

Learning Patterns of Tourist Movement and Photography from Geotagged Photos at Archaeological Heritage Sites in Cuzco, Peru

Abstract

The popularity of media sharing platforms in recent decades has provided an abundance of open source data that remains underutilized by heritage scholars. By pairing geotagged photographs from Flickr with machine learning and computer vision algorithms, we build upon the current theoretical discourse of anthropology associated with *visuality* and *heritage tourism* to identify travel patterns across a known archaeological heritage circuit, and quantify visual culture and experiences in Cuzco, Peru. Leveraging large-scale in-the-wild tourist photos, our goals are to (1) understand how the intensification of tourism intersects with heritage regulations and social media, aiding in the articulation of travel patterns across Cuzco's heritage landscape and to (2) assess how aesthetic preferences and visuality become entangled with the rapidly evolving expectations of tourists, whose travel narratives are curated on social media and grounded in historic site representations.

Keywords: geotagged photos, anthropology and computer vision, clustering, visuality, heritage tourism, visual recognition, Peru

1. Introduction

The rise of social media in recent decades has led to the large-scale influx of publicly distributed images. Photo-sharing websites like Flickr (est. in 2004) now house over 6 billion images generated by more than 40 million unique users (Crandall and Snavely 2012). Originally developed to aid in the organization of photos and enable sharing between users, media-sharing websites offer new and underutilized areas of study for cultural heritage researchers when paired with computer vision and machine learning algorithms. Heritage studies focused on visuality and tourism currently remain theoretically driven, with few moving beyond the conceptual stage to incorporate

methods that generate empirical evidence and quantify visual culture (i.e. photographs).

Using Cuzco, Peru, as a case study, we present an innovative application of computer vision and machine learning methods to understand archaeological heritage circuits, the evolution of their aesthetic legacies, tourist movement trends (and associated economic influences), and real-world visual experiences at heritage sites. To our knowledge, this is the first time that vision and machine learning techniques have been used to study the visual experiences of archaeological heritage tourism. Our goals are two-fold: to understand how the intensification of tourism intersects with heritage regulations and social media, aiding in the articulation of travel patterns across a known heritage landscape; and to assess how aesthetic preferences and visuality become entangled with the rapidly evolving expectations of tourists, whose travel narratives are curated on social media and grounded in historic site representations.

Located in the southwestern Andes mountains of Peru, Cuzco is the former capital of the Inca Empire (1438-1532 CE) and one of the world’s best-known tourist regions. Tourism contributes \$7.6 billion annually to the Peruvian economy and provides 3.9% of Peru’s GDP (Rice 2018). Much of Peru’s modern tourist economy centers on the Cuzco region and the idea of Andean “timelessness” which links the Inca past to the present (Covey 2017). The perception of timelessness depicted in Cuzco’s imagery aids in the creation of heritage hierarchies, aestheticizes places, and assigns economic and cultural values to heritage landscapes by selecting what is meant to be “seen” and preserved (Watson and Waterton 2010). In particular, imagery produced by 19th and 20th century explorers helped to popularize Cuzco as a tourist destination and critically influenced how archaeological sites were curated for mass consumption. The Inca sites first reported by the explorer Hiram Bingham have remained at the core of Cuzco’s tourism industry in part due to the sensationalism surrounding Bingham’s “discovery” of Machu Picchu in 1911.

Today, archaeological heritage tourism is structured by the *Boleto Turístico del Cuzco* (BTC), a multi-site pass that grants access to 10 sites in and around Cuzco. This includes: Sacsayhuaman, Tipón, Pikillacta, Tambomachay, Puca Pucara, Chinchero, Ollantaytambo, Pisac, Moray, and Qenqo. With the exception of the Wari (600-1000 CE) site of Pikillacta, all of the archaeological properties on the BTC are monumental Inca sites. Those 10 sites plus two others—the UNESCO World Heritage sites of Cuzco and

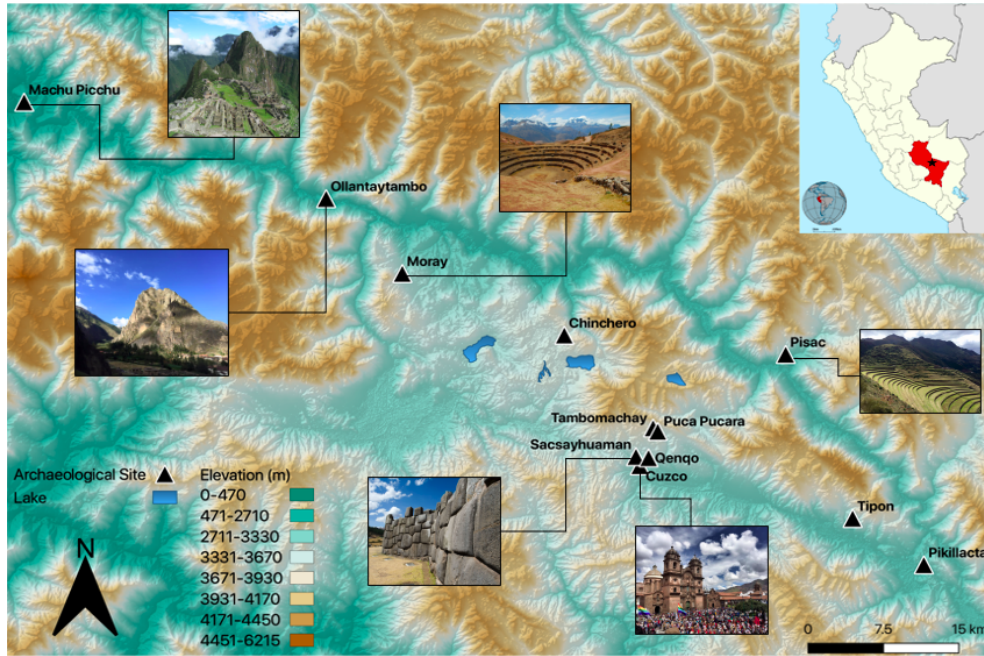


Figure 1: Multiscalar map of Cuzco, Peru with photographs representing the six most visited archaeological sites in our study.

Machu Picchu—comprise the 12 archaeological sites (Fig. 1) in the circuit we study.

Our work is interdisciplinary in nature, bridging the fields of Anthropology and Artificial Intelligence (AI) to build on the theoretical underpinnings currently driving studies of visual culture and heritage tourism. We approach the visual experiences, perceptions, and expectations of tourists traveling to the Cuzco region through a study of visibility. “Visibility” is a concept in anthropology that incorporates how vision is culturally and historically constructed and is defined by “how we see, how we are able, allowed, or made to see, and how we see this seeing” (Rose 2016). Heritage tourism is highly visual in nature, with archaeological sites being marketed and packaged to tourists as part of a national ‘brand’ that is produced and reinforced over time. The images that accompany heritage tourism and those that are shared on social media influence the commodification of landscapes, ruins, and native bodies and constrain what consumers see during short-term visits. They then in turn impact how tourists record, disseminate, and shape heritage

images. For this reason, visuality is a significant framework from which to critically assess the historic trajectory and conceptualization of archaeological heritage circuits generated through imagery.

We aim to quantify visuality and tourism patterns across Cuzco’s heritage circuit. To this end, we create a tourist photo dataset from Flickr, with photos from 2,261 Flickr users at 12 sites with associated geotags and timestamps spanning 2004-2019. We introduce a series of analysis using vision and learning to recover insights relevant to anthropology from this imagery. In particular, tourist movement across sites is fit to a Markov model, which is then applied to determine travel patterns between BTC sites via multiple mini-itineraries (day trips), as well as general sequences of movement between BTC and UNESCO sites. The density of unique users, total photos taken, and average time spent by users per site are analyzed to assess relative popularity among Cuzco’s archaeological sites, as well as patterns of movement *between* sites that reflect how multiple, rural heritage sites are bundled into scripted tour itineraries. Through our analyses, we compare the significance of BTC circuits and individual archaeological heritage sites on the landscape. Finally, to analyze the iconicity of modern tourist photos, we identify major photography themes at each site by clustering in the visual feature space of the photos, and we consider how canonical images define Cuzco’s tourist experience and continue to reinforce the region’s aesthetic legacy.

