Taylor & Francis
Taylor & Francis Group

Journal of Location Based Services

ISSN: (Print) (Online) Journal homepage: www.tandfonline.com/journals/tlbs20

Social media data for content creation in location-based games

Nancy Fazal & Pasi Fränti

To cite this article: Nancy Fazal & Pasi Fränti (30 Oct 2024): Social media data for content creation in location-based games, Journal of Location Based Services, DOI: 10.1080/17489725.2024.2414000

To link to this article: https://doi.org/10.1080/17489725.2024.2414000

JOURNAL OF LOCATION BASED SERVICES
https://doi.org/10.1080/17489725.2024.2414000

Social media data for content creation in location-based games

Nancy Fazal and Pasi Fränti

Machine Learning Group, School of Computing, University of Eastern Finland, Joensuu, Finland

ABSTRACT

Content in location-based games (LBGs) fundamentally depends on the location, and its global coverage is challenging. Typical content creation approaches involve crowd-sourcing, but recent research has also considered Web crawling and OpenStreetMap data as potential resources. A real-world location is a fundamental element in LBGs content. Media items such as images, short video clips, names, geocaches, and riddles can also be attached to the locations for games supporting sightseeing tours and treasure-hunts. In this paper, we evaluate the usability of social media data for content creation purpose in LBGs. We focus on three content items: *location*, *image*, and *name*. For this purpose, we retrieve geotagged images from services such as *Flickr*, *Foursquare*, *Yelp*, and *Google Places* from 2019 to 2022. Our experiments with six regions show that the Flickr images are representative of place with higher probability (32%) and location accuracy (96%), but relevant name extraction relied on the external method. Foursquare and Yelp images always have the correct location (100%), and the name is usually relevant (100%), but the image is rarely representative (22%, 4%). Google Places, despite providing large volumes of data and a variety of establishments, only 5% of representative images were found; however, location accuracy (75%) and relevant names (100%) are reasonable measures.

ARTICLE HISTORY
Received 3 January 2024
Accepted 2 October 2024

KEYWORDS
Social media data; location-based games; content creation; geotagged images; positional accuracy

1. Introduction

Location-based games (LBGs) are a special case of the wider class of *location-based services* (LBSs). The games can be played everywhere, where game information is coupled to a location. The real-world gameplay makes the experience fun for the players (Neustaedter, Tang, and Judge 2013). With the advancement of smartphone technologies, LBGs have been made available to a larger public. Some famous examples include *Pokémon Go*, *Pikmin Bloom*, *Jurassic World Alive*, *Ingress*, *Zombies Run*, *Draconius GO*, *The Walking Dread: Our World*, and *Minecraft Earth* (Laato et al. 2020).

CONTACT Nancy Fazal fazal@cs.uef.fi

Despite the massive success of Pokémon Go, they still face the challenge of content creation for their worldwide scalability. As a result, many such games limit themselves to specific regions only and do not try to reach an international audience. Different games have used different content creation approaches. Pokémon Go encourages the more experienced players to submit *place of interest* (POI) after they reach a certain level. Once submitted, the nominations are evaluated by Niantic's player community. Geocaching uses a similar approach by outsourcing content creation to its players. O-Mopsi allows the registered players to create new games, but quality control is monitored by the game administrators only (Fränti and Fazal 2023).

Fazal, Nguyen, and Fränti (2019) studied the efficiency of a web crawler for finding geotagged images to serve the content creation. The results were sparse; only 6845 retrieved images (<0.1%) contained geotags. They further studied the OpenStreetMap (OSM) data, where, for a given region, the available amenities with sufficient metadata were retrieved, and images were extracted from the given websites or Wikipedia links. This approach outperformed Web Crawling (Fazal, Nguyen, and Fränti 2019); however, OSM data was underrepresented in Asian countries, and good content was limited to urban areas only (Fazal, Mariescu-Istodor, and Fränti 2021).

In this paper, we evaluate the use of social media data for content creation and its challenges. To our knowledge, this is the first study on such data source in the context of LBGs. The research community has already found many motivations to work with the geo-referenced data (Jaffe et al. 2006). Memon et al. (2015) studied geotagged photos and historical weather data to recommend locations according to tourists' time and preferences. Cao et al. (2010) used geotagged images to suggest tourist destinations based on visual matching. A user can either enter a keyword describing the place of interest or provide a photo of the desired scenery, and the system checks into its database for places that match the visual characteristics as per the user's interest. Jaffe et al. (2006) propose a method to automatically select representative images as a summary set of photos for a given spatial region using an extensive collection of geotagged images. They further used visualisations from geotagged image datasets to generate *Tag Maps*.

Slava et al. Kisilevich et al. (2010) explored the attractive areas or significant events characterised by high photo activity in a specific area by collecting metadata of geotagged photos from Flickr using its publicly available API. Zheng et al. (2009) built a world-scale landmark recognition engine that recognises the landmarks on earth by using geotagged photos from photo-sharing websites Picasa and Panoramio and travel guide articles from wikitravel.com.

Lemieux (2015) proposed recording of criminological data using geotagged photos. The method is easy to use, not expensive, and readily available technology. Knowing that criminologists are often concerned with the location that hosts crime and contextual information about the crime, geotagged photos offer a unique way to collect such information. Huang (2016) provides a method to derive location recommendations using geotagged photos from Flickr that match a tourist's travel preferences and visiting context, such as time of day, weather, and season.

Nugraha and Damen (2013) and Guan and Chen (2014) investigate geotagged photos for post-disaster damage assessment. Sun et al. (2015) built a recommendation system using geotagged photos from Flickr to provide users with the most famous landmarks and the best travel routes between the landmarks. Research on geotagged photos has further focused on hotspot discovery (Yang, Gong, and U 2011), behaviour modelling (Zheng, Zha, and Chua 2012), place semantics extraction (Jaffe et al. 2006), community classification (H. Hu et al. 2015), the relationship between happiness and mobility patterns (Frank et al. 2013), transportation planning, targeted advertisements, and photo management applications (Crandall et al. 2009), personalised sightseeing tours by using content from Flickr, Wikipedia, and Google Maps (Brilhante et al. 2015), POI mining (Bui and Park 2017), Automatic region of interest detection (Belcastro et al. 2018) and artificial surface validation (Xing et al. 2017).

Social media data has been used for various LBSs, but the quality varies. One reason is the non-commercial contributors using different tools and technologies, various levels of precision, and most importantly, the lack of standardisation. Hence, the data is not error-free, and a quality assessment is therefore needed before using it (Senaratne et al. 2017). In gaming, the quality of images is crucial especially in treasure hunt type of games where players search the targets or perform some actions related to them. For example, Pokémon Go has a predefined set of guidelines like excluding people, body parts, live animals, and blurry photos (https://niantic.helpshift.com/hc/en/21-wayfarer/faq/2769-photo-guidelines/). To address the issue, we study the following three open questions to evaluate the usefulness of geotagged photos for content creation in LBGs:

(1) Is the image representative?
(2) Is the location accurate?
(3) How to extract relevant tag (a name that best describes the content)?

In treasure hunt games, a player is searching for the target and the associated image must therefore be representative so that the player can visually recognise the target she is looking for. Not every image taken at the place is suitable. Another difference to other location-based applications is that we need a precise point of interest. Larger regions may be problematic as game target as the player cannot know exactly where she is expected to go (Fränti and Fazal 2023). Many applications often focus more on the density of images in a certain area to identify the regions of interest. Technical issues like mismatch between the location of target and the location where picture is taken also matters.

The quality is not the only issue. Availability of data (quantity) matters even more. Games are needed everywhere, the players happen to be (demand), but data is mostly distributed in touristic areas where lots of people travel (supply). In this paper, we study the most relevant social media platforms to find out how well they cover (quantity) selected regions, and how useful is the data (quality) they provide.

Automatic extraction of points and regions of interest from large collection of photons has also been developed in (Cao et al. 2010) and (Kisilevich et al. 2010). However, we are more interested in the availability of the data (supply) but do not address how the selection should be automated. If there is no supply, there is nothing to be automated. We also do not aim at developing new quality assessment measures but study if the already existing measures (and how) can be applied for games. While our results are primarily meant for content creation in games, the results can generalise to other location-based applications as well, specifically supporting sight-seeing.

2. Content in location-based games

Content creation in LBGs is a huge challenge. Players can be anywhere, and content can be any real-world location, including monumental landmarks and smaller recognisable objects, such as benches, road signs, or post boxes. A playable location may require media items such as an image, a short video clip, and a name. Some games may also aim for players to perform actions at the locations, such as taking pictures of the target, solving a riddle, or logging findings.

