

A Deep Learning Ensemble Framework for Off-Nadir Geocentric Pose Prediction

Christopher Sun^{*1}, Jai Sharma¹, and Milind Maiti¹

¹Monta Vista High School

Abstract—Roughly 6,800 natural disasters occur annually, and this alarming number continues to grow due to climate change. Effective methods to improve natural disaster response include change detection, map alignment, and vision-aided navigation to allow for the time-efficient delivery of life-saving aid. Current software functions optimally only on near-nadir images taken around ninety degrees above ground level. The inability to generalize to oblique images increases the need to compute an image’s geocentric pose, which is its spatial orientation with respect to gravity. This Deep Learning investigation presents three convolutional models to predict geocentric pose using 5,923 nadir and off-nadir red, green, and blue (RGB) satellite images of cities worldwide. Prior to the ensemble, an autoencoder is assembled to condense the $256 \times 256 \times 3$ images to $32 \times 32 \times 16$ latent space representations, demonstrating the ability to learn useful features from the data. The first model in the ensemble is a U-Net Fully Convolutional Network (FCN) with skip connections used to predict each image’s corresponding pixel-level above-ground elevation mask. This FCN achieves a median absolute deviation of 0.335 meters and an R^2 of 0.865 on test data. Afterward, the elevation masks are concatenated with the RGB images to form four-channel inputs fed into the second model in the ensemble, which predicts each image’s rotation angle and magnification scale, the components of its geocentric pose. This Deep Convolutional Model achieves an R^2 of 0.943 on test data, significantly outperforming previous models designed by researchers. In addition to achieving superior model performance, outlier removal was performed through supervised interpolation, and the usefulness of data features was gauged through a sensitivity analysis of elevation masks, to target future avenues of feature engineering. The high-accuracy software built in this study contributes to mapping and navigation procedures to accelerate disaster relief and save human lives.

Index Terms—geocentric pose, convolutional neural network, ensemble learning

I. INTRODUCTION

With the number of natural disasters on the rise, there is an urgent need for land management technologies that provide disaster relief [9]. This growing frequency of natural disasters is paralleled by a rapid increase in the availability of large satellite imagery data sets, which has introduced the challenge of big data interpretation for practical disaster relief applications such as cartography [5]. For unmanned aerial vehicles (UAVs) to effectively respond to natural disasters, they must have adaptable and versatile software that perform a wide array of necessary functionalities, such as detecting objects, monitoring changes on land, and mapping surroundings

to allow for communication with other UAVs and humans. While researchers such as Mishra et al. [4] have implemented object detection algorithms to assist in search and rescue operations, there has been less progress in the tasks of vision-aided navigation and map alignment for disaster relief. The core technology behind these tasks is finding the geocentric pose of an aerial vehicle. Geocentric pose consists of three components: above-ground elevation, angle of orientation with respect to gravity, and scale between the actual and apparent sizes of objects on the ground. Currently, geocentric pose is easiest to predict when UAVs take near-nadir images from around ninety degree angles, almost directly above the ground. However, the geocentric pose of oblique, non-nadir images is harder to predict, creating the need for computational models to provide a more efficient, reliable, and accurate solution.

There has been little research conducted on satellite-aided geocentric pose detection using deep learning models. One of the few works that utilized this approach is that of Christie et al., who created an encoding for Geocentric Pose using monocular satellite imagery, attempting to rectify object parallax of monocular images and improving the accuracy of object localization for Earth observation tasks. Christie et al. present the idea of learning Geocentric Pose using oblique monocular images and state how this research can be expanded to other applications such as airborne or ground-based cameras [1] [2]. Shortcomings of this work include the accuracy of geocentric pose prediction.

Hence, we expand on this prior work by using a novel ensemble framework to achieve higher performance. Our research objective was to extract features from near-nadir and oblique RGB satellite images to predict each image’s geocentric pose using an ensemble methodology that consists of two convolutional models. First, a fully convolutional U-Net Elevation Model is used to predict the pixel-wise elevations of RGB satellite images. Then, a Geocentric Pose Model is used to predict the scale and angle of the satellite given the RGB image as well as the predicted pixel-wise elevations (Figure 1).