We first review related work in Sec. 2. Sec. 3 provides an overview of archaeological heritage tourism in Cuzco, Peru. Sec. 4 introduces our methodology for the case study. In Sec. 5, we present the statistical results and mined patterns of tourist movement and canonical view themes from our Flickr dataset. We conclude with a summary of our findings.

2. Related Work

Archaeological applications of machine learning have been limited to typological and stylistic analyses of artifacts (Hörr and Brunnett 2014, Gansell and Wiggins 2014), with the majority of studies concerning the representation of peoples and places centered around theoretical discussions of visuality (Garrod 2009, Burns and Lester 2010, Urry and Jonas 2011, Sand 2012).

Scholars in the Cuzco region have made significant theoretical contributions to studies of heritage and visuality (Poole 1992, Scorer 2014, Silverman 2013). This scholarship has focused on tracing the history of the visual economy and the emergence of new technologies (e.g., mass printing and

photography) that transformed the production, circulation, and interpretation of visual culture over the last three centuries (Poole 1992, Scorer 2014). Investigations yielded the development of a visual aesthetic for Andean ruins (and indigenous people), whose perceived “authenticity” became part of their value within national heritage programs and the tourist economy. This paper builds on these thematic narratives, applying machine learning methods to the current standard for visual dissemination—internet-based digital images—to generate empirical patterns that can link current practices to the overall historical arc of exploration and heritage tourism in Cuzco.

More recently, heritage studies have taken a top-down approach to visibility using aerial photography, satellite imagery, and drones. These remote sensing techniques have become popular archaeological tools for “seeing”, monitoring, and discovering archaeological heritage sites from afar (Casana and Laugier 2017, Bevan and Lake 2016, Parcak 2015). Remote sensing imagery tends to be collected at larger scales by governments, private companies, and scientific researchers and it is typically designed to capture landscape and site characteristics, not to see how tourists use those places. More relevant to our work, publicly posted images especially those shared on social websites like Instagram, Flickr, and Twitter, offer a distinct visual resource. These photos are accompanied by geotags, timestamps, and user IDs, which permit the direct observation of how tourists move within and between heritage sites (e.g., (Lee and Tsou 2018, Zheng et al. 2012, Hu et al. 2019) study tourist movement in national parks and urban cities), as well as which images they choose to capture and post online.

In the computer vision literature, the popularity of photo sharing platforms has opened up new possibilities to understand our world at a lower cost and in a more automated way. For example, at popular tourist sites like Trevi Fountain and Notre Dame, photos are taken densely from all viewpoints and locations around the sites, by large volumes of tourists. Pioneering work harnessed this open source data to apply structure from motion (Hartley and Zisserman 2003) and reconstruct 3D point clouds for entire sites (Snavely et al. 2006). Multi-modal analysis of large collections offers new ways to visualize tourist trends (Crandall et al. 2009). As many tourist photos are street scenes or self-portraits irrelevant to the sites themselves, automatic detection of *iconic images* can help capture key parts of a site, with applications in photo summarization and browsing (Zhang et al. 2018) and landmark segmentation (Simon and Seitz 2008). When geotags are not available, vision methods can estimate geolocation (latitude and longitude) from the images

themselves (Hays and Efros 2008, Kalogerakis et al. 2009, Chen and Grauman 2011). Timestamps and the photo owner’s ID provide temporal context that allows for travel sequence modeling for accurate geolocation (Kalogerakis et al. 2009, Chen and Grauman 2011, Crandall et al. 2009).

Pushing forward to combine these spatial and temporal findings to gain socio-cultural insights, recent studies have measured and tracked the occurrence of ecological phenomena (Zhang et al. 2012, Wang et al. 2013, 2016), while also successfully estimating the demography of 200 US cities by using 50 million images of street scenes gathered with Google Street View cars (Gebru et al. 2015, 2017). Prior to such applications, these results were only available through manual labor which does not scale well or update in real-time. Recently, scholars have turned toward social media and visual content analysis using geotagged photos to understand tourist behavior and perceptions at destination hotspots (Zhang and Li 2019, Vu and Ye 2015, Stepchenkov and Zhan 2013, Kaufman and Stettler 2019). At a high level, our work is in a related socio-cultural vein. However, we are the first to explore machine learning for heritage tourism studies, specifically for heritage circuits and visual culture in Cuzco, Peru. Our findings of tourist flows within a known circuit provide insights to heritage conservation, while our content analysis of tourist photos sheds light on how modern peoples visually experience Cuzco’s heritage landscape and reproduce the region’s aesthetic legacy.

3. Historic Arc of Inca Archaeological Tourism in Cuzco

Next we present historical context for our inter-disciplinary study. Following more than a century of foreign expeditions to the Cuzco region, archaeological tourism to the former Inca heartland took root in the early 1900s (Rice 2018). Hiram Bingham represents the nexus of the expeditionary aspiration to “discover” something unknown and introduce it to the developed world and the more scripted itinerary of the tourist. Through his relationship with *National Geographic Magazine* (Bingham 1913), Bingham shared his early encounters with the “lost city” of Machu Picchu, but his expeditions also helped to crystallize the modern touristic route in Cuzco by illustrating and describing key sites, including Sacsayhuaman, Tambomachay, Puca Pucara, Ollantaytambo, Pisac, Tipón, and Pikillacta (Bingham 1922). The vivid imagery of Bingham’s publications encouraged growing numbers of foreigners to visit the Cuzco region as tourists, with guidebooks appearing by

the early 1920s (Bauer 2004). Over the next decade, government officials recognized the economic potential of tourism and began to fund reconstruction projects at tourist sites, including Sacsayhuaman, Tambomachay, Pikillacta, and Pisac.

Following a powerful earthquake that destroyed most of Cuzco’s modern buildings in 1950, archaeological tourism reemerged with the help of the newly formed United Nations Educational, Scientific, and Cultural Organization (UNESCO). Annual tourist visits to the Cuzco region topped 100,000 by the mid-1970s (Steel 2008) and local archaeologists and officials sought to develop broader itineraries that would encourage tourists to spend more time (and money) in the Inca heartland. Established in 1978, Cuzco’s original Boleto Turístico (BTC) allowed 10-day visitor access to the region’s top archaeological attractions (Rice 2018). Tourist initiatives in the 1970’s also included the promotion of the Inca Trail, a heritage route to Machu Picchu (Maxwell 2012), which along with the monumental center of Cuzco became the first Peruvian properties to be added to UNESCO’s World Heritage List in 1983.

Today, roughly 90% of all international visits to Peru feature a stop in Cuzco, with half of these including trips to Machu Picchu (Larson and Poudyal 2012). Machu Picchu is one of the “New7Wonders” and annual visits to the Cuzco region now top 3 million, an increase of more than 500% since 2000 (Dircetur 2017). The influx of tourists to these two UNESCO sites has had the added effect of bolstering sales of the BTC while entrenching and perpetuating Cuzco’s established heritage narrative of Inca monumentality.

In 2008, the BTC circuit was expanded to include: a 1-day pass to Sacsayhuaman, Qenqo, Tambomachay, and Puca Pucara (BTC I); a 2-day pass to Cuzco’s six museums, and the archaeological sites of Tipón and Pikillacta (BTC II); and a 2-day pass to Pisac, Ollantaytambo, Chinchero, and Moray (BTC III). The addition of the 1 and 2-day passes has contributed to an uptick in visits to Cuzco’s more remote sites (e.g., Moray (Dircetur 2017)), and has also affected movement patterns between BTC sites, as we will examine quantitatively through image data below.