TidyCity uses riddles with a task to determine the real-world locations that each riddle describes and verify it by visiting the location (Wetzel, Blum, and Oppermann 2012). *Outcatch* is inspired by an old Finnish board game called Great Star of Africa. Its content also involves hiding or seeking caches. *Codecrackers* involve picking up hints that leads a player to the secret location (https://dasbox.be/encyclopedia-of-location-based-games/). *Huntzz* involves real-world treasure (scavenger) hunts and tour guides, which may involve attractions, events, and charities (http://www.huntzz.com/). *Loquiz's* content involves quizzes, tours, and treasure hunts. Besides that, it offers an extensive question and mission database (https://loquiz.com/). *Munzee* involves finding hidden QR codes in the public space called 'munzees'. Almost all the world countries have about 4 million Munzees (https://www.munzee.com/).

O-Mopsi (http://cs.uef.fi/o-mopsi) (Fränti, Mariescu-Istodor, and Sengupta 2017) is an orienteering game played outdoors. Finding real-world targets using a smartphone occurs in a forest lacking houses, roads, and other visible landmarks. A typical target comprises three elements: name, location, and photo, and finding it is mentally rewarding, which constitutes the primary motivation to play the game (see Figure 1). Similarly, a PokéStop in *Pokémon Go* is an in-game real-world location that encourages the players to go outside, discover new places and connect with people. It has a name, description, and photo attached to it (https://pokemongo.fandom.com/wiki/Pok%C3%A9Stop). In contrast, the locations in *Randonautica* contain no media items, and players are expected to have their adventure (https://www.randonautica.com/). *Geocaching* is the world's largest treasure hunt game with a big community. The content involves millions of geocaches located around the world. They are waterproof containers and come in different shapes, sizes, and difficulties. Once found, a player can log their findings online (https://www.geocaching.com/play) (see Figure 1).

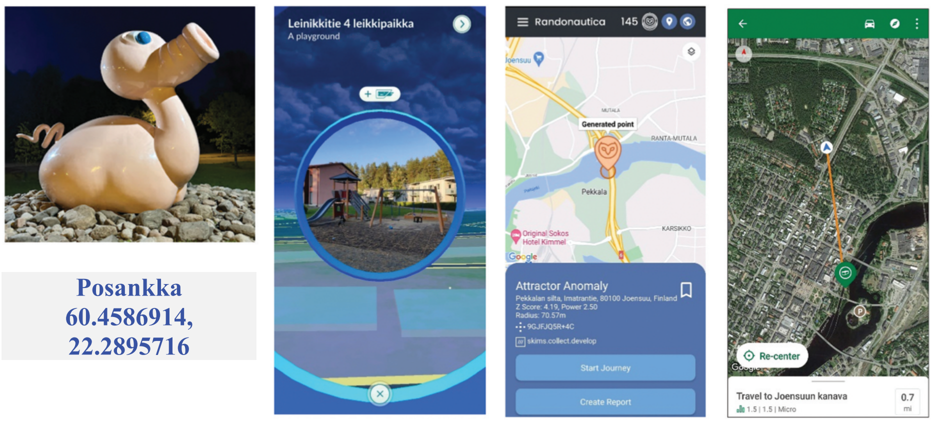

Figure 1. A target in O-Mopsi, Pokémon go, Randonautica and Geocaching.

The content in games like *Pokémon Go, O-Mopsi* and *Geocaching* are mainly crowdsourced by the players. One issue in crowdsourcing, however, is that the data quality can be affected by vandalism. However, community tends to detect and fix such problems quickly (Juhász et al. 2020).

3. Social media services

Lewis (2009) states that social media technologies are computer-mediated communication tools that connect people and allow them to share content. They have developed a trustworthy platform for information sharing, which leads to enormous amounts of data in the form of daily encounters, especially photos (Osatuyi 2013). The data is widespread to almost all the regions in world since it is a part of everyday life (Barbier and Liu 2011). Despite large volumes, most of the data is private. Reasons include licence terms and data crawlability, such as photos uploaded on Facebook and Instagram do not have an open access licence. Services like Getty Images, Shutterstock, and Bing Image Search API provide coarse location information, where precise location is the key for linking data with geographic applications (Zhu et al. 2022). We study four social media platforms: Flickr, Foursquare, Yelp, and Google Places, offering free data, official APIs (Application Programming Interface), and freedom to use for research purposes.

Flickr is an image and video hosting website created by a Vancouver-based organisation named *Ludicorp* in 2004 and currently owned by Yahoo! It allows its users to interact by sharing comments about photography and creating groups of specific interests, which makes it unique among other social networks. Each photo contains metadata added by the camera, such as camera model, camera settings, and date taken. Besides, photographers can add metadata to the images depicting their location, free text descriptions, or visual content.

The textual metadata associated with images not only represents the context of an image but also tells a lot about the social circle of a photographer. The images hosted on Flickr are available for both commercial and non-commercial purposes. It offers a powerful and completely free-to-use API, which is very generous regarding the total volumes of data provided (Spyrou and Mylonas 2016). The users can also manually annotate the image on a map if the location is not embedded in Exif. For every image missing location, Flickr offers an easy-to-use interface to add location data.

Foursquare (FSQ) is a location-based social media service that provides users with a personalised local search experience. As a geo-location application, it uses the global positioning system in mobile devices to generate locations and provide a list of places. 'Check-in' is the process of identifying the user's location. The platform is based on the idea that people can use mobile devices to interact with their environment, and it aims to provide recommendations of the best places around the user's current location (https://en.wikipedia.org/wiki/Foursquare_City_Guide). Foursquare provides Places API, which allows to bring location context into applications. The given endpoints further permit places to be searched in the FSQ database along with their details and photos (https://location.foursquare.com/developer/reference/search-data). It, however, does not allow the users to annotate location manually but encourages them to upload the photos that best describe the business.

Yelp was founded in 2004 and is used to publish crowdsourced reviews about businesses. It was founded by former PayPal employees Russel Simmons and Jeremy Stoppelman. As of December 2021, approximately 244.4 million reviews were available on Yelp. The company had 46 million unique visitors on its webpage and 56.7 unique visitors on its mobile sites as of 2021 (https://en.wikipedia.org/wiki/Yelp). The Yelp Fusion API provides rich business data in the form of names, addresses, photos, contact numbers, hours of operation, and price levels from millions of businesses worldwide (https://docs.developer.yelp.com/docs/fusion-intro). It does not provide an option to annotate the location manually and refrains

Table 1. Statistics of social media services.

Attributes	Flickr	Foursquare	Yelp	Places API
Photos	10 billion (2015)	-	-	-
Users	>112 million	Over 50 million active users	178 million visitors monthly	>1 billion
Growth rate	25 million photos per day	-	244 million reviews by 2021	-
Launched	2004	2009	2004	2005
Live browser	Yes	Yes	Yes	No
Access mechanism	Official API	Official API	Official API	Official API
API rate limit	3600 queries per hour	99,500 regular API calls	500 API call per 24 hours	100 requests per second
Response size	First 4000 results	2 photos per venue for non-commercial	-	up to 10 photo elements
License	Partly CC	Proprietary	-	Proprietary
Mapping Services	OpenStreetMap, Leaflet and Mapbox	Mapbox and OpenStreetMap	Google Maps	Google Maps

from posting shaky or out-of-focus photos. In the case of applying filters, Yelp encourages not to overdo it.

Google Places API is a service offered by Google that returns information about places using HTTP requests. *Places* are defined within this API as establishments, geographic locations, or prominent points of interest. It gives access to the millions of photos stored in the places database, while users cannot perform manual location annotation (https://developers.google.com/places/web-service/intro). See Table 1 for different attributes collected.

4. Data collection and preprocessing

The data collection process was performed for 2019–2022 within a radius of 3 km from a given location – assuming it to be the suitable gameplay area (Fränti and Fazal 2023). Data contains all the publicly available geotagged images and their associated textual tags/names. Flickr is very generous regarding the total volumes of data provided; thus, we collected only a sample of photos for which unique users were identified. Yelp, Foursquare, and Google Places APIs have a limited response, so we consider all the retrieved images. The location references used, and the total number of geotagged images found are shown in Table 2. The dataset is available for download in (https://cs.uef.fi/GeoSoMe/).

We selected six different locations for this study, including famous monuments and parks as follows:

Monuments: Helsinki Cathedral, Stonehenge, Leaning Tower of Pisa
Parks: Hyde Park, Mont des Arts Garden, Koli National Park

Table 2. Summary of the collected data.

