# An Automated Ground-to-Aerial Viewpoint Localization for Content Verification

Tania Sari Bonaventura, Luca Maiano, Lorenzo Papa and Irene Amerini
Sapienza University of Rome, Italy
Email: bonaventura.1916415@studenti.uniroma1.it, [amerini, maiano, papa]@diag.uniroma1.it

***Abstract*—In many forensic examinations and media authenticity verifications, it is essential to reconstruct the place where a photo was taken without having information such as GPS coordinates and other metadata available. In recent years, satellite imagery has been used in some of these reconstructions to map the image captured from the ground with the photos from above. This task, known as ground-to-aerial mapping, allows a very effective localization but remains a complex and time-consuming task. Deep learning-based methods allow for accurate automatic image matching; however, many of these solutions can be challenging to explain and, therefore, barely applicable in scenarios where it is necessary to justify the analysis. Consequently, in this paper, we propose a fully automated, explicable solution that is able to perform image-matching tasks based on a graph-based solution. Our proposed pipeline is composed of four stages in which we extract a graph representation of the images that we use for matching. Moreover, the designed pipeline improves previous related methods of $17.84\%$ for the mean IOU top-1 and of $32.71\%$ for the top-3.**

## I. Introduction

Satellite images have become an essential investigative tool in many journalistic analyses. Whether verifying the authenticity of facts, reporting on conflict zones, or reconstructing the location of a specific event due to partial video evidence leaked online, satellite imagery is widely adopted in newsrooms. The BBC was one of the first to use satellite images to report on the internment camp system used to imprison Muslim minorities in China's Xinjiang region in 2017 [1]. As soon as the first evidence of these events became public, the regime immediately censored all relevant documents like social media posts. The story would have been quickly suppressed if it were not for satellite imagery. Since then, many other analyses have made it possible to reconstruct where several events like brutal murders took place [2] thanks to the *cross-view* matching between ground evidence and aerial images. Similar considerations could be made on the current war in Ukraine, where several crimes and political stories have been spread.

In this paper, we examine the problem of automatic ground-to-aerial viewpoint localization from a *forensic perspective*. Although it is a very active area of research, many state-of-the-art techniques involve end-to-end deep learning-based methodologies that lack explainability. The lack of these characteristics represents a problem for media verification and authentication, as the matches obtained must be supported by clear and justifiable evidence [3]. Given these requirements, in this paper, we propose to address the problem based on graph matching as shown in Figure 1. First, the analyzed images are segmented to identify the objects of interest present within them. Subsequently, each entity becomes a node in a graph connected to the others based on a covisibility window. This approach offers a complete automation of the matching process without sacrificing the interpretation of the results, thus offering an advantage over deep learning-based models, which, although they have proven to be very accurate, are challenging to unfold and, therefore, hardly usable in investigations [4].

(a) Ground-view graph

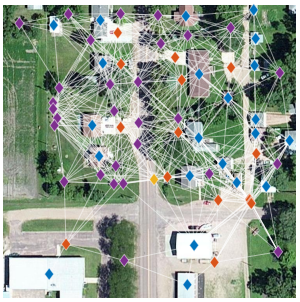

(b) Aerial-view graph

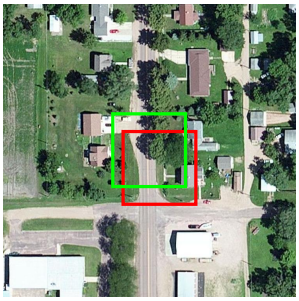

(c) Final matching

Fig. 1. An example application of our method. Figures 1-(a) and (b) depict the aerial and ground-level images, respectively. Our proposed method extracts a graph representation of the two images and matches the graphs to localize the ground view over the aerial one. In (c), we show our predicted location (in red) and the corresponding ground truth (in green).

