Replicating Human Perception of Proximity: Applying Neural Networks in the Monitoring of Suspicious Air Activities in Amazon

Gabriel Dietzsch
Instituto Nacional de Pesquisas Espaciais
Avenida dos Astronautas, 1.758 - Jd. Granja
São José dos Campos - SP - Brasil

Elcio Hideiti Shiguemori
Instituto Nacional de Pesquisas Espaciais
Avenida dos Astronautas, 1.758 - Jd. Granja
São José dos Campos - SP - Brasil

Abstract

This study investigates whether a neural network can approximate human perception of proximity in the context of air traffic surveillance involving suspicious aircraft near regular and irregular airstrips. Using a DenseNet model trained on a synthetic dataset of graphical representations, the study evaluates the network's ability to classify visual proximity relationships without explicit distance computation. While the model achieved moderate performance (F1-Score of 78.4%), results were limited by overfitting and the low variability of the data. Larger datasets did not improve performance, suggesting the importance of visual diversity over quantity. These findings validate the feasibility of modeling human-like spatial reasoning through neural networks in controlled environments. The research establishes an experimental baseline for future studies involving more complex data and architectures, such as EfficientNet or Transformers, to further improve model generalization and practical applicability.

Introduction

The Amazon, with its extensive biodiversity and natural wealth, has been a target of various illicit activities, such as deforestation, illegal mining, and drug trafficking (Pajolla 2022). These unlawful practices are often enabled by the use of aircraft, which facilitate the swift transport of supplies and resources necessary to sustain such activities in remote and inaccessible regions (Potter 2023),(Dietzsch and Shiguemori 2024).

The susceptibility of the Amazon rainforest was a key focus at the Amazon Summit, convened in Belém do Pará in August 2023. Participants underscored the critical need to protect the region and curb organized crime, emphasizing the role of air traffic control in halting the spread of these unlawful activities (Artaxo 2023). The Brazilian government, specifically, put forward a proposal for an Integrated Air Traffic Control System aimed at countering illegal air traffic, drug trafficking, and related crimes.

Robust air traffic control is essential for the detection and prevention of criminal activities (Potter 2022). Currently, identifying suspicious flights depends on analyzing tabular and informational data, including flight plans and radar information, handled by human operators (For 2020). Nevertheless, human limitations and operational complexities contribute to a higher likelihood of errors and missed detections.

Suspicious air traffic frequently utilizes clandestine airstrips, presenting significant monitoring challenges. Given this context, this study proposes an investigation into the ability of a neural network to understand the concept of proximity with respect to the distance of an aircraft from regular and irregular airports. The approach will employ a synthetic dataset of graphics showing examples of aircraft near both regular and irregular airports. The research seeks to determine whether a neural network can, much like a human, interpret whether an aircraft is near an irregular airport simply by 'looking at' a graphical representation.

Despite the extensive use of neural networks in remote sensing and classification tasks, few studies have addressed the possibility of modeling perceptual concepts, such as spatial proximity, based solely on visual input. Existing approaches typically rely on structured numerical data or explicit geometric computations. This work seeks to fill that gap by proposing a visual learning task that approximates how human operators might intuitively interpret relative positions in simplified airspace representations.

The central question of this research is: "Can a neural network learn to approximate proximity judgments between graphical elements in a way that resembles human intuition?" The general objective is to determine whether it is possible to replicate a human expert's perception of near and far in comparing an aircraft's position relative to a regular and an irregular airport, identifying which one is closer through classification.

Accurate and efficient flight classification is crucial for strengthening security and combating illicit activities in the Amazon (Valente 2023). This study will contribute to the development of innovative technological solutions, enhancing the monitoring and response capabilities of the relevant authorities.