II. MATERIALS AND METHODS

A. Data Used For Experimentation

The data used for this research was part of the Urban Semantic 3D Data Set, publicly available from the IEEE DataPort [3]. The data set contained 5,923 RGB images, their pixel-wise above-ground elevations, and their measurements of scale and angle. The images were taken from four locations in North and South America (Table I).

Elevation was given in centimeters, scale was given in pixels per centimeter, and angle was given in radians. Due to computational constraints, we resized the images from $2048 \times 2048 \times 3$ pixels to $256 \times 256 \times 3$.

B. Outlier Removal

The data set contained some elevation masks with invalid pixel values. For example, some pixel values were improbably large, higher than the tallest building in the city the data was acquired from. The solution to this problem was location-based thresholding, which involved finding pixel value thresholds for each city included in the data set. We found that images from Jacksonville and Omaha did not contain any outliers. For Atlanta images, we counted the number of pixels N that exceeded 4,000 centimeters for each image. If N exceeded 100, we determined that the image indeed contained a building taller than 4,000 centimeters. If N was less than 100, we interpolated all said pixels with the median elevation in the image. For San Fernando images, the outlier threshold was a hard cap at 3,000 centimeters, because the landscape was fairly even and did not contain tall buildings.

One limitation of location-based thresholding was that thresholds were experimentally determined, since there was no objective way to best deal with the outliers in the data set. In particular, the challenge was discriminating between large pixel values that were valid because of the presence of a tall building and those which were invalid.

C. Autoencoder Model

Before embarking on the ensemble model framework, we trained an autoencoder on the RGB images to gauge the practicality of our deep learning approach (Figure 2). Our reasoning was that if meaningful latent space representations could be extracted from the RGB images, the other convolutional models would likewise be able to extract useful features to predict geocentric pose.

The encoder used convolutional operations to condense the input images to matrices of shape $32 \times 32 \times 16$. These intermediate matrices were then flattened to a vector of length 16384, after which a dense layer was used to arrive at

the latent space representation of shape 512×1 . Finally, Multidimensional Scaling (MDS) was used to visualize these high-dimensional latent space representations on a 2D plane. An inverse architecture of the encoder was used to decode the latent space representations back into images of shape $256 \times 256 \times 3$. The autoencoder, made by combining the encoder and decoder into one model, was trained with the goal of retrieving decoded images as similar to the original input images as possible. The accomplishment of this task would confirm the existence of useful features that can be extracted using deep learning techniques.

D. U-Net Elevation Model

We trained a fully convolutional neural network to predict the elevation mask of each RGB image. The model was based on the U-Net architecture designed by Sun (2022) for a wildfire classification task, but included modifications such as having more convolutional blocks, and as a result, more skip connections between downsampling and upsampling layers [8] (Figure 3). The filter size was 5×5 throughout the model. The first convolution involved 32 filters, and this number was doubled for each layer during downsampling and halved for each layer during upsampling. Batch Normalization was employed between convolution layers for a regularization effect, and the ReLU activation function was used throughout the model. The model was configured with an 85-15% train-validation split and trained for around 700 epochs using Adam optimization, a constant learning rate of 0.001, and a mini-batch size of 32 images. Images with Not a Number (NaN) values were excluded from the training data set.

The U-Net Elevation Model was used to predict each RGB image's elevation mask containing pixel-wise heights above ground.

E. Geocentric Pose Model

The second model in the ensemble framework was the Geocentric Pose Model, a convolutional neural network with a series of five double convolution layers and maximum pooling in between each series (Figure 4). The resulting $7 \times 7 \times 256$ shape was flattened and fed through six fully connected layers to generate the predictions. Batch Normalization and the ReLU activation function were employed throughout the model. The model was configured with an 80-20% train-validation split and was trained for around 600 epochs using Adam optimization, a constant learning rate of 0.001, and a mini-batch size of 32 images.