Service	Helsinki cathedral	Koli national park	Leaning tower of pisa	Stonehenge	Hyde park	Mont des arts garden
	Location (Lat/Lon), radius = 3km					
	60.17050 24.95218	63.09730 29.80617	43.72297 10.39660	51.17902 -1.82620	51.50742 -0.16574	50.84543 4.35710
	Geotagged Photos retrieved					
Flickr	246	8	108	100	53	9
Foursquare	50	35	50	20	50	50
Yelp	50	1	50	5	50	50
Google Places	293	14	263	23	197	299

A brief introduction of the APIs and their endpoints used are as follows:

4.1. Flickr API

The Flickr API is comprehensive, well-documented, and available for non-commercial use. It provides API Explorer support, enabling API requests through the web browser. The requests and response formats include REST, XML-RPC, SOAP, JSON, and PHP. It returns, at most, the first 4,000 results for any search query. However, we must request the API key, a unique identifier, to authenticate an application or user. It also offers a wide range of API endpoints to collect data on photos, blogs, cameras, groups, and people. The images are stored with rich information, such as unique image IDs, owners, titles, tags, and locations. The location references in latitude/longitude format are manually added by the user on Flickr's interactive map or from external GPS, cameras, and phones with built-in GPS.

We use the *flickr.photos.search, flickr.photos.geo.getLocation* and *flickr.photos.getInfo* endpoints to fetch public geotagged photos with street-level accuracy and metadata. The current range Flickr offers is 1–16, where the world level is 1, country, region, city, and street levels are 3, 6, 11, and 16, respectively. We further limit our search query to outdoor images only. Flickr allow the users to add free text tags for the images and later standardises them by removing the space between words and converting them into lowercase. However, we rely on the raw tags (https://www.flickr.com/services/api/).

4.2. Foursquare API

The Foursquare API continuously updates and verifies points of interest with access to user-generated tips, tastes, and photos. It has over 100 million POIs in 247 countries and territories. To use the API, authentication via the API key is compulsory. Usage of Places API is, however, subject to the rate limits. It further provides live browser support to make API requests on the web browser. We first used the/*places/search* endpoint to retrieve all the places within a given area, and then, using the unique place ID, we retrieved the outdoor photos via/*places/{fsq_id}/photos* endpoint (https://location.foursquare.com/developer/reference/search-data). Foursquare, however, does not allow users to add free text descriptions to places or photos. We, therefore, rely on the given place name. The images retrieved do not have an individual location; instead, all the images for a given place have one general location assigned, which is the location of a place itself. For each place, ten photos are returned by default, and a maximum of 50 photos can be queried. Some places may not have any images.

4.3. Yelp fusion API

The Yelp Fusion API lets you get the best local content and reviews from millions of businesses. It also uses an API key to authenticate all the endpoints. It is subject to the daily rate limit and offers 5,000 API calls per 24 hours by default. We use the */businesses/search* to retrieve businesses and */businesses/{business_id_or_alias}* endpoint to retrieve images for each business ID (https://docs.developer.yelp.com/docs/fusion-intro). There could be more than one image without any individual location provided but a location of the business itself.

4.4. Google places API

Places API is a service that returns the formatted location data and imagery about establishments and prominent points of interest (POIs). Google Places API also requires an API key for authentication purposes. It accepts the request as a standard URL with specific endpoints such as */place* or */photo*. The response could be JSON or XML, as specified in the request. Resources available through the Places API contain Place search, Place details, Place photos, Place autocomplete, and Query autocomplete. It also offers Java, Python, Go, and Node.js client libraries to call this API. We use the */place/nearby search/* endpoint to retrieve places within the given area. The establishments of interest for this paper includes an *amusement_park, art_gallery, bakery, cafe, church, florist, university, zoo, supermarket, shopping_mall, stadium, school, restaurant, pharmacy, museum, theater, library*, and *city*.

Next, we use the Place Photo service, which is part of the Places API. This service gives access to the millions of photos stored in the place database. The photo references returned from the Places API were used to make Place photo requests as */place/photo?parameters* (https://developers.google.com/places/web-service/intro).

Table 3. Additional attributes for geotagged images.

Additional attributes	Description
Street View link	For given spatial reference, the street view link (http://www.google.com/maps?q=Your+Sign+Location+in+Street+View@"+latitude+," "+longitude+"&cbll="+latitude+," "+longitude+"&layer=c) was created to help verifying the location accuracy of the geotagged images.
Seasonal	To mark an object in the photo as seasonal or not to estimate the target's lifetime.
Scene type	To mark the scene type found in the photo, e.g. statue, building, sign, etc.
Comments	To record any additional comments by the authors.
Location source	For Flickr images, the location source was identified as if the location was recorded with GPS or manually annotated on a map.

4.5. Data preprocessing

The retrieved data primarily contains the photos and their metadata containing a rich amount of information such as location (latitude/longitude), views, user-contributed textual tags (for Flickr photos only), unique photo identifier, owner, date taken, and location. Foursquare, Yelp, and Google Places also provide additional information, such as the unique businesses/places IDs, place status (operational/closed/closed temporarily), and place categories. However, we must still record additional attributes for each photo, such as the *street view link*, *seasonal*, *scene type*, *analyst comments*, and *location source* (see Table 3).

We use *flickr.photos.getExif* and *flickr.photos.geo.getLocation* endpoints (https://www.flickr.com/services/api/) to examine if the device recorded the location or manually added by a user. The data is then imported into the PostgreSQL database.

5. Methodology

5.1. Evaluating the representative image

A *representative image* can be defined as one that best describes the essence of a given location. Whereas images containing people, animals, objects most likely not there, highly formatted, captured from far distance, paintings, portraits, and signs along the main road where pedestrian access might not be possible are termed as non-representative images (see Figure 2). In LBGs, primarily supporting sightseeing a location can be of a unique art, nature signs, public libraries, museums, local shops, parks, exercise equipment in public places, post boxes, landmarks, statues, buildings, benches, bird hide, fountains or other recognisable structures (see Figure 3). Indoor locations are acceptable too if they are publicly accessible. However, private residential properties, burial grounds, funeral homes, cemeteries, and graveyards should be avoided out of respect. Gravestones belonging to historical or other significant community figures can be used if accessible (https://niantic.helpshift.com/hc/en/21-wayfarer/faq/2775-content-guidelines/). Since social media users can link any image to a location, extracting a representative image for every location is

Figure 2. Examples of non-representative images.

Figure 3. Examples of representative images.

essential. In Figure 4, an image that best represents the Leaning Tower of Pisa is the last one on the right side, though all the images were tagged correctly.

Since, no image analysis and machine learning methods can answer if a given image is representative for a target location, we thus perform a manual inspection at this stage. Even widely popular games like Pokémon Go relies on the

Figure 4. Leaning Tower of Pisa images added by social media users.

human inspection for representative image selection. The process began by loading all the images on ArcGIS map viewer (https://www.arcgis.com/home/webmap/viewer.html), and for each image on the map, the attribute table contains the *URL* link, which, on click, opens a full-size image in a new browser window. The authors now make qualitative evaluation by carefully inspecting the image content and marks it as a representative (Yes) or not (No) in the attribute table. Other attributes, such as *seasonal*, *scene type*, and any additional comments, are also recorded at this stage.

5.2. Evaluating the location accuracy

For players to arrive at the desired location and to be able to achieve it, the given location accuracy is critical (Nicklas, Pfisterer, and Mitschang 2001). LBGs encouraging sightseeing tours, such as O-Mopsi and Pokémon Go have images attached to the given locations. Thus, a player must be at the approximate camera location from which the image was captured. The interaction radius can, however, be different from one game to another. In this section, we evaluate the location accuracy of representative images by comparing their published geographic location to the manually corrected approximate camera/photographer location based on the image content. The regions and their corresponding numbers of analysed representative images are listed in Table 2.

Hauff (2013). studied the location accuracy of Flickr images while using Wikipedia for ground truth locations of the venues. Locations were manually annotated inside, outside, or unknown to the venue. Results show that location accuracy highly depends on the location's popularity. Images/Videos taken at popular locations have a high accuracy of about 11–13 metres difference, whereas, for less popular venues, a difference of approximately 47–167 metres was reported on average.

