Through the Eyes of Instagram: Analyzing Image Content utilizing Meta's Automatic Alt-Text

João Francisco Hecksher Olivetti
joao.olivetti@ufv.br
Universidade Federal de Viçosa - UFV
Florestal, MG, Brasil

Philipe de Freitas Melo
philipe.freitas@ufv.br
Universidade Federal de Viçosa - UFV
Florestal, MG, Brasil

ABSTRACT

Multimedia communication has become an essential part of social media, with images representing a significant part of the content on most platforms. This study investigates image content on Instagram through Meta's internal image classification algorithm, Automatic Alt-Text (AAT). Our approach differs from research on data from comments and hashtags because of the use of actual visual descriptions as the means of understanding the kinds of the content published on the network. Our analysis of 200k posts reveals 1,471 unique tags being used to characterize image content on Instagram, representing mostly objects, food, animals, locations and other common components of social media photos. Notably, we found that content about personal aesthetics is highly popular on the platform, with person and selfie being respectively some of the top two most common tag and post categories, being also highly related to other tags such as makeup, lipstick and eyeliner. Furthermore, we explored the connections between tags, representing very popular content trends within the network. Finally, we uncover substantial differences in posting behavior of influencers and news pages when compared to regular users, observing they post more frequently and about more specific content, suggesting what may attract more engagement on Instagram.

KEYWORDS

Instagram, alt-text, social media, image classification, image tagging, complex networks, influencers

1 INTRODUCTION

The Web has become an integral part of daily life for a significant portion of the global population. With the increasing prevalence of online social networks and the multimedia nature of communication online, visual content understanding is crucial. To enhance accessibility, various tools have been implemented to assist disabled users in this task. One notable example is "Alt-text" (alternative text), a textual description of an image included in the HTML code of a webpage. It describes the appearance and function of an image to users who cannot see it, those using screen readers, or when the image fails to load.

In 2017, Instagram and Facebook implemented Meta's Automatic Alt-Text (AAT) algorithm, which automatically extracts features from images and photos and provides tags and textual descriptions of visual content. Meta's AAT has proven effective in aiding blind users to navigate these platforms and understand the image content [29]. Furthermore, recent extensions to the AAT system created a large-scale description of content on Instagram, as this feature provided valuable details to the entire platform, beyond just assisting visually impaired users. Image analysis poses significant challenges for researchers due to the inherent difficulty of both manual and automated processes to categorize its content. However, it is a key social currency on the Web [13], and substantial insights can be gained from analyzing images on Instagram, ranging from social behavior [5] to political campaigns [22] and even typification of fake profiles [30]. Most existing studies, though, rely on traditional analysis of textual comments or hashtags from posts, which may not accurately represent the actual content portrayed in pictures.

This study aims to analyze the structure of Meta's Automatic Alt-Text (AAT) system in categorizing Instagram images. We seek to answer the following research questions: **RQ1:** What are the most popular content tags on different Instagram user classes? **RQ2:** What is the relationship between these contents? **RQ3:** Is there a substantial difference between the content posted by different types of users?

To achieve this, we utilized a custom-built data scraper to collect 61,944 Instagram accounts and 198,623 posts. This effort extracted approximately 550k Alt-text tags describing the content of Instagram images. After that, by applying complex networks analysis, association rules and ranking metrics, we conducted a broad analysis to better understand and categorize the content posted on Instagram, as also the relation of these tags among the users and its popularity.

Our results indicate that the current state of the AAT algorithm shares similarities, but also has significant differences from what is described in the literature. Through 1,471 unique tags used by AAT that we found within our dataset, objects, foods, animals, locations are among the most common topics on Instagram posts. We found that personal aesthetic content is highly popular on Instagram, often linked to beauty tags such as "makeup", "lipstick", and "eyeliner". Lastly, we observed substantial differences in posting behavior between influencers, news pages, and regular users.

2 BACKGROUND AND RELATED WORK

This section shows similar, adjacent or related research to the present work. A discussion on what alt-text is also included as context on this papers main datapoint.

2.1 Automatic Alt-Text

Alt-text is the content present in the alt attribute of HTML documents. Its original usage was to show a brief description of an object that could not be loaded. Nowadays, it is used mainly for

In: Proceedings of the Brazilian Symposium on Multimedia and the Web (WebMedia'2024). Juiz de Fora, Brazil. Porto Alegre: Brazilian Computer Society, 2024.

ISSN 2966-2753

accessibility purposes, allowing screen readers to describe visual content to visually impaired users.

In social media, as content is user-uploaded, most images end up having no alt-text associated with them. This has led many Online Social Networks to allow users to add their own description of the image. This feature, however, remains very underutilized. There has been an extensive body of work focusing on alt-text, including analyses on its manual and automated construction[11], pre-processed image-description pair datasets[27], inquiries on what Blind or Low-Vision (BLVs) people want from these tags[28], as well as what harms can arise from alt-text tagging[16].

This has motivated Meta's (by the time called Facebook) researchers to develop the technology known as Automatic Alt-Text (AAT), with its original development being discussed in [29]. Their article focuses on the process of assuring the system was useful to people with vision disabilities.

2.2 Social Media Analysis

The analysis and extraction of content from social networks is increasingly relevant as these platforms have become integral to daily life. Studies in this field offer insights into social impacts and phenomena on these networks, as evidenced by studies like [2, 24]. One prevalent approach is sentiment analysis through various strategies, ranging from simplified methods like the influential study by [14] to more advanced neural network techniques [23, 25]. These studies aim to computationally classify the emotional content of sentences, a complex task with numerous applications [8, 9].

