use classification and region

	River Valleys	Titicaca	Overall
Active	12.5°	4.13°	5.8°
Abandoned	18.4°	13.4°	14.6°
Not Agricultural	14.9°	13.0°	14.7°

4.5.3.3 Geomorphology

Closely related to the slope is the geomorphology of the landscape. The distribution of agricultural lands with regard to geomorphons is shown in Figure 4.9. Currently in-use agricultural fields are much more likely to be located on flat, footslope (transition from slope to flat), or shoulder (transition from flat to slope) landforms, while abandoned files are much more likely to be situated on slopes.

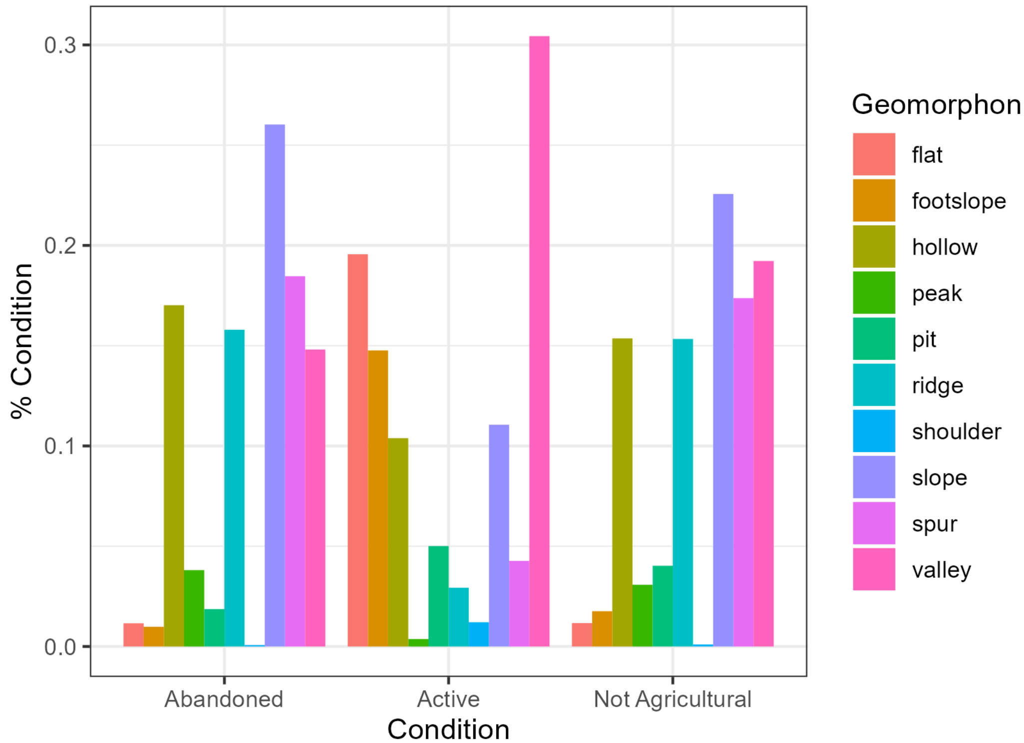

Figure 4.9: Percentage of features that fall in each Geomorphon class. Note that active fields are located on flatlands and valley bottoms while abandoned fields tend towards slopes. Both distributions are distinct from that of the background "Not Agricultural" class.

4.5.3.4 Aspect

Aspect (the direction the slope faces) is yet another environmental variable commonly expected to be associated with agricultural production. In the southern hemisphere, north facing slopes should receive more sunlight. Depending on context, this may be a benefit for plants who receive more opportunity for photosynthesis, or may be a burden as more direct sunlight requires higher water demands (Kohut 2023). However, generally, it does not appear from this data that Aspect has a substantial effect on the construction or maintenance of agricultural fields.

Removing flat areas (which by definition do not have an Aspect) and calculating the circular means of the angles show little difference between Active, Abandoned, or Non-Agricultural classes. In the River Valleys, the circular mean shows $\sim$220° from North, that is, the slopes on average face South West, this follows the trends as water flows from the

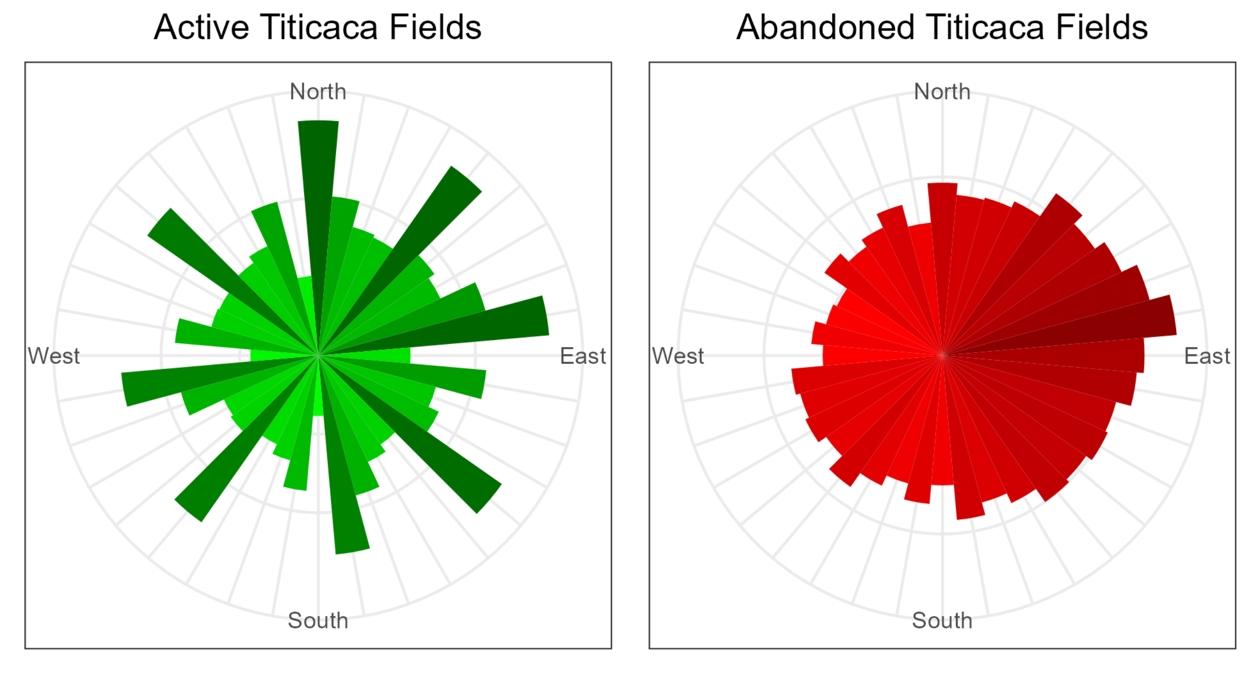

Figure 4.10: Distribution of Aspects for fields in the Titicaca Region. The active fields tend to be on land that is close to flat, resulting in more erratic distribution. Abandoned fields are much more frequently terraced, and reflect the distribution of the background.

mountain slopes to the Pacific Ocean. The survey area only includes land on the South Western side of Lake Titicaca, therefore for fields near Lake Titicaca, the average slope faces the North East, directing water towards the Lake. The one deviation from this pattern is the Abandoned fields in the Titicaca Region, which tend to be directed further East on average. This may be due to the valleys at the western boundary of the region which are more likely to be abandoned and are oriented more easterly.

4.5.3.5 Precipitation

Water is vital to agriculture, particularly in the arid environment of the study area. Therefore, it may be expected that precipitation amounts would be highly correlated with agricultural production, however, a comparison of fields to precipitation amounts in the wettest quarter of the year (during the agricultural season), shows that this relationship is more

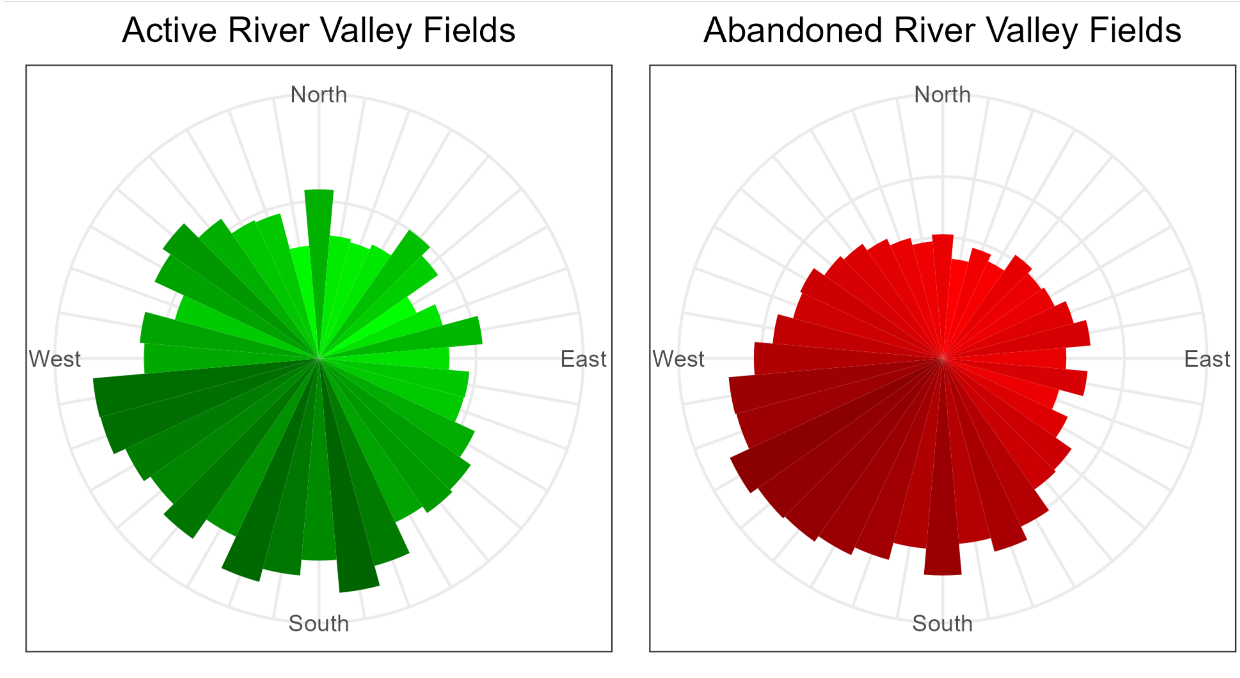

Figure 4.11: Distribution of Aspects for fields in the western river valleys. Note the slight south-western tendency for both active and abandoned fields. This reflects the background distribution as the Andes slope towards the Pacific.

complicated than expected. Indeed, in the River Valleys, it appears that agriculture is inversely related to precipitation amounts with active fields receiving the least water with approximately 208mm on average, abandoned fields receiving 272 mm and non agricultural lands receiving the most with 335mm.

This pattern is likely due to an inverse correlation between elevation and precipitation amounts. The highlands (where it is too cold to grow crops) receive significantly more rain than the valleys themselves, which rely on rivers, springs, and irrigation networks to water their crops. The Titicaca fields, being higher and close to a large body of water, receive significantly more precipitation with milder temperatures. As a result, the expected pattern of increased cultivation with increased precipitation holds true.

Table 4.4: Circular means of the aspect show that active and abandoned fields as well as non agricultural land tend towards the direction of the drainage basin, the Pacific Ocean for the river valleys, and Lake Titicaca for the land in its zone of influence.

	River Valleys	Titicaca
Active	218° (SW)	57° (NE)
Abandoned	220° (SW)	90° (E)
Not Agricultural	212° (SW)	46° (NE)

Table 4.5: Mean precipitation (mm) for the wettest quarter for each land-use classification and region

	River Valleys	Titicaca	Overall
Active	201	446	399
Abandoned	270	446	403
Not Agricultural	336	438	350

4.6 Validating the Results: Comparisons to previous survey efforts

4.6.1 GeoPACHA

The manual/visual survey for agricultural fields completed by the GeoPACHA project identified approximately 558,197 ha of land currently under cultivation in the study region. This includes terraces, such as those found in the river valleys on the western end of the study area, as well as waru waru and flat agricultural fields in valley bottoms and around Lake Titicaca. However, as discussed in the description of data sources, this number is a substantial overestimate as it also includes many areas that do not contain agricultural fields as a result of the lower-resolution "regionally accurate" survey methods required for covering such a large area by hand. Over the entire study region, the AI-assisted survey identified 391,158 ha of land currently under cultivation, a 30% reduction from the manual survey. Excluding the data around Lake Titicaca to focus on the river valleys the pattern remains the same, with 105,359.37 hectares of land under cultivation identified by the GeoPACHA manual survey and 74,763.61 ha by the AI-Assisted survey a 29% reduction. Visual comparison between the two datasets, as well as the data quality metrics, suggest that these differences are primarily due to the elimination of features that were incorrectly identified

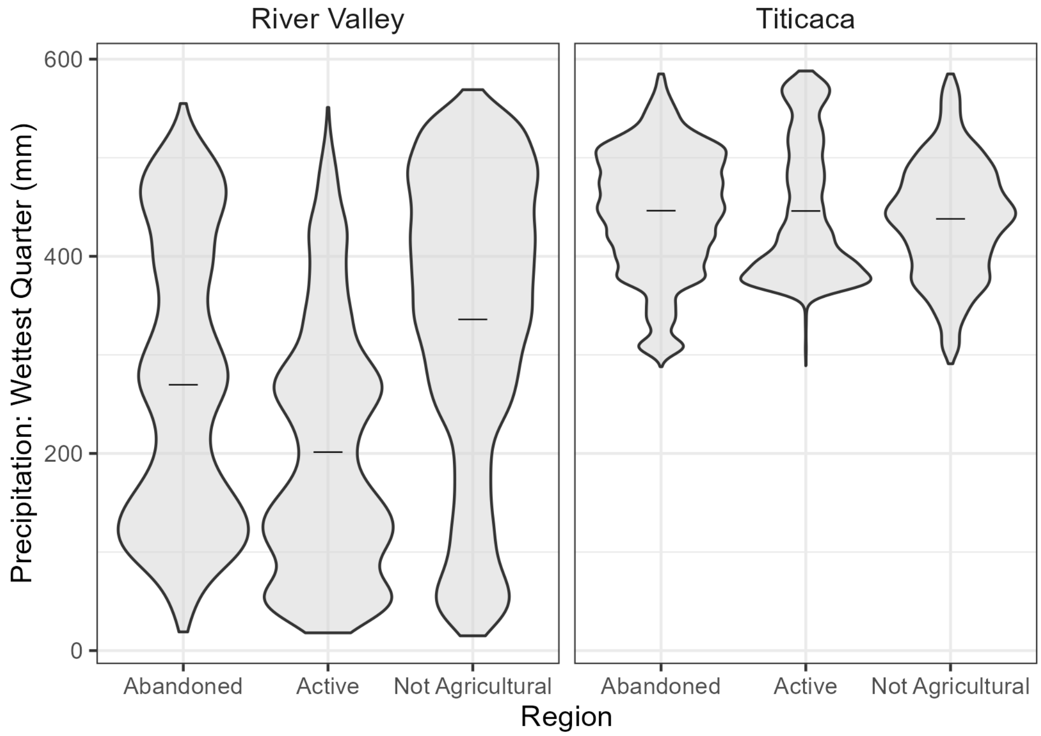

Figure 4.12: Violin plot showing the distributions of agricultural features across rates of precipitation in the wettest quarter for the different survey regions. Horizontal lines designate means.

in the manual survey, rather than features that were missed in the AI-Assisted survey. This result strongly speaks to the value of an AI-assisted approach for inventorying agricultural fields in satellite imagery. While the area surveyed by hand for the AI-assisted survey was very small in comparison to the full manual survey, the resulting data was much more precise. For analyses such as estimations of agricultural production capacity a 30% reduction in land area is substantial, fundamentally changing our understanding of the agricultural landscape.

4.6.2 Programa Andenes

The Programa Andenes study was conducted between 2012 and 2013 by the Consortium of Private Organizations for the Promotion of Small and Micro Enterprises (Copeme), the Center for Promotion and Development Studies (Desco), and the Promotion of Life

(Fovida). This research, published in 2021, is to the author's knowledge, by far the most extensive study of active and abandoned agricultural terracing in Peru and represents a monumental effort, identifying over 340,000 hectares of terracing. 140 of the districts reported as inventoried by Programa Andenes are included in the Southwest Highlands study area, with $\sim$125 districts falling mostly or entirely within the study area. Unfortunately, no GIS files of the data are included with the report, and the pdf maps included are very low resolution, making it impossible to compare the results of the surveys directly. Nevertheless, the report does include measurements of the total area of terracing for each district, and breakdowns between active and abandoned terracing for most (131) districts. These measurements can be compared to the results of the AI-assisted survey to examine their similarities and differences.

Given the results of the GeoPACHA survey and the large scale of the Programa Andenes project, one might expect that the AI-Assisted survey will identify less agricultural area than the Programa Andenes. However, this is not the case. Of the 140 Districts in the study region with data for the total area of terraces (currently in use or abandoned) in the Programa Andenes report, the AI-Assisted survey identified more agricultural fields in 110 of them, with over half of all districts showing more than a 50% increase over the Programa Andenes project. This is primarily due to the differences in definitions of targets of interest between Programa Andenes (which sought to only map terraced fields) and the AI-Assisted Survey (which sought to map all agricultural fields, including non-terraced agricultural land). This contrast is especially clear in the 35 districts around Lake Titicaca, where the AI-Assisted survey regularly identified 2 orders of magnitude more agricultural fields than Programa Andenes did terracing (Figure 4.13).

However, in other cases, the AI-assisted survey found agricultural terracing that was missed in the Programa Andenes survey. For example in the district of Andaray, where the steep slopes require terracing for all agricultural fields, Programa Andenes reports a total of 69.85 hectares of terracing. As shown in Figure 4.14 this is less than half the size of

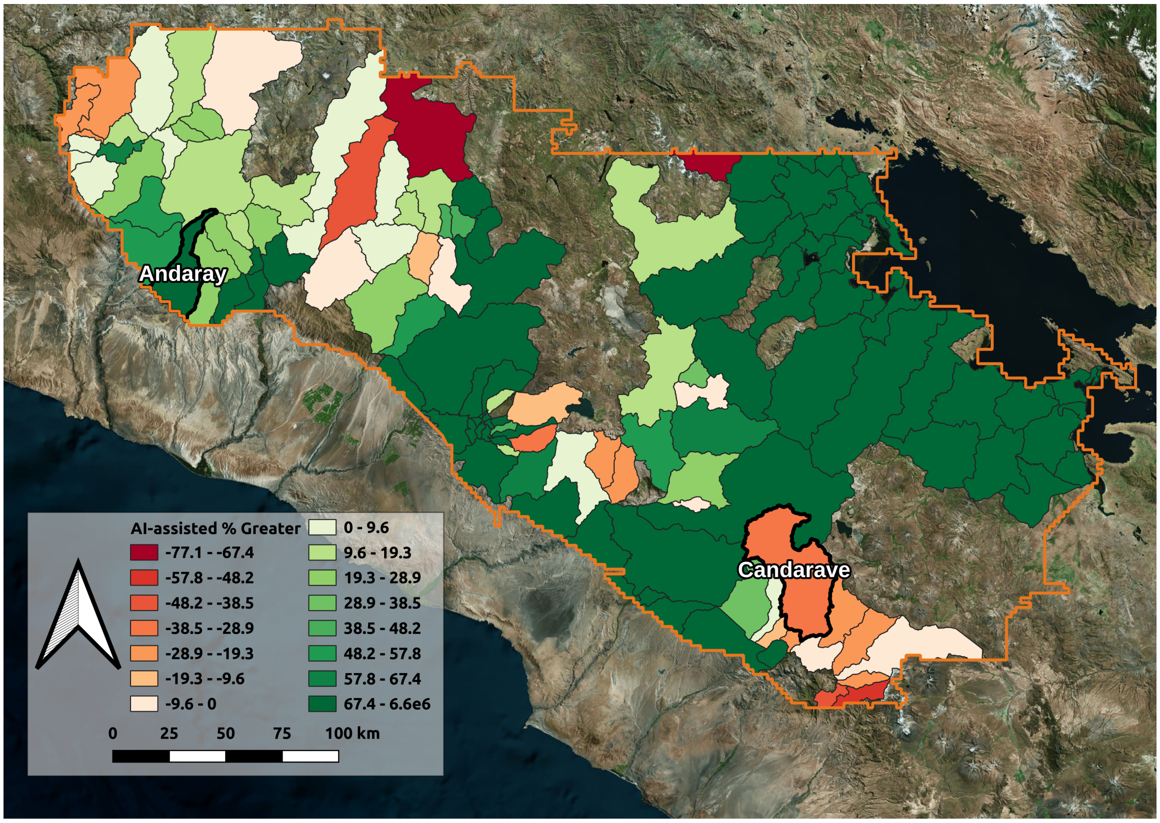

Figure 4.13: Percent difference between AI-assisted survey results and Programa Andenes results. Andaray and Candarave are highlighted for further discussion in the text.

a single active terrace complex identified by the AI-Assisted survey in Andary. In total, the AI-assisted survey identified over 1000 hectares of active and abandoned terracing in Andaray.

Given the tendency of the AI-Assisted survey to find more fields than the Programa Andenes project, it is important to examine the 26 districts where the AI-Assisted survey identified less agricultural land. Two of these, Candarave and Quilahuani included regions where the imagery used in the AI-Assisted survey had insufficient coverage, and so missed a section of Terracing (Figure 4.15). This accounts for less than 1300 hectares of primarily active agricultural fields and will be corrected in future work.