Cvetojevic, Juhasz, and Hochmair (2016) compared the annotated location information of Twitter and Instagram images with the photographer's position and the object being photographed. The photographer's position was estimated by 47 students who used their local knowledge and image content. Later, the authors verified those new estimated locations based on the satellite image view and Google Street view. The distances between objects and photographers follow a power law function with an exponent value of 1.31 and an R-squared of 0.89. Y. Li, Snavely, and Huttenlocher (2010) applied a prioritised feature-matching algorithm for location recognition tasks. Benford et al. (2003) report that network traffic logs from Rotterdam points and some locations, i.e., the narrow built-up streets at the centre, need better location accuracy, connectivity, or both. Thus, with careful scouting, game designers should focus on good coverage areas.

Hochmair and Zielstra (2012) and Zielstra and Hochmair (2013) measured the location accuracy of Flickr and Panoramio geotagged images by measuring the distance between the published image and the estimated camera position based on the image content. The geographic coordinates of the images were visualised as point features in ArcMap. Based on the image content, the analyst estimated and marked the camera position on the ArcMap aerial image background layer as a new point feature. The analysis revealed a better location accuracy for Panoramio images than the Flickr datasets.

Park et al. (2010) studied that the GPS data only identifies the camera location, but the viewing direction remains unknown. To produce more precise location information for the viewing direction of geotagged photos, they used Google Street View and Google Earth satellite images. Their method has two steps: 1) visual matching between a user photo and available street views to find the viewing direction, and 2) When only overhead satellite view is available, near orthogonal view matching between the photo and satellite imagery computes the viewing direction. Their experiments on a dataset of 55 images showed an average mean error of 11.1°.

Hochmair, Juhász, and Cvetojevic (2018) relied on the name tag and the author's local knowledge of the Salzburg downtown area to measure the distance offsets between mapped POIs and their actual location. Their results revealed that Google and OSM POIs do not have any positional errors, closely followed by Yelp, which had minor offsets in 12.7% of the cases. Foursquare had the next smallest error rate (32.5%), followed by Instagram (40.6%) and Facebook (43.6%).

We use a similar approach to Zielstra et al. (L. S. Kennedy and Naaman 2008; Moxley et al. 2009) and add images to the ArcMap as point features based on their published geographic coordinates. Each point feature contained an attribute with the URL and Google Street View Link. This allows the image content to be viewed in a web browser and the location to be verified with the help of Google Street View Link when available. The authors first test if the published location seems to be an approximate location of a photographer/camera by comparing the image content with Google Street View. If so, no other action is required, and the published location is accepted (see Figure 5).

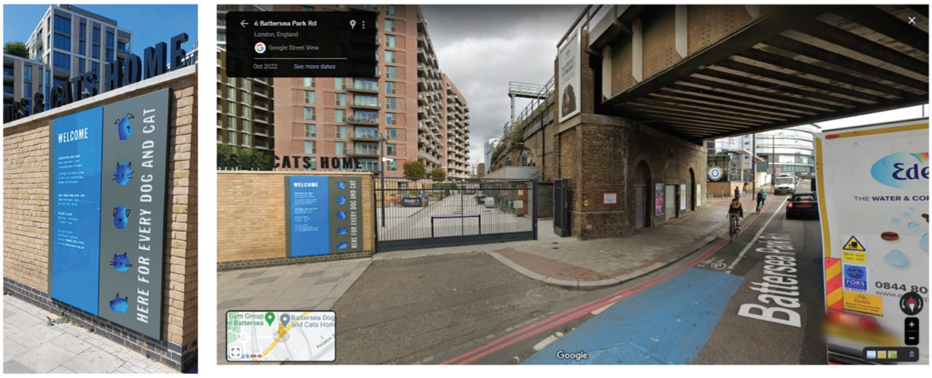

Figure 5. Published location verified using Google street view. Published location: 51.478091, -0.145306; Google Street View Link:http://www.google.com/maps?q=Your+Sign+Location+in+Street+View@51.478091,0.145306&cbll=51.478091,-0.145306&layer=c.

The authors used an imagery to estimate the photographer's location, where published location was not verified using Google Street View because either it was missing or did not correspond to the image content. The new location of the feature point is retrieved, and the Geodesic distance between the published and estimated location is computed. The following Arcade expression (see Figure 6) has been used to calculate the distance and to update the attribute table (https://developers.arcgis.com/arcade/).

```
var geom = Geometry($feature)
var fixed_latitude = geom.y
var fixed_longitude = geom.x
var orig_pt = Point ({ x:  Round($feature.longitude,6), y:
Round($feature.latitude,6), spatial reference: { kid: 4326 } });
var fix_distance = distanceGeodetic(Geometry($feature), orig_pt, 'kilometers')
$feature.latitude
return {
"result": {
    "attributes": {
      "fixed_latitude" : fixed_latitude,
      "fixed_longitude" : fixed_longitude,
      "distance": fix_distance
    }
  }
}
```

Figure 6. Arcade expression.

5.3. Evaluating the relevant tag

A relevant tag can be described as a textual name that captures the essence of a target, i.e., both the given location and the associated image. Fränti and Fazal (2023) emphasise that the name of a target should be short, preferably just one or two words. Statues and historical sites often have long names, such as *Terrace of the Lions* and *Buddha of Bamiyan*. Names of the commercial sites (e.g. *Starbucks*, *Mokkamaa*, *Houkutus*) are short and catchy, which is generally acceptable as it is.

Most social media services allow users to add unstructured textual tags to media objects, such as images, photos, and videos. They differ in semantic meaning and can explain different aspects of a media object, such as locations, dates, camera information, and people. This process is usually referred to as *tagging*. However, harvesting knowledge from these tags is challenging because of their free-form nature, noise, and semantic uncertainties (Moxley et al. 2009). A better understanding of the tag's semantics can help many applications such as automatic extraction of visual examples events/landmarks (L. S. Kennedy and Naaman 2008; Quack, Leibe, and Van Gool 2008), keyword search for images, tag-driven image annotations (Moxley, Kleban, and Manjunath 2008), evolving understanding of the world (Shirky 2005) and to generate the geographic and temporal labels (L. Kennedy et al. 2007).

In this section, we focus on selecting a relevant tag from a list of user-provided tags for Flickr images. Flickr standardises the tags by removing the space between words and converting the letters into lowercase. For example, a user tag 'My Helsinki' would become 'myhelsinki' (Y. Hu et al. 2015). We, however, rely on the raw tags for relevant tag extraction, which would later serve as the name of a target (see Figure 7). Among the selected services, only Flickr images had textual tags. Thus, we deal separately with Foursquare, Yelp, and Google Places. Note that Flickr has also unstructured textual descriptions, which could provide further insight for generating the text description of the POI. In Pokémon Go, targets usually have both the name and additional description.

Photo	Photo tags	Relevant Tag
	Belgium, Brussel, Bruxelles, Brussels, street stuff, street art, graffiti, urban art, canon, ptr, Peter Heuts	graffiti

Figure 7. Relevant Tag selection from given user-added tags.

Overell et al. (2009) used a generic method for classifying Flickr tags into semantic categories using third-party resources such as Wikipedia and Open Directory. Compared to Wordnet, their method improved the Flickr vocabulary coverage by 115%. Rattenbury and Naaman 2009; Rattenbury, Good, and Naaman (2007) assigned the place and event semantics to the Flickr tags using location and time metadata associated with photos. They further establish an entropy-based technique to automatically identify place and timed event tags. Moxley et al. (2009) classified 1.69 million Flickr images to find geographical areas with sufficient density to extract place and landmark semantics.

X. Li, Snoek, and Worring (2008) proposed a novel voting algorithm to compute tags' relevance concerning an image from the tagging behaviours of visual neighbours of that image. Their experiments on one million Flickr images verified their proposed algorithm. Compared with the baseline using original tags, retrieval using improved tags increases mean average precision by 24%. Begelman, Keller, and Smadja (2006) used clustering techniques to find semantically related tags. Sigurbjörnsson and Van Zwol (2008) introduced a novel and generic method for recommending tags for locations, artefacts, and objects. The method combines tag co-occurrence with tag aggregation strategies and promotion functions. The evaluation results with 200 Flickr photos showed that both tag aggregation strategies are effective, and it is crucial to consider the co-occurrence values of candidate tags. The promotion function is an effective way to incorporate the ranking of tags. Kennedy et al. (2007) used a tag-driven method to extract place and event semantics from Flickr tags based on each tag's metadata patterns.

Fazal and Fränti (2024) proposed a Tag-tag method for relevant tag extraction, which exploits the semantic relationship between objects detected on an image and its associated user tags. Their method is independent of visually similar images and is not restricted by tag groups. Further, they relied on state-of-the-art pre-trained machine learning models for object detection and pre-trained vectors on the part of the Google News dataset (word2vec-google-news-300) of about 100 billion words. Their experiments on Flickr images demonstrate the efficiency of the proposed method and report two possible reasons why the relevant tag was not correctly identified: 1) the Relevant tag itself was missing in the user-contributed list of tags, 2) the Relevant tag got skipped because it was either not present in Wordnet dictionary or had no pre-trained vector in word2vec-google-news-300 dataset.