Other approaches focus on understanding social media dynamics, such as Complex Network modeling, which provides deep explanatory insights into relationships. A notable study in this area by [7] used Twitter data to show that follower count alone is not a good metric for influence on the network. Other similar works include [18, 19, 26]. Twitter has been extensively studied due to its APIs and openness to academic research. However, with recent changes, there is a growing trend towards studying other social networks. Despite challenges in extracting data from less accessible networks, research on these platforms, such as the work by [21], highlights the value of image content and systematizes the investigation of fake news, revealing how messaging apps function as social networks.

The focus on images, which are computationally more costly to analyze than texts, partly explains the fewer studies on Instagram [13]. Nonetheless, significant studies on Instagram do exist. For instance, [12] used Complex Network modeling to analyze community interactions during elections. Other studies, like [10, 15] also use this approach.

Some research examines the social implications of image content across various platforms, such as [3, 4]. The foundational study by [13] is one of the few focusing on understanding Instagram's image content. Using early computer vision mechanisms, the authors identified challenges like defining relevant categories and the high computational cost. This study, by categorizing images from various users, directly connects to our work. Other studies, such as [22] and [20] also focus on the image content on Instagram

2.3 Research Gap

Instagram is predominantly a visual platform; however, most studies focus on the textual content of posts, such as comments, which may not accurately reflect the platform's full picture. By analyzing the alt-text descriptions of images, this study aims to bridge this gap, providing greater transparency and understanding of Instagram. It seeks to deepen the discussion about the topics present in this environment.

3 METHODOLOGY

In this section, we describe the approach and strategies used to gather process and analyze data from Instagram's webpage. It details the methodology used for data collection, analysis, and also relevant limitations encountered during the study. The methodology is divided into three main parts: Data Collection, Data Analysis, and Limitations.

3.1 Data Collection

To investigate Meta's Automatic Alt-Text (AAT) data on Instagram, we utilized a custom-built data scraper developed using Selenium[1], a tool that automates web browser navigation, enabling us to collect data without the use of APIs.

The time period of the data collection was between 09/2023 and 02/2024, about six months. This was the time we needed to have a sample close in size to the one Meta's researchers used to build the AAT algorithm[29], both having about 200K posts. It is important to note that the our scraper collected posts going from most recent to older, so some posts may be from dates prior to the start of the extraction, having been posted previously. Since Regular Users might not post anything for very long periods of time, their posts are from an irregular time period. One observation noted from checking these is that posts from before 2018 have no AAT tags, matching the time frame from the original paper, published in 2017.

Some of the challenges presented were dynamic content loading, frequent changes in Instagram's HTML structure, rate limiting and blocking, and obfuscated tags. Dynamic content loading requires waiting for the content to fully load before extracting alt text. Updates to Instagram's HTML structure break scraping scripts, requiring updates. Additionally, Instagram imposes rate limits on requests, which can lead to IP blocking or suspension. Moreover, Instagram's web version does not show posts if no account is logged in. Requests made in this context repeatedly return pages reporting an error. In order to collect Instagram data, one must be logged in to a valid, not-banned account. To manage these obstacles, 8 collection accounts were created. The crawler then cyclically changes the account used to scrape Instagram profiles, through Selenium, preventing individual accounts from exceeding limits without delaying extraction. We observed the most common way to get an account banned is to log in manually multiple times in a short period.

The data collection loop, which directly interacts with a user's page on the network, operates as follows: first, it checks if the current account to be extracted is already in the database. If not, a new data point is created, containing the username and the total number of posts from that profile. The algorithm then proceeds to the next step, scrolling and accessing each post on the account to

[1] Available at: https://selenium-python.readthedocs.io/

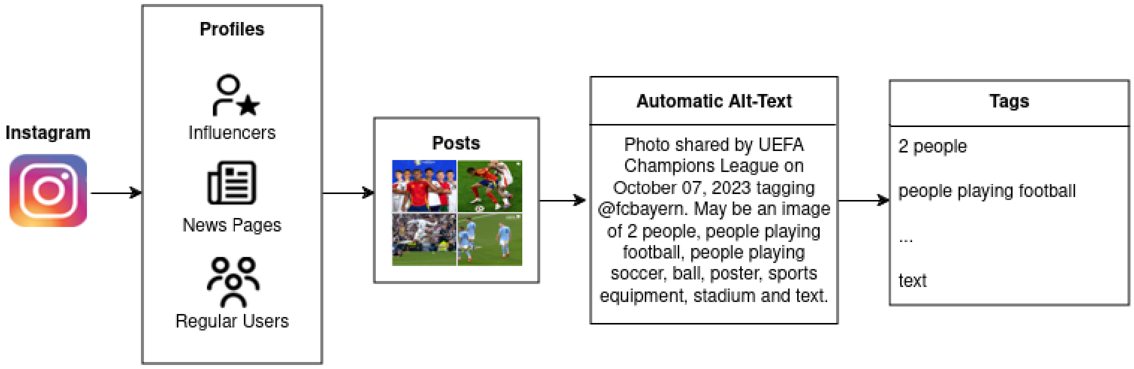

Figure 1: Flowchart of Instagram data collection methodology.

check if they contain alt-text, until the end of the page is reached or one of several code-imposed limits is hit to avoid unnecessary operations. With the web-scraper implemented, we targeted three main types of Instagram users for data collection:

Influencers: We start our profile's collection with the most popular accounts within the network. These influential users were obtained from HypeAuditor's ranking of the top 1000 accounts[2] with the most engagement, which calculates the ranking as a composite measure of total followers and the average number of likes and comments per post. These profiles mainly belong to celebrities (such as `@cristiano` and `@taylorswift`), but also include some politicians (e.g. `@narendramodi`) and a few "*meme pages*" (e.g. `@9gag`). 86 accounts were excluded from top 1000 for reasons as the profiles being deactivated, turning private, or archiving posts.