Another 7 districts fall on the edge of the survey boundary, suggesting that agricultural features outside of the AI-Assisted survey may account for the lower return rate. For the

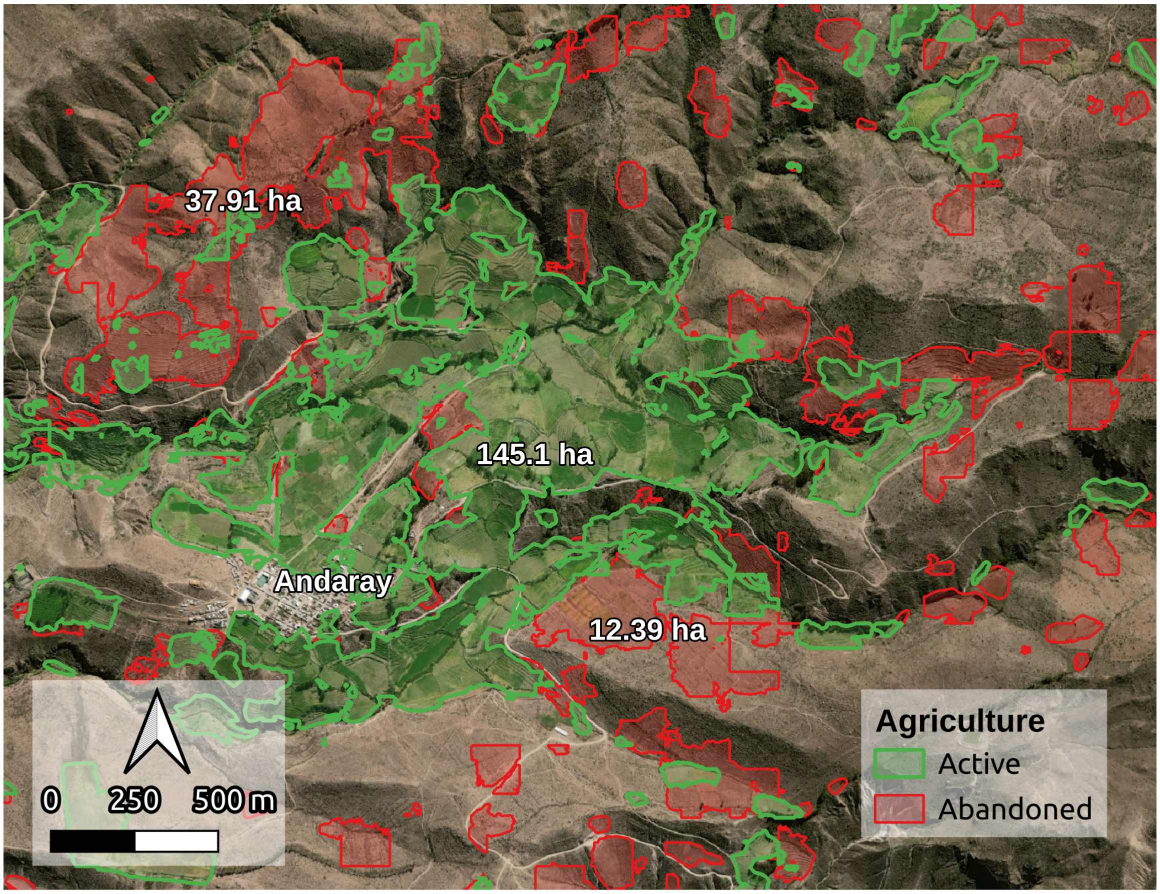

Figure 4.14: Agricultural fields in Andaray. These terraced fields cover much more area than is listed in the Programa Andenes results, suggesting that they may have been missed in the manual survey.

remaining 17 districts, it is difficult to determine why the AI-Assisted survey detected less than the Programa Andenes inventory without the original Programa Andenes data. It is possible that the AI-Assisted survey is achieving a higher level of detail than the Programa Andenes survey as it did for the GeoPACHA survey, thereby removing features that were incorrectly labeled as terracing. This hypothesis is supported in some districts, such as Achoma, Puyca, Susapaya, Sitajara, and Huanuara as a visual inspection suggests that the AI-Assisted dataset is thoroughly inventorying the currently in use fields, while reporting numbers substantially lower than those reported by Programa Andenes. In other cases, it is possible that the AI-Assisted survey missed features identified by the Programa Andenes project, particularly in the less reliable abandoned fields dataset. This seems to be the case

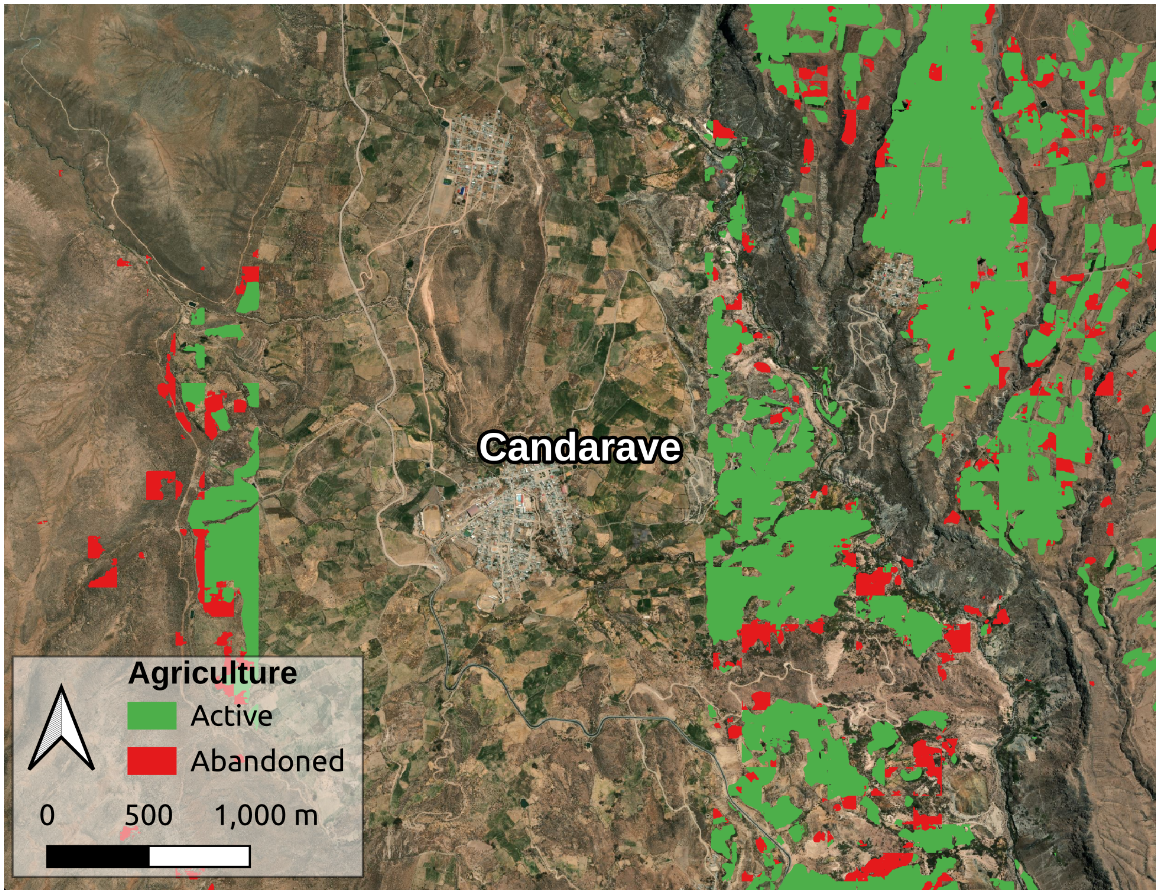

Figure 4.15: Fields in Candarave that were missed in the AI-assisted survey due to insufficient imagery coverage. One can clearly see the lines where the edges of the images stopped. This will be corrected for future analyses.

around the Pocsi District, where the model seems to have particularly struggled with identifying abandoned terracing. Further improving the abandoned field aspect of the model will be important in future research.

4.7 Discussion: The promise of AI-assisted survey

The results of the AI-Assisted survey of agricultural fields are highly promising for the future of understanding regional and supraregional patterns of agricultural deintensification. "Brute force" manual survey methods can be highly detailed at a local level, achieving very precise and highly accurate data. However, achieving such a level of detail at regional or trans-regional scales can be prohibitively time-consuming for manual surveyors. As a result, large-scale manual surveys are often forced to operate at lower spatial resolutions, resulting in the inclusion of nonagricultural features as shown by the GeoPACHA manual survey, or in missing features of interest as demonstrated in the Progama Andenes project.

Deep-learning offers a way to achieve both goals simultaneously by generating high-resolution data at very large spatial scales. As this research demonstrates, while the current deep-learning model may not surpass human capacity at the local scale, it far exceeds it at the trans-regional scale. The resulting extensive and high-resolution data will enable more accurate analyses of the capacity of agricultural production and the extent, distribution, and causes of agricultural deintensification and field abandonment. Furthermore, once trained, a deep learning model can be rapidly deployed on new imagery. Previously unsurveyed regions can be covered in a matter of hours, rather than months, and areas that have been previously examined can be monitored as new satellite imagery comes available, allowing us to track for the first time the expansion and contraction dynamics of agricultural infrastructure in the Andes at a regional scale. Such a dynamic perspective would shed light on the flexibility and stability of Andean agricultural infrastructure in light of ongoing political and climatological transformations, providing insight into how past processes may have shaped agricultural production, and how best policymakers may respond to them in the future.

Initial exploratory analyses of the data offer insights into how agriculture is distributed on the landscape in the southwestern highlands of Peru. As expected, environmental variables such as elevation and ground slope have significant impacts on the creation and maintenance of agricultural fields across the study region. Future work will seek to further refine these analyses and to tease apart the relationships between them. Elevation and precipitation, for example, likely interact to confine the zones in which agricultural production is practical. Crops at high elevations are likely to receive more water, but are also more susceptible to frost, while crops at lower elevations are safe from frost, but may have difficulty receiving enough water without extensive irrigation. Furthermore, aspect has reportedly had an impact on the suitability of land for agricultural production (Kohut 2023) in areas of the survey region such as the Colca Valley, however from the trans-regional perspective, it does not seem to have had a substantial impact on either the creation or maintenance of agricultural fields. Such regional variations bear further investigation.

Finally, this research offers a fresh perspective on the question of how much agricultural land has been abandoned in the western cordillera of the south-central Andean highlands. While estimates of the rate of field abandonment have ranged as high as 80% and archaeological projects to measure abandoned land have shown local abandonment rates as high as 40%, regional and transregional surveys by the Peruvian government have shown the average rate to be much lower, between 16% and 24%. The AI-assisted survey aligns well with the governmental surveys, showing estimations of an average abandonment rate between 21 and 23%, Indeed, in our estimation, the AI-assisted survey likely slightly overestimates the rate of agricultural abandonment due to the model's tendency to identify modern erosion control measures as abandoned agricultural fields. While every effort was made to remove these errors where they were observed, the large scale of the region makes it likely that some similar features were missed. Remembering that local rates of abandonment may be much higher,(or much lower) than the average, this research demonstrates the importance of developing regional or trans-regional data to understand broad patterns.

Much archaeological research on terracing and agriculture has been devoted to the Colca Valley, with limited examples from other regions such as Cusco (Kendall and Drew 2019) or the Titicaca Basin (Erickson 2000; Langlie 2018). While decades of intensive research has offered invaluable insights into the dynamics of agricultural production and abandonment in the Colca Valley, this paper suggests that these lessons cannot necessarily be taken as representative of the southwestern highlands as a whole. Rather, it appears likely that the Colca Valley may have especially high rates of abandonment, leading archaeologists to overestimate the amount of abandonment that occurred trans-regionally, at least in the South-western Highland region. This is perhaps not surprising. Archaeologists have traditionally focused our efforts on “archaeological sites,” that is, locations with a high density of human activity in the past that are often now abandoned. As we expand to examine the surrounding landscape and agricultural infrastructure that once supported the abandoned settlements we study, it is not unlikely that we are more likely to encounter abandoned agricultural features as well. We must remember that this experience does not necessarily represent the landscape as a whole, and seek to contextualize our investigations still more broadly in the regional and interregional context, a goal that may be assisted through AI-assisted surveys.

4.8 Conclusions

In spite of their importance, agricultural fields (particularly terracing) in the Andes have attracted less attention from archaeologists than other features of archaeological interest. This is largely due to the fact that information about fields is dispersed across the dozens or even hundreds or thousands of hectares of agricultural land that may surround prehistoric or modern towns, while archaeological information is much more concentrated in those settlements themselves. Agricultural fields also tend to be primarily of economic interest, perhaps making them less attractive as subjects of study than elaborate ceremonial centers or residential complexes where more aspects of daily life intersect (Donkin 1979, p. 25). Finally, fields located on terracing can be notoriously difficult to date archaeologically, due to the transportation of soil and embedded materials often required for construction and the continual remodeling necessary for maintenance. Nevertheless, these monumental-scale infrastructure projects were vital to the economic, social, and political lives of Andean people throughout prehistory and following the Spanish conquest and deserve further examination. This research represents the first step on this path, mapping where these features are, and their current state of use or abandonment.

4.8.1 Current and Future Work

There are many avenues for future research on agricultural features generated by the above methods and data. One dimension includes continuing to refine the AI models. The reliable identification of abandoned terracing is a major challenge even for human researchers. This makes it especially difficult to build a reliable AI model because the performance of the model will be limited by the quality of the data it is provided. However, there are methods for managing imprecise or unreliable data, known as "weak supervision," which may help improve the quality of the abandoned terrace predictions. We are currently exploring these methods as one potential avenue for refining the AI-Assisted survey. Furthermore, agricultural features identified by the model can now be refined by human surveyors and then fed back into the model as additional training data. This should result in substantial improvement in the model's performance.

Secondly, the inventory of agricultural features has shown that rates of field abandonment, and the relationship of field production and abandonment to environmental variables, is non-stationary across the study region. That is, local populations have made different choices, and possibly been offered different affordances, about how to manage agricultural resources. Mapping these variations and incorporating more social and human experiential variables into the modeling process will be vital to developing new, better, and more sophisticated hypotheses about how these variations arose. This research can then guide future studies in the field to test hypotheses and "ground truth" the data produced through trans-regional satellite surveys.

On a more local level, Denevan mapped in-use and abandoned agriculture of the Colca valley in 1931 using imagery from the Shippee Johnson expedition, and again in 1974 using aerial imagery acquired from the Servicio Aerofotográfico Nacional in Lima to trace the changes in agricultural production over a 43 year period (Denevan 1986). Digitizing Denevan's maps for direct comparison to the AI-Assisted survey would be instructive about the performance of the survey, but more importantly would allow us to track the changes that have occurred after another 50 years, thereby gaining a better understanding of the dynamics and transformations of agricultural production in the Colca valley over the course of a century. This and similar local projects, conducted in cooperation with local experts, will be vital to developing not only a broad, but deep understanding of the dynamics of agricultural infrastructure.

Another important dimension of future research is developing better understanding of the data produced by the AI-assisted survey itself. While the overall metrics of the survey were good, it was not equally reliable everywhere, with some regions appearing to contain higher rates of false positives than others. While much of this imbalance was removed through manual data cleaning, understanding and correcting for these patterns will be important as we deploy the model to new images and expand the inventory to new regions.

Finally, constructing further deep-learning models to further improve the quality and kinds of information available will allow researchers to ask new kinds of questions. Currently, the inventory has mapped the location of agricultural fields at the level of agricultural complexes. However, there is great variation in the types of fields that may occur within a complex. Denevan (1986) discusses 6 types of terraces in the Colca Valley alone, (largely shaped by their ecological context), and Langlie (2018) demonstrates that the morphology of terraces within a complex can reveal patterns of social resistance, in addition to those of social control commonly attributed to Inka terraces. Future models may be designed to map the boundaries of individual fields (Wang, Waldner, and Lobell 2022), rather than those of field complexes, allowing for this kind of morphological analysis at trans-regional scales. In each of these applications, AI-assisted survey methods promise transformational research at previously impossible archaeological scales.

CHAPTER 5

Agricultural Infrastructure and Reducciones in the River Valleys

Authors: James Zimmer-Dauphinee, Steven Wernke, Parker VanValkenburgh

5.1 Introduction

Prehispanic agricultural production in the highlands of the western cordillera of the southern Peruvian Andes was made possible through the construction of monumental-scale infrastructure projects by prehispanic andean farmers. Terracing covers much of the slopes of the highland, and are fed by lengthy canals, bringing water from streams and rivers to thousands of hectares of agricultural fields. Agricultural surpluses in these valleys helped to feed empires, allowing the Inkan Empire to reach its peak in population and agricultural production during the Late Horizon (Chavez 2019; Janusek and Kolata 2004; Wernke 2013). However, the indigenous population and the agricultural landscape underwent another transformation following the Spanish invasion. Under the command of the Spanish Crown, Viceroy Francisco de Toledo sought to increase political stability, evangelization, and economic extraction in the Viceroyalty of Peru. Toledo's reforms included the "General Resettlement of the Indians", which forcibly displaced over 1.5 million indigenous people from their traditional homes into compact, gridded towns with a devastating impact on the native Andean populace. It disrupted traditional communities and land-use arrangements, displaced people from their productive lands, and exacerbated the transmission of European-introduced diseases.

The dynamics of the Spanish invasion and its aftermath on agricultural production are still poorly understood. Historical data and archaeological research suggest that Spanish influence and power were differentially implemented across the landscape with varying results (Wernke 2007a, 2013; Davies 1984), yet most research into the particular dynamics of Reducción placement and agricultural deintensification in the western south-central Andean highlands has been restricted to a single valley. Over 30 years of research in the Colca Valley (Treacy 1987; Wernke 2010) have given key insights into the construction of reducciones, the abandonment of agricultural fields (Denevan 1988; Guillet et al. 1987; Treacy 1990b) and the complex and dynamic relationship between these historical facts (Wernke 2013). However, the extent to which the lessons learned from this research can be applied to a broader regional or trans-regional context has not been demonstrated. This research seeks to address this shortcoming by examining 3 key questions about the placement of reducción settlements and the effects of this on the process of agricultural deintensification that occurred following the fall of the Inka empire. The primary questions this research explores are:

1. Were reducción settlements strategically located to ensure accessibility to existing agricultural infrastructure?

2. Does a greater distance from a reducción settlement increase the odds of agricultural field abandonment?

3. Does the importance of proximity to reducción settlements in the process of agricultural deintensification vary across the study region?

Question 1 aims to improve our understanding of the street-level process of negotiation and creativity that was undertaken between Spanish officials and indigenous populations during the planning and construction of reducción settlements. Official Spanish colonial policy fundamentally conflicted with the prehispanic Andean economic system on which it was dependent. Official policy offered no solutions for this problem, leaving it to the discretion of, local inspectors, *encomenderos*, and *curacas* to resolve, including through the selection of reducción location (Hemming 1970, p. 394). This question, therefore, seeks to examine the importance of access to agricultural infrastructure in this process. Question 2 examines the dynamics of the agricultural systems once the resettlement had occurred, seeking to understand whether increased travel time to reach fields had a strong effect on the productivity and ultimately maintenance of those fields. In this phase, the negotiation between Spanish and indigenous populations continued, as reducción residents were forced to either abandon very distant fields or abandon the reducción to maintain them. The third question examines the region for variation in the results of these negotiations, seeking to understand comparatively where differences in the outcome may have occurred. Together, these questions lead to new hypotheses and productive directions for future investigation. While variation in the results can be mapped across the landscape, precisely what social, political, and economic factors shaped these results will require further archival and field-work. The results of this research, therefore, suggest new avenues for future research that the data suggest are most likely to provide new insights into the dynamics of resettlement and agricultural production.

5.2 Archaeological Context

5.2.1 The "Reducción General de Indios"

As discussed in previous chapters, the Reducción General de Indios, literally the "General Resettlement of Indians" was one of the largest forced resettlement programs in history, relocating over 1.5 million native Andean people into over 1,000 planned towns (Hemming 1970, p. 395). These towns were designed to bring Spanish conceptualizations of civilization to native Andean peoples through the imposition of Mediterranean-style urbanization as a part of Peruvian Viceroy Francisco de Toledo's plan in the 1570s to create order and control following decades of violence, and unrest in Peru (Stern 1993, p. 80). The design of the reducciones placed the spiritual, governmental, and social lives of its citizens in the physical center of town with a large rectangular plaza adjoined by the church and surrounded by gridded streets, the church, and the cabildo (town hall) (Mumford 2012; Wernke 2013, p. 214). In contrast to this Spanish ideal of urbanism, late prehispanic Andean settlements were generally smaller and dispersed throughout the valley to take advantage of the diverse, vertically-stacked, ecological zones of the Andean river valleys (Murra 1972). Distributing the population throughout the valley in this way, Andean people were able to leverage a wide array of agricultural and pastoralist resources, from cotton on the coast to coca and maize in the mid valleys, to potatoes and quinoa in the upper valleys and camelids in the high altitude puna (D'Altroy 2014; Davies 1995, p. 174; Mumford 2012; Murra 1972; Stern 1993, p. 4). Andean farmers constructed terracing on valley walls to utilize irrigation networks which provided water in the arid and semi-arid environment, allowing regions such as the Colca Valley to become known as a breadbasket of the prehispanic Andes.

Spanish efforts to disrupt this system through forced centralization into urban settlements endangered the pre-existing social and economic networks. This was a problem for the Spanish whose primary goals were the parasitic extraction of resources from Andean people, forcing the colonizers to balance their desire for control and surveillance with the viability and sustainability of their attempted imposed system (Mumford 2012; Stern 1993; Wernke 2007b, 2013). This offered substantial opportunities for Andean people to negotiate, manage, and shape the terms of the resettlement. Precisely how Spanish colonizers and indigenous Andean people negotiated this balance is unclear, and likely varied across communities and local officials. Toledan guidance about the process of selecting locations for reducciones emphasized level terrain, healthy climate, and minimizing disruption to the local economy by placing them on top of major pre hispanic settlements, while simultaneously locating them as far as possible from pre hispanic cemeteries and shrines, and abandoning and burning pre hispanic settlements (Mumford 2012, pp. 119–121; Wernke, VanValkenburgh, and Saito 2020, S66). In sum, Toledo's guidance was vague, contradictory, and left much to local magistrates. Therefore, the balance between the reducción ideal and practical reality was likely highly varied and contingent on local processes.

Stern writes that in Huamanga, Andean people often abandoned the reducciones with the support of local officials, returning to live closer to their dispersed landholdings (1993,

p. 90). Such a pattern has also been shown in Huarochirí, with the population quickly moving to new post-reducción towns in the decades that followed the resettlement (Oré Menéndez 2022). Meanwhile, Wernke's research on the Colca Valley has shed some light on how the process of negotiation played out there, where reducción residence was enforced. Drawing on visita documents as well as local knowledge and toponyms, Wernke was able to show that reducción settlements were placed closer to fields controlled by higher-ranking "right side" ayllu groups than the fields of the lower-ranking "left-side" (Wernke 2007b, 2013, p. 283) and that modern patterns of field abandonment are patterned with respect to reducción locations. This implies that, at least in the Colca Valley, local prehispanic social structure and associated agricultural landholdings played an important role in selecting the location of the reducciones.