6. Experiments and results

6.1. Representative image selection

The observations recorded by authors reveals that the Flickr provides most significant percentage of representative images, followed by Foursquare. As a business platform, Yelp had the smallest representative images found since most were of people, food, and indoor interiors. Google Places, on the other

Table 4. Representative images found (representative images/total images).

Region	Social media services: Flickr	Foursquare	Yelp	Google maps
Koli	1/8	6/35	0/1	1/14
Helsinki Cathedral	27/246	17/50	2/50	24/293
Leaning Tower of Pisa	52/108	14/50	2/50	16/263
Stonehenge	41/100	7/20	0/5	3/23
Mont des Arts Garden	7/9	4/50	2/50	13/299
Hyde Park	42/53	8/50	2/50	2/197
Total	170/524	56/255	8/206	59/1089
Percentage	32%	22%	4%	5%

hand, despite providing a large number of images and having a variety of establishments, only 5% of them were found to be representative (see Table 4).

6.2. Estimation of location accuracy

A total of 170 representative images were inspected from Flickr. We initially identify the source of location information for images, such as if the location was coming from the device or if the user had manually annotated. For this purpose, we used *flickr.photos.getExif* endpoint to retrieve the image's Exif data and compare the location values with *flickr.photos.geo.getLocation* endpoint. If they were identical, then Flickr probably got the location from Exif. However, the user could still have manually set the Exif data before uploading the photo. In this case, the authors double-check the location with the Google Street View link. This approach worked, and for a total of 138 images, the location was concluded to be taken from the device. For 25 images, authors manually tuned their location at approximate camera position based on the image content (see Figure 8). Images whose location could not be verified were excluded.

For Foursquare, 56 representative images were inspected, of which 28 were tagged correctly at the camera location, and the rest were fixed by authors. A total of 8 Yelp images inspected were found to be correctly tagged at the approximate camera location. For Google Places, 59 images were inspected, of which 14 place records were found to be duplicates and thus excluded. For the rest of the images, the location was correctly annotated at the camera location. Hence, no further action was required (see Table 5). Our results for Google

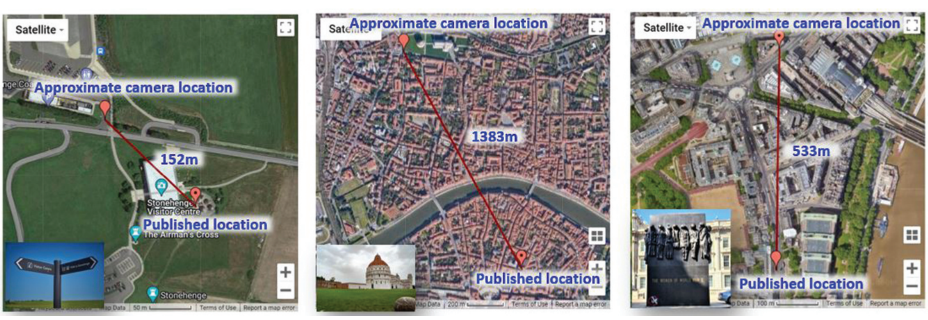

Figure 8. Locations manually fixed at the approximate camera location.

Table 5. Location accuracy estimation results.

	Flickr	Foursquare	Yelp	Google places
Images inspected	170	56	8	59
Camera location found	82%	50%	100%	76%
Location fixed	14%	50%	-	-
Location unverifiable/duplicate images	4%	-	-	24%

Table 6. Erroneous scene types of flickr and foursquare.

		Scene types						
Source	Distance error (m)	Landmark	Statue	Building	Shop	Road sign	Telephone box	Windmill
Flickr	MIN	72	51	37	–	153	–	38
	MAX	1383	533	164	–	–	–	–
Foursquare	MIN	–	41	55	29	–	76	–
	MAX	–	881	97	67	–	–	–

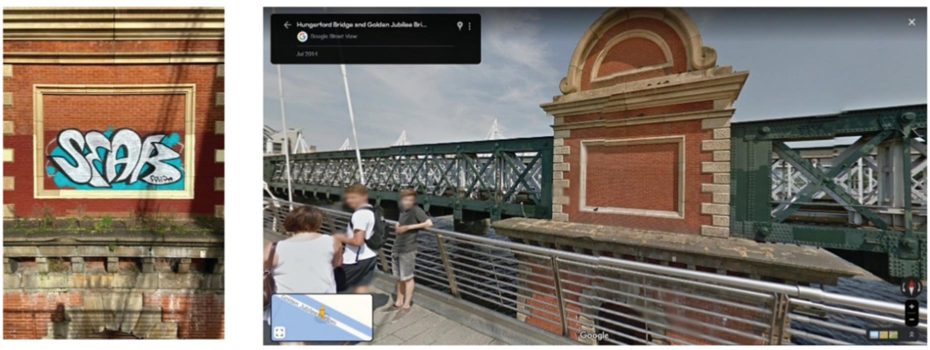

Figure 9. Street view needs to be updated to conclude whether graffiti remains. The latest street view is from 2014, whereas the graffiti image was uploaded in 2021.

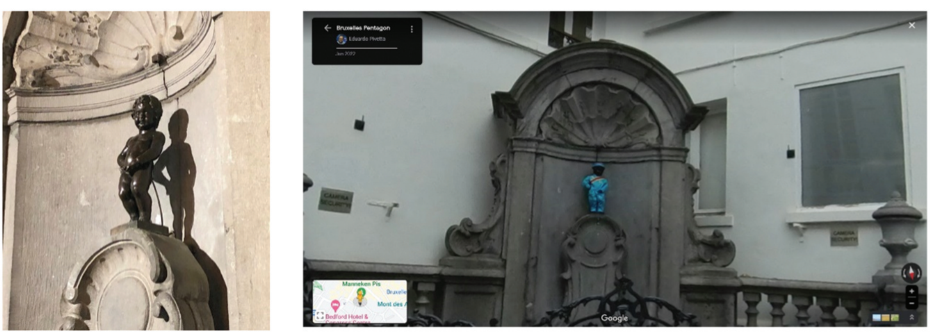

Figure 10. A statue subject to visual changes.

Places and Yelp, having no location accuracy errors, completely align with the results reported by Hochmair et al. (L. S. Kennedy and Naaman 2008).

Table 6 provides the observed distance errors for different scene types in Flickr and Foursquare images. We define the distance error as displacement when a user places the image at the location of an object being photographed instead of the location of a photographer/camera. For scene type *Landmark*s, it was observed that in the case of monumental landmarks, the manual tagging behaviour of a user is more inclined towards tagging the location at the landmark itself instead of the camera location. For *Statues*, a similar pattern is observed. The statue of *Achilles* was randomly tagged inside Hyde Park just because it exists there. In the case of large buildings, the location was sometimes tagged at the far back side of the building from where the correct view given in the image could not be seen. For shops, road signs, telephone boxes, and windmills, the uploader randomly tagged the location at the exact road instead of setting it up correctly to the actual camera location.

The images of art and small aesthetic objects could also be nice representative images, but measuring the location accuracy without visiting a published location is not trivial. We exclude such images. For example, in Figure 9, the image of graffiti was marked as representative, but the latest street view available was from the year 2014, which made it difficult to decide if it is still present or removed. Graffiti has a short life span and is subject to change (Fränti and Fazal 2023). Therefore, we include only those for whom Street View was available for the years we collected the images. This is one limitation of our approach; however, for recognisable and long-life structures, it does not hold.

Similarly, a small piece of art randomly tagged inside Hyde Park could not be verified for location accuracy using Street view or imagery. Hence, such cases were excluded, too. An image of a statue in Brussels was chosen as a representative image. It, however, is subject to visual changes and gets dressed up regularly. It is the only statue in the world with a wardrobe comprising 1094 outfits. We thus exclude this image (see Figure 10).

6.3. Relevant tag extraction

We opt for the same method as Fazal and Fränti (2024) for extracting the relevant tags for Flickr images. From our experiments on 170 Flickr representative images, 61 were discarded because of the missing tags. For the remaining 109 images, Wu-Palmer similarity and Cosine similarity measures correctly identified representative tags for 48% and 34% of the images, respectively (see Table 7).