News Pages: Profiles from this category represent respected news outlets, tabloids, gossip magazines and information-focused pages. These were selected starting from manually selected "seeds", mainstream news outlet users (e.g. `@nytimes`), which were then used to acquire more news pages through Instagram's recommendation system. A total of 59 pages were found and collected through this process. This fewer number of profiles was enough to reach a similar number of publications, as shown in Table 1, in which the 59 news pages have, in sum, more posts than the 914 Influencers.

Regular Users: Also included were ordinary Instagram users. Due to the vast size of the Instagram network and its restrictions, we developed a strategy to incorporate a diverse set of profiles to our data collection. To achieve this, we collected the usernames of followers of Influencer and News profiles. This aimed to increase the randomness and variation on the sample, as popular profiles usually have a set of followers with a wide range of interests and demographics. This method allowed us to include a sample of 60,971 general Instagram user base and 150,662 posts in our analysis.

3.2 Instagram Alt-text Structure

After collecting data from Instagram users and their posts, it is essential to understand the formatting of posts and the placement of alt text for effective data collection. This subsection explains the structure of a typical Instagram post, how alt text is integrated, and the challenges associated with extracting this information.

A typical Instagram post consists of visual content (images or videos), textual content (captions, comments, and hashtags), and

[2]Available at: https://hypeauditor.com/top-instagram/

Photo shared by CNN Brasil on January 22, 2024 tagging @cnneconomia, and @cnnpolitica. May be an image of 4 people, bureau and text that says 'CNN BRASIL Você por dentro de tudo. Igualdade salarial: empresas devem informar ao ao governo remuneração de empregados'.

Discarded Preamble
Automatic Alt-Text Data

Figure 2: Example of alt-text structure of Instagram post.

"May be a Twitter screenshot of 1 person, microphone and text."

"May be an image of text that says 'BREAKING UK and US launch new air strikes on Houthi targets in Yemen'."

"May be a black-and-white image of map, chandelier, globe, armchair, chaise lounge and indoors"

Image Class
Classification Tags
OCR Results

Figure 3: Features extracted from the Instagram alt-text.

alt-text, which is a hidden layer of textual information associated with the visual content. The alt-text is embedded within the `<img>` in the HTML structure of an Instagram post, specifically as the value of its alt attribute. After reaching the actual content of alt-text, it can be presented in different formats and variations. We outlined the alt-text structure to classify and utilize it effectively. One of the first observations was that not all images containing alt-text have image "classification tags". This aligns with Meta's information on [29], which explains that before the AAT technology was implemented, alt-texts were either manually added by users, which was very rare, or it had the generic "User's Photo" phrase attached to it. Another alt-text format containing no tags only includes the preamble, similar to that shown in Figure 2.

Finally, most of the AAT content has a structure similar to the format portrayed in Figure 3, highlighting how the alt text is located within the HTML code. The uncovered structure revealed three

	Accounts	Posts	Tags	Unique Tags
Regular Users	60.971	159.662	420.593	1470
News Pages	59	20.945	52.426	992
Influencers	914	18.016	73.133	1095
Total	61.944	198.623	546.152	1471

Table 1: Summary of the Instagram dataset collected.

main features in the Instagram alt-text: (1) An "**Image Class**" at the beginning of the AAT, represented by the description "*May be a* ****"; (2) a set of "**Description Tags**", which are class labels assigned to images with the aim of describing the visual content; and occasionally (3) "**OCR text**", which represent actual text present within the images.

With the AAT features defined, we implemented an algorithm to process this information. Using Python, we created the `TextContent` class. This class receives an alt-text string and divides it into the relevant features, accessed via the class' getters. All tag analysis operations were mediated by this structure, reducing errors and simplifying manipulation. Data collection results are summarized in Table 1. Our final dataset includes over 60,000 users and approximately 200,000 posts, from which we extracted around 550,000 alt-text tags that describe the content of the images on Instagram.

3.3 Limitations

There are some limitations associated with this collecting methodology and the resulting dataset. Beyond all restrictions imposed by Instagram to gather data from the platform, the main limitation is that, considering that Instagram has more than a billion users worldwide, our collection of 60k users and 200K posts represents a small sample, not necessarily representative of the entirety of the network. However other key studies in the field use datasets of similar sizes, with Meta using exactly 200K posts to build sample what tags they would need to initially build AAT algorithm[29]. Also, given the AAT system automatically assigns tags to the images, the data may present some mislabeled content, and we cannot ensure the quality of the image labeling. That said, this is the actual tool Instagram uses to provide information for the screen reader users, and thus our view of the content is the same a BLV user may have within the platform.

4 RESULTS AND DISCUSSIONS

In this section, we present our analysis of the data collected. We examine the structure of Meta's AAT system, using tools such as complex networks and association rules metrics, focusing on how it categorizes and describes visual content. Our findings include an analysis of tag popularity, relationships between alt-text tags, and a comparison of the content shared by regular users, news pages, and Influencers.