After elevation mask predictions were obtained from the U-Net Elevation Model, these predicted masks were concatenated with the original RGB images along the channels axis to produce four-channel images. These images were then fed as inputs to the Geocentric Pose Model for the prediction

of scale and angle. Together, the U-Net Elevation Model and the Geocentric Pose Model predicted all three components of geocentric pose.

III. RESULTS

A. Autoencoder: Latent Space Representations¹

a) *Visualization for Scale*: No clear association was observed from the visualization of scale encodings (Figure 5). This may have been due to the decrease in variance that comes with dimensionality reduction. Still, the autoencoder was able to reconstruct the original images to a large extent.

b) *Visualization for Angle*: Contrary to the results of the scale encodings, the encoding visualizations for the angle showed a clearer association. Specifically, images with higher angles were aggregated spatially to the left while images with lower angles were aggregated spatially to the right (Figure 5). The significance of this result was twofold. First, this visualization demonstrated that there is a relationship between an RGB image and its angle. Second, the workflow implemented revealed that this said relationship can be learned through a convolutional approach. Hence, we considered it feasible to use deep learning techniques to learn this relationship and predict the geocentric pose of RGB images.

B. U-Net Elevation Model Performance

The models in the ensemble framework were evaluated by the R^2 metric, defined as

$$R^2 = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2},$$

where n is the number of values, y_1 is the i th true value, $\hat{y}_1$ is the i th predicted value, and $\bar{y}$ is the mean of all true values.

The U-Net Elevation Model achieved an R^2 metric of 0.926 on the training data set and 0.865 on the validation data set. The model suffered from minimal overfitting and had a high degree of generalizability to unseen data, which was both numerically and visually confirmed (Figure 6).

1) *Interpolation of Invalid Data*: Recall that images containing NaN values were excluded from training. This was to allow NaN pixels to be replaced by predicted pixel heights after the U-Net Elevation Model converged on valid data. We performed interpolation on all NaN values, which served as a data cleansing and quality control mechanism that allowed all 5,923 images to be used for the Geocentric Pose Model.

¹It must be noted that these visualizations do not capture all nuances of the high dimensional data, so they are used as proof of concept.

2) *Elevation Mask Visualizations*: Notice the similarity between the true and predicted elevation masks (Figure 6). Furthermore, notice how the third example of invalid data displayed above contains a small yellow patch of NaN values near the top left of the elevation mask. For images like this, interpolation was only needed for a small number of pixels. On the other hand, for images containing mostly NaN values, such as the fourth example above, most of the elevation mask was reconstructed using the predictions of the U-Net Elevation Model.

C. Geocentric Pose Model Performance

The Geocentric Pose Model achieved an R^2 metric of 0.991 on the training data set and 0.943 on the validation data set, weighed equally between scale and angle. The following conclusions are derived from Figure 7.

a) *True vs. Predicted Value*: The predicted scale and angle are plotted against the true scale and angle, with linear regression coefficients of 0.995 and 0.994, respectively.

b) *Performance Distributions*: The distributions of the true values and predicted values have similar shape and spread for both scale and angle, showing the model's ability to capture trends between the RGB-Elevation inputs and geocentric pose.

c) *Distribution of Absolute Error Loss*: The distributions of absolute error loss for both scale and angle are heavily skewed right, which shows that the model rarely makes predictions that are faulty to a large degree. The distribution for angle is more skewed than that for scale.

d) *Sorted Absolute Error*: The graphs of sorted absolute error for both scale and angle take a skewed shape. On average, the model was able to predict the scale to within 0.386 pixels per decameter, and the angle to within 0.079 radians ($\approx 4.5^\circ$).

The final metrics for the geocentric pose ensemble framework are summarized in Table II.