These relationships between pre-existing agricultural infrastructure and reducción settlement locations are intriguing as they offer insight into the strategies and negotiations employed in the emplacement of reducción settlements in a way that can be measured at regional scales. In cases such as the Colca Valley, where Andean people did indeed resettle into new towns and the locations of those towns were the product of Andean and Spanish negotiations of ideals, one would expect the reducción town to be as close as possible to existing agricultural infrastructure. In contrast, the need for reducción settlements to be located near existing agricultural infrastructure would not be as great if indigenous people continued to maintain residences outside of the reducción (Oré Menéndez 2022; Stern 1993). As a result, one would expect the maintenance of agricultural fields to be less affected by the proximity of reducción settlements. Of course, the exploration of these patterns cannot offer any final answers to the modes of negotiation and management employed during the resettlement by the Spanish and Indigenous people. Many different patterns likely existed, including but not exclusive to the two models discussed above and equifinality suggests that many different processes may lead to similar outcomes. Nevertheless, a thorough examination of the relationship between agricultural infrastructure and reducción settlement location will allow us to generate hypotheses about what variations may exist in reducción placement strategies across the region which can be explored further in future research.

5.2.2 Agricultural Deintensification and Ruralization

While the distribution of agricultural infrastructure may give indications about the process of negotiation and decision-making that led to the placement of reducción settlements, the continued use or abandonment of said agricultural infrastructure may provide information about the ultimate results and implications of those strategies. In the Colca Valley, agricultural fields further from reducción settlement locations were more likely to be abandoned than those that were more proximate, with other variables such as elevation being held equal (Wernke 2013). This suggests that regions of once productive land were no longer maintained as the burden of traveling from centralized settlements to distant fields became too great to bear (Davies 1984; Mumford 2012; Stern 1993; Wernke 2007b). This process would also be self-reinforcing, as a lack of regular maintenance on distant fields and canals would make them less productive and reliable, and therefore less worth an investment of future labor. With canals in particular, this could become a problem for more proximate fields, if they are reliant on distant canals for reliable irrigation, thereby violating the arbitrary Spanish distinction between urban and rural. This is a process that we are terming "ruralization," signaling its dialectic production as a result of an enforced Mediterranean-style urbanization.

5.2.3 The Study Region: The Pacific Drainages of the Southern Peruvian Highlands

As described in the previous chapter, the region of interest for the AI-assisted agricultural inventory project covered both the river valleys of the western cordillera of the southern Peruvian Andes and the southwestern boundary of Lake Titicaca. The agricultural fields around Lake Titicaca compose the majority of agricultural land in the study region of Chapter 4, however, the processes of agricultural development and maintenance in the region immediately adjacent to Lake Titicaca are very different than those in the river valleys, due to the warming effects of the lake and the high availability of water. Modeling these two processes together does not make sense as the decision-making on whether to construct, maintain, or abandon a field is very different in each of these contexts. In contrast to the Titicaca fields, the river valleys are characterized by highly restricted availability of agricultural land, limited access to water, and dynamic temperatures across a wide variety of slopes and elevations. This combination of restriction and dynamism is particularly interesting, as it provides a greater need, and opportunity, for creativity in managing agricultural infrastructure, and may lead to more starkly differentiated distributional patterns. This research, therefore, chooses to focus analyses on the river valleys of the western cordillera of the south-central Andes, as shown in Figure 5.1.

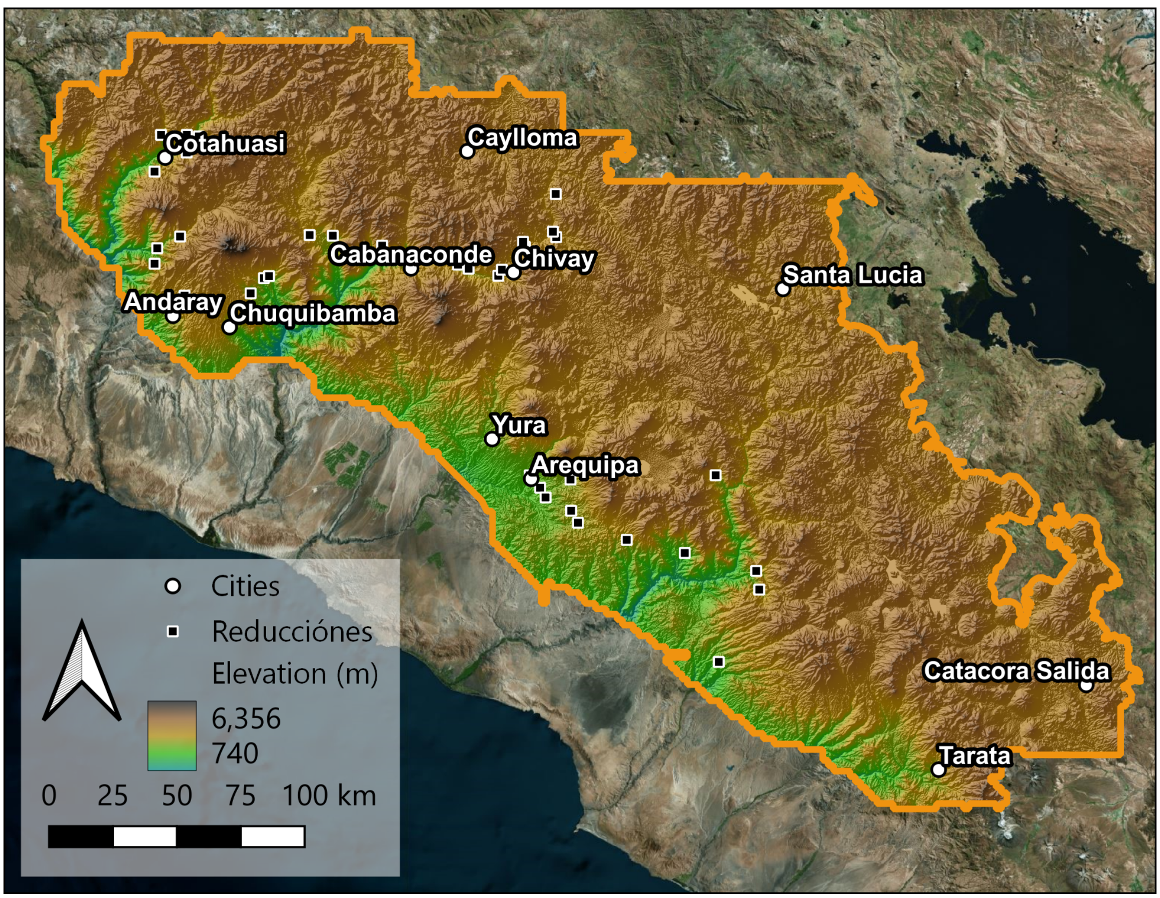

Figure 5.1: Western River Valleys Survey Region. This map also shows the locations of known reducción settlements used in this analysis.

5.2.4 Simplifying Assumptions

This analysis is premised on two major assumptions, namely: 1) The distribution of agricultural fields in the modern-day river valleys of the western cordillera of the southern Peruvian Andes is approximately equivalent to that at the time of the Spanish invasion. 2) The distribution of currently-maintained and abandoned agricultural fields is significantly shaped by the results of the Toledan resettlement program and the social, economic, and ecological restructuring that resulted.

These are significant assumptions that likely do not hold true everywhere in the study region, and may not be fully valid anywhere. For example, the Majes-Siguas irrigation project has generated approximately 15,950 ha of new agricultural land since the 1970s, drawing water from the Colca Valley to irrigate desert planes nearly 100 km away (Autodema 2016). Fortunately, these fields fall outside of our study region and, in any case, would be easy to exclude from any analysis. Excluding such major agricultural development projects, there is hope that assumption 1 is valid in most cases as the construction of agricultural fields, (and particularly agricultural terracing), is a highly labor-intensive process. Historical records show that Spanish colonists in Arequipa undertook limited efforts to produce new fields, preferring to use the extensive field systems already in existence (Davies 1984). The severe population decline during and following the conquest ensured that there were extensive fields available for farmers who sought land, and it required far less effort to maintain a field than to construct a new one (Davies 1984). Another point in favor of the relatively static distribution of agricultural fields is the success of efforts in the Colca Valley to match current field toponyms to the names found in the visita documents (Wernke 2013). This suggests that field distributions (and generational knowledge about them) have remained reasonably stable over the last 400 years. As fields are maintained through kinship networks across generations, they are not just locations of agricultural production but retain deeper meanings of social relation and reciprocity, making them resilient and stable places in the landscape.

The second assumption is harder to justify, given nearly 400 years of economic transformation, Peruvian Agrarian Reform, long-standing governmental and local interest in expanding agricultural production (Denevan 1988; Kendall and Drew 2019; Treacy 1987, 2019), and the documented dynamism of growth and contraction of the maintenance of agricultural fields through time (Denevan 1986). Nevertheless, this assumption once again appears to hold true in the Colca Valley, where there is a clear relationship between the proximity to reducción settlements and the active or abandoned condition of agricultural fields. In general, fields closer to the reducción are less likely to be abandoned, contingent on the elevation (risk of frost) and other social/environmental factors. It, therefore, appears that, at least in the Colca Valley, broad patterns of agricultural deintensification that were established during or as a result of the resettlement persist today, perhaps due to the substantial effort required to revitalize agricultural terracing and the canal systems that support them. This paper brings together two major sources of regional-scale archaeological data, the locations of reducciones and the locations of agricultural fields, along with an array of ecological and environmental variables to examine the process of agricultural deintensification as a result of the 16th-century forced resettlement of Andean people.

5.3 Data

5.3.1 Reducción Locations

The locations of a remarkable number of the more than 1000 reducciones estimated to have been constructed during the Toledan Resettlement are known (Mumford 2012; VanValkenburgh 2017; Wernke 2007b; Wernke, VanValkenburgh, and Saito 2020). As an ongoing collaborative project, the Linked Open Gazetteer of the Andean Region (LOGAR) (Wernke and Saito 2019) has been connecting Colonial-era documentation to modern cartographic features such as toponyms, and (primarily) modern towns, and villages. Through this work, the locations of approximately 673 reducciones across Peru have been identified, forming the most comprehensive database of reducciones compiled so far (Wernke, VanValkenburgh, and Saito 2020, p. 564). Initial analysis of this data by LOGAR researchers shows that reducción settlements tend to be either coastal settlements below 250m, or highland settlements, which have a peak in density between 3000 and 3500 meters (Wernke, VanValkenburgh, and Saito 2020, p. 567). LOGAR identifies 37 reducción features in our study region, which forms the bulk of the data used in this analysis. However, the LOGAR data's reliance on modern cartographic features means that it misses reducción settlements that are no longer occupied, or have changed their names. Therefore, the reducción location data was further augmented with data from GeoPACHA.

One of the principal survey objectives for GeoPACHA was the identification of missing relict reducciones, however, this data is still under review, making it difficult to evaluate how much confidence to place in the features identified. Therefore, this paper takes a conservative approach. Of the 47 features in GeoPACHA marked as reducciones within the study region, only 3, (Polobaya, Huamanmarca I, and a feature designated by Locus Id 5983) were deemed sufficiently likely to be reducciones to be included in the dataset, resulting in a total of 40 reducción locations to be used in modeling. We do not expect this to be a complete census of reducciones in the survey area (no such census exists), and many towns (particularly ones that are no longer occupied) may be missing from the dataset, however, we hope that this data nevertheless provides a sufficient cross-section of settlements to speak broadly about their distribution on the landscape. Furthermore, we acknowledge that reducciones were not the only settlement type that may have influenced agricultural deintensification during the late 16th and early 17th centuries. For example small rural hamlets, often populated by workers who left the reducción to care for fields or herds, were increasingly granted official legitimacy as an ”annex” in the 17th and 18th centuries (Abercrombie 1998; Oré Menéndez 2022, pp. 160–162).Nevertheless, as one of the largest and most immediate disruptions of settlement patterns and the social order that generated them, we have selected the reducciones as our unit of interest for this analysis.

5.3.2 Field Locations

The second major source of regional archaeological data is the results of the AI-assisted inventory of agricultural fields conducted in the previous paper. Some modifications have been made to the data described in the previous chapter for this analysis. Data cleaning has continued, further reducing the number of false positives in both the active field (fields currently under cultivation) and abandoned field (fields no longer under cultivation) datasets. The changes in this are marginal compared to the extent of active and abandoned agricultural fields identified in the region but may be relevant as this paper seeks to move between local and regional scales. In spite of this additional cleaning, the data still remains imperfect with occasional false positives and negatives. These problems are particularly prevalent in the Abandoned Field data set, as identifying abandoned fields is a much more challenging task for both AI and human surveyors as they more closely resemble other environmental features and their boundaries are often less well defined as terrace walls decay. These imperfections are important to remember in the modeling process and in model interpretations. Nevertheless, as shown in the previous chapter, this data remains more complete and more precise than any other survey at this scale known to the authors.

Secondly, in contrast to the previous chapter, this research will focus on the river valleys of the western cordillera of the south-central Peruvian Andes, excluding the region around Lake Titicaca as shown in Figure 5.1. This is because the following analyses, (which take into account elevation, slope, and walking distances from reducción settlements) are premised on verticality, aridity, micro-ecological diversity, and the constraints of cultivable land in the valleys. In contrast, the agricultural fields on the Lake Titicaca floodplain are more expansive, frequently designed to deal with excesses of water, and are ecologically moderated in temperature and precipitation by the environmental effects of the lake (Erickson 2000; Janusek and Kolata 2004; Kolata 1986; Stanish 2003). This forms a fascinating and distinct set of processes that must be understood in their own context, a project that we will undertake in the future, but is not the focus of this dissertation.

Table 5.1: Correlation matrix for a selection of environmental variables. Notice the high correlation with Elevation.

	Elevation	Mean Annual Temp	Mean Diurnal Temp Range	Precip Wettest Quarter
Elevation	–	–	–	–
Mean Annual Temp	-0.94	–	–	–
Mean Diurnal Temp Range	0.41	-0.54	–	–
Precip Wettest Quarter	0.67	-0.48	-0.04	–

5.3.3 Environmental Correlates

Several environmental correlates were selected for analysis, either due to apriori expectation or experimental demonstration that they are relevant to agricultural production/deintensification. Potentially relevant environmental correlates identified for these analyses were mean annual temperature, mean diurnal temperature range, mean precipitation of the wettest quarter, elevation, slope, aspect (direction of slope face), geomorphon classification (classification of local terrain shape), distance from streams, and walking cost-distance from reducción settlements. The first three of these correlates were obtained from the WorldClim database of weather and climate data (Fick and Hijmans 2017), and apriori are expected to have an effect on agricultural production. Frost risk is a major concern in the highlands, potentially damaging or destroying crops, and the arid-to-semi-arid environment demands that crops be grown in the wettest season. However, these variables were all highly correlated with each other and with Elevation (Table 5.1). As can be seen in Figures 5.2 and 5.3 the higher the elevation, the lower the temperature and the wetter the climate.

Such multicollinearity is a substantial problem for developing rigorous statistical models, making the results inaccurate and difficult to interpret. Therefore the 30m Elevation model from the Shuttle Radar Topography Mission (SRTM) (Center 2017) was selected as a proxy variable to simultaneously represent each of these biologically relevant climatic

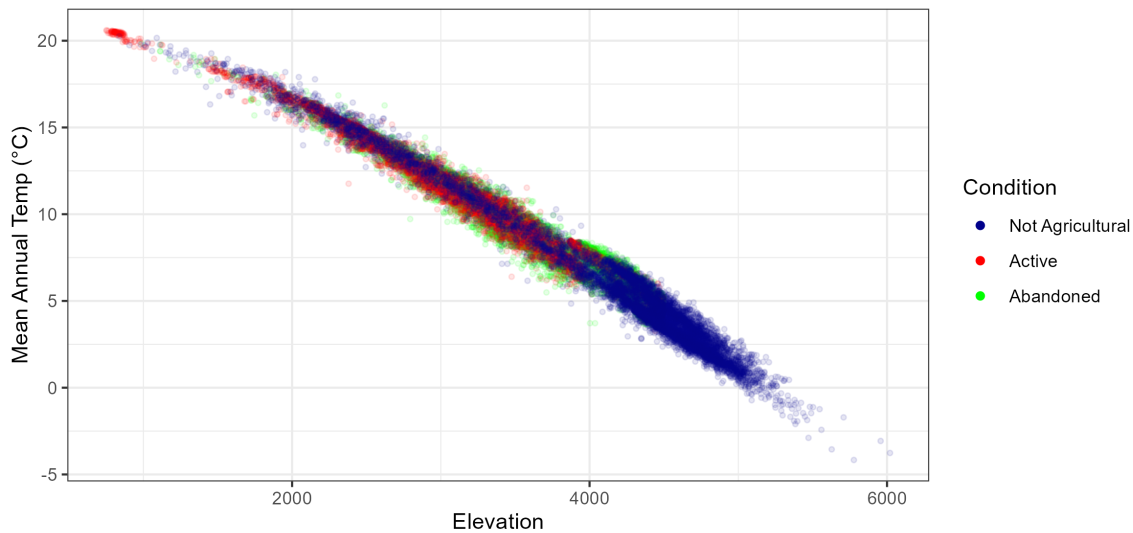

Figure 5.2: Scatter plot showing the high correlation between elevation and mean annual temperature

variables. This operationalization decision mirrors the way in which Elevation is often discussed in relation to agriculture in the academic literature. Though it is common to say "no agriculture occurs above 4500m," this observation is more a commentary on the covariate of temperature which results in temperatures too cold to support agriculture above 4500m, than a result of the inherent properties of elevation itself.

Slope (in degrees), aspect, and geomorphon classification were all derived from the SRTM Elevation model using the GRASS Geographic Information System as described in the previous chapter. However, the geomorphon variable was not used in the final modeling procedure because it increased the model's complexity without substantially improving the model's predictions, likely due to its correlation with the slope variable. The "Flat" geomorphon category, for example, is equivalent to a slope of 0. Aspect was similarly rejected for inclusion in the model as it did not contribute meaningfully to the model's predictions, as may be expected from the exploratory analysis in the previous chapter.

A linear distance from streams dataset was also generated for these analyses using GRASS GIS. Stream locations were calculated from the digital elevation model (DEM)

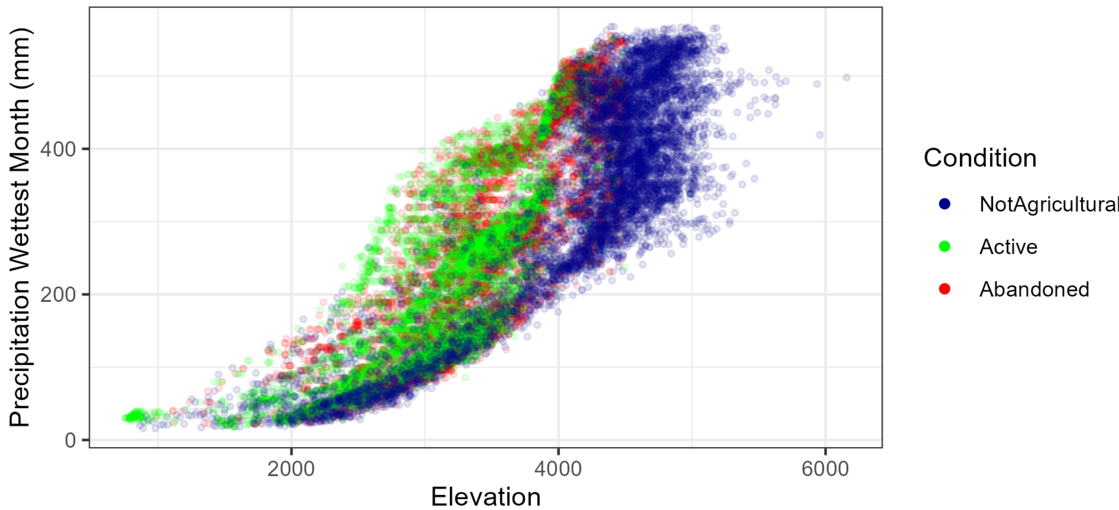

Figure 5.3: Scatter plot showing moderate correlation between precipitation and elevation.

using the GRASS r.stream.extract function with an accumulation threshold of 100,000. Visual inspection with satellite imagery shows that the resulting stream network corresponds well with the appearance of visible surface water in the imagery. The linear distance from these streams (in meters) was then calculated for the entire study region. However, this variable should be treated with caution, as streams used for irrigation in the river valleys are often fed by springs (Denevan 1988), which cannot be easily modeled from the elevation data. Furthermore, Andean people constructed canals to transport water great distances from the streams from which it originates. Indeed, it has been argued the primary purpose of terracing was to facilitate this irrigation (Guillet et al. 1987; Treacy 1990a). Therefore, the distance from a stream may not be a perfect predictor of the availability of spring water. However, it seems reasonable to assume that the longer the canal that is required to transport water, the more labor would be required to maintain that canal across the whole distance. Therefore, it is likely that fields which are reliant on the canal to transport water a great distance from the original stream are more likely to be abandoned than those that are more proximate to the stream and may have a shorter canal.

Finally, a cost-distance raster estimating walking times (in hours) from reducción settlements was calculated using GRASS's r.walk function which uses the Aitken (1977)/Langmuir (1984) hiking function to estimate the time it takes to walk a distance based on the ground's slope. Though not calibrated for the Andes specifically, this analysis is preferable to linear distance as it takes into account the effects of slope on human walking speed, making it a more reliable measure of human experience (White 2015). Additionally, the highly efficient implementation of this algorithm in GRASS makes it possible to calculate the walking time across the entire study area. Unfortunately, this dataset must be interpreted cautiously as well, as it is almost certain that our reducción dataset is incomplete. Locations which appear to be very far from a reducción may in fact be very close if there is an unidentified reducción or other settlement such as an annex or farmstead is nearby. Nevertheless, for short distances, where it is less likely that an unknown settlement will dramatically change the values, this data may provide crucial insights into the effects of distance from reducción on agricultural deintensification.

5.4 Methods

5.4.1 Modeling Field Locations

A first step to understanding the effect of reducción settlements on agricultural deintensification is to develop a model of where agricultural fields could be produced and maintained on the landscape. This model of what affects the distributions of agricultural infrastructure and deintensification helps us to understand what environmental variables are important to these processes and what their effects might be. As explored in the previous paper, for example, Elevation and slope both appear to significantly modify the odds of the presence of agricultural fields, or an existing field's abandonment, as do the mean diurnal temperature range, and average rainfall in the wettest month. Aspect (the direction the slope faces) in contrast, does not. This chapter seeks to more formally model these relationships using multinomial regression.