Generally speaking, the matching between images taken from an overhead point of view and the ground level is a fundamental task in computer vision applications due to the high amount of available information that can be extracted. Before tackling the ground-to-aerial problem, researchers focused their attention on ground-to-ground image matching. Hays et al. [5] proposed the first data-driven method that sorted out the problem of geo-localization from ground-level images. However, this solution relied more on scene categorization rather than localization retrieval. Another typical technique to comprehend relationships among images for data collection and geolocalization is based on 3D reconstruction [6]–[8] and geometric constraints, both in urban and natural environments. Baatz et al. [9] focused on mountainous areas to pull out the recognition of the skyline given a digital elevation model of a country. Following this concept, Lin et al. [10] proposed the first approach using ground and aerial images to retrieve

979-8-3503-3959-8/23/$31.00 ©2023 IEEE

geolocalization via a data-driven approach. Satellite images are now much more widespread and cover every region of the planet, offering a conspicuous advantage over terrestrial photos, which can be more challenging to collect. This offers a substantial advantage in forensic applications.

Recently, with the advent of deep learning, AI-based algorithms have been exploited to improve the matching performances of previous methods. Three are the main architectural structures employed to handle the task: (i) the siamese-like networks, (ii) the generative adversarial networks (GANs) and (iii) transformer networks. The goal of siamese-like architectures is to extract shared features between the ground point of view and the overhead one. Subsequently, the distance between extracted features is computed in order to understand if there is a relationship between the two or if there are some features that are immune to the significant shift in the point of view. Lin et al. [11] are the first to introduce the Where-CNN, a siamese network that achieves superior results when compared to traditional hand-crafted features. Subsequently, Vo and Hays [12] and Shi et al. [13] focused on recovering orientation besides location using a soft-margin triplet loss on top of a Siamese CNN network and a siamese-like network relying on polar coordinate mapping, respectively. Whereas, Hu et al. [14] insert the NetVLAD [15] layer in addition to a Siamese network to identify features that are consistent with large changes in perspective. Shi et al. [16] aims at transforming the features from the ground domain to the aerial one utilizing a novel feature transport module together with a Siamese network. Other methods exploit GANs to synthesize images related to the two viewpoints and use them as additional information useful for obtaining a better understanding of the scene. Deng et al. [17] propose a GAN-based methodology in which the generator produces a ground-level image from the aerial point of view that is then compared with the real ground query image to retrieve the matching. On the other hand, Regmi et al. [18] aim at generating an aerial viewpoint of a ground-level panorama query image in a way that the transformed image has scene representations similar to the images it is matched against. It is also worth mentioning more recent deep learning-based approaches that rely on vision transformers (ViT). For example, Tian et al. [19] suggest a conditional GAN combined with a Transformer to synthesize an aerial image that appears with the same style as the ground view. Whereas, Zhang et al. [20] work towards the use of limited field-of-view images captured from more common devices such as smartphones and digital cameras instead of panoramic ground images, grasping sequential spatiotemporal features via the implementation of a VGG16 and temporal feature aggregation module inspired by the ViT architecture. Even if these deep learning-based solutions achieve excellent performance, they are still difficult to explain and, therefore, difficult to apply in a forensic scenario in which it is necessary to justify the output of the analysis. Moreover, state-of-the-art deep learning methods are designed in a supervised setting, while here, we propose an unsupervised approach that leverages the view-independent adjacency properties of visible landmarks to create comparable graph structures.

Differently from the previous methods that require high computational costs for explicit feature extraction, we model this problem via a probabilistic framework that matches graph representations of the image obtained from connecting the salient objects of the image.

## II. Methodology

In this section, we present our pipeline for ground-to-aerial viewpoint localization. Our pipeline, shown in Figure 2, automates the entire task without sacrificing model explicability, making the solution applicable in forensic investigations. We achieve this goal by building upon the previous landmarks graph matching technique introduced in Verde et al. [21]. Differently from this paper, which required an extensive human labeling process, our proposed methodology automatically extracts the graph representation of the image. This graph representation has the advantage of being able to take into account image features that are not dependent on the viewpoint, therefore preserving the adjacency relationships among objects no matter the angle or point of view from where the image is taken. From here on, we introduce the stages of our pipeline, which is composed of four stages: *Stage-A* segments the image extracting the objects of interest in the satellite and panoramic images, *Stage-B* identifies the nodes that will compose the graph, *Stage-C* is used to generate the graph connecting all the nodes, and *Stage-D* computes the matches between the two-view images.