4.1 Post Types on Instagram

Once the data was pre-processed and ready, we began the analysis and understanding of Instagram content. Figure 4 illustrates the high-level classification of Instagram images by AAT. This classification reveals key differences in content types shared by various user groups. The "*Image Class*" feature is the first description the

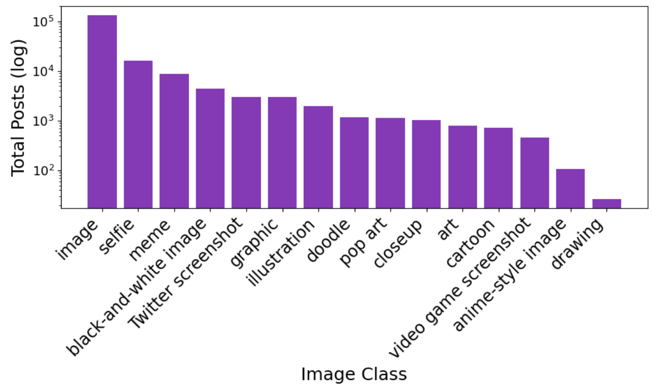

Figure 4: Types of posts present on Instagram data.

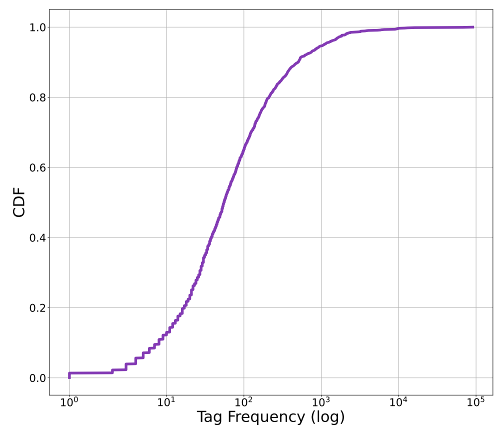

Figure 5: CDF of popularity of each tag.

AAT algorithm provides for users, serving as a general guide to what type of post the image could be. Although most content is classified as an image, Instagram differentiates whether the content is also a selfie, a meme, an illustration and other types. Unlike tags, a post can belong to only one ImageClass. When considering all 15 classes found, however, the methodology behind their definition is not clear. About half of the classes are art-related (illustration, doodle, pop art, art, cartoon, anime-style image, and drawing), with no obvious distinction among them. Although some classes have more well-defined content, such as "selfie", "back-and-white images" or "anime-style", manual observation within each class did not clarify the specific reasoning behind this labeling.

4.2 RQ1: Content Popularity by User Type

We also analyzed the popularity of the description tags provided by Meta's Automatic Alt-Text (AAT) system. The number of tags per post varies, with some posts having only one tag and others containing up to ten tags. Our findings show that only a few tags appear just once, and more than 95% of tags are present in at least 2 or more posts. This suggests Instagram has limited pre-built tag classes that they assign to describe images on the platform. Also, about 65% of the tags appear up to 100 times among the images and approximately 5% of the tags are highly popular, present in more than 1,000 different posts as shown in CDF of Figure 5.

	Regular Users	Influencers	News Pages
1.	text (56K)	text (13K)	text (20K)
2.	1 person (50K)	1 person (8K)	1 person (7K)
3.	one or more ppl (10K)	stadium (2K)	2 people (3K)
4.	2 people (10K)	2 people (2K)	poster (2K)
5.	smiling (8K)	poster (2K)	magazine (2K)
6.	poster (8K)	makeup (1K)	3 people (1K)
7.	top (8K)	people playing football (1K)	newspaper (738)
8.	hair (8K)	people playing soccer (1K)	4 people (735)
9.	beard (8K)	one or more people (1K)	people standing (520)
10.	makeup (7K)	playing football (1K)	5 people (469)

Table 2: Ranking of top 10 most popular tags on Instagram.

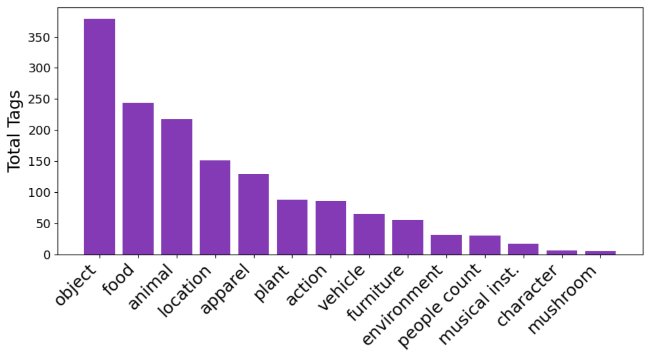

Figure 6: Manual classification of tag categories.

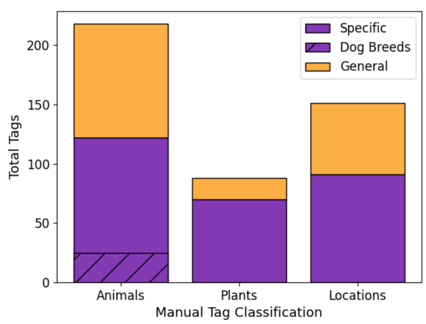

Figure 7: Tag specificity/generalization level per type.

Additionally, we examined the rankings of the most popular tags, as shown in Table 2. This table lists the most frequent tags for each user type. We observed that for all user types, tags like "*text*" and "*1 person*" are the most common. This aligns with findings from other studies, such as [5] and [13], indicating that individual photos of people with some captions are extremely popular on the Instagram. Other frequently tags, such as "*selfie*", "*hair*" and "*beard*" further highlights the prevalence of personal photos on the platform.