There are three main classifications of the landscape we are interested in understanding for this phase of the analysis:

1. Currently-in-use “Active” agricultural fields

2. “Abandoned” agricultural fields

3. Uncultivated or “Not Agricultural” land

Multinomial regression extends binomial regression (also known as logistic regression) to allow the estimation of the likelihood of belonging to one of three or more classes based on independent covariates. In doing so it defines one class (in this case “Not Agricultural”) as the reference class, and then calculates the likelihood of belonging to another class in reference to this class. This allows us to evaluate the effect each of the covariates has on the odds of creating an agricultural field, and how likely that field is to have been abandoned. We generated a stratified random sample of 25,534 points, 7967 located within active agricultural fields, 9567 in abandoned agricultural fields, and 8000 in the not cultivated background. The ecological variables were then sampled at each of these locations. Slight differences in the number of points for each category result from the randomness of sample generation, however, the approximate balance of the data should provide the best estimate of model coefficients. Of primary interest in this exploratory model is the sign and magnitude of the estimated coefficients, representing whether each covariate makes the cultivation and abandonment of a particular location more or less likely and the size of that effect relative to the other covariates respectively. This will enable us to include the most relevant variables in later analyses of agricultural deintensification in relation to reducción settlement location.

5.4.2 Locating Reducciones

The next phase in this analysis is to construct a formal model of the locations of reducción settlements on the landscape. The goal here is to generate a null model that is capable of simulating plausible distributions of reducción settlements. For example, it was very unlikely that a reducción would be constructed above 4500 meters, or on a very steep slope. Before we can understand other features of reducción distributions, such as whether there is evidence of clustering, or whether they are located especially close to agricultural fields, we must first take these environmental features into account.

The R Statistical package "SpatStat" was used to construct an inhomogeneous Poisson Point Process (PPP) model to describe the distribution of reducción settlements in the study area. Inhomogeneity, in the context of PPP modeling, means that the intensity of the point pattern (i.e. the frequency of reducciones) is affected by covariates, such as Elevation or Slope. Both a priori and empirically, this is clearly the case. As discussed by Wernke et al (2020), the density of reducción settlements is strongly related to elevation, with most highland reducciones occurring around 3000m in and few located below 2000 or above 4000 (Figure 5.4).

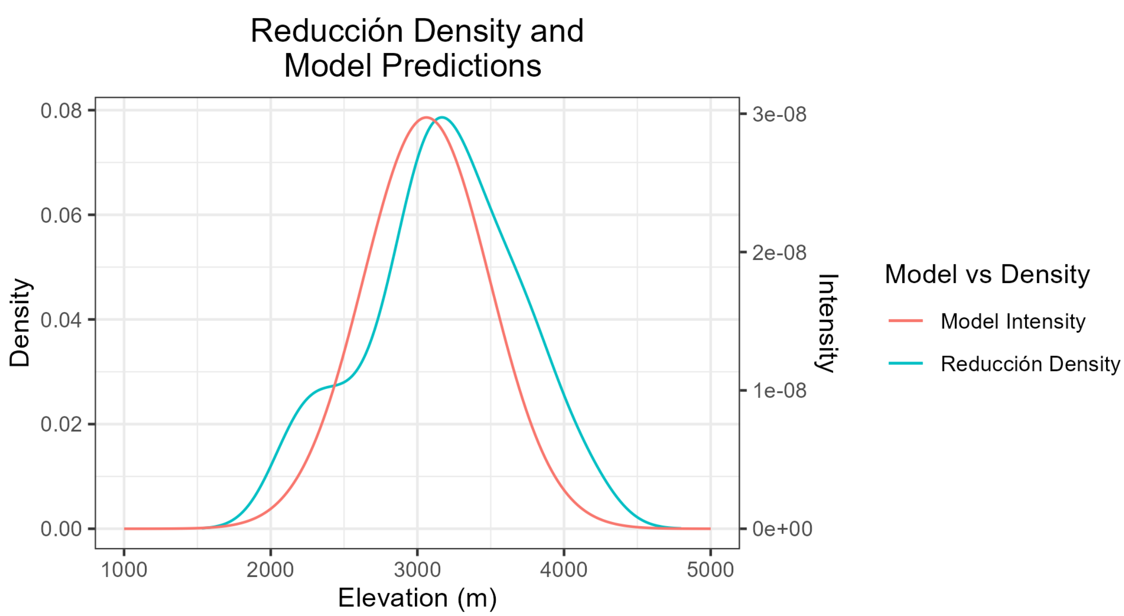

Figure 5.4: Plot of reducción density as a function of elevation and the corresponding intensity distribution estimated by the ppp model.

This is the most habitable zone of the highlands, as the environment becomes too arid at lower altitudes outside of the valleys, and too cold at higher ones. Exploratory analysis confirmed this non-linear relationship for the study area, with the density of reducción settlements gradually increasing between 2000 m and 3000 m and then gradually decreasing until 4500 meters. This relationship was therefore modeled as a quadratic function, shown in Figure 5.4. To assist in the model's numerical stability and assist in interpretation, the elevation data was divided by 100. The model coefficients therefore can be read as the change in the log odds for every 100 meters elevation change.

Toledo also advised that reducciones should be located on level terrain (Wernke, VanValkenburgh, and Saito 2020, s66). This is logical as it would be much more difficult to construct a large, gridded town on steep and uneven slopes, and appears to be supported by the data as the distribution of reducciones shows a clear tendency for lower slopes. While other considerations were certainly taken into account when locations were selected to construct reducciones, incorporating these two features into our analysis allows us to create a model of where it would be reasonable to consider constructing a reducción. The model is a Poisson point process because it does not model any interaction between the points, that is, the model assumes that the presence of one reducción does not make another reducción nearby more or less likely to occur. This is likely not the case in reality, as settlements may move from one location to another and be rebuilt as a part of ongoing negotiations between the ideals of urbanization and the realities of resource production and management in the Andes. These new settlements may be more opportunistically located but are likely to still be close to the old settlement location, resulting in a cluster of settlements. Alternatively, it may be unlikely that several reducción settlements would be constructed very close to each other simultaneously as this would force competition for already limited resources. This would result in an inhibition effect (more spaced than random) on the intensity of the point process. Nevertheless, it is useful to make the simplifying assumption of no interaction between settlement locations during their construction for the purposes of mitigating computational and mathematical complexity. While a log-Cox model (or other point process model) could theoretically capture such clustering and inhibition effects, the number of reducciones in the study region is limited (n=40), making it challenging to rigorously model such complex processes. Furthermore, the goal of this model is not to accurately capture the parameters of reducción construction, but rather to serve as a null model for the estimation of locations that would be suitable for the placement of a reducción. This goal can be achieved without taking into account the interaction effects between reducción settlements, as is demonstrated by the goodness-of-fit measures and tests discussed in the results section of this chapter. As this research expands to further regions and the size of the point pattern (number of reducciones) is greater, it may be possible to adopt more sophisticated modeling procedures, however, such complexity is not necessary at this time.

Having constructed a point process model for the intensity of construction of reducciones, it was possible to simulate point patterns that mirror the distribution of reducción settlements, resulting in patterns of reducción settlements that are plausible, but different from reality. Using the calibrated model, we simulated the construction of reducción settlements 1000 times, allowing us to compare the actual distribution of settlements to settlement locations simulated by the model.

5.4.3 Influence of agricultural infrastructure on reducción settlement location

A primary question of interest is whether access to pre-existing agricultural infrastructure was a major influence in the selection of reducción settlement locations. If this is the case, one would expect reducción settlements to have access to more agricultural fields (either active or abandoned) on average than the simulated settlements. For the purposes of modeling the accessibility of agricultural fields from settlement locations, we designated any field within a two-hour walk from the reducción settlements to be “accessible” to agriculturalists in that reducción. While it is possible that agriculturalists may occasionally have chosen to work fields more distant than this from the reducciones, in the authors’ estimation distances further than this introduces a serious burden on the potential productivity of the land as the laborers would have to spend significant portions of each day in transit, or (if possible) move away from the reducción. Rather, it seems likely that most full time-residents would have made great efforts to travel shorter distances than this. In Wernke's analysis of the fields of Corporaque, he found that the mean walking time to abandoned fields for "left side" ayllus was 109 minutes (1.8 hours), while the mean walking time to unabandoned fields was 46 minutes (.75 hours) for the same groups (Wernke 2013). A 2-hour cost basin, therefore, provides a generous outer bound for fields that may be considered "accessible."

The walking time from each reducción location (simulated and actual) was calculated using the SRTM DEM. First, each reducción point location was buffered with a radius of 10 kilometers, and the DEM was cropped to this extent, greatly reducing the computational load for calculating cost-distance and enabling the computation in R. Then the cost-distance was calculated using Tripcevich's hiking function in the R "movecost" package. Tripcevich's hiking function follows the form:

$$speed = \frac{4.028 \times 46}{((arctan(|slope|) \times \frac{180}{\pi}) + 4.127) + 46} * \frac{1}{N} \tag{5.1}$$

Where slope is expressed in degrees, thereby modeling travel speed as a function of slope. This model was selected for this analysis, (rather than other walking models such as Tobler's, Aitken 1977/Langmuir 1984, Garmy, Kaddouri, Rozenblat, and Schneider's hiking function, or others) as it was calibrated for movement in the Andean highlands by tracking the progress of a Llama caravan in 2007 as it carried salt to Chancara, just off the western edge of the study region. Though the dynamics of a long-distance llama caravan are likely different to those of agriculturalists traveling to their fields, Tripcevich describes the caravan's movement as steady travel all day, with short breaks and a lunch break. Such steady travel should well represent the movement of farmers to their fields.

The 2 hour walking time cost basin polygons for each simulated and real reducción settlement location were intersected with both the active and abandoned field polygons to calculate the total area of agricultural infrastructure within a 2 hour walk from the reducción location. The average accessible agricultural area was calculated for each of the 1000 simulated instances of reducción distribution and compared to the average accessible agricultural area of the 40 real reducción locations. This allows us to evaluate the access of reducción settlements to agricultural infrastructure in relation to the distribution produced by the null model of reducción locations.

5.4.4 Effects of Reducción Settlement Location on Agricultural Deintensification

Following previous research in the Colca Valley, we hypothesize that the forced resettlement of the population into reducciones in the river valleys shaped decisions to maintain or abandon extant agricultural fields, and the results of this process still visible on the landscape today (Wernke 2015). Fields located at great distances from reducción settlements would be very difficult or impossible to cultivate while maintaining a permanent residence in the reducción. As a result, we expect to see the rate of field abandonment increase with distance from reducciones across the study region. The dataset from the multinomial regression of field locations was filtered to only include examples of active or abandoned agricultural fields and a logistic regression model was constructed including the variables identified as important in the exploratory multinomial agricultural model, in addition to the walking-distance from reducciones variable.

However, it should be observed that in terms of an agriculturalist's ability to cultivate a field while maintaining a residence in the reducción, there is no meaningful difference between a field that is 6 hours away, to one that is more than 30 hours away. In either case, it is impossible to walk to the field, perform a full day's labor, and return to the reducción at night. Any variation in active and abandoned field distribution at this distance therefore must be the result of processes that this model does not attempt to evaluate, such as the maintenance of a remote residence. Furthermore, due to the incompleteness of the reducción dataset, it is impossible to evaluate whether a particular field is actually several days' walk from a reducción, or whether a much more proximate one is simply unknown. Therefore, only sample locations that were within 6 hours of a reducción were included in this analysis to evaluate the effects of distance from reducciones on deintensification.

5.4.5 Mapping Local Variations

The previous analysis assumes that maintaining a permanent and full-time residence in the reducción was mandatory and common practice. However, it is possible that this restriction was less robust in some instances than in others, with farmers allowed to maintain part or even full-time residence outside of town. This would make it much more possible to continue to cultivate and maintain distant fields, and we would expect the distance from reducciones to be much less powerful as a predictor of agricultural field abandonment. It is, therefore, useful to look for variations in the importance of this variable at local levels, rather than focusing only on the trans-regional pattern. One way to do this is to repeat the analyses many times while limiting the extent of the analyses to smaller areas (Kvamme 2005). We therefore divided the study region into 103 overlapping regions of 0.5°. Any data that was further than 6 hours from a reducción, any region that included less than 100 sample points total, or included less than 30 sample points representing either active or abandoned agricultural features, was filtered out. This helps to lower the noise resulting from outliers (false positives or negatives in the terracing dataset), resulting in a consistent measure of the importance of the distance from the closest reducción in the likelihood of agricultural field abandonment. We then repeatedly ran a logistic regression using elevation, slope, distance from streams, and distance from reducciones as covariates, (latitude was dropped as it is not expected it will have substantial effects for such small regions) using only the sample data points contained in that region, and mapped the coefficient for the distance from reducción variable.

5.5 Results

5.5.1 Modeling Field Locations

The exploratory multinomial model showed latitude, elevation, slope, and distance from streams to all be significant contributors to agricultural field construction and abandonment. Their coefficients (representing the change in log odds of being active or abandoned agricultural fields in comparison to the non-cultivated background) and their 95% confidence intervals are reported in 5.2.

Table 5.2: Multinomial model coefficients, odds-ratios, and confidence intervals.

Variable	Units	Model Coefficient	Δ Odds/Δ Unit	Odds 95% CI
Latitude:Active	Degrees	0.4919	1.635	1.548 - 1.728
Latitude:Aband	Degrees	0.2742	1.315	1.257 - 1.377
Elevation:Active	100 m	-0.2131	0.808	0.803 - 0.814
Elevation:Aband	100 m	-0.1297	0.878	0.873 - 0.884
Slope:Active	10°	-0.0585	0.943	0.940 - 0.947
Slope:Aband	10°	0.0058	1.006	1.003 - 1.009
Distance to Stream:Active	1km	-0.3396	0.712	0.699 - 0.725
Distance to Stream:Aband	1km	-0.0881	0.916	0.905 - 0.927

These values can be converted to the change in odds by exponentiating them, which is also reported for convenience and ease of interpretation. In Table 5.2 therefore, the Latitude change in odds is 1.6 for active fields, meaning there is a 1.6 times increase in the odds (or, equivalently, there is a 60% increase in odds) for an active field for every degree further North in the study region. Similarly, the Elevation (100m change) for active fields value of 0.808 shows that for every 100m increase in elevation, there is a 20% decrease in the odds of an active field. Caution must be used when interpreting these values, as they are contingent on the arbitrary selection of units used for this analysis. For example, the values would change significantly if one were to model changes in elevation by the meter (the change in odds will be much lower per meter change in elevation) or by the kilometer (the change in odds will be much higher per kilometer change in elevation).

It would not be correct, therefore, to suggest that in general a change in elevation has a greater effect than a change in slope. Rather, the model suggests that a 100m change in elevation has a greater effect than a 10° change in slope, the implications of which are not immediately clear. Rather, it is more useful to look at the direction of change (does field construction/abandonment become more or less likely), and to compare coefficients between the Active and Abandoned fields for a single variable. For a 100-meter increase in elevation, the model shows a 20% decrease in the odds of an active field, but only a 13% decrease in the odds of an abandoned field, if all other variables are held constant. This suggests that higher-elevation fields are less likely to be constructed and are more likely to be abandoned than fields at lower elevations.

Similar logical patterns hold for distance from streams and slopes, with both active and abandoned fields becoming increasingly less likely with distance from water, and the odds of active fields decreasing with an increase of slope. However, it is interesting to note that the odds of an abandoned field are nearly unaffected by increases in slope, and may even be very slightly increased. This reflects the high frequency of terracing on the valley slopes, and the high rates of abandonment of terracing in the valleys today, mirroring the observations made in the Colca Valley (Denevan 1986, 1988). The strongest effect across the region comes from latitude, which shows an approximately 60% increase in the odds of active fields and an approximately 30% increase in the odds of an abandoned field for every degree further north. This may be due to the increase in warmth as one moves closer to the equator, but likely is mostly the result of the selection of the survey area in relation to the agricultural infrastructure. The southern portion of the survey area is very dry and contains little to no agricultural fields, resulting in a marked increase as one moves north toward less arid climates.

With these four covariates, the model achieves a McFadden's pseudo R^2 value of 0.21.

Table 5.3: PPP model coefficients, odds-ratios, and confidence intervals.

Variable	Units	Model Coefficient	Δ Odds/Δ Unit	Odds 95% CI
Elevation	100m	1.6437	5.1744	2.816 - 9.509
Elevation^2	100m	-0.0268	0.9735	0.965 - 0.983
Slope	Degrees	-0.1888	0.8280	0.775 - 0.885

Given the complexity of the phenomenon, the restriction of this model to purely environmental variables, and the limitations of the data, this is a moderately good fit. More importantly, the model highlights which environmental variables are important to agricultural production and field abandonment and produces a baseline understanding of the direction and strength of these effects which can be used to better interpret the effects of other factors in agricultural deintensification.

5.5.2 PPPm

The next stage of analysis was the construction of the inhomogeneous Poisson point process model. This model shows that slope and elevation also have statistically significant effects on the frequency of reducción settlements in the highlands, at a less than 0.01 level (model coefficients summarized in 5.3).

However, of greater interest to our research is how well the model does at estimating the intensity distribution of reducciones across the landscape contingent on elevation and slope. A visual inspection of the intensity map produced by the model in Figure 5.5, showing that the model expects higher intensities in river valleys, particularly the Arequipa, Colca, and Cotahuasi valleys. This aligns well with the known distribution of settlements and matches intuitively our expectations about where reducciones may be located.

Another check of model fit is the graph of the envelope for the K statistic for data simulated by the model (Figure 5.6). This graph shows the distribution of the K statistic (which tests for clustering and dispersion) for 100 simulated point patterns, the statistic for the true data in black, and the theoretical value for the statistic in red. A poorly fitted model would differ widely between the observed and the theoretical data, with the observed

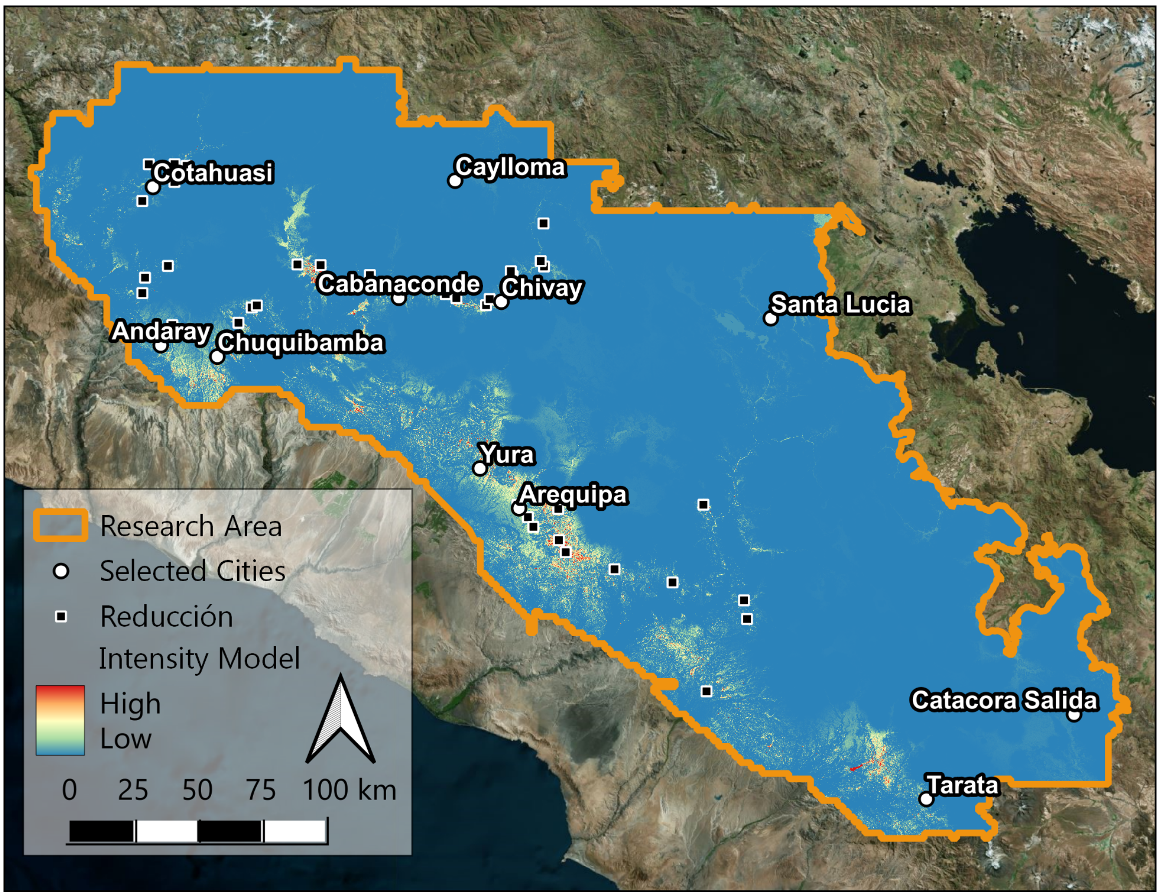

Figure 5.5: Estimate of Intensity of Reduccion settlements by inhomogeneous PPP model accounting for elevation and slope. Note the correlation with known point distribution.

data landing outside of the simulated envelope for much of the graph. Instead, our model aligns the theoretical and observed data nearly perfectly and never deviates outside of the envelope.

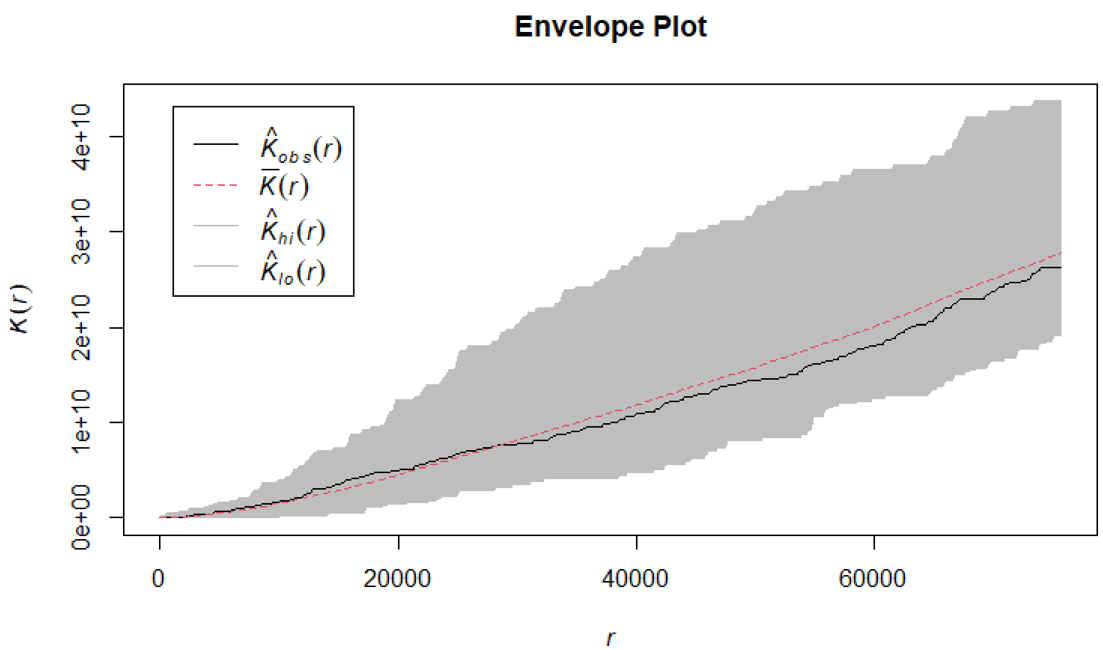

Figure 5.6: Plot of PPP model envelope for K statistic, suggesting a high quality of model fit.