***Stage-A*: Semantic Segmentation.** Our pipeline takes as input a panoramic ground viewpoint image and a wide aerial perspective from a satellite photograph. These images are initially processed with common AI-based semantic segmentation [23] techniques in order to extract a labeled image where each pixel corresponds to a given class. We select this approach with the aim of building a pipeline that is automated without the need to extract the significant landmarks a priori. To demonstrate the effectiveness of the proposed pipeline, we segment the image with respect to generic entities that can appear in any viewpoint image, such as *buildings*, *pavements*, *roads*, and *trees*. Figure 2 reports an example of the output of the semantic segmentation task, both for the satellite and the panoramic ground image. We report each label in the image with a different color, that is, (1) buildings are colored in blue, (2) pavements are depicted in light blue, (3) roads are yellow, and (4) trees are pictured in green. All the irrelevant elements are blacked out in the segmented image.

**Stage-B: Node Generator.** Once the objects in the images have been labeled, we want to convert the segmented image into a graph. In this stage, we represent each object in the image as a separate labeled node. To do so, the segmented image obtained from the previous step is filtered to handle one class at a time. Then, we convert the segmented objects into a binary mask where white pixels represent the relevant class, and black pixels depict the background, as illustrated in Figure 3. The binary images are finally analyzed by the Spaghetti labeling algorithm introduced in Bolelli et al. [24] using 4-way connectivity. An example of this process is shown in Figures 3(a)-(d). Each color represents a group of pixels belonging to the same connected component, i.e., different objects within a class are depicted in different colors. The same framework is applied to both aerial and ground images. The last step of this phase is the extraction of the centroid of each of the connected components. This point will represent the coordinates of the node.

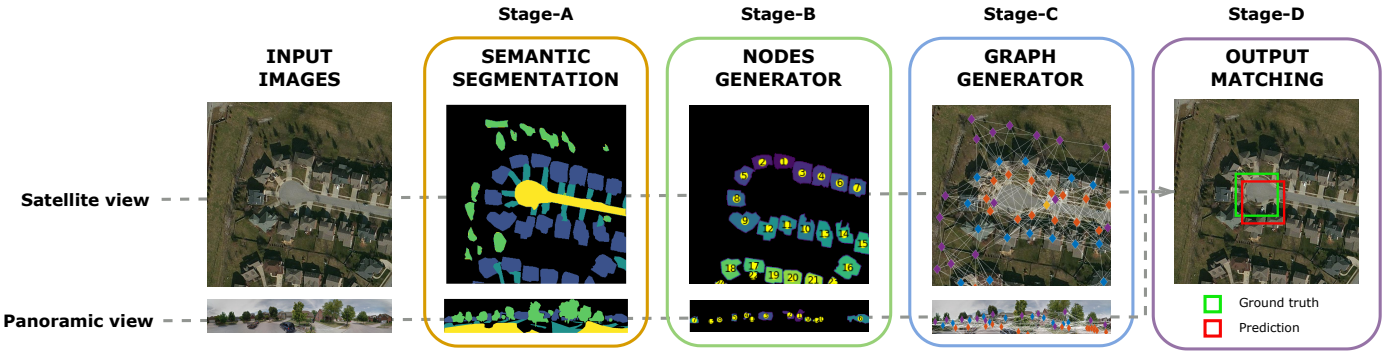


Fig. 2. Overview of our proposed pipeline based on semantic graph for ground-to-aerial image matching. **Stage-A** depicts the semantic segmentation operation applied to the images taken from CVUSA dataset [22], both to a panoramic image and the corresponding aerial view. **Stage-B** represents the connected components algorithm and the corresponding nodes generator. **Stage-C** shows the graph generator for both panoramic and overhead images. Finally, **Stage-D** represents the final matching retrieval between the two points of view.