IV. DISCUSSION

A. Insights from Deep Learning Models

a) *Autoencoder Model*: The latent space representations of the RGB images were twelve times smaller in storage than the original images, yet preserved important features within the data. A visualization of these encodings through MDS revealed a learnable correlation between an image and its angle of capture. Though no clear correlation was determined when observing scale encodings, possibly due to low variance retained when reducing dimensionality, the performance of the autoencoder motivates future avenue of exploration, which involve utilizing autoencoders as feature extractors that can feed

high-level characteristics of RGB images into the ensemble framework.

b) Geocentric Pose Ensemble Model: After the Geocentric Pose Model achieved an optimal performance, we sought to understand the usefulness of the features extracted by the U-Net Elevation Model. We proceeded to test the degree of importance of the elevation masks in regards to geocentric pose prediction by training three different version of the Geocentric Pose Model, as follows.

- 1) **RGB-Only Model:** The RGB-Only Model received only the RGB images as inputs and acted as a negative control.
- 2) **RGB-Elevation Model** The RGB-Elevation Model was the standard Geocentric Pose Model.
- 3) **Elevation-Only Model** The Elevation-Only Model received only the elevation masks as inputs and acted as a baseline performance standard.

These three models were trained with the same hyperparameter specifications as described in Section II-E, with the exception of the input shape, which differed along the channels axis.

The R^2 of each of the three defined models is listed in Table III, with the models numbered according to the order they were introduced. Furthermore, the P -values between adjacent R^2 metrics calculated as a difference in proportions is shown. The results show that the presence of the elevation mask is integral for the Geocentric Pose Model to achieve high performance metrics. The R^2 of the RGB-Elevation Model exceeded that of the RGB-Only Model for both scale and angle with a statistically significant P -value. The R^2 of the RGB-Elevation Model exceeded that of the Elevation-Only Model for scale with a statistically significant P -value. However, the Elevation-Only Model surprisingly outperformed the RGB-Elevation Model in angle predictions, but by an extremely small margin.

The results of these models show that elevation masks are a crucial feature for the prediction of geocentric pose. Elevation masks could specifically be beneficial in predicting the angle of satellites, as knowledge of relative building heights can help in determining the orientation of satellites. Likewise, processing RGB images in conjunction with elevation masks could provide information about the scale, which relies on a measure of the conversion factor between pixels in the image and distances in the real world.

c) Limitations and Future Improvements: When analyzing specific predictions of low accuracy, we were surprised to find that the Geocentric Pose Model struggled the most for predicting the angle of some near-nadir images, likely because the height of buildings are difficult to perceive at these angles, making it difficult to distinguish between angles in this range.

As stated previously, computational constraints obligated us

to downsize the input images, which lowered the resolution, making it more difficult to predict geocentric pose. Instead, improvements can be made by employing random cropping, which would allow us to reduce the number of parameters while keeping the original resolution. Performance can further be enhanced through the use of image augmentation, particularly for cities such as Atlanta and Jacksonville, to increase the resilience of the model towards distortions of the image and to prevent the model from overfitting on the city with the most labeled data. Additionally, the ensemble framework may benefit from the usage of the Ground Sample Distance (GSD) feature included in the meta-data set.

B. Comparison to Prior Research

While Christie et al. achieved a cumulative R^2 metric of 0.799, the models described in this study achieved a cumulative R^2 metric of 0.943, albeit on different testing examples [2]. Nevertheless, this demonstrates a major increase in accuracy and reliability of geocentric pose predictions. We attribute this higher performance to our ensemble methodology and the data preprocessing that served as quality control.

V. CONCLUSION

With a reliable ensemble model to identify the geocentric pose of both near-nadir and oblique images, the versatility of software fueling natural disaster response can be enhanced. We have designed an accurate framework that has applications in disaster response through vision-aided navigation, change detection, map alignment, and search-and-rescue. When unmanned aerial vehicles learn the physical features of their surroundings, they can transfer information between each other to allow for specialized systems and more effectively summon human users when assistance is needed.

Need for Further Research: Future research should explore optimal outlier removal methods for this data set. Further research is also necessary to investigate how the model architectures described in this paper could be improved. Changing the size of the bottleneck layer in the U-Net Elevation Model or using two separate, simpler models to predict scale and angle are avenues that should be explored. In addition, further testing needs to be performed to understand how well both the Geocentric Pose and the U-Net Elevation Model generalize to new data from cities around the world, since different city architectures could pose a potential challenge for the models.