Finally, a Balanced Independent Two-Stage Monte Carlo Test (designed to test the goodness-of-fit for spatial patterns) fails to reject the null hypothesis that 100 model simulations and observed data come from different distributions with a p-value of 0.99. Together these tests suggest that the inhomogeneous PPPm fits the data well and can be used to simulate plausible distributions of reducción settlements.

5.5.3 Reducción Locations and Agricultural Infrastructure

1000 simulated point patterns representing plausible distributions of reducción settlements were generated. These simulations produced between 10 and 42 locations each, with an average of 25.4 points per simulation for a total of 25,401 plausible reducción locations. Cost basins representing a 2-hour walking time were generated for each of these settlements, covering 433,407 ha per simulation on average. For each simulated point pattern, the agricultural area within its cost basin was calculated and divided by the number of points in that simulation. So, for example, the first simulation had 21 points whose cost basins covered a total of 18052.5 ha of active and abandoned agricultural fields, for an average area of 18052.5 ha / 21 sites = 859.6 ha of agricultural land. Repeating this calculation for each simulated point pattern produced 1000 average measures of access to agricultural land for different plausible distributions of reducción settlements.

Following the central limit theorem, this distribution of sample averages produces an approximately normal distribution with a mean of 717.7 ha and a standard deviation of 171.2 ha. Repeating the process once more for the locations of actual reducción settlements collected from LOGAR and GeoPACHA databases showed that the true reducción settlements had an average of 1241 ha of accessible land with agricultural fields, placing it on the far-right tail of the distribution (Figure 5.7), more than 3 standard deviations above the average.

Only 5 of the 1000 simulations had access to more agricultural land on average. Taking as a null hypothesis that the locations of reducción settlements are not related to the accessibility of agricultural infrastructure, this result strongly rejects this hypothesis, suggesting instead that reducción settlements in the river valleys were carefully and intentionally placed with greater than average access to agricultural infrastructure.

5.5.4 Effects of Reducción Settlement Location on Agricultural Deintensification Global and Local

To better understand the effects of distance from a reducción on the likelihood of field abandonment, we constructed a binomial regression model to predict whether a field was active or abandoned, using the environmental variables identified as important in the multinomial model and the walking cost distance from a reducción in hours.

Latitude no longer appears to have a large effect, as it did in the multinomial model.

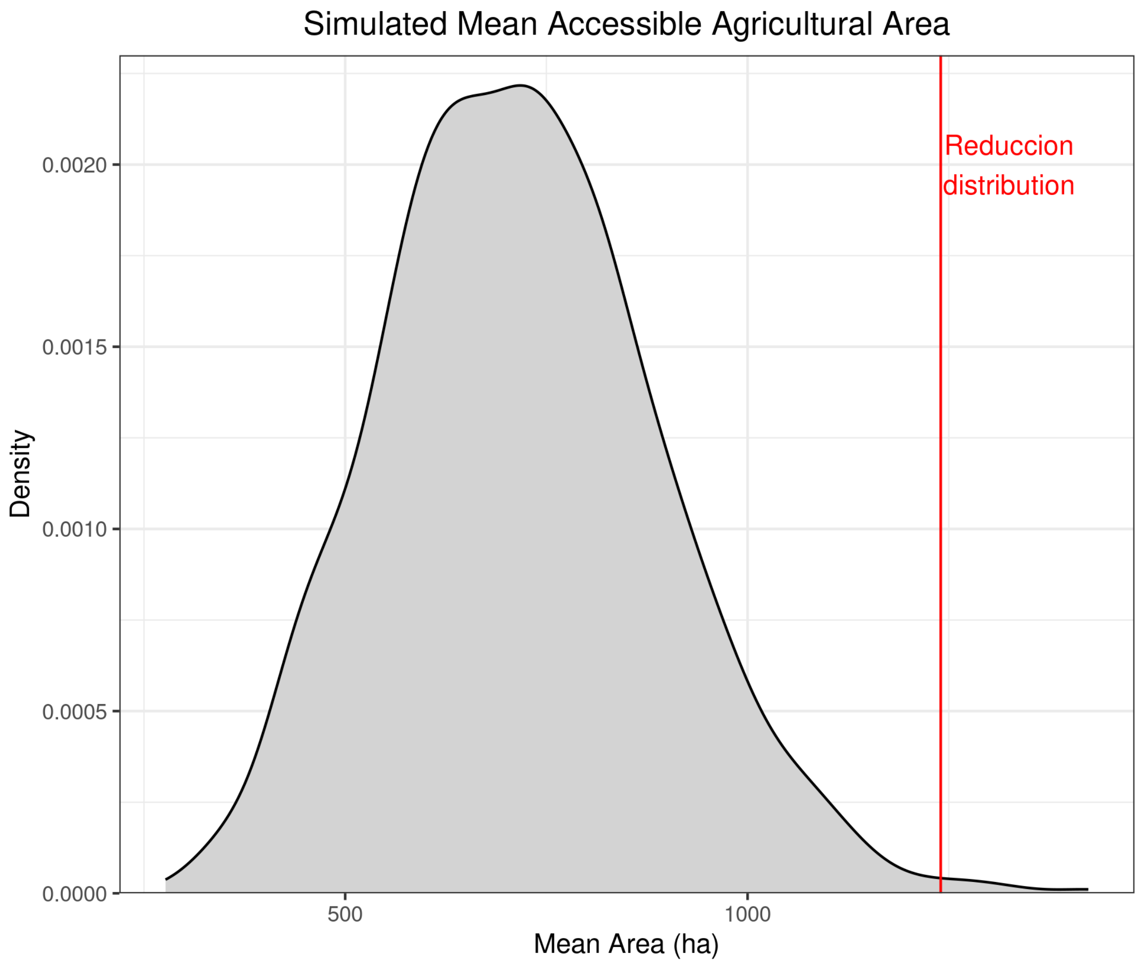

Figure 5.7: Distribution of mean area of agricultural infrastructure within 2 hour walk of 1000 simulated reducción patterns. Note mean for true reducción distribution is in the 99.5th percentile.

This is likely in large part due to the restriction of the model to the relatively narrow bands of latitude where agricultural fields are actually located. Distance to streams, in contrast, appears to have the largest effect on the odds of abandonment. While this is suggestive that the distance water must be transported by canals is predictive of field abandonment, the relatively large scale of the units (kilometers) may overemphasize the importance of this variable. When taken with the above caveats about this variable, the significance of this result should be treated with caution. Of greater interest, the elevation variable shows an approximately 6.6% increase in the odds of field abandonment with a 100 meter gain

Table 5.4: Binomial regression model coefficients, odds, and confidence intervals in the prediction of whether an agricultural field will be maintained or abandoned.

Variable	Units	Model Coefficient	Δ Odds/Δ Unit	Odds 95% CI
Latitude	Degrees	-0.0472	0.9539	0.876 - 1.038
Elevation	100 m	0.0637	1.0658	1.056 - 1.076
Slope	Degrees	0.0826	1.0861	1.081 - 1.092
Distance to Streams	Kilometers	0.2719	1.3124	1.278 - 1.348
Time to Reducción	Hours	0.1082	1.1143	1.080 - 1.150

in elevation, matching the result from the multinomial model, and slope shows approximately an 8% increase in odds for a 1° increase in slope angle. Finally, the model also shows a greater than 11% increase in the odds of abandonment for every hour walking time further away from known reducciones, all other factors held constant, suggesting that for most of the study region, distance from reducciones substantially influences agricultural deintensification.

However, this is a global model, summarizing the effects across the entire study region. Repeating the modeling procedure for data sampled at smaller scales across the landscape allows us to map the variations in covariate importance across the landscape. For example, the distribution of coefficients for the elevation covariate shows that it is highly important in predicting the abandonment of fields in the western highlands, but it is less important, or even a negative predictor in the lower river valleys to the east (Figure 5.8).

This makes sense as farmers at lower elevations do not have to be as concerned about frost risk, and indeed for crops suited to colder environments may see benefits to maintaining fields higher in the valley. Similarly, mapping the coefficient of the distance from reducción settlements can suggest where proximity to reducciones was important for maintaining fields, and where there may have been more flexibility. In the Colca Valley, and fields to the east of Arequipa, the coefficient for distance from reducciones is very high,

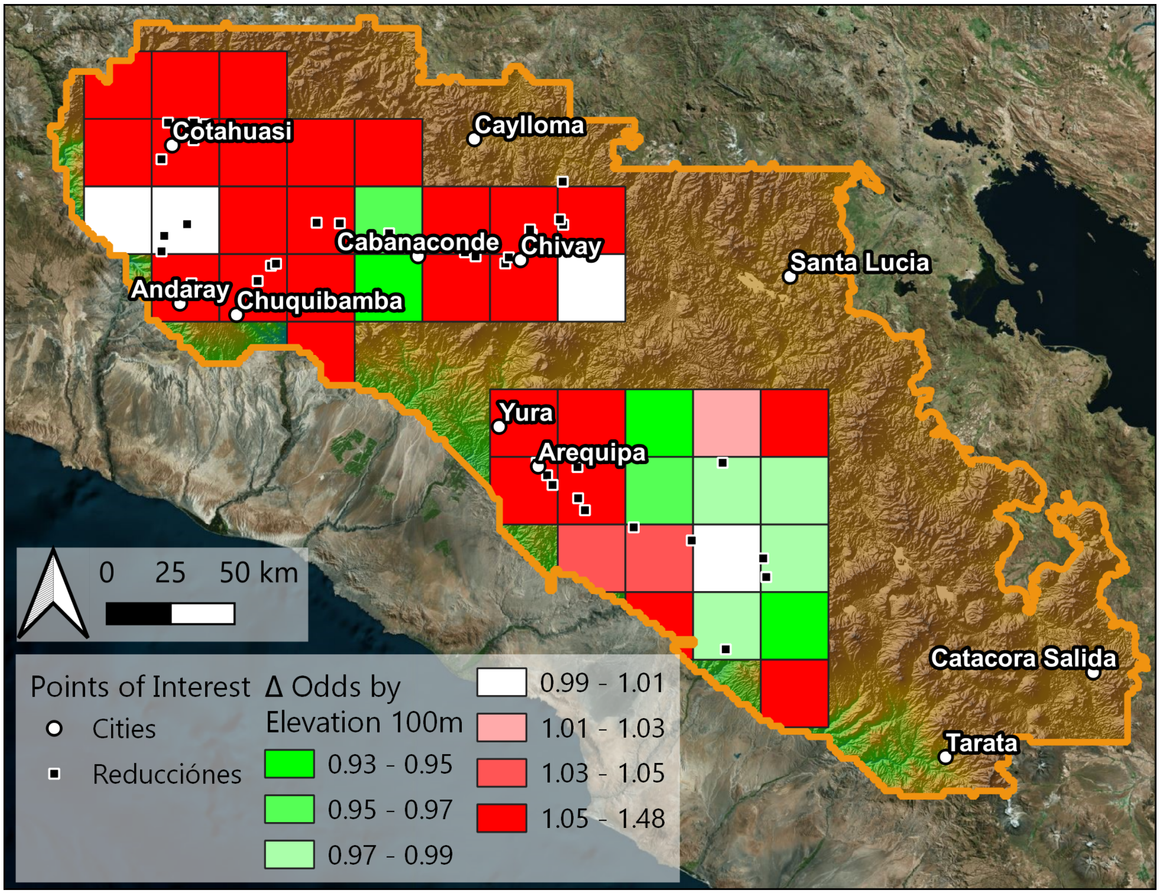

Figure 5.8: The variation in model estimates of the odds of field abandonment due to elevation. Noise resulting from small sample size is likely responsible for some of the variation, however a general trend of elevation being especially important in regions where average elevation is higher is apparent.

suggesting that the distance from reducción fields substantially increases the rate of abandonment (Figure 5.9).

Intriguingly, fields in the north western portion of the study region seem to be less dependent on proximity to reducciones, showing much lower or even negative coefficients, suggesting that perhaps fields further away from the reducción are more likely to be maintained.

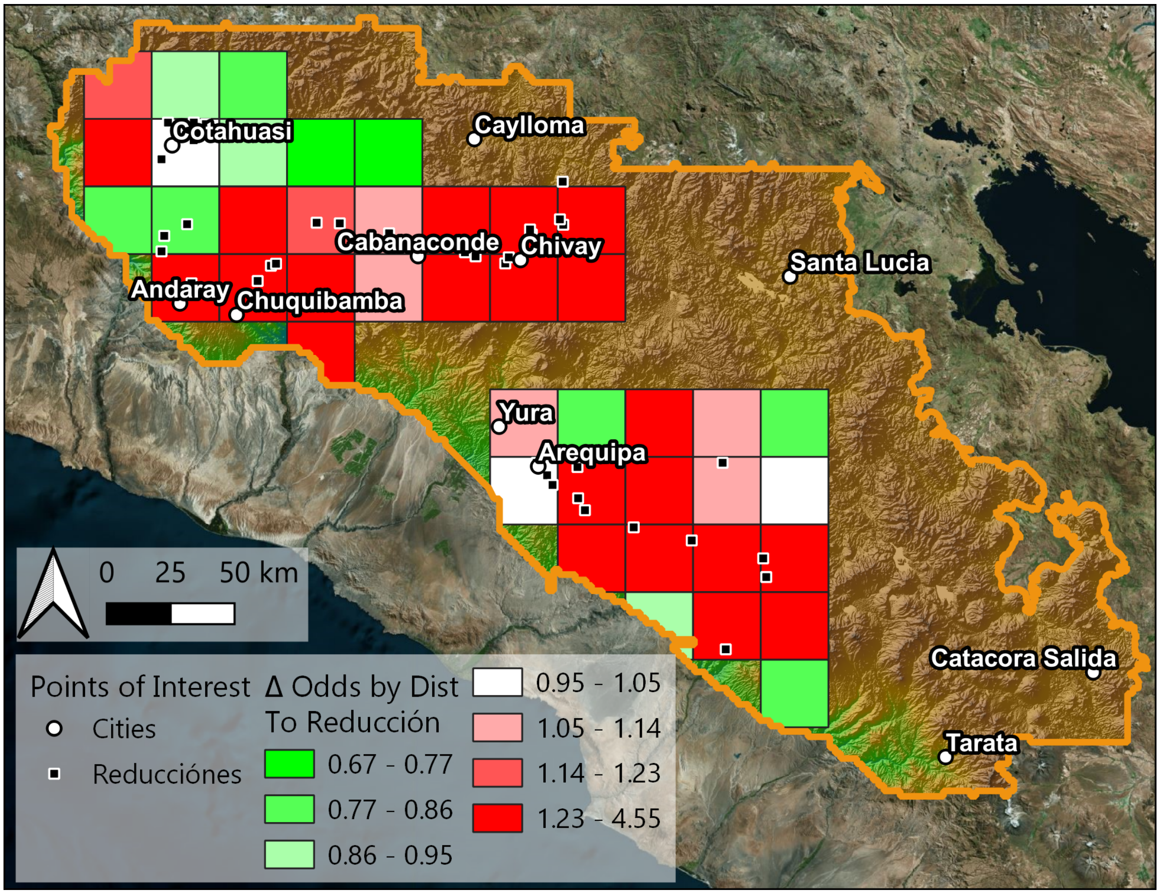

Figure 5.9: The variation in model estimates of the odds of field abandonment due to distance from reducción. Noise resulting from small sample size is likely responsible for some of the variation, however it appears that travel time to a reducción is an important factor in the Arequipa and Colca Valleys, while it is not further to the west near Cotahuasi.

5.6 Discussion

5.6.1 Reducción Placement

The idea that reducción settlements were placed strategically close to important agricultural resources is neither new nor particularly surprising. Toledo himself recommended that the resettlement cause as little disruption as possible, and both Spanish and Indigenous actors had many incentives to maintain as much agricultural production as possible, the Spanish for exploitation and extraction, the indigenous for survival and economic independence. However, the methods of limiting this disruption are not a given. A model wherein the resettlement towns were constructed but the population continued to primarily live outside of town, or at least maintained active secondary residences near their fields, would limit economic disruption without necessitating the placement of the new towns close to readily accessible agricultural fields. Indeed, it may be advantageous to place such a town further away, following (from the Spanish perspective) Toledo's other directives of "abandoning" pre hispanic towns, cemeteries, and ritual places, while (from the Indigenous perspective) limiting the day-to-day contact with the place of Spanish power. Broadly, such a model is not supported by this analysis, which shows that most reducción settlements were placed close to agricultural fields.

An alternative model, that reducción settlements were placed with relatively easy access to pre-existing agricultural infrastructure in an effort to mitigate some of the massive disruptions of the resettlement, finds more support in this analysis. The extreme difference between the actual and the simulated distribution of reducciones in relation to access to fields suggests that it would be difficult to find locations for reducciones that have better access to fields while simultaneously meeting the slope and elevation requirements. This likely implies that reducciones were likely placed quite near prehispanic settlements, cemeteries or ritual places, in spite of the Toledan admonishment to the contrary. The placement of reducciones therefore represents the results of a negotiation between ideals, practicality, and social, political, economic context, rather than a strict adherence to the plans of the resettlement.

We do not mean to imply, however, that Spanish and indigenous actors were working together to maximize economic gain. As actors in an ongoing and empowered (though imbalanced) negotiation, differing goals, incentives, and power structures shaped the outcomes. This is clear from Wernke's work on field abandonment around the reducción of Corporaque. The goal may not have always been to optimize the area available for agriculture, but rather to optimize access for particular people, access to particularly productive lands, or access to land that could grow socially valued crops such as maize. Sometimes, reducciones may have indeed been placed with little regard for access to existing agricultural infrastructure. Indeed, a closer look at the reducción data shows an interesting pattern in this respect. While most of the settlements have access to very large amounts of agricultural fields within a 2 hour walk, a few, such as Huamanmarca I, Machu Llacta, and "Boginoni_200420-203445" have dramatically less access to agricultural infrastructure than even the average for the simulated settlement locations. Also in contrast to most of the reducciones in our dataset, these reducciones are abandoned today.

There are a few possible ways to interpret this pattern. One is that reducciones which were placed without careful regard for the availability of agricultural infrastructure were abandoned to create new towns with better access to such important resources. In this scenario, negotiations about the placement of reducción settlements broke down to such a degree that the town's location was untenable after construction. Often, the town was then moved to a nearby and better-suited location. Alternatively, it is possible that such abandoned settlements are the result of changes in local social, political, or economic conditions that caused a town that was thriving without agricultural infrastructure to no longer persist. One example of this may be to revive the model of periodic and temporary occupation of reducción settlements with primary residences located closer to the fields. Such a town may have been constructed, maintained, and survived as long as there was a need to return to it periodically to fulfill obligations to the Spanish, or as a meeting place for festivals. However as Spanish control waned or became focused elsewhere it would no longer have been necessary to support the town which could then be abandoned.

The differences in agricultural access between abandoned reducción settlements and settlements that remain occupied to this day form a pattern that also serve as a reminder that the above results and discussion are preliminary and require deeper investigation. The reducción data included in this research primarily comes from the LOGAR database as this is the most complete, accurate, and peer-reviewed collection of reducción settlements. However, the LOGAR data was primarily collected by matching the names of reducciones in Spanish colonial documents to those of modern towns and place names today, suggesting that the sample of reducciones is strongly biased away from reducciónes that have been abandoned. This bias in our data was somewhat mitigated through the inclusion of settlements identified in GeoPACHA such as, Huamanmarca, Polobaya, and Locus 5983 which are abandoned. However, if it is the case that a lack of access to agricultural infrastructure is a strong predictor of the abandonment of a reducción, and there are many reducción settlements in the study region who's locations remain unknown, then this bias may go a long way towards explaining why the reducción locations in our data appear to have so much higher access to agricultural fields on average. Future research to confirm the locations of other abandoned reducciones and efforts to evaluate the completeness of known reducción data would do a great deal to improve confidence in these analyses.

5.6.2 Agricultural Deintensification

It has previously been shown that agricultural field abandonment in the Colca Valley is correlated with distance from the reducción of Coporaque when elevation and other environmental factors are held constant (Wernke 2013), but the generality of this pattern has not previously been investigated across the western cordillera of the south-central Andean highlands. This research successfully expands this finding, encompassing several river valleys across the western south-central highlands and thereby demonstrating the lasting impact of the Toledan resettlement on agricultural deintensification. However, a global summary of the effect reducciones have had on agricultural fields over the survey region is limited in the insights it can provide. On the ground, variations in negotiation, land management, ecological and political context, and individual farmer decisions open up a wide variety of potential outcomes that are likely to go unfortunately unnoticed and unconsidered when merely considering global statistics. In contrast, one of the greatest potential benefits of regional and supraregional studies such as this one is the ability to comparatively evaluate differences in strategy between different localities. One may imagine several ways in which local dynamics may differ from the global trends of reducción settlements and field abandonment. One anomaly may be the wholesale abandonment of a region, such as occurred at Churajon, where nearly 35,500 ha of fields were completely abandoned in the early 17th century (Szykulski 2008; Szykulski et al. 2000). No settlement, reducción or otherwise remained to continue to maintain and cultivate these fields. A visual inspection of the agricultural data suggests that this is a rare occurrence, with some level of cultivation persisting in most regions where abandoned fields are present.

Another potential local divergence from the global pattern is the presence of agricultural fields with no nearby reducción. Fields such as those south of Ampato and Sabancaya (Figure 5.10) are the most apparent of these, with no known reducciones within 40 kilometers.