Figure 11 demonstrates some example images for which Tag-tag correctly identified relevant tags. However, from the sightseeing tour perspective, a tag , Corgi, would have been more appealing than just a 'dog'. The ground truth tags for Figures 11 and 13 are underlined. Figure 12 has some good representative images for which tag extraction got restricted because a user who uploaded an

Table 7. Representative tag extraction using the tag-tag method (Fazal and Fränti 2024).

Monuments	Images inspected	Average number of tags	Images with missing tags	Wu-Palmer correct predictions	Word2Vec correct predictions
Helsinki Cathedral	27	7	14	38%	23%
Stonehenge	41	6	11	40%	3%
Leaning Tower of Pisa	51	6	22	48%	17%
Koli	2	9	0	100%	100%
Hyde Park	42	7	14	46%	32%
Mont des Arts Garden	7	6	0	14%	28%




Belgium, art, canon, Brussel, Brussels, street stuff, graffiti, urban ptr, Peter Heuts, Bruxelles	Westminster, England, London, CLOG, Central London Outdoor Group, sculpture, corgi, dog, Victoria, UK	Leake street, art, street, Iphone, Leake Street tunnel	Women, WW2, UK, Monument, London, England
Wordnet/Word2Vec			
graffiti/ graffiti	dog/dog	art/art	monument/monument

Figure 11. Relevant tag correctly identified.



Stonehenge, Photographie, wbayer.com	England, London, United Kingdom UK, museum, Great Britain, Vi&A, Victoria and Albert, arts culture, Europe, West London, Structures & Architecture	Хельсинки, Finland, helsinki	europe, nicolas, mk, travel, trip
Ground truth			
Graffitied stone	Sculpture	Pohjola	Stonehenge
Wordnet/Word2Vec			
Stonehenge/Stonehenge	museum/museum	Helsinki/Finland	europe/trip

Figure 12. The relevant tag itself was missing.

Angleterre, Camden, Camden Town, England, London, London Borough Of Camden, Royaume-Uni, The United Kingdom, UK, Art De Rue, Photo De Rue, Photography, Street, Street Art, Street Photography, Divers, Focal Length - 88 Mm, Focal Length In 35mm Format - 88 Mm, High Iso, Ilce-7c, Iso 2500, Sony, Sony Ilce-7c, Sony Ilce-7c E 28-200mm F2.8-5.6 A071, Created By Dxo, Dxo, Dxo Photolab, Dxo Photolab 6, Edited Photo	Stonehenge, Wiltshire, Neolithic Village	Piazza dei Miracoli, Italy, Italien, Pisa, Stadt, City, Tuscany, Toskana, Outdoor, ThroughHerLens, Gebäude, Building, Architektur, Architecture, Domplatz, Piazza del Duomo, Baptisterium, Baptistery, Leaning Tower of Pisa, Schieferturm, Turm, Tower	London, Camden, Regent's Park, London Parks, Frieze Sculpture, Frieze Sculpture 2022, Public Sculpture, Tim Etchells
Wordnet/Word2Vec			
Street/Photography	Stonehenge/Stonehenge	Building/Tower	Camden/Camden

Figure 13. The relevant tag went missing; it either needed to be present in the WordNet dictionary or had no pre-trained vector in the word2vec-google-news-300 dataset.

image did not bother to add a tag defining the image content, e.g., an image of the historical landmark 'Stonehenge' is missing from the list. The images shown in Figure 13 have representative tags in their user-provided list. However, the opted method could not identify them because they were missing in the Wordnet dictionary or the word2vec-google-news-300 dataset as per Tag-tag limitations.

Next, we carefully inspected the representative images from Foursquare, Yelp, and Google Places for which textual tags were unavailable. Therefore, we relied on the given names. We manually inspected and found them all (100%) relevant, descriptive, and capturing the essence of a place, thus accepted as it is (see Figure 14).

7. Discussion

The use of social media data for content-creation purpose in location-based games supporting treasure hunt/sightseeing tours is evaluated for the first time. Three essential game elements are studied, i.e., representative image selection, location accuracy verification, and relevant tag/name extraction. Experimental results showed that Flickr had the most representative images, followed by Foursquare. In contrast, Yelp and Google Places had a minimal

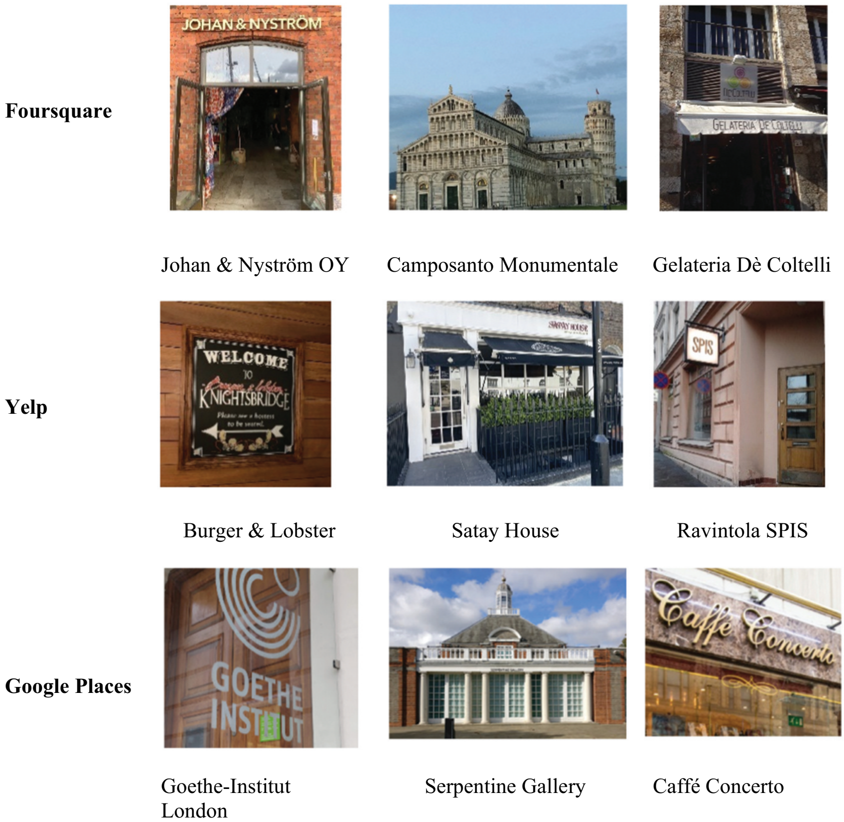

Figure 14. Example images from Foursquare, Yelp, and Google places with their names accepted as it is.

number because Yelp, a business platform, mainly offered indoor images. Google Places, despite having various establishments, provided only 5% representative images.

For the location accuracy aspect, Foursquare and Yelp achieved 100% of the correct camera location for inspected images. In contrast, Flickr and Google Places achieved 96% and 75% accuracy, respectively. The possible reasons why the camera location could not be verified include missing or too old Street View and required in-person visit. In rare cases where a target was subject to the visual changes, e.g., a statue of Manneken Pis, was excluded from the experiments. The distance offset in various scene types in Flickr and Foursquare images showed that the vast Landmarks and statues are subject to significant displacement, followed by the buildings, road signs, telephone boxes and shops, etc.

For relevant tag/name selection, the given names for POIs were accepted as they are for all the services, excluding Flickr. Since the user-provided tags in Flickr were unstructured and noisy, there was a need to select one relevant tag. We applied Tag-tag (Fazal and Fränti 2024) to make such

Table 8. Overall representative images, tags, and location accuracy.

	Flickr	Foursquare	Yelp	Google places
Representative Images	32%	22%	4%	5%
Location Accuracy	96%	100%	100%	75%
Relevant name	41%	100%	100%	100%

a selection, which correctly identified representative tags for 41% of the Flickr images (see Table 8).

8. Conclusion

Content creation in location-based games (LBG) is challenging, as content inherently depends on location. Even players of commercially successful games like Pokémon Go and Ingress often complained about the lack of content worldwide. As a result, many games limit themselves to a specific region and do not attempt to appeal to a worldwide audience. Approaches like Web Crawling, OpenStreetMap (OSM), Crowdsourcing, Wikipedia, and relying on the game administrators have been explored in the past. However, localising content to every possible game area remains a concern.

Social media data has been profoundly used in recent years in different applications. However, the quality of data varies and needs thorough assessment before use. This paper exploits social media data for content creation in treasure-hunt games and sightseeing. A target comprises three entities: an image, location, and name. We selected four social media services: Flickr, Foursquare, Yelp, and Google Places. A dataset was collected from 2019 to 2022 comprised geotagged images with their textual tags or names. The region of interest includes three famous landmarks and three Parks named *Helsinki Cathedral*, *Leaning Tower of Pisa*, *Stonehenge*, *Koli National Park*, *Hyde Park*, and *Mont des Arts Garden*, respectively.