Neural networks are widely used for classification across various fields. They optimize control in oil production (Mikhaylov et al. 2023), recognize emotions in images with models like ResNet50, VGG16, and VGG19 (Matijašević, Samcovic, and ogatović 2022), classify road types in autonomous driving (Perić et al. 2023), and automate weed and seed classification in agriculture (Ganesan et al. 2023). These applications highlight their versatility and real-time effectiveness in decision-making and automation. In satellite communication traffic classification, Graph Convolutional Networks (GCN) and Graph Attention Networks (GAT) have proven more accurate than traditional MLP methods (Hao et al. 2023).

In aviation, machine learning techniques, including an adapted k-means for classifying aircraft by position and velocity (Vincent-Boulay and Marsden 2024), CNNs for aircraft model identification by noise characteristics (Morinaga, Mori, and Yamamoto 2023), and a combination of ResNet, GCN, and LSTM for airport flight flow prediction under environmental influences (Zang and Zhu 2022), demonstrate the impactful use of neural networks in air traffic management and flow forecasting.

Predicting human perceptual similarity remains a challenging area of research. This aspect of human vision is believed to involve various levels of visual analysis, including shapes, objects, texture, layout, color, and more (Rosenfeld, Zemel, and Tsotsos 2019).

This study aims not to compute distances explicitly but to assess whether a neural network can identify spatial proximity patterns from visual data in a manner consistent with human perception. While computing distances may appear straightforward, this study seeks to validate if neural networks can deduce visual proximity patterns in cases where geometrical arrangements may challenge human perception.

Method

While the problem may appear trivial in two-dimensional Euclidean space, human judgment is not always based on exact geometric calculation, especially in scenarios involving visual ambiguity, occlusion, or overlapping elements. This study intentionally uses simplified graphical representations to isolate the visual proximity factor and test whether neural networks can learn to classify such scenes in a way that aligns with human labeling—without relying on hard-coded distance formulas. This design allows us to test the network's ability to learn spatial reasoning from data, not rules.

This work was developed using a synthetically generated dataset of positional graphs. Each graph includes a symbol for an airplane, a symbol representing a regular airport, and a symbol representing an irregular airport. Initially, 1,000 examples were generated, containing the x, y position of the airplane symbol closest to the regular airport symbol (Figure 1), followed by 1,000 examples where the airplane symbol was closest to the irregular airport symbol (Figure 2).

To assess the possibility of replicating human judgment regarding the proximity of the graphical elements involved, DenseNet was used: each layer receives information from all previous layers (Figure 3), which can help the model learn more complex features (Kim, Heo, and Han 2024), potentially useful for distinguishing between images that should be classified as regular or irregular based on the proximity of the elements involved.

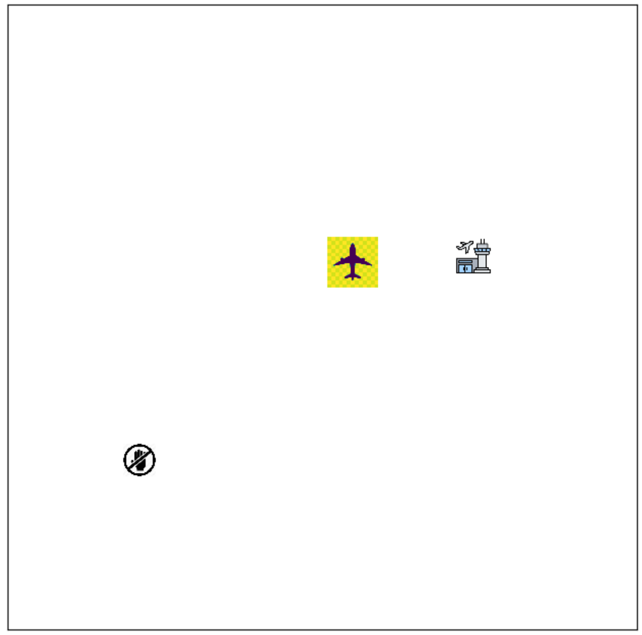

Figure 1: Example of a graph with an aircraft closer to a regular airport (regular classification).