4. Methodology

Having provided the historical context for our study, we now present our technical methodology to elicit a quantitative story for this heritage circuit,

and to connect the last 15 years of tourism to visuality established by early explorers.

We collected data from Flickr to obtain images taken at the 12 archaeological sites in our study (Sacsayhuaman, Tipón, Pikillacta, Tambomachay, Puca Pucara, Chinchero, Ollantaytambo, Pisac, Moray, Qenqo, Machu Picchu, Cuzco) and to create user albums for each unique user ID. To detect patterns of tourist movement and analyze their change over time, we model tourists’ transition sequences with a Markov model (Kemeny and Laurie Snell 1960). To identify scenes and objects photographed by users at each site and detect patterns of iconicity, a convolutional neural network is used to extract semantic features from images. A clustering algorithm is then used to discover common image themes at each site, such as mountains, stonework, alpaca, terracing, etc. We detail each of these components in the following section.

Dataset Creation Images for all BTC and UNESCO sites were collected using Flickr’s API for the 15-year period since Flickr’s creation, from 2004 to 2019. Google Earth was used to establish a central point of latitude and longitude for each site, as well as an estimated spatial extent (km) and buffer.

We downloaded the data in two stages. In Stage 1, we collected photos taken at all 12 archaeological sites (listed in Sec. 1) by querying with the site’s GPS coordinates and taking the top 4,000 retrieved images from each site. Images with GPS coordinates beyond the estimated buffer from the central point of the site were later eliminated. In Stage 2, we then expanded this collection to form an album for each user by querying with all unique user IDs from Stage 1, and specified to retrieve photos taken within the time the user visited the sites. Each photo includes meta-data like the image’s url, owner’s ID, geotag, and timestamp. In total, 57,804 images were collected from 2,261 users. Table 1 shows the total number of photos per site and the number of unique visitors who traveled to the 12 sites. These metrics were used to infer a site’s popularity ranking within our dataset. Machu Picchu was found to be the most popular site visited (as it is advertised as the penultimate destination in the Cuzco region), while the Wari site of Pikillacta (the only non-Inca site included on the BTC) was found to be the least popular.

Mining Tourist Traveling Patterns We model tourists’ traveling patterns with a Markov chain, a stochastic model for a sequence of events that assumes that the probability of each event depends only on the state at-

tained in the previous event (Géron 2017). In our study, an event is a site (e.g., Cuzco), and a sequence is the order of sites a tourist traveled (e.g., start in Cuzco and go to Sacsayhuaman and finally to Machu Picchu). With photos associated with each unique user (tourist) and site, we represent the tourist traveling pattern among all sites with a Markov matrix P . It is a square matrix that describes the transitions of a Markov chain, where the ij -th entry P_{ij} is the probability of moving from site i to site j . Each row thus sums to 1. We calculate the transition frequencies from site to site to obtain these probabilities. This Markov matrix allows us to discover popular transitions between sites, how transitions are affected by policies regulating heritage landscapes (e.g., the release of the BTC ticket package (Sec. 5)), and how transitions change over time, as we will discuss below.

Discovering Canonical View Themes Deep neural networks trained on large scale data provide a powerful and compact representation. In vision, features extracted from deep convolutional neural networks (CNN) are widely used to capture objects and scenes in an image (Donahue et al. 2014, Sharif Razavian et al. 2014, Girshick et al. 2014). Here, we adopt a successful CNN called ResNet50 (He et al. 2016), which is pre-trained on ImageNet (Deng et al. 2009), a database with more than 10 million images labeled with 1K common object categories. We extract our image features using its penultimate layer, since higher layers are known to capture semantic features such as object categories. Pre-training the visual encoder enriches the features to capture common low – and high – level visual patterns, beyond mere pixels. By utilizing a CNN pre-trained on ImageNet for our t-sne and clustering analysis, we are able to identify multiple real-world objects simultaneously without deploying different algorithms for each object and without the added cost of manual labor (for which the task would be impossible).

With this rich representation, our goal is to find common themes of canonical views taken at each archaeological site. *Clustering* is a data mining technique that groups objects in a way such that those in the same group (called a cluster) are more similar (in some sense) to each other than to those in other groups (clusters). As such, each cluster corresponds to a discovered theme, and all photos taken at a site can be captured by the representative images, i.e. *exemplars*, from all themes.

To this end, we adopt affinity propagation (AP) (Frey and Dueck 2007) clustering. AP views data points as nodes in a network and exchanges messages between nodes until a set of good quality exemplars emerges. It has

the benefits of automatically deciding the number of clusters in a data set, and also representing clusters as *exemplars* instead of a mean/median of the cluster. Fig. 7 shows examples of our discovered themes.

Aside from discovering canonical views with clustering, we also explore the statistics of which types of scenes are observed across the photos at each site. Specifically, to compute our scene-site occurrence matrix (which we will present below in Fig. 5) we use ResNet50 pre-trained on MIT Places (Zhou et al. 2017). MIT Places consists of over 10 million images and 400 unique scene categories. Because ImageNet is an object-centric database, it is not possible to extract scene labels, making it necessary to also train on MIT Places for this specific task. While some of the general AI classifications associated with the Places database reflect non-Andean categories (e.g. medina, rice paddy, catacomb, etc.), shared features remain high and do correlate to Andean life-ways (e.g. market, agriculture, tombs, etc.).

5. Results and Analysis

We analyze (1) tourists’ statistics at each site, (2) whether tourist transition patterns may be attributable to available BTC tickets, and (3) the discovered canonical view themes.

Table 1 shows the total number of photos taken per site per year and Table 2 shows site popularity based on the total number of photos taken per site and the total number of unique visitors. While our best efforts were made to avoid sampling bias and statistical oversight, not all tourists traveling to BTC and UNESCO sites take the same amount of photos, geotag photos, or even upload photos to Flickr. For these reasons, biases may exist when calculating the average time a user spent at each site since our calculations are dependent on photo timestamps. This may also create biases when generating site popularity rankings across Cuzco’s heritage landscape. Despite these factors, our quantitative historical study is valuable for addressing tourist practices at a key moment of economic (BTC elaboration and mass tourism), social (share-ability of personal imagery), and technological (ubiquity of smart phones with cameras) change in the Cuzco region.

Flickr’s rising popularity since its inception in 2004 has led to an increase in data (moving toward the present) as more users signed up to organize, store, and share their images. As seen in Table 1, the number of photos per year ranged significantly across sites over time. The least amount of data collected for our study corresponded to the launch of the Flickr platform in

Table 1: Total number of photos taken per site from 2004-2019. The most popular photo years for each site are shown in bold.

(a)

Year	Machu Picchu	Cuzco	Sacsayhuaman	Ollantaytambo	Pisac	Chincho
2019	293	71	38	6	4	0
2018	1768	675	96	232	172	57
2017	2697	789	326	369	410	205
2016	2365	719	330	409	224	149
2015	2903	1296	434	268	160	79
2014	2601	1245	542	229	308	74
2013	3376	1157	609	515	530	150
2012	4050	1852	362	439	399	50
2011	2649	1005	186	164	132	34
2010	1899	756	221	230	201	87
2009	1678	775	225	123	123	43
2008	2794	848	327	174	292	13
2007	1073	380	188	78	42	29
2006	619	131	53	21	57	14
2005	651	124	8	28	15	0
2004	134	19	13	3	8	0

(b)

Year	Moray	Qenqo	Puca Pucara	Tambomachay	Tipón	Pikillacta	Total
2019	0	0	13	0	0	0	425
2018	75	9	21	15	0	1	3121
2017	122	62	72	80	61	12	5205
2016	73	31	33	36	46	11	4426
2015	76	40	42	33	46	100	5477
2014	50	57	34	42	7	7	5196
2013	161	84	116	55	54	43	6850
2012	35	48	59	30	13	1	7338
2011	64	31	49	39	77	38	4468
2010	67	13	19	12	8	18	3531
2009	35	47	48	51	28	0	3176
2008	36	108	23	54	1	3	4673
2007	11	17	10	1	82	50	1965
2006	25	7	0	0	5	4	936
2005	10	1	1	1	0	0	839
2004	0	0	0	1	0	0	178

Table 2: Total number of photos taken per site and total number of unique visitors per site. Site popularity (ranked) is inferred from these numbers.