REFERENCES

- [1] Christie, G., Abujder, R. R. R. M., Foster, K., Hagstrom, S., Hager, G. D., & Brown, M. Z. (2020). Learning geocentric object pose in oblique monocular images. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 14512-14520).

- [2] Christie, G., Foster, K., Hagstrom, S., Hager, G. D., & Brown, M. Z. (2021). Single view geocentric pose in the wild. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 1162-1171).
- [3] Foster, K., Christie, G., & Brown, M. (2020). Urban Semantic 3D Dataset. IEEE Dataport. <https://dx.doi.org/10.21227/9frn-7208>
- [4] Mishra, B., Garg, D., Narang, P., & Mishra, V. (2020). Drone-surveillance for search and rescue in natural disaster. Computer Communications, 156, 1-10.
- [5] Muruganandham, S. (2016). Semantic Segmentation of Satellite Images using Deep Learning (thesis). Faculty of Electrical Engineering Department of Cybernetics.
- [6] Ronneberger, O., Fischer, P., & Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In International Conference on Medical image computing and computer-assisted intervention (pp. 234-241). Springer, Cham.
- [7] Sharma, J., Sun, C., & Maiti, M. (2022). Jaisharmz/geocentric-pose: Github repository for the research paper "A deep learning ensemble framework for off-nadir geocentric pose prediction." GitHub. Retrieved from <https://github.com/jaisharmz/geocentric-pose>
- [8] Sun, C. (2022). Analyzing Multispectral Satellite Imagery of South American Wildfires Using Deep Learning. In 2022 International Conference on Applied Artificial Intelligence (ICAPAI) (pp. 1-6). IEEE.
- [9] Van Aalst, M. K. (2006). The impacts of climate change on the risk of natural disasters. Disasters, 30(1), 5-18.

Data Availability Statement

The data used for this research was part of the Urban Semantic 3D Data Set, publicly available from the IEEE DataPort [3].

A. Additional Information

The authors have made the code for the research in this paper open source and available on GitHub (click) [7].

C.S, J.S, and M.M contributed equally to the research described in this manuscript and the preparation of the manuscript, including figures and tables. The authors declare no competing interests.

Figures




Fig. 1: Ensemble Model to predict Geocentric Pose.




Fig. 2: Shown above is the Autoencoder Model trained with 16,823,107 parameters, consisting of an encoder (left) and a decoder (right).




Fig. 3: The U-Net Elevation Model contained 23,966,337 parameters.




Fig. 4: The Geocentric Pose Model contained 7,076,370 parameters (Adapted from Sun (2022)).




Fig. 5: MDS was employed to visualize the latent spaces representations of the encoded images. Above, each point represents an encoded image, which is colored by the magnitude of its scale and angle, respectively.




Fig. 6: The elevation masks predicted by the U-Net model are compared to the true elevation masks (left). The U-Net interpolations are displayed next to their pre-interpolation counterparts containing large amounts of invalid pixel values.




Fig. 7: The above plots show the resemblance of true and predicted values as well as the distribution of error of the Geocentric Pose Model.

Tables

TABLE I: Data Acquisition

Location	No. of Images
San Fernando, AR	2325
Atlanta, USA	704
Jacksonville, USA	1098
Omaha, USA	1796

TABLE II: Geocentric Pose Ensemble Model Metrics

	U-Net Model		Geocentric Pose Model			
	Elevation (m)		Scale (px/dam)		Angle (rad)	
Metric	Train	Val	Train	Val	Train	Val
MAD	0.263	0.335	0.192	0.386	0.063	0.079
MAE	0.847	1.049	0.243	0.582	0.091	0.160
R ²	0.926	0.865	0.990	0.924	0.991	0.963

TABLE III: Geocentric Pose Model Varying Input Trials

	1) R ²	P-value	2) R ²	P-value	3) R ²
Scale	0.591	8.55E-80	0.924	2.24E-4	0.881
Angle	0.733	7.30E-55	0.963	4.84E-1 ²	0.963

²The Elevation-Only Model outperformed the RGB-Elevation Model in angle predictions.