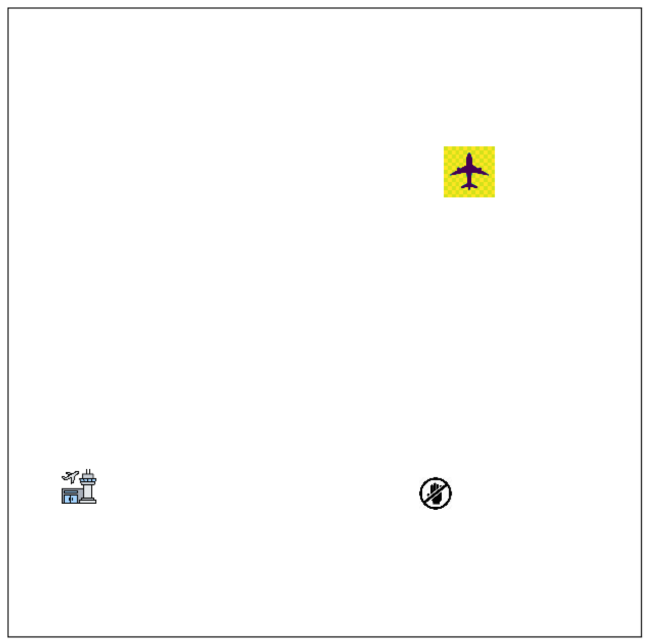

Figure 2: Example of a graph with an aircraft closer to a irregular airport (irregular classification).

DenseNet (Huang et al. 2018) was selected due to its efficient reuse of features and its ability to mitigate vanishing gradient issues. The architecture used consists of 5 layers with varying filters (conv base layers, global average pooling, dense layers with 256 and 1 neuron) and 1 batch normal-

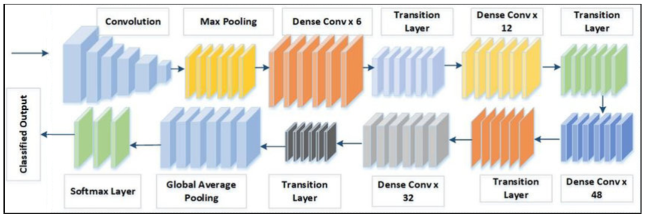

Figure 3: DenseNet architecture, adapted from (Mahum et al. 2022)

ization layer indirectly through the pre-trained conv base.

The training and validation procedures followed the steps for the graphical data dataset:

a) **Data Preparation**: The images were divided into training and validation sets. Samples were balanced to ensure the same number of images (1,000) in the "regular" and "irregular" classes. Data generators with data augmentation were used to increase the variability of the data set;

b) **Configuration of Data Generators**: ImageDataGenerator was used to normalize pixel values and apply data augmentation (only horizontal and vertical flip);

c) **Training**: The model was trained using the fit method, with steps per epoch and validation steps calculated as the total number of images divided by the batch size. The EarlyStopping condition was used to stop training if the validation loss did not improve after 15 consecutive epochs; and

d) **Validation**: The performance of the model was evaluated in each epoch using a separate validation set. Accuracy and loss metrics were monitored in both the training and validation sets.

To evaluate the performance model, the following metrics were used:

- **Accuracy**: Measure of the proportion of correct predictions in relation to the total number of predictions.
- **Loss**: Calculated using the binary cross-entropy loss function, providing an indication of how well the model fits the training and validation data.
- **Validation accuracy**: Accuracy calculated on the validation set to monitor model performance on unseen data.
- **Validation loss**: Loss calculated on the validation set to monitor whether the model is suffering from overfitting.

In the end, to evaluate the model's classification capability, another 59 graphical examples with proximity classified as regular and 57 graphical examples with proximity classified as irregular were generated. These examples were never seen by the model resulting from the training.