Site	# photos	# unique visitors	popularity rank
Machu Picchu	31550	1498	1
Cuzco	11842	1273	2
Sacsayhuaman	3958	524	3
Ollantaytambo	3288	381	4
Pisac	3077	348	5
Moray	840	156	6
Chincho	984	143	7
Qenqo	555	139	8
Puca Pucara	540	125	9
Tambomachay	454	108	10
Tipón	428	34	11
Pikillacata	288	27	12

2004, while the year with the greatest amount of data available *across* sites was 2013. The largest number of *total* photos per year in our dataset can be found in 2012. There is also the question of whether Flickr users choose to make their photos publicly available, further influencing sampling strategies and the overall results of our analysis. Although it is not a perfect vehicle to approach how *all* tourists move across Cuzco’s heritage sites, the Flickr archive reflects a dynamic reality of publicly available images.

Popularity of Sites Site popularity was determined by comparing the total number of Flickr images gathered per site versus the total number of unique visitors per site. Moreover, the average amount of time a user stayed at a site was calculated from user albums collected in Stage 2. The earliest and latest time a photo was taken at each site were identified for each user. Time elapsed from the earliest to the latest time serves as the total amount of time a user spent at a given site.

A density map was generated (using ArcGIS software) from the total number of photos taken by users at all 12 sites to visualize landscape hotspots (Fig. 2, Tab. 2). The number of unique visitors (traveler icon) and average time spent (clock icon) at each site are also shown. Machu Picchu was unsurprisingly the most popular site (1,498 visitors and 31,550 photos taken, followed by Cuzco (1,273 visitors and 11,842 photos taken), as the former is one of the most iconic heritage sites globally and the latter is the major point of entry to the region. Of the sites visited, Pikillacta and Tipón were the

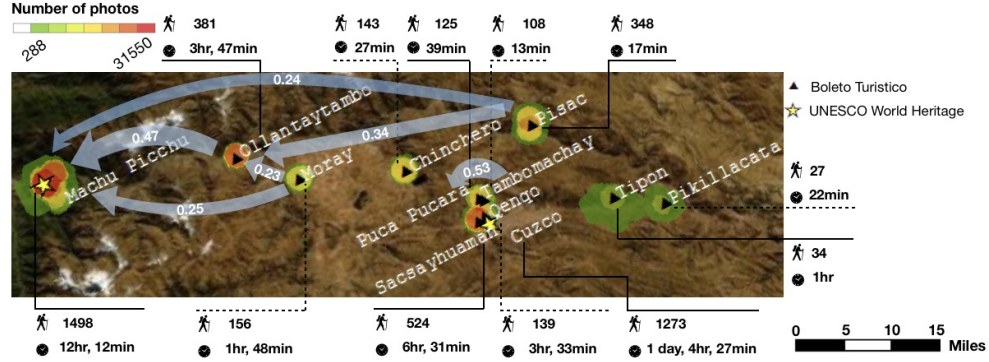


Figure 2: Overview of tourist movement patterns discovered in the image data. Popularity of sites measured by the unique number of tourists (shown by traveler icons) and number of photos taken per site (color coded in heatmap). The average time tourists spent per site is shown by clock icons. Top transitions between sites are on arrows labeled by probabilities.

least popular, likely due to their remote location to the southeast of Cuzco.

Popularity rankings were generally consistent between the two variables (visitor count and photo frequency), with slight differences between Chinchero and Moray. Tourists spent the longest time at Cuzco, which serves as the region’s only transportation hub and is where most hotel and guesthouses are located. Tourists were found to spend upwards of 10 hours at Machu Picchu. Travel to Machu Picchu is often packaged as a day trip or overnight stay via rail, allowing for more time to be spent at the site. Tourists may also choose to trek the Inca Trail, which involves a multi-day excursion through the Andes mountains to reach Machu Picchu. By contrast, tourists spent less than an hour at Tambomachay, Pisac, Chinchero, and Puca Pucara. Tour groups with limited time schedules may account for the amount of time spent at Pisac and Chinchero, whereas Tambomachay and Puca Pucara are small and can be covered quickly before moving to more popular sites like Sacsayhuaman and Qenqo.

These statistics offer a window into the cultural heritage circuit over the last 15 years in ways that are not possible with traditional methods. For example, ticket purchase rates or manual surveys do not capture durations of visits, photographic preferences, etc., but the large-scale tourist photos do.

Transition Patterns Changing over Time Next we examine the movement of tourists. We refer to our Markov matrix over all 12 archaeological sites as the *transition* matrix. First, we compute the transition probabilities for the entire 2004 to 2019 time range. The top 6 transitions with probabilities greater than 0.2 are plotted as arrows in Fig. 2. The most frequent transition is from Tambomachay to Puca Pucara as they are only a 3 minute walking distance. The other most frequent transitions moved toward Machu Picchu. In most cases, visitors traveled to one of four sites prior to departing for Machu Picchu: Ollantaytambo, Moray, Pisac, and Chinchero. Their proximity to Machu Picchu may account for this pattern.

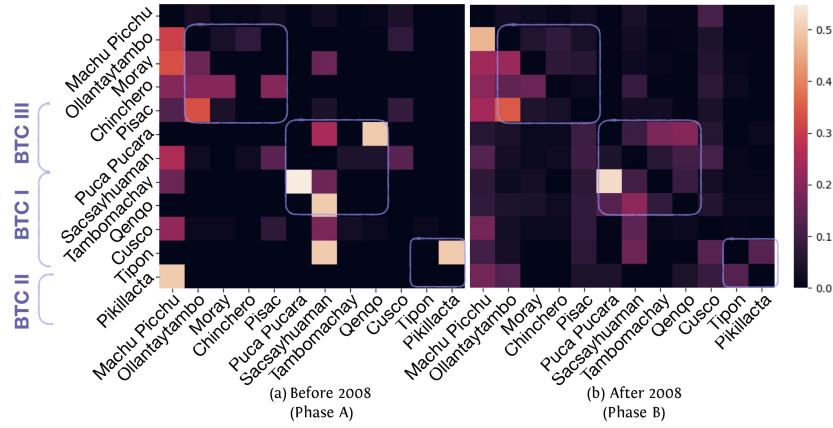


Figure 3: Transition matrix comparison of before and after the introduction of the BTC tourist package (2008). In general, transitions after 2008 become denser. Three outlined boxes represent site transition within BTC I, II, III respectively. Note how transitions within each group become more dispersed after the ticket packages were released.

Next, we compute a separate transition matrix for travels made before (Phase A) and after 2008 (Phase B), which marks the introduction of additional BTC tourist packages: a 1-day pass to Sacsayhuaman, Qenqo, Tambomachay, and Puca Pucara (BTC I); a 2-day pass to the archaeological sites of Tipón and Pikillacta (BTC II); and a 2-day pass to Pisac, Ollantaytambo, Chinchero, and Moray (BTC III). Fig. 3 shows the matrices plotted as heatmaps, where the i^{th} row and j^{th} column element contains the probability of moving from site i to site j . Phase A displays higher site specific transition sequences, like Pikillacta to Machu Picchu and Puca Puca to Qenqo,

and low regional diversity of movement between sites. The site specific transitions in Phase A decrease in Phase B as regional transition probabilities increase across the matrix. This suggests that there is a greater range and dispersal of site transitions in Phase B, weakening the stronger site-to-site sequences that previously existed. The three outlined boxes represent BTC I, II, and III groupings respectively.