In total, we extracted 546.152 tags among the posts within our dataset; however, we found those tags comprise a total of only 1,471 unique tags used by Instagram to describe the image content. To have a wide view of what those tags are and the actual content of Instagram images, we manually labeled them into broader categories, following a similar approach of authors from [13]. This categorization was done through iterative labeling, in which each step the categories were grouped, resulting in 14 classes by the end, as illustrated in Figure 6. Interestingly, we also observed and labeled a different degree of specificity/generalization of each tag. While some tags represent very general concepts, such as "*text*", some are very specialized. For example, Instagram differentiates a lot of specific cat and dog breeds, as well as famous real-world locations. Figure 7 represents the portion of tags defined as general specific. For example, "*duck*" is considered a general tag, whereas "*mallard*" is specific. Following this process, we discovered that objects, food and animals are the categories with most tag variation on Instagram. An interesting observation from this is that the categories with more tags are not necessarily the most popular among posts. Furthermore, we find a redundancy among tags. For instance, tags such as "*soccer*" is accompanied by tags such as "*playing soccer*" and "*people playing soccer*". These specific action tags could reflect results from [29] in which impaired users indicated that people are the most interesting part of an image, especially their mood and what they were doing (action).

Additionally, we found typos in tags like "*acquatic animal*" and "*riding a crousel*", in which all occurrences have the same error, but this did not repeat in tags such as "*carousel*" and "*people riding a carousel*".

4.3 RQ2: How AAT Content Interacts on Instagram

Next, we focused on the relation between kinds of image content posted on Instagram by investigating common co-occurrence of tags within our dataset. To consider what users usually put together on their photos on Instagram, we firstly use complex network analysis to develop the graph of tags presented in Figure 8 and assess common paths that posts follow on the network with community algorithms. In this graph, the nodes represent each of the tags collected, and each edge represents the co-occurrence of two tags in a post; edges are weighted by how many times the tags appeared together. The size of a node is determined by its degree, with the most well-connected nodes being the largest. As expected, these are nodes such as "*text*' and "*1 person*", which are the top 1 and 2 most frequent within the dataset.

The grouping of nodes was made using the Community Detection algorithm described by [6], as implemented in Gephi. It resulted in uncovering the 10 communities, 8 of which had a relevant enough size, differentiated by color in the graph. It is interesting to observe how some types of content are grouped on Instagram and tend to appear together in a post. Examples such as "*cup*", "*bottle*" and "*tea*" linked together show how objects that one would expect to be together actually appear in images. In the same direction we also have "*lipstick*", "*makeup*", "*hair*", "*eyeliner*", all around "*1 person*" suggesting the popularity of such kind of selfies within the network. In general, our findings on this network analysis revealed linked groups of content representing as following: **Red**: Generally contains objects associated with apparel. Mirrors the "Fashion" category from [13]; **Pink**: includes most of the "People Count" mentioned in the original AAT paper[29]. Also includes objects, notably things that usually appear alongside people; **Cyan**: Represents landscapes, their components and some outdoor activities; **Light Green**: Contains food and kitchen utensils. Also, a category in [13]; **Dark Green**: Species of animals, plants and types of furniture; **Deep Blue**: Includes only flowers and four other tags, those being "*bouquet*", "*vase*" and "*flower pot*"; **Grey**: Vehicles, animals used in transportation and some other apparently unrelated critters; **Brown**: Eight out of the nine items in this category are containers, such as "*bag*", "*wallet*", "*purse*" and "*suitcase*".

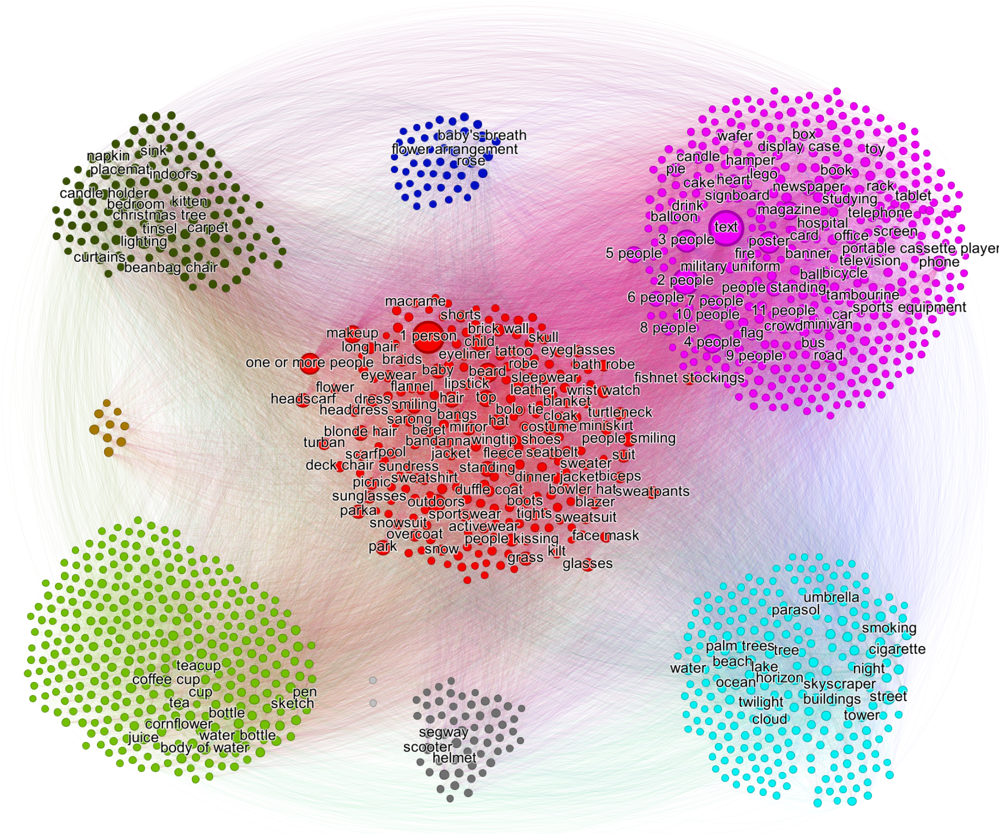

Figure 8: Network of tag relationships. Nodes are tags and edge means co-ocurrence in same post.