Results and Discussion

In an attempt to mitigate overfitting, basic data augmentation techniques (horizontal and vertical flips) were applied, along with the use of an EarlyStopping callback during training. Additionally, larger synthetic datasets comprising 5,000 and 10,000 images were generated and tested. These variations, however, did not result in any significant improvement in validation performance or reduction of the performance gap, indicating that the overfitting observed may not be solely attributable to dataset size. It is likely that the simplicity and low variability of the synthetic images contributed to the model's tendency to overfit. Future work should consider the application of more advanced regularization strategies, such as dropout and L2 penalties, as well as the evaluation of alternative neural network architectures with reduced parameter counts, in order to improve model generalization.

The result of neural network training is shown in Table 1:

Table 1: Trainning and Validation Metrics

Metric	Trainning	Validation
Accuracy	0.9062	0.7500
Loss	0.3257	0.3777

The evaluation metrics for the test set are presented in the table 2:

Table 2: Evaluation results for target

AUC	F1	Precision	Recall
0.883	0.784	0.785	0.784

The confusion matrix for the test set can be seen in Figure 4.

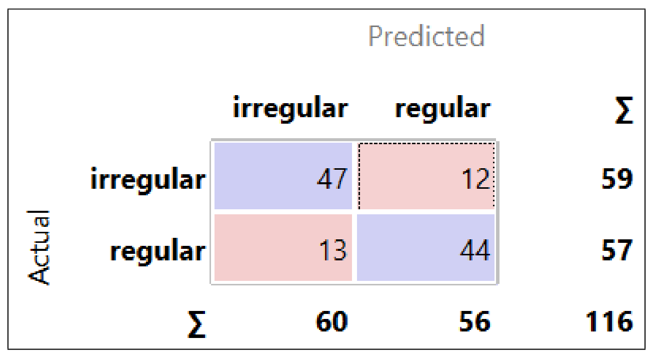

Figure 4: Confusion matrix for the test set

With an accuracy of approximately 90.62% on the training set, the model has learned the characteristics of this set well. This value suggests that the model can effectively identify patterns in the training samples.

The accuracy of 75% on the validation set is significantly lower than that of the training set, indicating that the model is not generalizing well. This gap between training and validation accuracy is a sign of overfitting, where the model is "memorizing" the training data rather than learning patterns applicable to new data.

The difference in loss between training (0.3257) and validation (0.3977) further reinforces the sign of overfitting. Generally, a higher validation loss suggests that the model is struggling to make correct predictions on the validation set and may be becoming reliant on specific training examples.

Regarding the evaluation metrics of the test set, it appears that an AUC value of 0.883 indicates that the model has a high ability to discriminate between the ”regular” and ”irregular” classes. The F1 Score of 0.784 indicates that the model has a good balance between precision and recall (or sensitivity). An accuracy of 0.785 indicates that 78.5% of the model’s positive predictions are actually positive. With a recall of 0.784, the model is correctly identifying about 78.4% of true positives.

The F1-Score of 78.4% demonstrates the model’s moderate success in understanding visual proximity patterns.. This result underscores the inherent challenge of teaching neural networks to replicate human-like intuition when interpreting synthetic images. The discrepancy between training and validation performance highlights overfitting as a key limitation, suggesting the need for improvements in data diversity and augmentation. Future iterations could explore deeper architectures, such as EfficientNet or ResNet, which may enhance the network’s capacity to capture subtle spatial relationships. Additionally, increasing the complexity and diversity of the data set could provide a more robust basis for training, further enhancing generalization to unseen data.

Regarding the confusion matrix, the model has a balanced performance between both classes, with precision and recall metrics close to ”regular” and ”irregular”. The model is treating the classes well, without presenting a strong bias towards one of them, which is positive for generalization. Incorrectly classified graphs have characteristics that can confuse the model, as seen in Figure 5.

Interestingly, several misclassified examples also posed significant challenges for human experts, particularly in cases where the aircraft was equidistant from both airports or where visual patterns were ambiguous. This alignment of human and model difficulties validates the complexity of the task and reinforces the relevance of this research. By focusing on these challenging scenarios, future studies can better explore how neural networks can approximate human decision-making processes under similar conditions.