For the BTC I grouping, Phase A transitions are concentrated around nearby sites, i.e., Tambomachay to Puca Pucara or Qenqo to Sacsayhuaman. During Phase B, transition patterns become more symmetric and diffused among all four sites. This is true for transitions in BTC II and III as well, though slightly less obvious. The change in BTC transition patterns may indicate how the ticket packages encourage tourists to explore more sites through condensed day-trip itineraries. In Phase B, travel patterns shift away from BTC I sites and toward BTC III sites probably due in part to increased hotel construction in the Sacred Valley, which has intensified the use of Ollantaytambo as the point of departure to Machu Picchu, as well as access to Chinchero, Pisac, and Moray.




Figure 4: Comparison of transition patterns considering 1-day and 2-day passes. Significant changes are noted in transitions from BTC III sites to Machu Picchu while observed transition probabilities of tourists traveling to Machu Picchu from BTC I sites remained low.

Movements between sites on the BTC III ticket (Chinchero, Pisac, Ollantaytambo, Moray) were the most complex across Phase A and Phase B. A popular circuit between BTC III sites was discernible during Phase A and appeared to be the antecedent for the BTC III (2-day) pass: Following the

Table 3: Transition trends from BTC I and BTC III sites to Machu Picchu during Phase B are compared. A significant change in transition rates is seen when moving from BTC III sites to Machu Picchu (10-day pass vs. 2-day pass), while travel from BTC I sites to Machu Picchu remains relatively low regardless of 10-day vs. 1-day pass in Phase B. For the BTC I site of Sacsayhuaman, this is a departure from transition popularity in Phase A.

(a) Transition trends from BTC I sites to Machu Picchu	
Site Transition	10 Day v.s. 1 Day Pass
Sacsayhuaman → Machu Picchu	13% vs 13%
Qenqo → Machu Picchu	8% vs 6%
Puca Pucara → Machu Picchu	6% vs 7%
Tambomachay → Machu Picchu	7% vs 8%
(b) Transition trends from BTC III sites to Machu Picchu	
Site Transition	10 Day v.s. 2 Day Pass
Ollaytantambo → Machu Picchu	45% vs 35%
Pisac → Machu Picchu	13% vs 21%
Chincherro → Machu Picchu	5% vs 12%
Moray → Machu Picchu	24% vs 22%

addition of the BTC III (2-day) pass, visits to BTC III sites and transition patterns shift significantly. The most significant of these impacts affected travel to Pisac. During Phase A, visitors to Pisac accessed the site from two main points: Chincherro or Sacsayhuaman. Phase B saw the disappearance of these two transition points to Pisac, with all sites similarly likely to access it, while the main transit *out* from Pisac, i.e., heading towards Ollantaytambo, remained the same in both Phases.

Overall, the transition matrix indicates that changes to the BTC reflect existing touristic access practices as well as the altered ways that tourists might move among the sites on a given BTC ticket and how they might transition between those tickets and other heritage sites.

BTC 1 and 2-day Passes Next we explore Phase B transition trends by isolating tourist movement within a 1 and 2-day time frame. Figure 4 represent heatmaps for 1-day and 2-day periods corresponding to the maximum time duration allowed for travel to BTC I, II, III sites during Phase B and transition probabilities greater than 10%. Our results show that the introduction of the 2-day pass in Phase B led to a significant shift in transitions

from BTC III sites to Machu Picchu when compared to the full 10-day pass. We observed a ten percent decrease in transitions from Ollantaytambo to Machu Picchu (10-day vs. 2 day) as other points of departure corresponding to BTC III (Chincho and Pisac) saw significant increases (Tab. 3). This also corresponds to similar changes observed when comparing the BTC transitions in Phase A and Phase B above (Fig. 3).

We also observed that movement from BTC I sites to Machu Picchu compared to both the Phase B 10-day and 2-day passes is significantly lower than movement from BTC III sites (Tab. 3). The low transition probability of BTC I sites to Machu Picchu in Phase B mainly impacted the site of Sacsayhuaman, which saw a relatively high rate of transition to Machu Picchu during Phase A and a severe decrease in transition probabilities in Phase B. While the changes to the BTC have not affected how BTC I sites near Cuzco are accessed, the addition of the 1 and 2-day passes in Phase B created strong differences in transitions to Machu Picchu from associated BTC III sites. This is attributed in part to an increase in hotel beds in the Sacred Valley allowing tourists greater flexibility when starting and ending their BTC III days at different sites (depending on whether they are staying in Cuzco, Ollantaytambo, or Urubamba).

Again, these findings are valuable to understanding the heritage circuit for anthropologists, and they become significantly more accessible by detecting patterns from in-the-wild photos, as we propose.

Canonical View Analysis Next we examine how tourists are photographing these sites, and how those patterns interact with heritage-based operations. We first detect the scene types present in each photo to obtain the distribution of captured scenes across all sites (Fig. 5). We then break this distribution down to first visualize which scenes dominate a specific site (Fig. 6), and then determine what the scenes look like at each site (Fig. 7). Finally, we quantitatively complement our above analysis with statistics of discovered photo clusters at each site (Tab. 4), and discuss our findings in detail.

What do tourists notice at the sites and consider worthwhile to photograph and post online? Do tourist images converge around a set of canonical views, some of which were first introduced to popular media more than a century ago? To answer these questions, we quantify the contents of discovered canonical views per site. We use a ResNet50 (He et al. 2016) pre-trained on scene categories in the MIT Places (Zhou et al. 2017) dataset to predict scene




Figure 5: Scene-site occurrence matrix showing distribution of photographed scenes across sites. The most common scene labels were: mountain path and valley (i.e. mountains), archaeological excavation and ruins (i.e. stonework), amphitheater (i.e. terracing). These statistics reveal what tourists saw and photographed most often at each of the 12 sites.

labels in our dataset. Predicted scenes were aggregated on photos in each site to retrieve the most frequent scene labels. We take the most frequent 10% of scenes from each site as a representative scene category, and plot this scene-site occurrence matrix (Fig. 5). The scenes are sorted from most natural (left) to most artificial (right). Common scenes occurring across all BTC and UNESCO sites were: ‘Mountain Path’ and ‘Valley’ (i.e. mountains), ‘Archaeological Excavation’ and ‘Ruins’ (i.e. stonework), and ‘Amphitheater’ (i.e. agricultural terracing). Labels such as ‘Medina’, ‘Kasbah’, and ‘Bazaar’ also corresponded well to market scenes captured in Cuzco, Pisac, and Chinchero. Roadside markets at the entrance of archaeological sites are common in the Cuzco region and account for similar labels associated with Pikillacta and Tipón. Our detected scene-labels show a strong correlation to major scene themes identified within our t-sne and cluster results below.

To visualize the photos at each *scene*, we sample photos from each site, extract their CNN features using ResNet50 and ImageNet, and run a t-sne (Maaten and Hinton 2008) visual embedding, as shown in Fig. 6.