Association Rules: The second approach to understand relationships between content is extracting the association rules through Apriori Mining algorithm [1], which can be used to point pairs or sets of tags that have high probabilities to appear together within a post. Since the large differences between user types were shown, the rule mining was applied separately on news pages, regular users and Influencers.

On the news pages set, there are many high-confidence association rules on pairs, especially those following the pattern `(people count tag) -> (text)`, with a confidence above 0.9 This helps to explain posts frequently seen in news pages, in which people involved in some news-worthy project are depicted along with text that gives context. Other similar rules are `(1 person, poster, magazine) -> (text)` and `(2 people, poster, magazine) -> (text)`. These show a very similar pattern, but they also include other common tags for news publications. In general, those rules suggest news pages have a much more homogeneous posting format, depicting usually a person or other object in a specific format of poster accompanied by some text describing the news headline.

The influencer users are those with the most association rules. A very representative and high confidence rule is `(stadium, playing soccer) -> (text, 1 person, playing football)`, which includes most of the tags associated with the sport. One reason for this may come from the AAT redundancy in action tags and the absence of difference between "*soccer*" and "*football*". Additionally, several famous football players are among the top 1000 most influential users in our dataset due to the list used to collect them, contributing to the popularity of this type of content within this user category. However, these also evidence that thematic posts of soccer are prevalent for those users. On the other hand, few rules have the tag "*ball*", suggesting that, despite the theme, the game itself is not the focus of the post. One of the main rules that is not related to sport is `(makeup, 1 person) -> (text)`, a rule seemingly more representative of the artists and other influencers of this category. Note the presence of textual content in these influencers rules, which is a common format for events, advertises, sponsored posts and other communication tools that celebrity accounts are used to post on Instagram.

Finally, the tags associated with the regular user type are much more diverse than those of the other categories, but generally have lower confidence levels. For instance, the rule `(smiling) -> (1 person)` has a confidence of 0.96, which is higher than most others. We also observe pairs such as `(ocean) -> (beach)`, with a confidence of 0.62. Interestingly, rules related to makeup appear more frequently and with higher confidence. For example, `(lipstick) -> (makeup)`, `(eyeliner) -> (makeup)`, `(makeup, eyeliner) -> (1 person)`, and `(makeup, blonde hair) -> (1 person)`, all of which have confidence levels close to 0.7. Additionally with

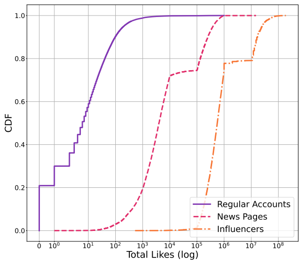

(a) Likes per Post

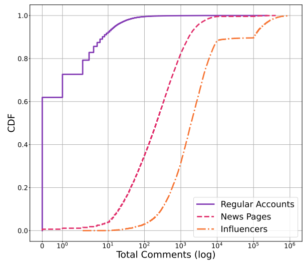

(b) Comments per Post

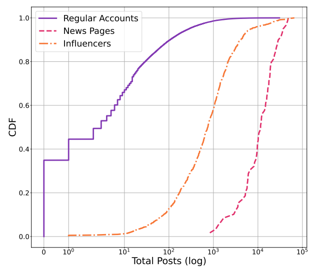

(c) Posts per User

Figure 9: Differences in engagement for different user types.

"person" and "selfie" being the second most common tags and post categories, it is evident that content about personal aesthetics is highly popular on Instagram. These findings indicate that users are highly engaged with content related to personal appearance and self-expression.

4.4 RQ3: Posting Behavior by User Type

Next, we explore the differences in posting behavior and engagement across different kinds user classes. We analyze the frequency of posts, the number of tags per post, as well as engagement metrics such as likes and comments from regular users, news pages Influencer profiles. These insights help us understand how different user types utilize the platform and how their content resonates with their audience.

Differences in Engagement: First, we examined some engagement metrics such as likes and comments per post, as shown in Figure 9. Influencers generally receive more likes and comments compared to regular users and news pages, reflecting their higher engagement levels due to their follower base and the nature of their content. It is notable that for regular users, about 20% and 60% of their posts have no likes and no comments at all, respectively. On the other hand, almost all influencer's posts reach more than 100k likes and 100 comments. For comparison, only 25% of news posts have similar likes. However, looking at the number of posts per type of user (Figure 9c, we can note that influencers are not the users with the most posts. News pages have a much higher number of posts than other types, with all news pages collected having more than 1,000 posts in their profiles, while less than 1% of regular users and just 40% of influencers have such numbers. News pages post most frequently, often several times a day. This contrasts with Regular users, where almost 40% have zero posts and about 90% have up to 100 posts. Since Regular users typically have no economic or professional reasons to post frequently, their totals are usually much lower.

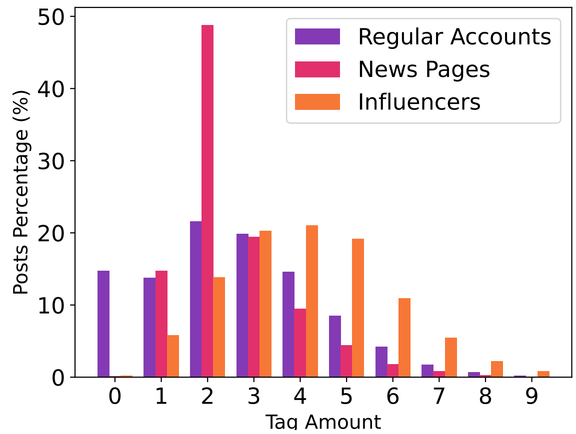

Figure 10: Number of tags per post for each type of user.