This occurrence made it necessary to limit the distance from reducción variable to only include fields less than 6 hours away from a reducción for the global analysis because the existence of a large quantity of active and abandoned fields an estimated 30-hour walk from the nearest reducción so severely skewed the distribution of the data that it obscured the more local effects of interest. Perhaps the most intriguing interpretation of this is the potential for communities that were able to avoid or resist resettlement and create or maintain systems of fields outside of Spanish control. Given the limitations of the reducción dataset, it is wise to be cautious in this interpretation. Rather, there are probably unknown or unrecorded reducciones located in this region that were not included in the dataset. There are several modern towns in the area that were likely reducciones that did not appear in the surviving fragments of the *Tasa de la visita general de Francisco de Toledo*, including Huanca, Taya, Huambo, and Lluta. Only further archival and field investigation will clarify the prehistoric, colonial, and modern contexts that give rise to these fields.

A more subtle form of avoidance or resistance to the resettlement may be as Stern proposed for Huamanga, the construction of a reducción settlement but the maintenance of households and fields outside and separate from Spanish control. The map showing the

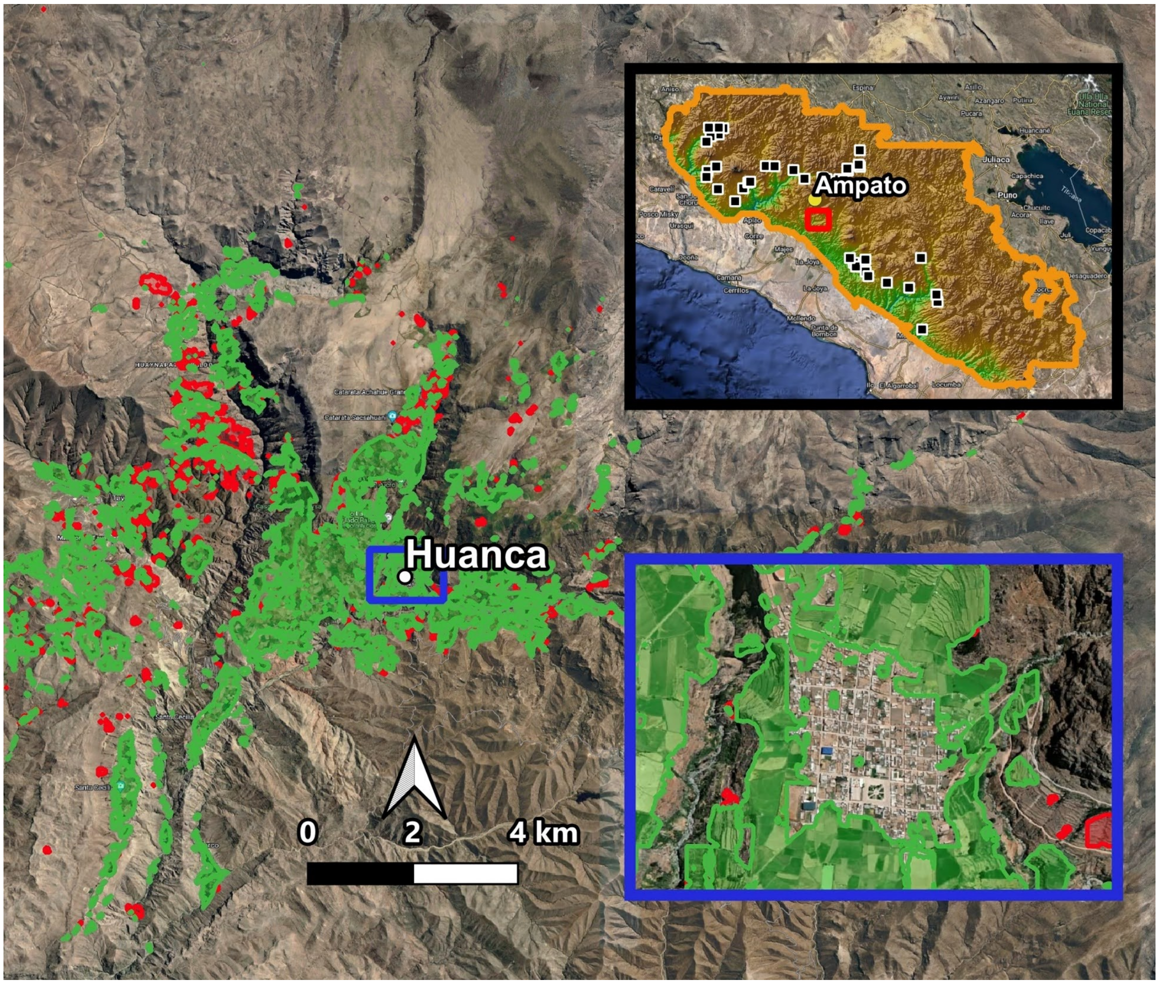

Figure 5.10: Fields to the south of Ampato have no known reducciones nearby. However, the modern town of Huanca has gridded streets and is located in a similar context to reducciones in other areas, suggesting that it may have been an as-yet unrecorded reducción.

local variation of field abandonment as a function of distance from reducciones does not generally support this description for the Arequipa Valley or Colca Valleys, which show high coefficients, suggesting a strong effect. However, such a pattern is shown in the westernmost valleys of the study area. Here the coefficients for distance from reducciones are mostly close to zero, suggesting limited influence on field abandonment as a result of distance from reducciones. In fact, what effect does appear to exist is actually slightly negative, suggesting that in this region proximity to reducción is predictive of field abandonment. Once again, this must be interpreted with caution. It is possible that an unknown reducción could change this dynamic. For example, if there is a band of abandoned fields halfway between a known and unknown reducción, both of which have active fields the model would measure substantial active fields at a greater distance than the abandoned ones, resulting in the lowering of this coefficient. However, unlike the fields to the south of Ampato, many reducciones have been identified in this region, and the pattern appears to exist in the center of this cluster of reducciones, as well as to their north, east, and south. This analysis, therefore, suggests this region as another fruitful area for additional archival and field study to understand what dynamics may have encouraged the abandonment of fields close to reducciones, while stimulating their growth and maintenance further away.

5.7 Conclusions

This research set out to examine trans-regional patterns and regional variations in agricultural deintensification as a result of the colonial forced urbanization of Andean people and the corresponding ruralization of non-urban land. From these analyses, it seems clear that reducciones were placed, on average, closer to existing agricultural fields than would be expected if their proximity to such infrastructure was not being considered. While the guidance of Toledo may not have prioritized such considerations, local Spanish and indigenous populations were strongly incentivized to maintain access to as many fields as possible. Therefore, street-level policy prioritized access to agricultural infrastructure over the official policies of the Viceroyalty. This research also suggests that reducciones that did not prioritize agricultural access may have been less likely to survive, though the sample size of abandoned settlements in this analysis is not large enough to come to any conclusions.

Once the reducciones were constructed, this analysis suggests that for most regions in the study area, agriculturalists were forced to make difficult decisions about which fields to maintain, and to sacrifice those fields that were too far away to work effectively. This is shown by the increase in odds of field abandonment for fields that are further away from the location of known reducciones, holding other variables constant. Such patterns are reflective of the Spanish de-prioritization of non-urban landscapes, a process we have designated ruralization. However, there are regional deviations from this pattern, suggesting that enforcement of reducción living may have varied across the study area, or the presence of as-yet unidentified reducciones. Further investigation will be required to understand these variations in their particular contexts. Nevertheless, these results show the power of trans-regional analyses for identifying local and regional patterns, and for suggesting potential important avenues for future investigation.

CHAPTER 6

Conclusions

Following the Spanish invasion, Viceroy Francisco de Toledo sought to transform indigenous people into easily exploited and obedient Christian subjects by transforming where they lived into highly organized and centralized towns (Mumford 2012, p. 120). This forced resettlement program, known as the Reducción General de Indios, was expected to instill ”policía” and facilitate Christian conversion through urbanization (Kagan 2000). However, the Spanish colonists were entirely dependent on the Andean production of goods for their own prosperity and survival, and existing Andean economic and social systems of production were reliant on networks of smaller and more dispersed settlements. Any disruption to this system endangered the extractive goals of the empire. This fundamental conflict between official colonial policy and street-level reality was managed at local and regional levels by street-level bureaucrats (both Spanish and Indigenous), who used their considerable discretion to implement the Toledan Reforms in ways that made sense in their own contexts. For example, while Toledo complained about reducciónes being constructed near prehispanic settlements ”where they have their idolatries and the graves of their ancestors” (Mumford 2012, p. 120) the practice appears to have been a common one, maintaining vital access to existing agricultural infrastructure.

Nevertheless, in seeking to urbanize the people, the Spanish formed what they saw as a complementary rural landscape. Agricultural fields located great distances from reducción settlements could not be maintained while the workers who claimed it were restricted to living in a reducción. As opposed to the dispersed settlements that characterized the prehispanic settlement pattern, the resettlement created these non-urban spaces as ”rural” and therefore lacking in the power, structure, and morality that the Spanish believed to define urban life (Kagan 2000). This dissertation has shown that this effort had lasting impacts on the agricultural landscape, with large tracts of once-productive agricultural land abandoned and decaying across the landscape. However, it also shows that this effort was deeply dependent on local circumstances and street-level policies. Most of the agricultural land identified in this paper was within a few hours' walk of a known reducción settlement, and distance to fields was just one of the many considerations reducción farmers had to consider, complicating the calculations, but not necessarily fundamentally changing them. In some regions, even this effect seems to have been limited, with distance to reducciones having little to no influence on field abandonment, suggesting that perhaps farmers did not remain restricted to the reducciones for long (Oré Menéndez 2022).

Rather than the uniformity one would expect from a top-down imposition of iron-clad law, the resulting picture is one of immense variation. However, such variation is only visible when looking at the broader trans-regional context. Such a perspective was enabled in this research through the application of a novel methodology – AI-assisted survey. The various merits and drawbacks of automated survey of remote sensing data compared to brute-force manual survey have been a topic of periodic interest and debate in archaeology for decades (Casana 2014; Comer and Harrower 2013; Custer et al. 1986). This dissertation eschews this divide by suggesting that AI is merely another tool for researchers to employ as they work to develop a better understanding of large-scale phenomena. As such it can be used to expand and improve the researcher's work, rather than replace it.

6.1 Summary of Work

Chapter 2 laid the foundation for an AI-assisted survey by establishing protocols for the creation of training and validation data to accurately evaluate automated survey performance in archaeological structure detection. However, it also made clear that the automated survey results alone did not produce data that is sufficiently precise or complete to be directly used in archaeological analyses. Too many archaeological features were missed, and far too many locations without archaeological structures were erroneously marked as containing them. What was not clear was how the model results compared to more traditional ”brute force” methods such as the manual inspection of satellite imagery. In previous survey projects in the middle east, researchers have claimed the ability to capture nearly 100% of visible archaeological features of interest through the manual inspection of CORONA satellite imagery (Casana 2014, 2020). This success was largely due to the rigorous design of the survey project and extensive training of its surveyors, but also is a result of the large size and distinctive visibility of the archaeological features of interest in this survey. The detection of archaeological features such as buildings, towns, or estancias in the south-western Peruvian highlands is likely to be an inherently more challenging task, as the features are smaller and less easily visually distinguished from the surrounding landscape or modern features. In this case, how well do trained surveyors perform in detecting archaeological features in the landscape, and what might automated techniques have to offer?

Chapter 3 addressed these questions through a comparison of the results between the automated survey and GeoPACHA, a survey project in Peru modeled on the CORONA success described above. GeoPacha undertook a ”brute-force” manual survey in Peru using publically available Google, Bing, and other commercial satellite images, and overlaps with approximately 3,000 square kilometers of the region autonomously surveyed using the imagery provided by the Digital Globe Foundation. Overcoming differences in image alignment, research design, and data representation, the results of these two surveys were compared, demonstrating that both the ”brute-force” and automated surveys had similar rates of feature recall. That is, they each missed approximately 30% of features on the landscape. Importantly, however, they did not miss the same features, and when taken together, the two surveys found about 30% *more* features than either survey alone. The biggest advantage in feature detection of the manual survey was therefore not its recall, but its precision. In our estimation, very few of the features identified manually are false positives, probably less than 5%. In contrast, nearly half of the features identified by the automated model were false positives, highlighting the need for manual review of any automated survey technique.

In addition to the automated model's relatively poor precision, an automated survey that is designed to detect the presence or absence of features fails to capture other important information such as the age or use of the buildings, or contextual data such as whether a structure is isolated, part of a larger settlement, or related to other archaeological features. In short, the model provides a filter, which tells researchers where to look, while the human collaborators are required to provide rich contextual information to make that location meaningful. This information can then be fed back into the AI to further improve its predictions in the future. This is the fundamental idea of an AI-assisted survey, which combines the speed of an automated survey with the depth of a brute-force one. It is my assertion AI-assisted surveys promise to be faster, more complete, and more accurate, all while maintaining the important metadata produced by human researchers.

Chapter 4 carries this idea forward conceptually, while applying it to a new problem, mapping the distribution of active and abandoned agricultural infrastructure. Large-scale agricultural deintensification in the Andes has long been a subject of archaeological interest (Denevan 1986; Guillet et al. 1987; Treacy 1990a; Wernke and Whitmore 2009), as researchers seek to understand the causes of the mass abandonment of agricultural infrastructure that occurred following the Spanish invasion. By measuring the scale and patterning of this abandonment at trans-regional scales we can examine regional variation in the colonial processes that affected agricultural production in the Andes. However, doing so using satellite imagery provides unique challenges for both human and automated surveyors. Prehispanic Andean people constructed terracing at an enormous scale, and digitizing its extent by hand is a slow process. It is also unclear when such digitizing is ”complete” as agricultural fields appear to have a nearly fractal level of complexity as one seeks to map them with increasing levels of precision. As one draws a boundary around a field system, one must decide how much time and effort to expend to exclude rivers, valley cliffs, unterraced hills, or towns that are situated amongst the fields. However, if one wishes to be more precise, one could also exclude roads that cross through and between fields, isolated buildings, and trees. Perhaps, rather than mapping systems of fields, one could map the extent of individual fields or terraces within a field system. Each level of precision requires exponentially more time to complete.

Therefore, the question is not how precise one can make a manual survey, but rather how precise can it be *at what scale* and *within what timeframe*. A manual approach to cover the 81,148 km² study area required months of work to produce a map of currently-in-use agricultural systems that we describe as ”regionally accurate”. That is, it tends to show where active fields are present, but the digitized boundaries may also include many non-agricultural features such as rivers or towns. This is a necessity for such large-scale manual surveys as a finer-grained approach that excludes these features takes far too long to practically complete at a large scale. We refined a sample of this data in 51 areas covering 1500 km² seeking to exclude non-agricultural features to the extent possible. While excluding the bulk of the towns, hills, and valleys, we did not exclude all small isolated structures, roads, and a vast array of other non-agricultural features. Even still, a 30 km² region with complex systems of fields could easily take between 5 and 8 hours of labor to complete. A survey of 100,000 km² could take decades. Using an AI-assisted approach, we can use our sample of higher-resolution data to map agricultural fields in a matter of weeks or months more precisely than would be feasible to do by hand in years.

Such a promise of speed and precision demands deeper evaluation. Chapter 4, therefore, includes a comparison of the data results to two existing manual surveys, one a preliminary GeoPACHA dataset that is still under editorial review, and the other a dataset collected by the Peruvian Department of Agrarian Development and Irrigation known as Programa Andenes. The AI-assisted survey results are shown to be markedly better than the preliminary data produced for GeoPACHA, due largely to the opportunistic sampling of abandoned fields in GeoPACHA and the lower resolution of the data produced. The comparison to Programa Andenes is more challenging, however, as Programa Andenes was interested solely in agricultural terracing, while the survey conducted for this dissertation was concerned more broadly with agricultural fields in general. As a result, the AI-assisted survey identified much more land area as agricultural than did Programa Andenes on average. However, in some instances, it was possible to show that the AI-assisted survey identified terracing that was missed in the Programa Andenes manual survey.

Finally, Chapter 4 begins an exploratory analysis of the distribution of active and abandoned agricultural fields with respect to environmental covariates. Elevation, Slope, and Geomorphology are all shown to affect the odds that Andean people would construct and maintain agricultural infrastructure. Simultaneously, aspect seems to have little to no meaningful impact on field construction and maintenance. This analysis is formalized in a multinomial model in Chapter 5. Ultimately, the AI-assisted survey showed an agricultural field abandonment rate of 24%, much lower than is commonly estimated or measured in the archaeological literature. This suggests that archaeological sampling may be biased towards regions where field abandonment is especially high, a logical conclusion given usual archaeological focus on abandoned settlements. The results of these studies must therefore be evaluated trans-regionally as one seeks to take the lessons learned from them to understand historic agricultural transformations outside of their particular contexts.

Chapter 5 begins this expansion, relying heavily on previous research in the Colca Valley and Huamanga to inspect the plausibility of extrapolating those lessons to the other river valleys in the western south-central Andean highlands. In the Colca Valley, agricultural deintensification was strongly related to the late 16th century forced resettlement of the indigenous population into centralized towns known as reducciónes (Denevan 1988; Wernke and Whitmore 2009; Wernke 2013). Chapter 5 attempts to characterize the distribution of these towns as an inhomogenous Poisson Point Pattern contingent on elevation and ground slope. This model was then used as a null hypothesis to test whether reducción settlements were constructed to prioritize access to agricultural infrastructure by simulating plausible alternative distributions of reducciones and measure how much agricultural land was within an estimated 2-hour walk from the reducciones. The simulation showed that the locations of known reducciónes have access to much more agricultural land than would be expected by the simulation, suggesting that as indigenous people constructed the towns, they did so with access to agricultural infrastructure in mind. This suggests a process of negotiation such as the one proposed by Wernke (Wernke 2013) in the selection of sites for reducción settlements, where the bureaucratic Spanish mandates of distance from prehispanic settlements and culturally important sites were subsumed by practical realities and indigenous influence.

Reversing the question, Chapter 5 also examines the effects of the Resettlement on patterns of agricultural deintensification and field abandonment. Work in the Colca Valley has previously shown that increased distance from the reducción increases the odds of field abandonment, contingent on environmental factors such as elevation and slope (Wernke 2013). This makes sense as distance increases the difficulty in performing agricultural labor, but for moderate distances this challenge may be balanced by differences in the suitability of land for particular crops. Both Spanish colonizers and indigenous residents of the reducciones would have been concerned about these challenges, as their wealth and survival depended on increased agricultural production, however, the ways in which these challenges were managed may have varied. Therefore the relationship between distance from reducción and agricultural abandonment is modeled in Chapter 5 in two different ways, globally and locally. Globally, the pattern observed in the Colca Valley appears to hold true, with a greater rate of abandonment at greater distances from reducciones after accounting for covariates such as slope, elevation, and distance from water. Locally, however, it appears that there is variation across the landscape, with the westernmost valleys showing little or even negative preferences for fields closer to reducción settlements. This suggests that different resettlement strategies may have been at play in this region, perhaps allowing farmers to maintain residences away from the reducción to enable greater agricultural productivity, though further research will be necessary to confirm this hypothesis.

Ultimately, this research shows that the forced resettlement of the indigenous Andean population had a lasting impact on agricultural infrastructure in the south-central Andes. As the Spanish sought to create urban communities, they also produced rural land which was outside the easy access of those forced to live in reducción settlements. Farmers were forced to balance ecological concerns such as elevation and ground slope with the distance they had to travel to maintain their fields, thereby limiting the variety of agricultural lands they could bring to bear. This additional burden would have made these communities more vulnerable to external forces such as poor crop yields or natural disasters. Nevertheless, Andean people took active roles in shaping the new agricultural landscape. Reducción settlements were located close to existing agricultural infrastructure, maintaining access to resources that had been constructed over generations in spite of the Spanish Crown's commitment to destroying Andean ties to their past. Furthermore, the effects of the resettlement were not uniform across the south-central Andes, with some regions seeing much higher rates of abandonment than others. As local Andean leaders and Spanish administrators implemented the reforms described by Toledo, they used their discretion to shape the resettlement in ways that they saw fit. We must therefore be careful when describing "The Resettlement," as the resettlement was different things in different places, and the local contexts shaped it to achieve different outcomes.

6.2 Future Research

6.2.1 Returning to structure detection

There are many aspects of this work that can be expanded or improved in the future. With regard to archaeological structure detection, further research has already begun as we work to move away from chip-based classification techniques and towards object detection methods, which not only predict the presence or absence of a structure in a particular ”chip” but also draw a box around individual features it identifies as archaeological structures. This will aid reviewers in identifying what features the model believes to be of archaeological relevance and, if sufficiently accurate, would allow the rapid counting of buildings across large areas, with the bounding box providing an approximate measure of building dimensions. Such detailed information will be valuable for calculating population dynamics and scale, and can be difficult to achieve even with the high-quality GeoPACHA metadata due to inconsistencies in how and whether counts of archaeological features were recorded during the brute force survey. Preliminary research with collaborators from the Vanderbilt Data Science Institute has shown this approach to be highly promising for detecting and measuring corrals and estancias and I look forward to applying it to residential structures as well. We are also working to expand the Deep Learning methods used to include continual learning, which will allow us to continue to improve the model as new data is added without having to fully retrain the model each time.

6.2.2 Dating Fields

Turning to the agricultural survey, the weakest aspect of this data is its lack of available dating information. To some extent, this is unavoidable, as dating agricultural fields, especially terracing, can be challenging on the ground, much less from satellite imagery. However, there may be ways of approaching this issue. Schreiber (1987) has used eight lines of evidence to estimate the dates of terrace construction during the Wari occupation of the Carahuarazo Valley. Among these were the relationships between settlements with known or estimated construction dates and the terracing. If, for example, a terrace is constructed around, but not through a site, then this suggests the site was in existence when the terrace was constructed. Such complex relationships are precisely the kind of information an AI-assisted may be able to determine where an automated survey would fail. Furthermore, local research questions are better elucidated, it may be possible to target particular analyses such as OSL dating to further estimates of terrace dates.

6.2.3 Modeling Canals

The lack of explicit consideration of canals is another challenge in the analyses of this dissertation. Distance from streams is taken as a proxy for the availability of water, but prehispanic canal systems can stretch for kilometers, bringing water to even distant fields (Dillehay, Eling, and Rossen 2005). However, given the distribution of agricultural infrastructure and a high resolution elevation model, it would likely be possible to extrapolate the locations of canals. Visually, it is often easy to intuit the existence of a canal, even if it is not visible in the imagery, because terracing will begin as a few narrow fields near the start of the canal and expand out and down as the canal moves downhill away from the stream. Understanding this process is particularly important as the maintainence or abandonment of fields fed by a canal may be determined by the maintenance or abandonment of the canal. I.e. Fields that are proximate to a reducción may be abandoned if the canal that feeds them brings water from a distant stream and is not sufficiently maintained. Determining the presence and path of canals, and combining it with the agricultural data produced in this dissertation, may therefore provide further insights into agricultural deintensification.