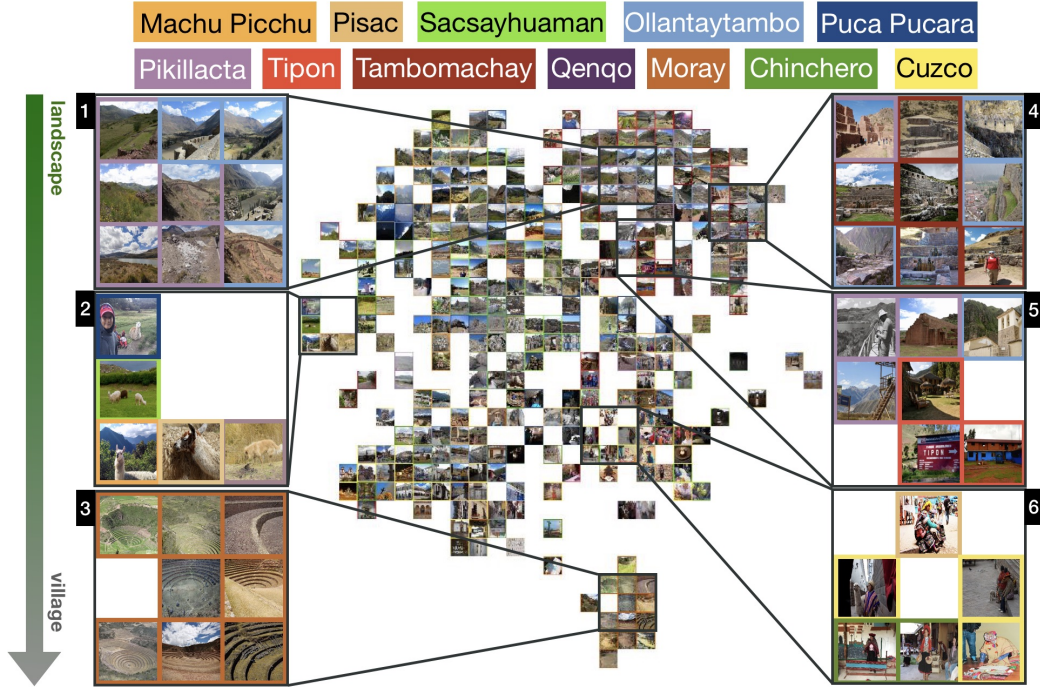


Figure 6: t-sne (Maaten and Hinton 2008) visualization of a subset of photos from all sites. Images are framed by colors corresponding to sites shown in top legend. Visualization shows a gradual theme variation from landscape views (top) to village-life (bottom). Bounding boxes show zoom-ins of discovered example themes (e.g. mountains, stone work, agriculture terracing). Similar architectural (Box 4) and terrain aesthetics (Box 1, 3) across sites reveals shared visuality in Cuzco’s heritage circuit.

Vizualization of shared themes (i.e., predicted scenes) across all sites include mountain landscapes, alpaca, and stone architecture (Box 1, 2, 4) which are similar to our common scene-labels. Photos of Peruvian peoples wearing traditional dress are mostly found in Cuzco, Chinchero, and Pisac (Box 6). At these locations, large open-air markets attract tourists wishing to purchase Peruvian textiles, pottery, and souvenirs. Cuzco also hosts many festivals and parades to which traditional dance groups from across the country are invited. A high percentage of images also correspond to colonial architecture, particularly in Cuzco and Chinchero where colonial churches are a major attraction (Box 5).

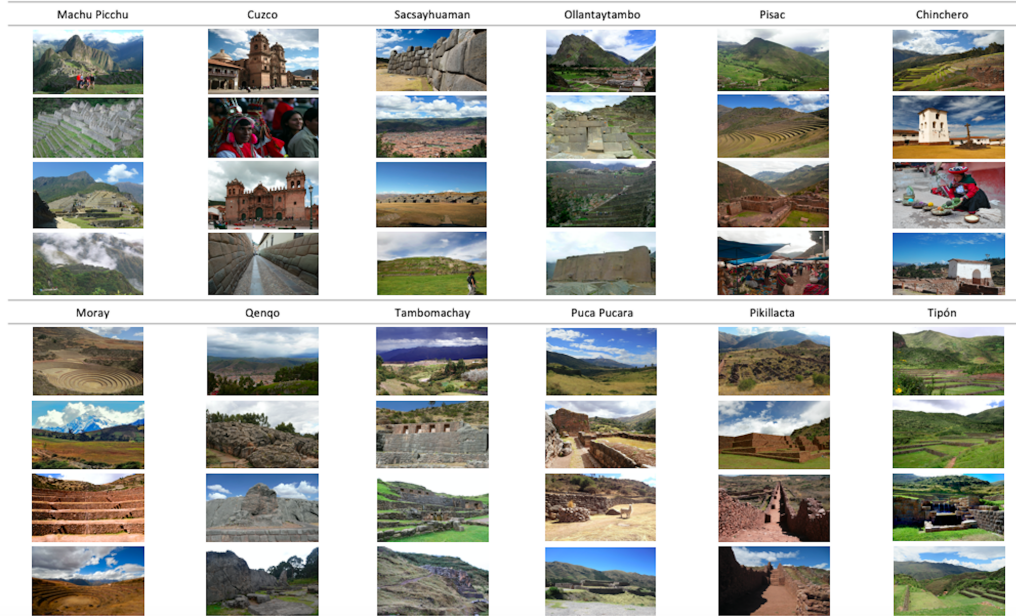


Figure 7: Representative cluster images are pictured for all 12 sites for all 57,804 photos captured by more than 2,000 unique travelers. Not only do popular themes (e.g. stone architecture, mountains, agricultural terracing) correspond to our distribution of detected scenes in Figure 5, but also images of site-specific scenes (e.g. colonial architecture, indigenous peoples, market scenes) are depicted for the sites of Cuzco and Chinchero.

To visualize representative photos at each *site*, we run affinity propagation (AP) algorithm per site to discover clusters, and take the center images from the largest clusters as representative images (Fig. 7). Although tourist interests leaned slightly toward scene-captures associated with artificial scenes containing some manmade elements, our results showed a near-equal split between natural and artificial scene labels corresponding with our cluster themes. Scene labels also showed a strong correlation with our t-sne visualizations. Our t-sne and cluster analysis indicates that modern tourists continue to be drawn to the same exotic view-sheds depicted in early expeditionary images: Andean landscapes (e.g. terracing), Inca ruins, indigenous people, and Colonial architecture (Fig. 8).

To quantitatively understand the silhouette of clusters discovered at each site, we look at three measurements: number of clusters discovered, separa-

tion score between clusters, and the compactness score within a cluster. The first two measure the *diversity* of clusters, as the more clusters discovered and the more separated clusters are, the more different clusters are from each other—and hence the more variety is present in how tourists photographed the given site. The compactness score measures the *homogeneity* of clusters, as the tighter or denser a cluster is, the more similar data points are to each other within that cluster—and hence the less variation in tourist photos of a particular landmark. The separation score is computed as the visual distances between cluster exemplars, averaged over all pairwise clusters. The compactness score is computed as the maximal pairwise visual distance between data points in the cluster, averaged over all clusters.¹

Table 4 presents the quantitative metrics. Cluster sizes range from 364 to 1,425, with 114 being the average for BTC sites and 1,034 for UNESCO World Heritage sites. Overall, UNESCO sites have more diverse clusters, as measured by the number of clusters and separation score, reflecting the larger transition time at these sites, whereby people are spending more time, and taking more varied images. Conversely, Machu Picchu and Ollantaytambo have the highest compactness scores, revealing that the reproduction of popular photo themes and viewpoints are more likely at large sites.

These patterns shed light on the impact of cultural heritage management decisions in recent years. Due to the growth of mass tourism, several of Cuzco’s most popular archaeological sites are no longer ‘free-range’. Instead, movement *within* these sites is restricted to a predefined path. For example, at Machu Picchu visitors must purchase their site permits months in advance to guarantee entry before they travel. Upon arrival to the site, tourists are guided through a set route to avoid damage to sensitive areas. In fact, the governing body of UNESCO continues to express concerns over the volume of tourists accessing the site. In 2019, new regulations were adopted limiting tourists to a maximum of 4 hours at the site, with three daily entry shifts scheduled between 6 am and 3 pm. Efforts to prevent the deterioration of heritage spaces and restrict movements have also been initiated at Ollantaytambo where tourists are directed up the Terraces of Pumatallis toward the Sun Temple and Wall of Six Monoliths before moving toward the site’s funerary complexes. Our results suggest that these conservation initiatives

¹As AP clustering provides few large clusters and many small clusters, we compute separation and compactness using the top 10% largest clusters.