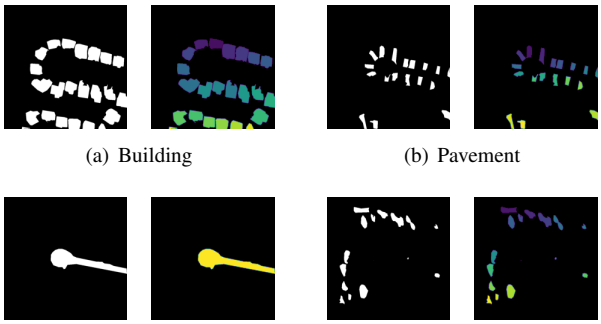

(a) Building (b) Pavement

(c) Road (d) Tree

Fig. 3. Example of connected component labeling algorithm applied to the binary mask (left side) of each label of a satellite image, i.e., building (a), pavement (b), road (c), and tree (d). On the right side of each label is illustrated the output of the connected component labeling algorithm: each color represents a connected component, i.e., different instances within a class are depicted in different colors.

**Stage-C: Graph Generator.** Now that nodes have been created, this stage connects them based on a *covisibility window* using the strategy proposed by Verde et al. [21]. The landmark graph is obtained by sliding a window through the image with a stride of one pixel. If two or more nodes are detected within the covisibility window, they are connected and stored as a *clique*. The covisibility window size is variable and adjusted based on the distance of the nodes with the aim of obtaining a connected graph, both for the aerial and the ground image. The output of this stage will be a labeled graph that represents all the spatial connections between the objects in the image.

**Stage-D: Graph Matching.** At this stage, we can finally compare the satellite image with the ground image by matching their graphs. Once the pictures are matched, the result gives an estimate of the location of the query ground panoramic image in the aerial view. The matching of the graphs is conducted by considering just cliques of the previous Stage-C that are considered *relevant* [21]; that is, the same clique in the aerial view must at least once occur on the query image. Given the drastic viewpoint change and possible element occlusions in the query image, traditional keypoint matching techniques would not work in this scenario, as some nodes visible in the satellite image may not be present in the ground image. Consequently, we cast the problem to a probabilistic framework by obtaining a collection of candidate locations $\mathcal{L}_x$ that represent different subgraphs of the satellite image and corresponding possible matches in the ground image $\mathcal{Z}$. Following Stumm et al. [25], we assume that the sparse normalized cross-correlation between location adjacency matrices represents the observation likelihood $P(\mathcal{Z}|\mathcal{L})$ of the ground query $\mathcal{Z}$ given a location $\mathcal{L}$ in the satellite image. Denoting the class adjacency matrix between classes $i$ and $j$ in $\mathcal{Z}$ and $\mathcal{L}$ by $W_{ij}^{\mathcal{Z}}$ and $W_{ij}^{\mathcal{L}}$, the likelihood is computed as shown in Equation 1.

$$P(\mathcal{Z}|\mathcal{L}) = \frac{\sum_{ij} W_{ij}^{\mathcal{Z}} \cdot W_{ij}^{\mathcal{L}}}{\sqrt{\sum_{ij} (W_{ij}^{\mathcal{Z}})^2 \cdot \sum_{ij} (W_{ij}^{\mathcal{L}})^2}} \tag{1}$$

At this point, we can apply Bayes' rule to derive the posterior probability of being in a location given the observation as follows in Equation 2.

$$P(\mathcal{L}_x|\mathcal{Z}) = \frac{P(\mathcal{Z}|\mathcal{L}_x)P(\mathcal{L}_x)}{P(\mathcal{Z}|\mathcal{L}_x)P(\mathcal{L}_x) + P(\mathcal{Z}|\overline{\mathcal{L}}_x)P(\overline{\mathcal{L}}_x)} \tag{2}$$

In conclusion, the candidate location that satisfies the maximum a posteriori (MAP) criterion in Equation 3 is the best possible matching for the present ground query.

$$\mathcal{L}_{MAP} = argmax_{\mathcal{L}_x} P(\mathcal{L}_x|\mathcal{Z}) \tag{3}$$

The output of the prediction of the matching pipeline is the covisibility window corresponding to the best candidate location $\mathcal{L}_{MAP}$ (i.e., the green square in Stage-D of Figure 2) to be compared to the ground truth (i.e., the red square in Stage-D of Figure 2). As said, the dimension of the covisibility window is variable, and as a consequence the dimension of the bounding box of the prediction will vary accordingly. This aspect will be deepened in Section III.