The geotagged images retrieved were evaluated for their representativeness, location accuracy, and the relevance of the tag/name when available. Experimental results showed that Flickr provided the most representative images (32%), whereas Foursquare had significantly less (22%). Yelp, being a business platform, has most of the indoor images, and Google Places has the fewest representative images despite having a variety of establishments.

All Foursquare and Yelp images had the correct camera location, whereas Flickr and Google Places achieved 96% and 75% of the images, respectively. Besides Flickr, we manually inspected the relevance of a target's name and accepted them since they were all concise and captured the essence of a target. Flickr images, however, had only a list of user-provided tags, which were noisy and unstructured. We thus used an external relevant tag extraction method called Tag-tag, which worked for 41% of the images.

Social media data has the potential to be used, but none is perfect. Flickr is the most promising regarding the large number of representative images available. However, it requires an additional tag extraction component, and about 4% of the image's location could not be verified. All sources require image content analysis to conclude whether the image represents a location. The percentage of valuable images in this regard was relatively modest: Flickr 32%, Foursquare 22%, Google Places 5%, and Yelp 4%. Besides the deficits mentioned earlier, copyright issues should also be considered, which may limit the unrestricted use of a material.

Our main conclusion is that even if the richest social media platforms may have data, their use has many challenges. Use of the data would require automated solutions for detecting location accuracy, selection of representative image, and selection of representative tag. For example, even the best image analysis and machine learning methods cannot answer whether a given image is representative for a target location in gaming context. These are challenges that should be studied further.

Disclosure statement

No potential conflict of interest was reported by the author(s).

ORCID

Nancy Fazal http://orcid.org/0000-0002-7099-1327
Pasi Fränti http://orcid.org/0000-0002-9554-2827

References

Barbier, G., and H. Liu. 2011. "Data Mining in Social Media." *Social Network Data Analytics*: 327–352. https://doi.org/10.1007/978-1-4419-8462-3_12.

Begelman, G., P. Keller, and F. Smadja. 2006. "Automated Tag Clustering: Improving Search and Exploration in the Tag Space." In C*ollaborative web tagging workshop at WWW2006,* Edinburgh, Scotland, 15–33.

Belcastro, L., F. Marozzo, D. Talia, and P. Trunfio. 2018. "G-RoI: Automatic Region-Of-Interest Detection Driven by Geotagged Social Media Data." *ACM Transactions on Knowledge Discovery from Data (TKDD)* 12 (3): 1–22. https://doi.org/10.1145/3154411.

Benford, S., R. Anastasi, M. Flintham, C. Greenhalgh, N. Tandavanitj, M. Adams, and J. Row-Farr. 2003. "Coping with Uncertainty in a Location-Based Game." *IEEE Pervasive Computing* 2 (3): 34–41. https://doi.org/10.1109/MPRV.2003.1228525.

Brilhante, I. R., J. A. Macedo, F. M. Nardini, R. Perego, and C. Renso. 2015. "On Planning Sightseeing Tours with TripBuilder." *Information Processing & Management* 51 (2): 1–15. https://doi.org/10.1016/j.ipm.2014.10.003.

Bui, T. H., and S. B. Park. 2017. "Point of Interest Mining with Proper Semantic Annotation." *Multimedia Tools & Applications* 76 (22): 23435–23457. https://doi.org/10.1007/s11042-016-4114-7.

Cao, L., J. Luo, A. Gallagher, X. Jin, J. Han, and T. S. Huang. 2010. "A Worldwide Tourism Recommendation System Based on Geotagged Web Photos." *2010 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2274–2277. IEEE. March.

Crandall, D. J., L. Backstrom, D. Huttenlocher, and J. Kleinberg. 2009. "Mapping the world's Photos." *Proceedings of the 18th international conference on World wide web*, 761–770. April.

Cvetojevic, S., L. Juhasz, and H. Hochmair. 2016. "Positional Accuracy of Twitter and Instagram Images in Urban Environments." *GI_Forum* 4 (1): 191–203. https://doi.org/10.1553/giscience2016_01_s191.

Fazal, N., and P. Fränti. 2024. "Relevant Tag Extraction Based on Image Visual Content." In *Communications in Computer and Information Science*, 283–295. Vol. 2025. Singapore: Springer. https://doi.org/10.1007/978-981-97-0827-7_25.

Fazal, N., R. Mariescu-Istodor, and P. Fränti. 2021. "Using OpenStreetmap for Content Creation in Location-Based Games." *2021 29th Conference of Open Innovations Association (FRUCT)*, Tampere, Finland, 109–117. IEEE.

Fazal, N., K. Q. Nguyen, and P. Fränti. 2019. "Efficiency of Web Crawling for Geotagged Image Retrieval." *Webology* 16 (1): 16–39. https://doi.org/10.14704/WEB/V16I1/a177.

Frank, M. R., L. Mitchell, P. S. Dodds, and C. M. Danforth. 2013. "Happiness and the Patterns of Life: A Study of Geolocated Tweets." *Scientific Reports* 3 (1): 2625. https://doi.org/10.1038/srep02625.

Fränti, P., and N. Fazal. 2023. "Design Principles for Content Creation in Location-Based Games." *ACM Transactions on Multimedia Computing, Communications and Applications* 19 (5s): 1–30. https://doi.org/10.1145/3583689.

Fränti, P., R. Mariescu-Istodor, and L. Sengupta. 2017. "O-Mopsi: Mobile Orienteering Game for Sightseeing, Exercising, and Education." *ACM Transactions on Multimedia Computing, Communications and Applications* 13 (4): 1–25. https://doi.org/10.1145/3115935.

Guan, X., and C. Chen. 2014. "Using Social Media Data to Understand and Assess Disasters." *Natural Hazards* 74 (2): 837–850. https://doi.org/10.1007/s11069-014-1217-1.

Hauff, C. 2013. "A Study on the Accuracy of Flickr's Geotag Data." *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval*, 1037–1040. July.

Hochmair, H. H., L. Juhász, and S. Cvetojevic. 2018. "Data Quality of Points of Interest in Selected Mapping and Social Media Platforms." In *Progress in Location Based Services 2018 14*, 293–313. Springer International Publishing.

Hochmair, H. H., and D. Zielstra. 2012. "Positional Accuracy of Flickr and Panoramio Images in Europe." In P*roceedings of the geoinformatics forum,* Salzburg, Austria, 3–6.

http://www.huntzz.com/.

https://loquiz.com/.

https://www.munzee.com/.

https://pokemongo.fandom.com/wiki/Pok%C3%A9Stop.

https://www.randonautica.com/.

https://www.geocaching.com/play.

https://en.wikipedia.org/wiki/Foursquare_City_Guide.

https://location.foursquare.com/developer/reference/search-data.

https://en.wikipedia.org/wiki/Yelp.

https://docs.developer.yelp.com/docs/fusion-intro.

https://developers.google.com/places/web-service/intro.

https://www.flickr.com/services/api/.

https://niantic.helpshift.com/hc/en/21-wayfarer/faq/2775-content-guidelines/.

https://developers.arcgis.com/arcade/.

https://niantic.helpshift.com/hc/en/21-wayfarer/faq/2769-photo-guidelines/.
https://cs.uef.fi/GeoSoMe/.

Hu, H., Y. Wen, T. S. Chua, J. Huang, W. Zhu, and X. Li. 2015. "Joint Content Replication and Request Routing for Social Video Distribution Over Cloud CDN: A Community Clustering Method." *IEEE Transactions on Circuits and Systems for Video Technology* 26 (7): 1320–1333. https://doi.org/10.1109/TCSVT.2015.2455712.

Hu, Y., S. Gao, K. Janowicz, B. Yu, W. Li, and S. Prasad. 2015. "Extracting and Understanding Urban Areas of Interest Using Geotagged Photos." *Computers, Environment and Urban Systems* 54:240–254. https://doi.org/10.1016/j.compenvurbsys.2015.09.001.

Huang, H. 2016. "Context-Aware Location Recommendation Using Geotagged Photos in Social Media." *ISPRS International Journal of Geo-Information* 5 (11): 195. https://doi.org/10.3390/ijgi5110195.

Jaffe, A., M. Naaman, T. Tassa, and M. Davis. 2006. "Generating Summaries and Visualization for Large Collections of Geo-Referenced Photographs." *Proceedings of the 8th ACM international workshop on Multimedia information retrieval*, 89–98. October.