Table 4: Clustering results on all sites. # clusters show the number of clusters discovered at each site by AP clustering. (Inter-cluster) separation shows how well separated a cluster is from other clusters. These together measure how different the clusters are; higher means more different. (Intra-cluster) compactness shows the average homogeneity of a cluster; lower means more homogeneous. UNESCO sites have the greatest amount of clusters (meaning more varied photos taken) and reflect how tourists spent longer periods of time at these sites. Large sites (Machu Picchu and Ollantaytambo) have the highest compactness scores, revealing the reproduction of popular photo themes, while smaller sites (Chincherro, Puca Pucara, and Qenqo) have lower compactness, likely due to unrestricted tourist movement there. See text for detailed discussion.

(a)						
site	Machu Picchu	Cuzco	Chincherro	Moray	Ollantaytambo	Pisac
# clusters \uparrow	1425	642	75	67	279	216
separation \uparrow	23.24	25.03	25.08	25.39	21.96	22.33
compactness \downarrow	20.67	23.47	26.03	21.80	20.46	22.42

(b)						
site	Puca Pucara	Qenqo	Sacsayhuaman	Tambomachay	Pikillacta	Tipón
# clusters \uparrow	48	59	257	47	36	56
separation \uparrow	21.26	23.05	23.68	20.92	21.93	22.39
compactness \downarrow	24.06	24.31	21.19	19.84	22.07	21.37

and entry regulations have increased photo homogeneity at larger sites as tourists are made to follow predefined routes through heritage attractions allowing for the reproduction of objects and scenes. At smaller sites like Chincherro, Puca Pucara, and Qenqo, photo homogeneity remained lower as tourist movement is unrestricted, allowing a wider variety of images.

Finally, our results hint at the influence of historical images in guiding how current tourists experience a heritage site. Notably, many of the representative images identified in the top 10% of clusters share aesthetic qualities with historic photographs published by Hiram Bingham in the early 20th Century (Fig. 8). The continued reproduction and distribution of these scenes on social media platforms creates a visual heritage narrative that influences the expectations of future travelers on the landscape. When a particular image or scene achieves dominance, alternative ways of experiencing landscapes are in-

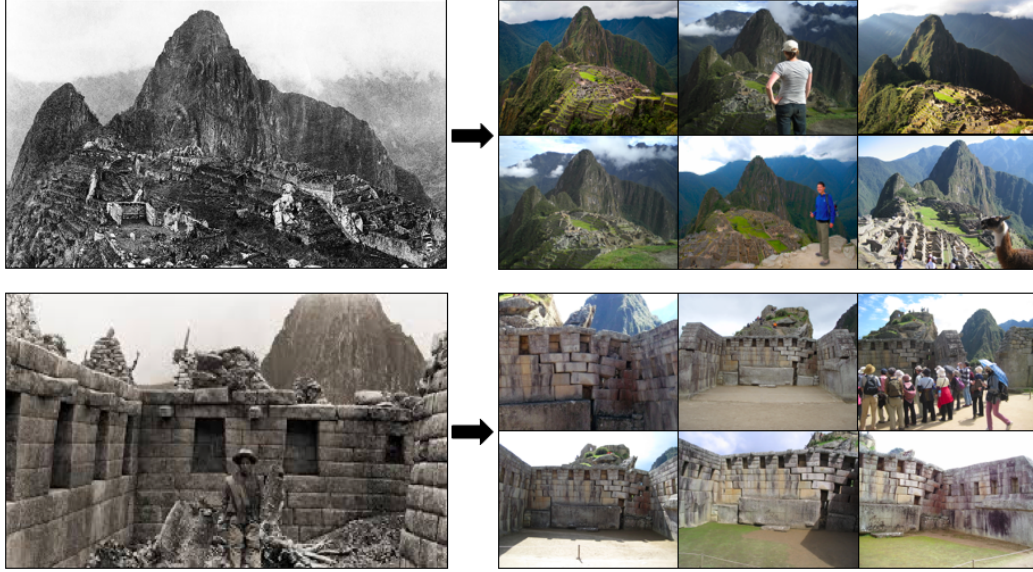


Figure 8: Left: Historic photographs taken by Hiram Bingham during the Yale Peruvian Expedition (1911-1915). Right: Flickr user images within the last 15 years (2004-2019). Images taken by Bingham show a strong visual influence on preferred scene-captures and objects photographed by tourists.

evitably obscured (Waterton 2009). Therefore, heritage discourses, like that generated by early explorers and perpetuated by Cuzco’s modern tourism industry, generate aesthetic legacies over time by reproducing that which is meant to be seen based on perceived experiences and values extracted from audience consumption (aesthetic, economic, cultural, etc. (Watson and Waterton 2010)).

6. Conclusion

Our study provides an innovative (and first of its kind) application of computer vision and machine learning algorithms to quantify the visibility of heritage landscapes and analyze the influence of heritage regulations on tourist circuits in the Cuzco region. By utilizing publicly available, geotagged source data from Flickr we are able to analyze tourist movement amongst BTC and UNESCO sites. Knowledge of travel patterns across heritage landscapes provides vital information not only for heritage conservation and the management of archaeological sites, but also for assessing the economic impacts that new regulations may have on local communities who depend on

the recurrent influx of visitors.

Through pattern recognition tools like feature extraction, image recognition, clustering, and t-sne embeddings, we are able to show a broad overview of the diversity and homogeneity of tourist generated images and identify how travelers are visually experiencing Cuzco’s heritage landscapes. These techniques also provide opportunities for qualitative and quantitative comparisons with historic imagery to address the reproduction of aesthetic legacies. The power of convolutional neural networks (CNNs) to extract image features and compare tens of thousands of photos within a short period, as well as the ability of affinity propagation (AP) to automatically cluster site images based on their visual characteristics discovers trends at a scale that would be extremely difficult and time consuming, if not impossible, through human labor alone.

The methods and techniques that we have adapted for our regional study of Cuzco’s heritage circuit can be generalized and scaled globally to study different heritage sites and time periods. The interdisciplinary nature of our approach and the unique insights we are able to generate with machine learning are significant for developing sustainable heritage practices as mass tourism increasingly threatens heritage sites worldwide. Future heritage research that could benefit from our ideas include: examining how tourist scene-captures have evolved over time, studying movement patterns and hotspots *within* heritage sites, and identifying outlier archaeological sites ‘on-the-rise’ through social media platforms and their potential economic influences.

Learning Patterns of Tourist Movement and Photography from Geotagged Photos at Archaeological Heritage Sites in Cuzco, Peru

References

- Bauer, B. S. (2004). Ancient cuzco: Heartland of the inca.
- Bevan, A. and Lake, M. (2016). *An Examination of Automated Archaeological Feature Recognition in Remotely Sensed Imagery*, chapter 3. Taylor & Francis.
- Bingham, H. (1913). *In the Wonderland of Peru: The Work Accomplished by the Peruvian Expedition of 1912, Under the Auspices of Yale University and the National Geographic Society*. National geographic magazine. Vol. 24 nr 4. April 1913. National Geographic Society.
- Bingham, H. (1922). *Inca land: explorations in the highlands of Peru*. Boston: Houghton Mifflin.
- Burns, Peter M., C. P. and Lester, J.-A. (2010). *Tourism and Visual Culture: Theories and Concepts*. Cambridge: CAB International.
- Casana, J. and Laugier, E. J. (2017). Satellite imagery-based monitoring of archaeological site damage in the syrian civil war. *PLoS ONE*, 12(11).
- Chen, C.-Y. and Grauman, K. (2011). Clues from the beaten path: Location estimation with bursty sequences of tourist photos. In *Computer Vision and Pattern Recognition*.
- Covey, R. A. (2017). Legacies of spanish colonialism in the inca heartland (cuzco, peru). *Brown J. World Aff.*, 24:185.