## III. Results

In this section, we report our experiments and implementation details. All the pipeline has been implemented in Python. In particular, we use the OpenCV[1] and NetworkX [26] libraries for extracting the connected components and generating the

[1] https://opencv.org/

nodes corresponding to each object as described in Section II. We evaluate our pipeline on the CVUSA [22] dataset, which contains more than $44k$ image pairs taken from the aerial viewpoint at a resolution of $750 \times 750$ and from ground level at a resolution of $1232 \times 224$. The dataset is assembled by downloading images depicting locations in the United States from Google Street View and Flickr. In all our experiments, we evaluate viewpoint localization in terms of Intersection Over Union (IOU), which measures the alignment between the ground truth and predicted bounding boxes.

As mentioned in Section II, we use a variable size covisibility window as in Verde et al. [21]; we evaluate our method with two other window sizes: $128 \times 128$, and $256 \times 256$. The ground truth bounding boxes are always placed at the center of the satellite image because of how the dataset was assembled.

**Experiment 1.** The first experiment investigates the trade-off between the variable window and a fixed strategy. We report, in Table I, the top-1 and top-3 mean IOU with respect to the estimated windows. As can be seen from the reported values, the proposed solution introduces a $96.56\%$ gain in terms of mean top-1 IOU when compared to the $256\times256$ fixed window and $390.37\%$ when compared to the $128\times128$ one. Moreover, on the same settings, we achieve a top-3 accuracy improvement of $105.33\%$ with respect to the fixed window $256 \times 256$ and $483.17\%$ for the $128 \times 128$ window. Based on those findings, we assume that the use of a variable size window allows the pipeline to better generalize over different scenarios and obtain more accurate matching.

TABLE I. PERFORMANCES OBTAINED WITH THE PROPOSED METHOD VARYING THE WINDOWS DIMENSION.

| Ground truth size | IOU top-1 | IOU top-3 |
|---|---|---|
| 128x128 | 0.0948 | 0.0766 |
| 256x256 | 0.2366 | 0.2175 |
| variable window | **0.4650** | **0.4467** |

**Experiment 2.** In the second experiment, we compare our proposed automated pipeline with respect to the interpretable method introduced by Verte et al. over the same test set considered in their paper, which is composed of 15 images extracted from the CVUSA dataset. We report in Table II the obtained results over different methodologies and window sizes. Based on the reported values, the proposed pipeline achieves an improvement on the top-1 and top-3 accuracy equal to $17.84\%$ and $32.71\%$, respectively, with the variable window configuration. We can also notice that the proposed method obtains slightly worse results with respect to Verde et al. in the case of fixed windows. However, our method is fully automated, and probably extracts better graph representations that lead to the best matching accuracy than the other method when it is considered a variable configuration for the covisibility window.

TABLE II. QUANTITATIVE EVALUATION OF THE PROPOSED AUTOMATED PIPELINE WITH RESPECT TO VERDE ET AL. THE BEST RESULTS ARE IN BOLD.

| Ground truth size | Verde et al. [21] | | Proposed method | |
|---|---|---|---|---|
| | IOU top-1 | IOU top-3 | IOU top-1 | IOU top-3 |
| 128x128 | 0.2215 | 0.1058 | 0.1503 | 0.0527 |
| 256x256 | 0.2789 | 0.2057 | 0.2535 | 0.1147 |
| variable window | 0.2724 | 0.2305 | **0.3210** | **0.3059** |