Differences in Content: After evaluating some posting behaviors, we take our attention to the actual content of the publication of each types of users. First, we see the tag frequency of their posts in the histogram of Figure 10. While about 60% of regular users' posts have up to 3 tags, we observe influencers generally have more tags within their publications. This suggests influencers have a higher variety of content in a single post. Also, it can be attributed to product placement, showing various objects, and the diverse lifestyles of celebrities and politicians. No posts from news pages have zero tags, but almost 50% have exactly two tags. This is because most images from news pages include a text, thus heavily increasing the chances of including the "*text*" tagging, consequently also have the OCR property, which is enough to describe the image without additional tags required.

To further measure the differences between content published by user types, we applied Kendall's Tau rank correlation metric [17] to compare rankings of the top 10, 100, and 992 tags of each account class, as shown in Table 3. The number 992 was used instead of 1000 because that was the total number of tags in the news pages dataset. Considering the top 10 tags of each set, the Kendall's Tau results show weak negative associations ($\tau < 0$) between Regular/Influencer and Influencer/News pages. The strongest negative correlation is between News pages and Regular users, indicating a moderate but significant negative association, suggesting News accounts are the most distinct user category.

The correlation results from Kendall's Tau show that there are substantial differences between all the types of users. Even when considering most of the tags, the positive correlation is at best moderate, suggesting some content is always popular. As comparison, we also apply Kendall Tau it to another set: two randomly distributed, equally sized partitions of the "Regular user" dataset. The results, shown in Table 4, indicate a strong positive correlation and that random samples would present similar tags ($\tau > 0.87$) among their content. This confirms that the negative correlations observed between user types are not random, and makes it evident that they have different behavior on Instagram.

Account Types	Top10 τ	Top100 τ	Top992 τ
Regular/Influencer	-0.077	0.092	0.528
Influencer/News	-0.167	-0.211	0.543
News/Regular	-0.446	-0.194	0.501

Table 3: Kendall's Tau Applied between different User Types

Account Types	Top10 τ	Top100 τ	Top1000 τ
Random A / B	0.918	0.911	0.872

Table 4: K τ of two regular user sets randomly partitioned.

5 CONCLUSION

In this study, we conducted an extensive investigation of image content on Instagram through AAT system. By analyzing alt-text data from nearly 200,000 user posts collected from 60,000 Instagram users, we provided significant insights into the visual content shared on the platform.

Our findings reveal 1,471 unique tags used by Instagram to categorize content, composed mostly of objects, food, animals, and locations. These show that content related to personal aesthetics, such as makeup, lipstick, and eyeliner, are highly popular on Instagram. Additionally, we identified several tags that are frequently connected, forming popular content trends within the network. Moreover, we highlight differences in the content shared by news pages, Influencers, and regular users. Notably, news pages tend to post images containing text, while Influencers often share content related to personal aesthetics and activities, such sports. Regular users, on the other hand, post a diverse range of content, although having fewer tags per post.

Our research also highlighted the evolution of the AAT algorithm, noting the introduction of many new tags, including some seemingly redundant. Future studies could expand on this work by incorporating larger samples and applying AAT context to studies on political campaigns, mental health, and disinformation on Instagram. Our approach offers a fresh perspective that complements existing research, enhancing our understanding of the diverse and dynamic nature of content shared on this social network.

REFERENCES

[1] Rakesh Agrawal, Ramakrishnan Srikant, et al. 1994. Fast algorithms for mining association rules. In *Proc. 20th int. conf. very large data bases, VLDB*, Vol. 1215. Santiago, 487–499.

[2] Wasim Ahmed, Josep Vidal-Alaball, Joseph Downing, and Francesc López Seguí. 2020. COVID-19 and the 5G Conspiracy Theory: Social Network Analysis of Twitter Data. *Journal of Medical Internet Research* 22, 5 (2020), e19458.

[3] Camila S Araujo, Luiz P D Correa, Ana P C Silva, Raquel O Prates, and Wagner Meira. 2014. It is Not Just a Picture: Revealing Some User Practices in Instagram. In *9th Latin American Web Congress*. 19–23.

[4] Camila S Araújo, Wagner Meira, and Virgilio Almeida. 2016. Identifying Stereotypes in the Online Perception of Physical Attractivenes. In *Proc. of the 8th Intl. Conf. of Social Informatics*. Springer, 419–437.

[5] Saeideh Bakhshi, David A. Shamma, and Eric Gilbert. 2014. Faces engage us: photos with faces attract more likes and comments on Instagram. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems* (Toronto, Canada) *(CHI '14)*. ACM, NY, 965–974.

[6] Vincent Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. 2008. Fast Unfolding of Communities in Large Networks. *Journal of Statistical Mechanics Theory and Experiment* 2008 (04 2008).

[7] Meeyoung Cha, Hamed Haddadi, Fabrício Benevenuto, and Krishna Gummadi. 2010. Measuring User Influence in Twitter: The Million Follower Fallacy. *Proc. of the International AAAI Conf. on Web and Social Media* 4, 1 (2010), 10–17.

[8] Koyel Chakraborty, Siddhartha Bhattacharyya, and Rajib Bag. 2020. A Survey of Sentiment Analysis from Social Media Data. *IEEE Transactions on Computational Social Systems* 7, 2 (2020), 450–464.