6.2.4 Alternate Regions: Lake Titicaca

Chapter 5 excluded fields around Lake Titicaca from this analysis due to the major differences between agricultural processes between these fields and those in the river valleys. However, the majority of agricultural fields identified in the AI-assisted survey of the GeoPACHA southwestern study area were located in the vicinity of the lake, including systems of abandoned terracing that far outstrip in extent those at Churajon. Therefore, any deeper understanding of agricultural deintensification in the south-central Andes demands further investigation of the Titicaca fields. Future research will seek to understand the relationships between the Spanish resettlement and patterns of agricultural deintensification in this region relate, and how they are similar or different to those in the river valleys.

6.2.5 Ground Truthing/Field Testing

Finally, the results of this dissertation suggest several locations for deeper investigation. While agriculture and reducciones mostly co-occur on the landscape, there are a handful of locations where this dynamic appears to not hold true in our data. The regions surrounding Huanca, Huambo, Taya, and Lluta in Caylloma Province contain extensive terracing that is not accompanied by any known reducción settlements. However, the town of Huanca itself appears very similar in form to a reducción. It is possible, and in my opinion likely, that Huanca is simply a reducción that is not listed in the surviving tasa fragments, and so was not identified in the LOGAR or GeoPACHA data. Other similar reducciones may also be in the area. Archival research may be able to confirm this, though I have not yet been able to track down sufficient evidence to be certain. On the ground archaeological investigation of the town may also be able to confirm this hypothesis. Other regions will likely require further investigation. Agriculturalists in the districts of Cotahuasi, Huaynacotas, Tomepampa, and others near the border of Arequipa and Ayacucho appear to have placed a lower priority on proximity to reducciones in determining whether to maintain a field. This may be the result of an unrecorded reducción that is skewing the data, or due to the high density of reducciones in the area making it so that very little cultivable land is far from a reducción. Alternatively, it may be the result of differences in colonial era prioritization of agricultural production over resettlement confinement. Addressing this question will likely require the production of a local history and archaeological investigation of the reducciones to search for departures between bureaucratic documentary evidence and local realities. Ultimately, the most powerful demonstration of the value of AI-assisted surveys is its ability to generate new questions and suggest new hypotheses, allowing us to consider new possibilities that we could never before explore.

Bibliography

Abercrombie, Thomas Alan (1998). *Pathways of memory and power : ethnography and history among an Andean people*. In collab. with Internet Archive. Madison, Wis. : University of Wisconsin Press. 638 pp. ISBN: 978-0-299-15314-4. URL: http://archive.org/details/pathwaysofmemory0000aber (visited on 07/13/2023).

Agapiou, Athos et al. (2013). "Optimum temporal and spectral window for monitoring crop marks over archaeological remains in the Mediterranean region". In: *Journal of Archaeological Science* 40.3, pp. 1479–1492. DOI: 10.1016/j.jas.2012.10.036.

AgroRural, MDAR: and CERA, eds. (2021). *Andenes para la vida: inventario y caracterización de andenes en los Andes tropicales del Perú*. Primera edición. Serie Ecología y Desarrollo 16. OCLC: on1298287920. Lima, Perú: Ministerio de Desarrollo Agrario y Riego : AgroRural ; Centro de Estudios Regionales Andinos. 319 pp. ISBN: 978-612-4121-35-7.

Aitken, R. (1977). "Wilderness areas in Scotland". Unpublished Ph.D. thesis. PhD thesis. University of Aberdeen.

Alcock, S. and J. Cherry (2016). *Side-by-Side Survey: Comparative Regional Studies in the Mediterranean World*. Oxford: Oxbow Books.

Alexakis, Dimitrios et al. (Feb. 23, 2009). "Detection of Neolithic Settlements in Thessaly (Greece) Through Multispectral and Hyperspectral Satellite Imagery". In: *Sensors* 9.2, pp. 1167–1187. ISSN: 1424-8220. DOI: 10.3390/s90201167. URL: http://www.mdpi.com/1424-8220/9/2/1167 (visited on 07/20/2022).

Autodema (Jan. 2016). *PROYECTO ESPECIAL MAJES SIGUAS – AUTODEMA*. Published: Retrieved 5 December 2018.

Azavea (July 19, 2022). *Raster-Vision*. Version 0.13. URL: https://github.com/azavea/raster-vision (visited on 07/21/2022).

Banning, E.B. (2002). *Archaeological Survey*. New York: Kluwer Academic/Plenum Publishers.

Bethell, Leslie, ed. (1984). *The Cambridge history of Latin America*. Cambridge [England] ; New York: Cambridge University Press. 11 pp. ISBN: 978-0-521-23223-4 978-0-521-24516-6 978-0-521-23224-1 978-0-521-23225-8 978-0-521-24517-3 978-0-521-23226-5 978-0-521-24518-0 978-0-521-26652-9 978-0-521-39524-3 978-0-521-49594-3 978-0-521-39525-0.

Bewes, James et al. (Feb. 1, 2019). "Artificial intelligence for sex determination of skeletal remains: Application of a deep learning artificial neural network to human skulls". In: *Journal of Forensic and Legal Medicine* 62, pp. 40–43. ISSN: 1752-928X. DOI: 10.1016/j.jflm.2019.01.004. URL: https://www.sciencedirect.com/science/article/pii/S1752928X18304219 (visited on 04/05/2023).

Boza Cuadros, Maria Fernanda (May 5, 2021). "Bridge and Boundary: The Maritime Connections of Colonial Arequipa, Peru". In: *International Journal of Historical Archaeology*. ISSN: 1573-7748. DOI: 10.1007/s10761-021-00599-3. URL: https://doi.org/10.1007/s10761-021-00599-3 (visited on 06/11/2021).

Buetti-Dinh, Antoine et al. (June 1, 2019). "Deep neural networks outperform human expert's capacity in characterizing bioleaching bacterial biofilm composition". In: *Biotechnology Reports* 22, e00321. ISSN: 2215-017X. DOI: 10.1016/j.btre.2019.e00321. URL:

https://www.sciencedirect.com/science/article/pii/S2215017X18301954 (visited on 07/20/2022).
Byeon, Wonmin et al. (Mar. 1, 2019). "Automated identification and deep classification of cut marks on bones and its paleoanthropological implications". In: *Journal of Computational Science* 32, pp. 36–43. ISSN: 1877-7503. DOI: 10.1016/j.jocs.2019.02.005. URL: https://www.sciencedirect.com/science/article/pii/S1877750318310226 (visited on 04/05/2023).
Canuto, Marcello A. et al. (2018). "Reading on LiDAR: Ancient lowland Maya complexity as revealed by airborne laser scanning of northern Guatemala". In: *Science* 361.6409, eaau0137.
Casana, Jesse (Aug. 2014). "Regional-Scale Archaeological Remote Sensing in the Age of Big Data: Automated Site Discovery vs. Brute Force Methods". In: *Advances in Archaeological Practice* 2.3, pp. 222–233. ISSN: 2326-3768. DOI: 10.7183/2326-3768.2.3.222. URL: https://www.cambridge.org/core/product/identifier/S2326376800000292/type/journal_article (visited on 03/19/2021).
— (Feb. 20, 2020). "Global-Scale Archaeological Prospection using CORONA Satellite Imagery: Automated, Crowd-Sourced, and Expert-led Approaches". In: *Journal of Field Archaeology* 45 (sup1). Publisher: Routledge, S89–S100. ISSN: 0093-4690. DOI: 10.1080/00934690.2020.1713285. URL: https://doi.org/10.1080/00934690.2020.1713285 (visited on 08/10/2022).
Casana, Jesse and Jackson Cothren (2013). "The CORONA Atlas Project: Orthorectification of CORONA satellite imagery and regional-scale archaeological exploration in the Near East". In: *Mapping archaeological landscapes from space*. Springer, pp. 33–43.
Casana, Jesse, John Kantner, et al. (May 1, 2014). "Archaeological aerial thermography: A case study at the Chaco-era Blue J community, New Mexico". In: *Journal of Archaeological Science* 45, pp. 207–219. DOI: 10.1016/j.jas.2014.02.015.
Center, Earth Resources Observation {And} Science (EROS) (2017). *Shuttle Radar Topography Mission (SRTM) Void Filled*. DOI: 10.5066/F7F76B1X. URL: https://www.usgs.gov/centers/eros/science/usgs-eros-archive-digital-elevation-shuttle-radar-topography-mission-srtm-void?qt-science_center_objects=0#qt-science_center_objects (visited on 05/09/2023).
Chase, Arlen F. et al. (2012). "Geospatial revolution and remote sensing LiDAR in Mesoamerican archaeology". In: *Proceedings of the National Academy of Sciences of the United States of America* 109.32, pp. 12916–12921.
Chavez, De la Vera Cruz (2019). "La sociedad Churajón: Un proceso de desarrollo local". Accepted: 2019-10-23T17:23:11Z Publisher: Universidad Nacional de San Agustín de Arequipa. PhD thesis. Arequipa, Peru: Universidad Nacional de San Agustín de Arequipa. URL: http://repositorio.unsa.edu.pe/handle/UNSA/9686.
Chen, Ting, Simon Kornblith, et al. (Oct. 25, 2020). "Big Self-Supervised Models are Strong Semi-Supervised Learners". In: *arXiv:2006.10029 [cs, stat]*. arXiv: 2006.10029. URL: http://arxiv.org/abs/2006.10029 (visited on 05/30/2021).
Chen, Xinlei and Kaiming He (Nov. 20, 2020). "Exploring Simple Siamese Representation Learning". In: *arXiv:2011.10566 [cs]*. arXiv: 2011.10566. URL: http://arxiv.org/abs/2011.10566 (visited on 12/11/2021).

Clement, Christopher Ohm and Michael E. Moseley (1991). “The spring-fed irrigation system of Carrizal, Peru: A case study of the hypothesis of agrarian collapse”. In: *Journal of Field Archaeology* 18.4, pp. 425–443.

Comer, Douglas C. and Michael J. Harrower (Jan. 10, 2013). *Mapping Archaeological Landscapes from Space*. SpringerBriefs in Archaeological Heritage Management. Google-Books-ID: yWJDAAAAQBAJ. New York: Springer Science & Business Media. 276 pp. ISBN: 978-1-4614-6074-9. URL: https://www.springer.com/gp/book/9781461460732.

Cook, Noble David (1992). “Impact of disease in the sixteenth-century Andean world”. In: *Disease and Demography in the Americas*. Ed. by John W. Verano and Douglas H. Ubelaker. Washington, DC: Smithsonian Institution Press, pp. 207–213.

— (1998). *Born to die: disease and New World conquest, 1492-1650*. New approaches to the Americas. New York: Cambridge University Press. xiii, 248. ISBN: 0521622085 (hc) 0521627303 (pbk.)

— (June 7, 2004). *Demographic Collapse: Indian Peru, 1520-1620*. Google-Books-ID: zBnzya3kDKAC. Cambridge University Press. 325 pp. ISBN: 978-0-521-52314-1.

Cook, Noble David and Alexandra Parma Cook (2007). *People of the volcano: Andean counterpoint in the Colca Valley of Peru*. OCLC: ocm74987603. Durham: Duke University Press. 319 pp. ISBN: 978-0-8223-3988-5 978-0-8223-3971-7.

Custer, Jay F. et al. (July 1986). “Application of Landsat Data and Synoptic Remote Sensing to Predictive Models for Prehistoric Archaeological Sites: An Example from the Delaware Coastal Plain”. In: *American Antiquity* 51.3. Publisher: Cambridge University Press, pp. 572–588. ISSN: 0002-7316, 2325-5064. DOI: 10.2307/281753. (Visited on 07/12/2023).

D’Altroy, Terence N. (Apr. 30, 2014). *The Incas*. Google-Books-ID: 8JmAAwAAQBAJ. John Wiley & Sons. 685 pp. ISBN: 978-1-118-61059-6.

Davies, Keith A. (1984). *Landowners in Colonial Peru*. Austin: University of Texas Press. ISBN: 978-0-292-74639-8 978-0-292-76621-1 978-0-292-76622-8 978-0-292-76623-5. URL: https://utpress.utexas.edu/books/davlap (visited on 11/26/2019).

Davies, Nigel (June 15, 1995). *The Incas*. Google-Books-ID: OnYaAAAAYAAJ. University Press of Colorado. 280 pp. ISBN: 978-0-87081-360-3.

Davis, Dylan S. (Sept. 1, 2020). “Defining what we study: The contribution of machine automation in archaeological research”. In: *Digital Applications in Archaeology and Cultural Heritage* 18, e00152. ISSN: 2212-0548. DOI: 10.1016/j.daach.2020.e00152. URL: https://www.sciencedirect.com/science/article/pii/S2212054820300515 (visited on 04/04/2023).

Davis, Dylan S. et al. (Aug. 1, 2021). “Deep learning reveals extent of Archaic Native American shell-ring building practices”. In: *Journal of Archaeological Science* 132, p. 105433. ISSN: 0305-4403. DOI: 10.1016/j.jas.2021.105433. URL: https://www.sciencedirect.com/science/article/pii/S0305440321001035 (visited on 04/04/2023).

Denevan, William M. (1986). “The cultural ecology, archaeology, and history of terracing and terrace abandonment in the Colca Valley of southern Peru. Volumes 1 and 2”. In: *Unpublished Manuscript submitted to the National Science Foundation and National Geographic Society*.

— (1988). “Measurement of Abandoned Terracing from Air Photos: Colca Valley, Peru”. In: *Yearbook. Conference of Latin Americanist Geographers* 14, pp. 20–30. ISSN: 1054-3074. URL: https://www.jstor.org/stable/25765694 (visited on 09/02/2019).

Denevan, William M. (Jan. 9, 2003). *Cultivated Landscapes of Native Amazonia and the Andes*. Google-Books-ID: 9cGtcuzTshEC. Oxford University Press. 436 pp. ISBN: 978-0-19-925769-0.

Deng, Jia et al. (June 2009). “ImageNet: A large-scale hierarchical image database”. In: *2009 IEEE Conference on Computer Vision and Pattern Recognition*. 2009 IEEE Conference on Computer Vision and Pattern Recognition. ISSN: 1063-6919, pp. 248–255. DOI: 10.1109/CVPR.2009.5206848.

Dillehay, Tom D., Herbert H. Eling, and Jack Rossen (Nov. 22, 2005). “Preceramic irrigation canals in the Peruvian Andes”. In: *Proceedings of the National Academy of Sciences* 102.47. Publisher: Proceedings of the National Academy of Sciences, pp. 17241–17244. DOI: 10.1073/pnas.0508583102. URL: https://www.pnas.org/doi/10.1073/pnas.0508583102 (visited on 06/06/2023).

Dodge, Samuel and Lina Karam (May 6, 2017). *A Study and Comparison of Human and Deep Learning Recognition Performance Under Visual Distortions*. arXiv: 1705.02498[cs]. URL: http://arxiv.org/abs/1705.02498 (visited on 07/20/2022).

Donkin, R. A (1979). *Agricultural terracing in the aboriginal New World*. OCLC: 470288398. Tucson: Publ. for the Wenner-Gren Foundation for Anthropological Research [by] the Univ. of Arizona Press. ISBN: 978-0-8165-0453-4.

Doutriaux, Miriam (2004). “Imperial conquest in a multiethnic setting: the Inka occupation of the Colca valley, Peru”. PhD thesis. Berkeley: University of California, Berkeley. URL: file:///G:/My%20Documents/references/articles_proposals%20sent%20to%20me/Dissertations/Doutriaux-2004.pdf.

Erickson, Clark L. (Sept. 1999). “Neo-environmental determinism and agrarian ‘collapse’ in Andean prehistory”. In: *Antiquity* 73.281, pp. 634–642. ISSN: 0003-598X, 1745-1744. DOI: 10.1017/S0003598X00065236. URL: https://www.cambridge.org/core/product/identifier/S0003598X00065236/type/journal_article (visited on 05/09/2023).

— (Dec. 31, 2000). “12. The Lake Titicaca Basin: A Precolumbian Built Landscape”. In: *Imperfect Balance*. Ed. by David L. Lentz. Columbia University Press, pp. 311–356. ISBN: 978-0-231-11157-7. DOI: 10.7312/lent11156-015. URL: https://www.degruyter.com/document/doi/10.7312/lent11156-015/html (visited on 05/05/2023).

Fick, Stephen E. and Robert J. Hijmans (2017). “WorldClim 2: new 1-km spatial resolution climate surfaces for global land areas”. In: *International Journal of Climatology* 37.12. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/joc.5086, pp. 4302–4315. ISSN: 1097-0088. DOI: 10.1002/joc.5086. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/joc.5086 (visited on 05/09/2023).

Fisher, John (June 1, 1998). *The Economic Aspects of Spanish Imperialism in America, 1492-1810*. Google-Books-ID: 4iO9DgAAQBAJ. Oxford University Press. 256 pp. ISBN: 978-1-78138-645-3.

Garrison, Thomas G. et al. (Oct. 2008). “Evaluating the use of IKONOS satellite imagery in lowland Maya settlement archaeology”. In: *Journal of Archaeological Science* 35.10, pp. 2770–2777. ISSN: 03054403. DOI: 10.1016/j.jas.2008.05.003. URL: https://linkinghub.elsevier.com/retrieve/pii/S0305440308001088 (visited on 07/20/2022).

Gelles, Paul H. (2000). *Water and power in highland Peru: the cultural politics of irrigation and development*. New Brunswick, N.J: Rutgers University Press. xvii+233. ISBN: 978-0-8135-6018-2.

Gil, Yolanda et al. (Mar. 17, 2019). "Towards human-guided machine learning". In: *Proceedings of the 24th International Conference on Intelligent User Interfaces*. IUI '19: 24th International Conference on Intelligent User Interfaces. Marina del Ray California: ACM, pp. 614–624. ISBN: 978-1-4503-6272-6. DOI: 10.1145/3301275.3302324. URL: https://dl.acm.org/doi/10.1145/3301275.3302324 (visited on 07/22/2022).

GRASS Development Team (2012). *Geographic Resources Analysis Support System (GRASS GIS) Software*. USA: Open Source Geospatial Foundation. URL: http://grass.osgeo.org.

Grizonnet, Manuel et al. (2017). "Orfeo ToolBox: Open source processing of remote sensing images". In: *Open Geospatial Data, Software and Standards* 2.1. Publisher: SpringerOpen, p. 15.

Guillet, David (1992). *Covering ground: communal water management and the state in the Peruvian highlands*. Linking levels of analysis. Ann Arbor: The University of Michigan Press. xviii+250. ISBN: 978-0-472-09474-5.

Guillet, David et al. (1987). "Terracing and Irrigation in the Peruvian Highlands [and Comments and Reply]". In: *Current Anthropology* 28.4, pp. 409–430. ISSN: 0011-3204. DOI: 10/fdjw38. URL: https://www.jstor.org/stable/2743482 (visited on 09/02/2019).

Guyot, Alexandre et al. (2021). "Combined Detection and Segmentation of Archeological Structures from LiDAR Data using a Deep Learning Approach". In: *Journal of Computer Applications in Archaeology* 4.1, pp. 1–19.

Harris, Douglas H. (Jan. 2002). "How to *Really* Improve Airport Security". In: *Ergonomics in Design: The Quarterly of Human Factors Applications* 10.1, pp. 17–22. ISSN: 1064-8046. DOI: 10/cqxq29. URL: http://journals.sagepub.com/doi/10.1177/106480460201000104 (visited on 04/27/2020).

Hauberg, Søren et al. (June 30, 2016). *Dreaming More Data: Class-dependent Distributions over Diffeomorphisms for Learned Data Augmentation*. DOI: 10.48550/arXiv.1510.02795. arXiv: 1510.02795[cs]. URL: http://arxiv.org/abs/1510.02795 (visited on 07/11/2023).

He, Kaiming et al. (Dec. 10, 2015). *Deep Residual Learning for Image Recognition*. DOI: 10.48550/arXiv.1512.03385. arXiv: 1512.03385[cs]. URL: http://arxiv.org/abs/1512.03385 (visited on 07/21/2022).

Hemming, John (1970). *The conquest of the Incas*. A Harvest book 258. New York: Harcourt Brace Jovanovich. 641 pp.

Hoberman, Louisa Schell and Susan Migden Socolow (1986). *Cities & society in colonial Latin America*. 1st ed. Albuquerque: University of New Mexico Press. xi+350. ISBN: 978-0-8263-0844-3.

Janusek, John Wayne and Alan L. Kolata (Dec. 1, 2004). "Top-down or bottom-up: rural settlement and raised field agriculture in the Lake Titicaca Basin, Bolivia". In: *Journal of Anthropological Archaeology* 23.4, pp. 404–430. ISSN: 0278-4165. DOI: 10.1016/j.jaa.2004.08.001. URL: https://www.sciencedirect.com/science/article/pii/S0278416504000479 (visited on 05/09/2023).

Kagan, Richard L. (2000). *Urban images of the Hispanic world, 1493-1793*. In collab. with Fernando Marías. New Haven: Yale University Press. x+235. ISBN: 978-0-300-08314-9.

Kendall, Ann and David Drew (Jan. 11, 2019). "The Rehabilitation of Pre-Hispanic Agricultural Infrastructure to Support Rural Development in the Peruvian Andes: The Work of the Cusichaca Trust". In: *The Oxford Handbook of Historical Ecology and Applied Archaeology*. Ed. by Christian Isendahl and Daryl Stump. Oxford University Press. ISBN: 978-0-19-967269-1. DOI: 10.1093/oxfordhb/9780199672691.013.24. URL: https://doi.org/10.1093/oxfordhb/9780199672691.013.24 (visited on 05/10/2023).

Kohut, Lauren E. (Mar. 30, 2023). *Farming Landscapes under Stress: Modeling Access to Pastures and Fields in the Late Intermediate Period Colca Valley (1100– 1450 CE, Arequipa, Peru)*. personal communication.

Kolata, Alan L. (1986). "The Agricultural Foundations of the Tiwanaku State: A View from the Heartland". In: *American Antiquity* 51.4. Publisher: Society for American Archaeology, pp. 748–762. ISSN: 0002-7316. DOI: 10.2307/280863. URL: https://www.jstor.org/stable/280863 (visited on 07/12/2023).

— (June 1991). "The Technology and Organization of Agricultural Production in the Tiwanaku State". In: *Latin American Antiquity* 2.2. Publisher: Cambridge University Press, pp. 99–125. ISSN: 1045-6635, 2325-5080. DOI: 10.2307/972273. (Visited on 07/13/2023).