**Experiment 3.** Our final experiment analyzes the robustness of the proposed pipeline with respect to every single object class for estimating their importance when constructing and matching the graphs. We report a graphical comparison in Figure 4, where the blue line shows the accuracy trend of our method while the dotted one reports the performances of Verde et al. [21] while evaluating the 15 image subset from CVUSA with the variable window setting. Figure 4 also reports the standard deviation of our method (in purple) compared to the one from Verde et al. (in red). Based on the reported values and previous findings, we can notice that our method outperforms the baseline when all four classes are present. Moreover, this experiment measures the impact of removing one single class at a time. From the figure, we can observe that removing the *pavement* or *tree* classes can lead to a performance gain with respect to operating on all four classes. Therefore, when using a set of three classes, our proposed pipeline achieves an average IOU improvement equal to $22.15\%$ with respect to the method from Verde et al.

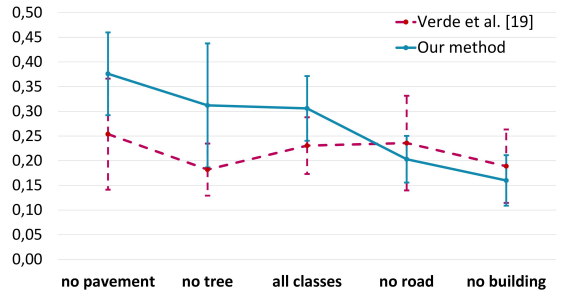


Fig. 4. Change in IOU in terms of mean and standard deviation when considering a subset of only three classes. As you can see, removing floors and trees improves performance.

## IV. CONCLUSION

In this paper, we introduced a fully automated pipeline for the ground to the aerial viewpoint localization. The proposed method allows for easy interpretation with respect to fully deep learning-based methodologies, making our strategy suitable for forensic investigations and media authenticity verification. Our method is fully automated compared to the work of Verde et al. [21], and to obtain an average IOU improvement of $25.3\%$. In future works, we plan to employ and study the explainability of graph neural networks to solve this task with a trainable but interpretable method. Finally, we plan to extend our study to different datasets.

## ACKNOWLEDGMENT

This study has been partially supported by SERICS (PE00000014) under the MUR National Recovery and Resilience Plan funded by the European Union - NextGenerationEU and Sapienza University of Rome project 2022–2024 “EV2” (003_009_22).

## REFERENCES

[1] A. Datta. (2023) Satellite imagery is taking journalism to new orbits. [Online]. Available: https://www.geospatialworld.net/prime/special-features/taking-journalism-to-new-orbits/

[2] F. Architecture. (2023) Geolocation: Forensic architecture. [Online]. Available: https://forensic-architecture.org/methodology/geolocation

[3] I. Amerini, A. Anagnostopoulos, L. Maiano, and L. R. Celsi, "Deep learning for multimedia forensics," *Foundations and Trends® in Computer Graphics and Vision*, vol. 12, no. 4, pp. 309–457, 2021. [Online]. Available: http://dx.doi.org/10.1561/0600000096

[4] S. W. Hall, A. Sakzad, and K.-K. R. Choo, "Explainable artificial intelligence for digital forensics," *WIREs Forensic Science*, vol. 4, no. 2, p. e1434, 2022. [Online]. Available: https://wires.onlinelibrary.wiley.com/doi/abs/10.1002/wfs2.1434

[5] J. Hays and A. A. Efros, "im2gps: estimating geographic information from a single image," in *Proceedings of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2008.

[6] D. M. Chen, G. Baatz, K. Köser, S. S. Tsai, R. Vedantham, T. Pylvänäinen, K. Roimela, X. Chen, J. Bach, M. Pollefeys, B. Girod, and R. Grzeszczuk, "City-scale landmark identification on mobile devices," in *CVPR 2011*, 2011, pp. 737–744.

[7] S. Agarwal, N. Snavely, I. Simon, S. M. Seitz, and R. Szeliski, "Building rome in a day," in *2009 IEEE 12th International Conference on Computer Vision*, 2009, pp. 72–79.

[8] Q. Shan, C. Wu, B. Curless, Y. Furukawa, C. Hernandez, and S. M. Seitz, "Accurate geo-registration by ground-to-aerial image matching," in *2014 2nd International Conference on 3D Vision*, vol. 1, 2014, pp. 525–532.