[9] P Chauhan, N Sharma, and G Sikka. 2020. The emergence of social media data and sentiment analysis in election prediction. *Journal of Ambient Intelligence and Humanized Computing* 12, 2 (2020), 2601.

[10] Rion Brattig Correia, Lang Li, and Luis M Rocha. 2016. Monitoring potential drug interactions and reactions via network analysis of instagram user timelines. In *Biocomputing 2016: Proc. of the Pacific Symposium*. World Scientific, 492–503.

[11] Emory J. Edwards, Michael Gilbert, Emily Blank, and Stacy M. Branham. 2023. How the Alt Text Gets Made: What Roles and Processes of Alt Text Creation Can Teach Us About Inclusive Imagery. *ACM Transactions on Accessible Computing* 16, 2, Article 18 (jul 2023), 28 pages.

[12] Carlos H G Ferreira, Fabrício Murai, Ana P C , Jussara M Almeida, Martino Trevisan, Luca Vassio, Idilio Drago, and Marco Mellia. 2020. Unveiling Community Dynamics on Instagram Political Network. In *Proc. of the 12th ACM Conf. on Web Science* (NY). ACM, 231–240.

[13] Yuheng Hu, Lydia Manikonda, and Subbarao Kambhampati. 2014. What We Instagram: A First Analysis of Instagram Photo Content and User Types. *Proc. of the International AAAI Conf. on Web and Social Media* 8, 1 (2014), 595–598.

[14] C. Hutto and Eric Gilbert. 2014. VADER: A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text. *Proc. of the International AAAI Conf. on Web and Social Media* 8, 1 (2014), 216–225.

[15] Simona Ibba, Matteo Orrù, Filippo Eros Pani, and Simone Porru. 2015. Hashtag of Instagram: From Folksonomy to Complex Network.. In *KEOD*. 279–284.

[16] Jared Katzman, Angelina Wang, Morgan Scheuerman, Su L Blodgett, Kristen Laird, Hanna Wallach, and Solon Barocas. 2023. Taxonomizing and Measuring Representational Harms: A Look at Image Tagging. *AAAI Conf. on Artificial Intelligence* 37, 12 (2023), 14277–14285.

[17] Mirella Lapata. 2006. Automatic evaluation of information ordering: Kendall's tau. *Computational Linguistics* 32, 4 (2006), 471–484.

[18] Sebastián Lozano and Ester Gutiérrez. 2018. A complex network analysis of global tourism flows. *International Journal of Tourism Research* 20, 5 (2018), 588–604.

[19] Hafiz A M Malik. 2022. Complex Network Formation and Analysis of Online Social Media Systems. *Computer Modeling in Engineering & Sciences* 130, 3 (2022), 1737–1750.

[20] Breno Matos, Fabrício Benevenuto, and Ana P C da Silva. 2021. A Nudez como Estereótipo de Beleza através das Lentes do Instagram. In *X Brazilian Workshop on Social Network Analysis and Mining*. 163–168.

[21] Philipe Melo, Fabrício Benevenuto, Daniel Kansaon, Vitor Mafra, and Kaio Sá. 2021. Monitor de WhatsApp: Um sistema para checagem de fatos no combate à desinformação. In *Anais do XXVII Simpósio Brasileiro de Sistemas Multimídia e Web (WebMedia'21)*. SBC, 79–82.

[22] Caroline Muñoz and Terri Towner. 2017. The Image is the Message: Instagram Marketing and the 2016 Presidential Primary Season. *Journal of Political Marketing* 16 (06 2017).

[23] László Nemes and Attila Kiss. 2021. Social media sentiment analysis based on COVID-19. *Journal of Information and Telecommunication* 5, 1 (2021), 1–15. Taylor & Francis.

[24] Serge Nyawa, Dieudonné Tchuente, and Samuel Fosso-Wamba. 2022. COVID-19 vaccine hesitancy: a social media analysis using deep learning. *Annals of Operations Research* (2022).

[25] Abubakr H. Ombabi, Wael Ouarda, and Adel M. Alimi. 2020. Deep learning CNN–LSTM framework for Arabic sentiment analysis using textual information shared in social networks. *Social Network Analysis and Mining* 10, 1 (2020), 53.

[26] Khubaib A Qureshi, Rauf A S Malick, Muhammad Sabih, and Hocine Cherifi. 2021. Complex network and source inspired COVID-19 fake news classification on Twitter. *IEEE Access* 9 (2021), 139636–139656.

[27] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. Conceptual Captions: A Cleaned, Hypernymed, Image Alt-text Dataset For Automatic Image Captioning. In *Proc. of the 56th Annual Meeting of the Association for Computational Linguistics*, Vol. 1. ACL, 2556–2565.

[28] Abigale Stangl, Meredith Ringel Morris, and Danna Gurari. 2020. "Person, Shoes, Tree. Is the Person Naked?" What People with Vision Impairments Want in Image Descriptions. In *Proc. of the 2020 CHI Conf. on Human Factors in Computing Systems*. ACM, 1–13.

[29] Shaomei Wu, Jeffrey Wieland, Omid Farivar, and Julie Schiller. 2017. Automatic Alt-text: Computer-generated Image Descriptions for Blind Users on a Social Network Service. In *Proc. of the 2017 ACM Conf. on Computer Supported Cooperative Work and Social Computing (CSCW '17)*. ACM, NY, 1180–1192.

[30] Koosha Zarei, Reza Farahbakhsh, and Noël Crespi. 2019. Typification of Impersonated Accounts on Instagram. In *IEEE 38th International Performance Computing and Communications Conf. (IPCCC)*. 1–6.