Jaffe, A., M. Naaman, T. Tassa, and M. Davis. 2006. "Generating Summaries for Large Collections of Geo-Referenced Photographs." *Proceedings of the 15th international conference on World Wide Web*, 853–854. May.

Juhász, L., T. Novack, H. H. Hochmair, and S. Qiao. 2020. "Cartographic Vandalism in the Era of Location-Based Games—The Case of OpenStreetmap and Pokémon GO." *ISPRS International Journal of Geo-Information* 9 (4): 197. https://doi.org/10.3390/ijgi9040197.

Kennedy, L., M. Naaman, S. Ahern, R. Nair, and T. Rattenbury. 2007. "How Flickr Helps Us Make Sense of the World: Context and Content in Community-Contributed Media Collections." *Proceedings of the 15th ACM international conference on Multimedia*, 631–640. September.

Kennedy, L. S., and M. Naaman. 2008. "Generating Diverse and Representative Image Search Results for Landmarks." *Proceedings of the 17th international conference on World Wide Web*, 297–306. April.

Kisilevich, S., M. Krstajic, D. Keim, N. Andrienko, and G. Andrienko. 2010. "Event-Based Analysis of people's Activities and Behavior Using Flickr and Panoramio Geotagged Photo Collections." *2010 14th International Conference Information Visualisation*, 289–296. IEEE. July.

Laato, S., T. Pietarinen, S. Rauti, and E. Sutinen. 2020. "Potential Benefits of Playing Location-Based Games: An Analysis of Game Mechanics." In *Computer Supported Education: 11th International Conference, CSEDU 2019*, 557–581. Vol. 11. Heraklion, Crete, Greece: Springer International Publishing. May 2–4, 2019.

Lemieux, A. M. 2015. "Geotagged Photos: A Useful Tool for Criminological Research?" *Crime Science* 4 (1): 1–11. https://doi.org/10.1186/s40163-015-0017-6.

Lewis, B. K. 2009. *Social Media and Strategic Communication: Attitudes and Perceptions Among College Students* Doctoral dissertation.

Li, X., C. G. Snoek, and M. Worring. 2008. "Learning Tag Relevance by Neighbor Voting for Social Image Retrieval." *Proceedings of the 1st ACM international conference on Multimedia information retrieval*, 180–187. October.

Li, Y., N. Snavely, and D. P. Huttenlocher. 2010. "Location Recognition Using Prioritized Feature Matching." In Computer Vision–ECCV 2010: 11th European Conference on Computer Vision, Heraklion, Crete, Greece, 791–804. Vol. 11. Berlin Heidelberg: Springer. September 5–11, 2010, *ProceEdited byings Part II*.

Memon, I., L. Chen, A. Majid, M. Lv, I. Hussain, and G. Chen. 2015. "Travel Recommendation Using Geo-Tagged Photos in Social Media for Tourist." *Wireless Personal Communications* 80 (4): 1347–1362. https://doi.org/10.1007/s11277-014-2082-7.

Moxley, E., J. Kleban, and B. S. Manjunath. 2008. "Spirittagger: A Geo-Aware Tag Suggestion Tool Mined from Flickr." *Proceedings of the 1st ACM international conference on Multimedia information retrieval*, 24–30. October.

Moxley, E., J. Kleban, J. Xu, & B. S. Manjunath. 2009. "Not All Tags are Created Equal: Learning Flickr Tag Semantics for Global Annotation." *2009 IEEE International Conference on Multimedia and Expo*, 1452–1455. IEEE. June.

Neustaedter, C., A. Tang, and T. K. Judge. 2013. "Creating Scalable Location-Based Games: Lessons from Geocaching." *Personal and Ubiquitous Computing* 17 (2): 335–349. https://doi.org/10.1007/s00779-011-0497-7.

Nicklas, D., C. Pfisterer, and B. Mitschang. 2001. "Towards Location-Based Games." *Proceedings of the international conference on applications and development of computer games in the 21st century: ADCOG*, 61–67. 21.

Nugraha, S., and M. Damen. 2013. "Utilization of Geotagged Photograph, Remote Sensing, and Gis for Post-Disaster Damage Assessment." *The Indonesian Journal of Geography* 45 (2): 187–204.

Osatuyi, B. 2013. "Information Sharing on Social Media Sites." *Computers in Human Behavior* 29 (6): 2622–2631. https://doi.org/10.1016/j.chb.2013.07.001.

Overell, S., B. Sigurbjörnsson, and R. Van Zwol. 2009. "Classifying Tags Using Open Content Resources." *Proceedings of the Second ACM International Conference on Web Search and Data Mining*, 64–73. February.

Park, M., J. Luo, R. T. Collins, and Y. Liu. 2010. "Beyond GPS: Determining the Camera Viewing Direction of a Geotagged Image." *Proceedings of the 18th ACM international conference on Multimedia*, 631–634. October.

Quack, T., B. Leibe, and L. Van Gool. 2008. "World-Scale Mining of Objects and Events from Community Photo Collections." *Proceedings of the 2008 international conference on Content-based image and video retrieval*, 47–56. July.

Rattenbury, T., N. Good, and M. Naaman. 2007. "Towards Automatic Extraction of Event and Place Semantics from Flickr Tags." *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, 103–110. July.

Rattenbury, T., and M. Naaman. 2009. "Methods for Extracting Place Semantics from Flickr Tags." *ACM Transactions on the Web (TWEB)* 3 (1): 1–30. https://doi.org/10.1145/1462148.1462149.

Senaratne, H., A. Mobasheri, A. L. Ali, C. Capineri, and M. Haklay. 2017. "A Review of Volunteered Geographic Information Quality Assessment Methods." *International Journal of Geographical Information Science* 31 (1): 139–167. https://doi.org/10.1080/13658816.2016.1189556.

Shirky, C. 2005. Ontology is Overrated: Categories, Links, and Tags. *http: //Www. ShirkyCom/Writings/Ontology_overrated.Htm*.

Sigurbjörnsson, B., and R. Van Zwol. 2008. "Flickr Tag Recommendation Based on Collective Knowledge." *Proceedings of the 17th international conference on World Wide Web*, 327–336. April.

Spyrou, E., and P. Mylonas. 2016. "A Survey on Flickr Multimedia Research Challenges." *Engineering Applications of Artificial Intelligence* 51:71–91. https://doi.org/10.1016/j.engappai.2016.01.006.

Sun, Y., H. Fan, M. Bakillah, and A. Zipf. 2015. "Road-Based Travel Recommendation Using Geo-Tagged Images." *Computers, Environment and Urban Systems* 53:110–122. https://doi.org/10.1016/j.compenvurbsys.2013.07.006.

Wetzel, R., L. Blum, and L. Oppermann. 2012. "Tidy City: A Location-Based Game Supported by in-Situ and Web-Based Authoring Tools to Enable User-Created Content." *Proceedings of the international conference on the foundations of digital games*, 238–241. May.

Xing, H., Y. Meng, D. Hou, F. Cao, and H. Xu. 2017. "Exploring Point-Of-Interest Data from Social Media for Artificial Surface Validation with Decision Trees." *International Journal of Remote Sensing* 38 (23): 6945–6969. https://doi.org/10.1080/01431161.2017.1368101.

Yang, Y., Z. Gong, and L. H. U. 2011. "Identifying Points of Interest by Self-Tuning Clustering." *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval*, 883–892. July.

Zheng, Y. T., Z. J. Zha, and T. S. Chua. 2012. "Mining Travel Patterns from Geotagged Photos." *ACM Transactions on Intelligent Systems and Technology (TIST)* 3 (3): 1–18. https://doi.org/10.1145/2168752.2168770.

Zheng, Y. T., M. Zhao, Y. Song, H. Adam, U. Buddemeier, A. Bissacco, and H. Neven. 2009. "Tour the World: Building a Web-Scale Landmark Recognition Engine." *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 1085–1092. IEEE. June.

Zhu, X. X., Y. Wang, M. Kochupillai, M. Werner, M. Häberle, E. J. Hoffmann, and K. Abdulahhad. 2022. "Geo-Information Harvesting from Social Media Data. *IEEE Geoscience and Remote Sensing Magazine* 10 (4): 150–180. https://doi.org/10.1109/MGRS.2022.3219584.

Zielstra, D., and H. H. Hochmair. 2013. "Positional Accuracy Analysis of Flickr and Panoramio Images for Selected World Regions." *Journal of Spatial Science* 58 (2): 251–273. https://doi.org/10.1080/14498596.2013.801331.