[9] G. Baatz, O. Saurer, K. Köser, and M. Pollefeys, "Large scale visual geo-localization of images in mountainous terrain," in *Computer Vision – ECCV 2012*, A. Fitzgibbon, S. Lazebnik, P. Perona, Y. Sato, and C. Schmid, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 517–530.

[10] T.-Y. Lin, S. Belongie, and J. Hays, "Cross-view image geolocalization," in *2013 IEEE Conference on Computer Vision and Pattern Recognition*, 2013, pp. 891–898.

[11] T.-Y. Lin, Y. Cui, S. Belongie, and J. Hays, "Learning deep representations for ground-to-aerial geolocalization," in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 5007–5015.

[12] N. N. Vo and J. Hays, "Localizing and orienting street views using overhead imagery," in *Computer Vision – ECCV 2016*, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds. Cham: Springer International Publishing, 2016, pp. 494–509.

[13] Y. Shi, X. Yu, D. Campbell, and H. Li, "Where am i looking at? joint location and orientation estimation by cross-view matching," in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 4063–4071.

[14] S. Hu, M. Feng, R. M. H. Nguyen, and G. H. Lee, "Cvm-net: Cross-view matching network for image-based ground-to-aerial geo-localization," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 7258–7267.

[15] R. Arandjelovic, P. Gronat, A. Torii, T. Pajdla, and J. Sivic, "Netvlad: Cnn architecture for weakly supervised place recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 5297–5307.

[16] Y. Shi, X. Yu, L. Liu, T. Zhang, and H. Li, "Optimal feature transport for cross-view image geo-localization," in *arXiv preprint arXiv:1907.05021*, 2019.

[17] X. Deng, Y. Zhu, and S. Newsam, "Using conditional generative adversarial networks to generate ground-level views from overhead imagery," *ArXiv*, vol. abs/1902.06923, 2019.

[18] K. Regmi and M. Shah, "Bridging the domain gap for ground-to-aerial image matching," in *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 470–479.

[19] X. Tian, J. Shao, D. Ouyang, A. Zhu, and F. Chen, "Smdt: Cross-view geo-localization with image alignment and transformer," in *2022 IEEE International Conference on Multimedia and Expo (ICME)*, 2022, pp. 1–6.

[20] X. Zhang, W. Sultani, and S. Wshah, "Cross-view image sequence geo-localization," in *2023 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2023, pp. 2913–2922.

[21] S. Verde, T. Resek, S. Milani, and A. Rocha, "Ground-to-aerial viewpoint localization via landmark graphs matching," *IEEE Signal Processing Letters*, vol. 27, pp. 1490–1494, 2020.

[22] S. Workman, R. Souvenir, and N. Jacobs, "Wide-area image geolocalization with aerial reference imagery," in *IEEE International Conference on Computer Vision (ICCV)*, 2015, pp. 1–9, acceptance rate: 30.3%.

[23] E. Xie, W. Wang, Z. Yu, A. Anandkumar, J. M. Alvarez, and P. Luo, "Segformer: Simple and efficient design for semantic segmentation with transformers," *Advances in Neural Information Processing Systems*, vol. 34, pp. 12 077–12 090, 2021.

[24] F. Bolelli, S. Allegretti, L. Baraldi, and C. Grana, "Spaghetti labeling: Directed acyclic graphs for block-based connected components labeling," *IEEE Transactions on Image Processing*, vol. 29, pp. 1999–2012, 2020.

[25] E. Stumm, C. Mei, S. Lacroix, and M. Chli, "Location graphs for visual place recognition," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 5475–5480.

[26] A. A. Hagberg, D. A. Schult, and P. J. Swart, "Exploring network structure, dynamics, and function using networkx," in *Proceedings of the 7th Python in Science Conference*, G. Varoquaux, T. Vaught, and J. Millman, Eds., Pasadena, CA USA, 2008, pp. 11 – 15.