- Crandall, D., Backstrom, L., Huttenlocher, D., and Kleinberg, J. (2009). Mapping the world’s photos. In *International World Wide Web Conference*.
- Crandall, D. and Snavely, N. (2012). Modeling people and places with internet photo collections. *Queue*, 10(5).
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition*.
- Dircetur (2017). Cusco: Dirección regional de comercio exterior y turismo-cuzco. <http://www.dirceturcusco.gob.pe/>. Accessed: 2019-07-07.
- Donahue, J., Jia, Y., Vinyals, O., Hoffman, J., Zhang, N., Tzeng, E., and Darrell, T. (2014). Decaf: A deep convolutional activation feature for generic visual recognition. In *International conference on Machine Learning*.
- Frey, B. J. and Dueck, D. (2007). Clustering by passing messages between data points. *Science*, 315(5814):972–976.
- Gansell, Amy Rebecca, J.-W. v. d. M. S. Z. and Wiggins, C. H. (2014). Stylistic clusters and the syrian/south syrian tradition of first-millennium bce levantine ivory carving: a machine learning approach. *Journal of Archaeological Science*, 44:194–208.
- Garrod, B. (2009). Understanding the relationship between tourism destination imagery and tourist photography. *Journal of Travel Research*, 47(3):346–358.
- Gebru, T., Krause, J., Wang, Y., Chen, D., Deng, J., Aiden, E. L., and Fei-Fei, L. (2017). Using deep learning and google street view to estimate the demographic makeup of neighborhoods across the united states. *Proceedings of the National Academy of Sciences*, 114(50):13108–13113.
- Gebru, T., Krause, J., Wang, Y., Chen, D., Deng, J., and Fei-Fei, L. (2015). Visual census: Using cars to study people and society. *Bigvision*.
- Girshick, R., Donahue, J., Darrell, T., and Malik, J. (2014). Rich feature hierarchies for accurate object detection and semantic segmentation. In *Computer Vision and Pattern Recognition*.

- Géron, A. (2017). *Hands-on Machine Learning with Scikit-Learn & TensorFlow*. O'Reilly Media Inc., Sebastopol, CA.
- Hartley, R. and Zisserman, A. (2003). *Multiple view geometry in computer vision*. Cambridge University Press.
- Hays, J. and Efros, A. A. (2008). Im2gps: estimating geographic information from a single image. In *Computer Vision and Pattern Recognition*.
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Computer Vision and Pattern Recognition*.
- Hu, F., Li, Z., Yang, C., and Jiang, Y. (2019). A graph-based approach to detecting tourist movement patterns using social media data. *Cartography and Geographic Information Science*, 46(4):368–382.
- Hörr, Christian, E. L. and Brunnett, G. (2014). Machine learning based typology development in archaeology. *Journal on Computing and Cultural Heritage (JOCCH)*, 7(1):2.
- Kalogerakis, E., Vesselova, O., Hays, J., Efros, A. A., and Hertzmann, A. (2009). Image sequence geolocation with human travel priors. In *International Conference on Computer Vision*.
- Kaufman, Michael, S. P.-H. L. and Stettler, J. (2019). Analysis of tourism hotspot behaviour based on geolocated travel blog data: The case of qyer. *ISPRS International Journal of Geo-Information*, 8(11):493.
- Kemeny, J. and Laurie Snell, J. (1960). *Finite Markov chains*. New York Springer-Verlag, Princeton, NJ.
- Larson, L. R. and Poudyal, N. C. (2012). Developing sustainable tourism through adaptive resource management: a case study of machu picchu, peru. *Journal of Sustainable Tourism*, 20(7).
- Lee, J. Y. and Tsou, M.-H. (2018). Mapping spatiotemporal tourist behaviors and hotspots through location-based photo-sharing service (flickr) data. In *International Conf. on Location Based Services*.
- Maaten, L. v. d. and Hinton, G. (2008). Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605.

- Maxwell, K. (2012). Tourism, environment, and development on the inca trail. *Hispanic American Historical Review*, 92(1):143–171.
- Parcak, S. (2015). Archaeological looting in egypt: A geospatial view (case studies from saqqara, lisht, and el hibeh). *Near Eastern Archaeology*, 78(3):196–203.
- Poole, D. (1992). *Vision, Race, and Modernity A Visual Economy of the Andean Image World*. Princeton: Princeton University Press.
- Rice, M. (2018). *Making Machu Picchu: The Politics of Tourism in Twentieth-century Peru*. UNC Press Books.
- Rose, G. (2016). *Visual methodologies: An introduction to researching with visual materials*. London: Routledge.
- Sand, A. (2012). Visuality. *Studies in Iconography*, 33:89–95.
- Scorer, J. (2014). Andean self-fashioning: Martín chambi, photography and the ruins at machu picchu. *History of Photography*, 38(4):379–397.
- Sharif Razavian, A., Azizpour, H., Sullivan, J., and Carlsson, S. (2014). Cnn features off-the-shelf: an astounding baseline for recognition. In *Computer Vision and Pattern Recognition workshops*.
- Silverman, H. (2013). *Cuzcotopia: imagining and performing the Incas*, chapter 10. Taylor & Francis.
- Simon, I. and Seitz, S. M. (2008). Scene segmentation using the wisdom of crowds. In *European Conference of Computer Vision*.
- Snively, N., Seitz, S. M., and Szeliski, R. (2006). Photo tourism: exploring photo collections in 3d. In *ACM Transactions on Graphics*.
- Steel, G. (2008). *Vulnerable careers: Tourism and livelihood dynamics among street vendors in Cusco, Peru*. Rozenberg Publishers.
- Stepchenkov, S. and Zhan, F. (2013). Visual destination images of peru: Comparative content analysis of dmo and user-generated photography. *Tourism Management*, 36:590–601.

- Urry, J. and Jonas, L. (2011). *The Tourist Gaze 3.0*. London:SAGE Publications.
- Vu, Huy Quan, L.-G. L. R. and Ye, B. H. (2015). Exploring the travel behaviors of inbound tourists to hong kong using geotagged photos. *Tourism Management*, 46:222–232.
- Wang, J., Korayem, M., Blanco, S., and Crandall, D. J. (2016). Tracking natural events through social media and computer vision. In *ACM International Conference on Multimedia*.
- Wang, J., Korayem, M., and Crandall, D. (2013). Observing the natural world with flickr. In *International Conference on Computer Vision Workshops*.
- Waterton, E. (2009). Sights of sites: picturing heritage, power and exclusion. *Journal of Heritage Tourism*, 4(1).
- Watson, S. and Waterton, E. (2010). *Culture, Heritage and Representation: Perspectives on Visuality and the Past*. Surrey: Ashgate Publishing.
- Zhang, Kun, C. Y. and Li, C. (2019). Discovering the toursits’ behaviors and perceptions in a tourism destination by analyzing photos’ visual content with a deep learning model: The case of beijing. *Tourism Management*, 74:595–608.
- Zhang, H., Korayem, M., Crandall, D. J., and LeBuhn, G. (2012). Mining photo-sharing websites to study ecological phenomena. In *International World Wide Web Conference*.
- Zhang, W., Tan, G., Lei, M., Guo, X., and Sun, C. (2018). Detecting tourist attractions using geo-tagged photo clustering. *Chinese Sociological Dialogue*, 3(1):3–16.
- Zheng, Y.-T., Zha, Z.-J., and Chua, T.-S. (2012). Mining travel patterns from geotagged photos. *ACM Intelligent Systems and Technology*, 3(3):56.
- Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., and Torralba, A. (2017). Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.