Körber, Moritz et al. (2015). "Vigilance Decrement and Passive Fatigue Caused by Monotony in Automated Driving". In: *Procedia Manufacturing* 3, pp. 2403–2409. ISSN: 23519789. DOI: 10.1016/j.promfg.2015.07.499. URL: https://linkinghub.elsevier.com/retrieve/pii/S2351978915005004 (visited on 07/20/2022).

Kvamme, Kenneth L. (2005). "There and Back Again: Revisiting Archaeological Locational Modeling". In: *GIS and Archaeological Site Location Modeling*. CRC Press. ISBN: 978-0-429-21178-2.

Lambers, Karsten, Wouter B. Verschoof-van der Vaart, and Quentin P. J. Bourgeois (Jan. 2019). "Integrating Remote Sensing, Machine Learning, and Citizen Science in Dutch Archaeological Prospection". In: *Remote Sensing* 11.7. Number: 7 Publisher: Multidisciplinary Digital Publishing Institute, p. 794. ISSN: 2072-4292. DOI: 10.3390/rs11070794. URL: https://www.mdpi.com/2072-4292/11/7/794 (visited on 08/15/2022).

Langlie, BrieAnna S. (Dec. 1, 2018). "Building ecological resistance: Late intermediate period farming in the south-central highland Andes (CE 1100–1450)". In: *Journal of Anthropological Archaeology* 52, pp. 167–179. ISSN: 0278-4165. DOI: 10.1016/j.jaa.2018.06.005. URL: https://www.sciencedirect.com/science/article/pii/S0278416518300151 (visited on 05/05/2023).

— (Apr. 1, 2022). "Locating Agricultural Terraces in the Southern Peruvian Andes". Conference. Conference. Society for American Archaeology. Chicago, IL. URL: https://documents.saa.org/container/docs/default-source/doc-annualmeeting/final-program/87th-annual-meeting-(chicago-il---2022).pdf?sfvrsn=1fe4b065_2 (visited on 05/09/2023).

Langmuir, E. (1984). *Mountaincraft and leadership*. Leicester: The Scottish Sports Council/MLTB. Cordee.

Lasaponara, Rosa and Nicola Masini (2007). “Detection of archaeological crop marks by using satellite QuickBird multispectral imagery”. In: *Journal of archaeological science* 34.2, pp. 214–221. DOI: 10.1016/j.jas.2006.04.014.

Latour, Bruno and Steve Woolgar (1979). *Laboratory life: the social construction of scientific facts*. Sage library of social research ; v. 80. Beverly Hills: Sage Publications. 272 pp. ISBN: 978-0-8039-0993-9 978-0-8039-0994-6.

Lecun, Y. et al. (Nov. 1998). “Gradient-based learning applied to document recognition”. In: *Proceedings of the IEEE* 86.11. Conference Name: Proceedings of the IEEE, pp. 2278–2324. ISSN: 1558-2256. DOI: 10.1109/5.726791.

Lin, Albert Yu-Min et al. (Dec. 30, 2014). “Crowdsourcing the Unknown: The Satellite Search for Genghis Khan”. In: *PLOS ONE* 9.12. Publisher: Public Library of Science, e114046. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0114046. URL: https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0114046 (visited on 08/10/2022).

Linares Delgado, Lucy Roxana (1989). “Cronología y relaciones culturales pre-hispánicas del Valle del Chili - Arequipa”. Licenciado. Arequipa: Universidad Católica Santa María. 77 pp.

Lipsky, Michael (2010). *Street-level bureaucracy: dilemmas of the individual in public services*. 30th anniversary expanded ed. OCLC: ocn503596064. New York: Russell Sage Foundation. 275 pp. ISBN: 978-0-87154-544-2.

Málaga Medina, Alejandro (1975). “Las reducciones en el virreinato del Peru (1532-1580)”. In: *Revista de Historia de América* 80. Publisher: Pan American Institute of Geography and History, pp. 9–42. ISSN: 0034-8325. URL: https://www.jstor.org/stable/20139179 (visited on 03/25/2021).

Matienzo, Juan De (1967). *Gobierno del Perú (1567)*. Ed. by Guillermo Lhomann Villena. Institut français d’études andines. ISBN: 978-2-8218-4567-1. DOI: 10.4000/books.ifea.3104. URL: http://books.openedition.org/ifea/3104 (visited on 07/12/2023).

Miller, Harvey J. (June 1, 2004). “Tobler’s First Law and Spatial Analysis”. In: *Annals of the Association of American Geographers* 94.2, pp. 284–289. ISSN: 0004-5608. DOI: 10.1111/j.1467-8306.2004.09402005.x. URL: https://doi.org/10.1111/j.1467-8306.2004.09402005.x (visited on 04/18/2023).

Montavon, Grégoire et al. (2019). “Layer-wise relevance propagation: an overview”. In: *Explainable AI: interpreting, explaining and visualizing deep learning*. Publisher: Springer, pp. 193–209.

Morlon, Pierre. (1996). “Comprender la agricultura campesina en los Andes Centrales : Perú-Bolivia / compilación y coordinación, Pierre Morlon ; traducción al castellano, Edgardo Rivera Martínez ; cuidado de la edición, Anne-Marie Brougère.” In: *Comprender la agricultura campesina en los Andes Centrales : Perú-Bolivia*. Travaux de l’Institut français d’études andines t. 96. Lima, Perú: Institut français d’études andines. ISBN: 84-8387-041-X.

Moseley, Michael E. (1983). “The good old days were better: Agrarian collapse and tectonics”. In: *American Anthropologist* 85.4, pp. 773–799. DOI: 10/bp38mq.

Mumford, Jeremy Ravi (2012). *Vertical empire: The general resettlement of Indians in the colonial Andes*. Duke University Press.

Murra, John V. (1972). “El ”control vertical” de un máximo de pisos ecológicos en la economía de las sociedades andinas.” In: *Visita de la Provincia de Leon de Huánuco en*

1562. Ed. by John V. Murra. OCLC: 252974767. Huánco, Peru: Universidad Nacional Hermillo Valdizan: Univ. Hermilio Valdizan, pp. 427–476.
Neira Avendaño, Máximo (1990). "Arequipa prehispánica". In: *Historia general de Arequipa*. Ed. by Máximo Neira Avendaño et al. Arequipa: Fundación M.J. Bustamante de la Fuente, pp. 5–213.
Oré Menéndez, Gabriela de los Ángeles (2022). "Legacies of Colonial Dislocation: Resettlement, Agricultural Deintensification and Infrastructural Reclamation in Huarochirí, Peru".
Parcak, Sarah H. (2009). *Satellite remote sensing for archaeology*. Routledge.
— (Mar. 6, 2017). *GIS, Remote Sensing, and Landscape Archaeology*. Vol. 1. Oxford University Press. DOI: 10.1093/oxfordhb/9780199935413.013.11. URL: http://oxfordhandbooks.com/view/10.1093/oxfordhb/9780199935413.001.0001/oxfordhb-9780199935413-e-11 (visited on 04/07/2022).
— (July 9, 2019). *Archaeology from Space: How the Future Shapes Our Past*. Illustrated edition. Henry Holt and Co. 283 pp.
Pour, Amin Beiranvand et al. (Sept. 1, 2021). *Multispectral and Hyperspectral Remote Sensing Data for Mineral Exploration and Environmental Monitoring of Mined Areas*. Google-Books-ID: DkBAEAAAQBAJ. MDPI. 416 pp. ISBN: 978-3-0365-1264-8.
QGIS Development Team (2009). *QGIS Geographic Information System*. URL: http://qgis.osgeo.org.
Rai, Arun (Jan. 1, 2020). "Explainable AI: from black box to glass box". In: *Journal of the Academy of Marketing Science* 48.1, pp. 137–141. ISSN: 1552-7824. DOI: 10.1007/s11747-019-00710-5. URL: https://doi.org/10.1007/s11747-019-00710-5 (visited on 07/21/2022).
Sandor, J.A. (1992). "Long-term effects of prehistoric agriculture on soil: examples from New Mexico and Peru". In: *Soils in archaeology: landscape evolution and human occupation*. Ed. by V.T. Holliday. Washington, D.C.: Smithsonian Institution Press, pp. 217–245.
Sandor, Jonathan A. (1987). "Initial investigation of soils in agricultural terraces in the Colca Valley, Peru". In: *Pre-Hispanic agricultural terraces in the Andean region. British Archaeological Reports, International Series* 359.1, pp. 163–192.
Saturno, William et al. (2007). "Putting Us on the Map: Remote Sensing Investigation of the Ancient Maya Landscape". In: *Remote Sensing in Archaeology*. Ed. by James Wiseman and Farouk El-Baz. Interdisciplinary Contributions To Archaeology. New York, NY: Springer, pp. 137–160. ISBN: 978-0-387-44455-0. DOI: 10.1007/0-387-44455-6_6. URL: https://doi.org/10.1007/0-387-44455-6_6 (visited on 08/17/2022).
Schreiber, Katharina J. (1987). "Conquest and Consolidation: A Comparison of the Wari and Inka Occupations of a Highland Peruvian Valley". In: *American Antiquity* 52.2. Publisher: Society for American Archaeology, pp. 266–284. ISSN: 0002-7316. DOI: 10.2307/281780. URL: https://www.jstor.org/stable/281780 (visited on 06/05/2023).
Selvaraju, Ramprasaath R. et al. (Feb. 2020). "Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization". In: *International Journal of Computer Vision* 128.2, pp. 336–359. ISSN: 0920-5691, 1573-1405. DOI: 10.1007/s11263-019-01228-7. arXiv: 1610.02391[cs]. URL: http://arxiv.org/abs/1610.02391 (visited on 07/21/2022).

Sevara, Christopher et al. (Feb. 1, 2016). “Pixel versus object — A comparison of strategies for the semi-automated mapping of archaeological features using airborne laser scanning data”. In: *Journal of Archaeological Science: Reports* 5, pp. 485–498. ISSN: 2352-409X. DOI: 10.1016/j.jasrep.2015.12.023. URL: https://www.sciencedirect.com/science/article/pii/S2352409X15302248 (visited on 04/09/2023).

Shingledecker, Clark et al. (2017). “Measuring vigilance abilities to enhance combat identification performance”. In: *Human factors issues in combat identification*. CRC Press, pp. 47–66.

Simard, P.Y., D. Steinkraus, and J.C. Platt (2003). “Best practices for convolutional neural networks applied to visual document analysis”. In: *Seventh International Conference on Document Analysis and Recognition, 2003. Proceedings.* Seventh International Conference on Document Analysis and Recognition. Vol. 1. Edinburgh, UK: IEEE Comput. Soc, pp. 958–963. ISBN: 978-0-7695-1960-9. DOI: 10.1109/ICDAR.2003.1227801. URL: http://ieeexplore.ieee.org/document/1227801/ (visited on 07/11/2023).

Smith, C. T., W. M. Denevan, and P. Hamilton (1968). “Ancient Ridged Fields in the Region of Lake Titicaca”. In: *The Geographical Journal* 134.3. Publisher: [Wiley, Royal Geographical Society (with the Institute of British Geographers)], pp. 353–367. ISSN: 0016-7398. DOI: 10.2307/1792964. URL: https://www.jstor.org/stable/1792964 (visited on 07/13/2023).

Somrak, Maja, Sašo Džeroski, and Žiga Kokalj (Jan. 2020). “Learning to Classify Structures in ALS-Derived Visualizations of Ancient Maya Settlements with CNN”. In: *Remote Sensing* 12.14. Number: 14 Publisher: Multidisciplinary Digital Publishing Institute, p. 2215. ISSN: 2072-4292. DOI: 10.3390/rs12142215. URL: https://www.mdpi.com/2072-4292/12/14/2215 (visited on 04/04/2023).

Stanish, Charles (1997). “Nonmarket Imperialism in the Prehispanic Americas: The Inka Occupation of the Titicaca Basin”. In: *Latin American Antiquity* 8.3. Publisher: Society for American Archaeology, pp. 195–216. ISSN: 1045-6635. DOI: 10.2307/971652. URL: https://www.jstor.org/stable/971652 (visited on 07/13/2023).

— (Mar. 12, 2003). *Ancient Titicaca: The Evolution of Complex Society in Southern Peru and Northern Bolivia*. University of California Press. 380 pp. ISBN: 978-0-520-23245-7.

Stern, Steve J (1993). *Peru’s Indian peoples and the challenge of Spanish conquest: Huamanga to 1640*. OCLC: 939053242. Madison, Wis.: University of Wisconsin Press.

Szykulski, Józef (1998). “Churajón, nuevos datos sobre el Periodo de Transición en el Sur del Peru: BAS 30”. In: *BAS - Bonner Amerikanistische Studien Estudios Americanistas de Bonn*, p. 31. URL: https://www.academia.edu/13840323/Churaj%C3%B3n_nuevos_datos_sobre_el_Periodo_de_Transici%C3%B3n_en_el_Sur_del_Peru_BAS_30 (visited on 07/31/2020).

— (2008). “Arquelogía de Churajón, Sur del Peru”. In: *Tambo. Boletin de Archeologia*. (Visited on 01/17/2019).

Szykulski, Józef et al. (2000). *Investigaciones arqueológicas en Churajón, Departamento de Arequipa - Perú Informe de los trabajos de 1998/1999*. Informe. Arequipa: Universidad Católica Santa María” de Arequipa, p. 57.

Tan, Jimin et al. (June 8, 2021). *A critical look at the current train/test split in machine learning*. DOI: 10.48550/arXiv.2106.04525. arXiv: 2106.04525[cs]. URL: http://arxiv.org/abs/2106.04525 (visited on 04/18/2023).

Tjoa, Erico and Cuntai Guan (Nov. 2021). "A Survey on Explainable Artificial Intelligence (XAI): Towards Medical XAI". In: *IEEE Transactions on Neural Networks and Learning Systems* 32.11, pp. 4793–4813. ISSN: 2162-237X, 2162-2388. DOI: 10.1109/TNNLS.2020.3027314. arXiv: 1907.07374[cs]. URL: http://arxiv.org/abs/1907.07374 (visited on 06/30/2022).

Treacy, John M. (1987). "Building and rebuilding agricultural terraces in the Colca Valley of Peru". In: *Yearbook. Conference of Latin Americanist Geographers*. Vol. 13. JSTOR, pp. 51–57.

— (1990a). "Agricultural Terraces in Peru's Colca Valley: Promises and Problems of an Ancient Technology". In: *Fragile Lands Of Latin America: Strategies For Sustainable Development*, pp. 209–229. ISBN: 978-0-429-04280-5. DOI: 10.4324/9780429042805-14.

— (1990b). "The fields of Coporaque: Agricultural terracing and water management in the Colca Valley, Arequipa, Peru." In.

— (Apr. 11, 2019). "Agricultural Terraces in Peru's Colca Valley: Promises and Problems of an Ancient Technology". In: *Fragile Lands Of Latin America*. DOI: 10.4324/9780429042805-14. URL: https://www.taylorfrancis.com/ (visited on 10/01/2019).

Treacy, John M. and William Denevan (1994). "The creation of cultivable land through terracing". In: *The archaeology of garden and field*. Ed. by Naomi Miller and Kathryn Gleason. Philadelphia: University of Pennsylvania Press, pp. 91–110.

Trier, Øivind Due, David C. Cowley, and Anders Ueland Waldeland (Nov. 29, 2018). "Using deep neural networks on airborne laser scanning data: Results from a case study of semi-automatic mapping of archaeological topography on Arran, Scotland". In: *Archaeological Prospection*. ISSN: 10752196. DOI: 10/gf2b68. URL: http://doi.wiley.com/10.1002/arp.1731 (visited on 04/30/2019).

Ur, Jason A. (2013). "CORONA satellite imagery and ancient near eastern landscapes". In: *Mapping Archaeological Landscapes from Space*. Springer, pp. 21–31.

Van Buren, Mary (1996). "Rethinking the Vertical Archipelago: Ethnicity, Exchange, and History in the South Central Andes". In: *American Anthropologist* 98.2. Publisher: [American Anthropological Association, Wiley], pp. 338–351. ISSN: 0002-7294. URL: https://www.jstor.org/stable/682892 (visited on 07/12/2023).

VanValkenburgh, Parker (Mar. 1, 2017). "Unsettling Time: Persistence and Memory in Spanish Colonial Peru". In: *Journal of Archaeological Method and Theory* 24.1, pp. 117–148. ISSN: 1573-7764. DOI: 10/gfw2mv. URL: https://doi.org/10.1007/s10816-017-9321-7 (visited on 03/15/2019).

Verschoof-van der Vaart, Wouter B. and Karsten Lambers (2019). "Learning to Look at LiDAR: The Use of R-CNN in the Automated Detection of Archaeological Objects in LiDAR Data from the Netherlands". In: *Journal of Computer Applications in Archaeology* 2.1, pp. 31–40.

Voulodimos, Athanasios et al. (Feb. 1, 2018). "Deep Learning for Computer Vision: A Brief Review". In: *Computational Intelligence and Neuroscience* 2018. Publisher: Hindawi, e7068349. ISSN: 1687-5265. DOI: 10.1155/2018/7068349. URL: https://www.hindawi.com/journals/cin/2018/7068349/ (visited on 07/11/2023).

Wang, Sherrie, François Waldner, and David B. Lobell (Jan. 2022). "Unlocking Large-Scale Crop Field Delineation in Smallholder Farming Systems with Transfer Learning and Weak Supervision". In: *Remote Sensing* 14.22. Number: 22 Publisher: Multidisciplinary Digital Publishing Institute, p. 5738. ISSN: 2072-4292. DOI: 10.3390/rs14225738. URL: https://www.mdpi.com/2072-4292/14/22/5738 (visited on 05/10/2023).

Wernke, Steven A. (2006). "The politics of community and Inka statecraft in the Colca valley, Peru". In: *Latin American Antiquity* 17.2, pp. 177–208.

— (2007a). "Analogy or erasure? Dialectics of religious transformation in the early doctrinas of the Colca Valley, Peru". In: *International Journal of Historical Archaeology* 11.2, pp. 152–182. ISSN: 1092-7697. DOI: 10/ftpv6k.

— (2007b). "Negotiating Community and Landscape in the Peruvian Andes: A Transconquest View". In: *American Anthropologist* 109.1, pp. 130–152. ISSN: 1548-1433. DOI: 10/b5g3gh. URL: https://anthrosource.onlinelibrary.wiley.com/doi/abs/10.1525/aa.2007.109.1.130 (visited on 10/22/2019).

— (2010). "A reduced landscape: toward a multi-causal understanding of historic period agricultural deintensification in highland Peru". In: *Journal of Latin American Geography*, pp. 51–83. DOI: 10/d7dwm6.

— (2013). *Negotiated settlements: Andean communities and landscapes under Inka and Spanish colonialism*. 1. paperback printing. Gainesville: Univ. Press of Florida. 371 pp. ISBN: 978-0-8130-6093-4 978-0-8130-4249-7.

Wernke, Steven A. and Akira Saito (2019). *LOGAR Linked Open Gazetteer of the Andean Region*. LOGAR Linked Open Gazetteer of the Andean Region. URL: http://www.logarandes.org/place/391.html (visited on 11/12/2019).

Wernke, Steven A., Parker VanValkenburgh, and Akira Saito (Feb. 20, 2020). "Interregional Archaeology in the Age of Big Data: Building Online Collaborative Platforms for Virtual Survey in the Andes". In: *Journal of Field Archaeology* 45 (sup1), S61–S74. ISSN: 0093-4690, 2042-4582. DOI: 10/ggkwz2. URL: https://www.tandfonline.com/doi/full/10.1080/00934690.2020.1713286 (visited on 02/15/2020).

Wernke, Steven A., Parker VanValkenburgh, James Zimmer-Dauphinee, et al. (in press). "Large-scale imagery survey in archaeology and the Geospatial Platform for Andean Culture, History, and Archaeology (GeoPACHA)". In: *Antiquity*. inpress.

Wernke, Steven A. and Thomas M. Whitmore (Aug. 1, 2009). "Agriculture and Inequality in the Colonial Andes: A Simulation of Production and Consumption Using Administrative Documents". In: *Human Ecology* 37.4, pp. 421–440. ISSN: 1572-9915. DOI: 10.1007/s10745-009-9261-2. URL: https://doi.org/10.1007/s10745-009-9261-2 (visited on 05/27/2023).

White, Devin Alan (Nov. 2015). "The Basics of Least Cost Analysis for Archaeological Applications". In: *Advances in Archaeological Practice* 3.4. Publisher: Cambridge University Press, pp. 407–414. ISSN: 2326-3768. DOI: 10.7183/2326-3768.3.4.407. URL: https://www.cambridge.org/core/journals/advances-in-archaeological-practice/article/abs/basics-of-least-cost-analysis-for-archaeological-applications/DE502C37794C0E200AE7FA6A7529E25E (visited on 07/13/2023).

Wittfogel, Karl August (1957). *Oriental despotism [electronic resource]: a comparative study of total power / by Karl A. Wittfogel.* In collab. with American Council of Learned Societies. ACLS Humanities E-Book. New Haven: Yale University Press. xix+556.

Wright, A. C. S. (1962). “Some Terrace Systems of the Western Hemisphere and Pacific Islands”. In: *Pacific Viewpoint* 3.1. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/apv.31005, pp. 97–101. ISSN: 2638-4825. DOI: 10.1111/apv.31005. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/apv.31005 (visited on 12/14/2021).

— (1963). “The Soil Process and the Evolution of Agriculture in Northern Chile”. In: *Pacific Viewpoint* 4.1, pp. 65–74. ISSN: 2638-4825. DOI: 10.1111/apv.41006. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/apv.41006 (visited on 02/15/2023).

Yang, Naisen and Hong Tang (Jan. 2020). “GeoBoost: An Incremental Deep Learning Approach toward Global Mapping of Buildings from VHR Remote Sensing Images”. In: *Remote Sensing* 12.11. Number: 11 Publisher: Multidisciplinary Digital Publishing Institute, p. 1794. ISSN: 2072-4292. DOI: 10.3390/rs12111794. URL: https://www.mdpi.com/2072-4292/12/11/1794 (visited on 07/22/2022).

Zimmer-Dauphinee, James, Parker VanValkenburgh, and Steven A. Wernke (in press). “Eyes of the Machine: Building Toward AI-assisted Satellite Archeological Survey in the Andes”. In: *Antiquity*. in press.

Zingman, Igor et al. (Aug. 2016). “Detection of Fragmented Rectangular Enclosures in Very High Resolution Remote Sensing Images”. In: *IEEE Transactions on Geoscience and Remote Sensing* 54.8. Conference Name: IEEE Transactions on Geoscience and Remote Sensing, pp. 4580–4593. ISSN: 1558-0644. DOI: 10.1109/TGRS.2016.2545919.