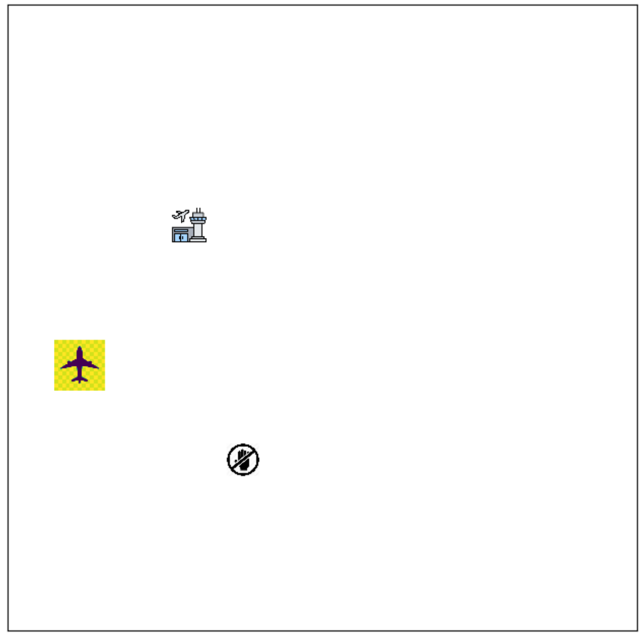

Figure 5: Example of a graph with equal distances between objects.

Conclusions

This study demonstrates the potential of neural networks, specifically DenseNet, to approximate human intuition in evaluating proximity relationships between graphical elements.

While the model achieved a moderate F1-Score of 78.4%, the results highlight both the promise and the challenges of using artificial intelligence to replicate human-like perception in synthetic spatial tasks.

The reliance on a single architecture and the use of synthetic datasets are notable limitations, potentially constraining the model’s generalizability and real-world applicability. Future research should investigate alternative architectures, such as ResNet, EfficientNet, or Transformers, to identify configurations better suited for this task. Incorporating real-world data and increasing dataset complexity could further improve the robustness and utility of the model. Additionally, exploring transfer learning and human-in-the-loop approaches may enhance the practical relevance of such methods in airspace monitoring and other domains requiring proximity-based decision making.

By addressing these limitations, future studies can build on the foundation established here, contributing to more effective applications of AI in air traffic surveillance and broader remote sensing challenges.

Although this research represents an initial step toward modeling human-like proximity perception through artificial neural networks, the findings contribute meaningfully by confirming the feasibility of training models to capture relative spatial arrangements from visual input alone. The moderate performance observed—especially when contrasted with human difficulties in certain ambiguous cases—provides valuable insight into the limits and potential of current architectures under simplified, controlled conditions. This study therefore establishes a validated experimental baseline, upon which more complex, diverse, and realistic datasets and models can be tested in future investigations.

Due to ongoing research and the use of synthetic datasets still under expansion, the code and data are not publicly available at this time. However, a public version is being prepared for future dissemination.

Acknowledgments

This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001. This study is a contribution to the C2 Project, of the Institute for Advanced Studies, of the Brazilian Air Force.

References

Artaxo, P. 2023. Cúpula da amazônia: frustações, avanços e perspectivas. *Plenamata*.

Dietzsch, G., and Shiguemori, E. H. 2024. Integrating artificial intelligence for air traffic surveillance: A case study using tabular and graphical data. *Associação de Especialistas Latinoamericanos em Sensoriamento Remoto*.

Força Aérea Brasileira, Brasil. 2020. *Instrução do Comando da Aéronáutica 100-18*.

Ganesan, R.; Manjula, G.; Chandan, M.; Thamarai, I.; Rajavelu, S.; Bhatt, U. M.; and Sarkar, S. 2023. Application of deep neural networks for weed detection and classification. In *2023 2nd International Conference on Applied Artificial Intelligence and Computing (ICAAIC)*, 1014–1024.

Hao, D.; Le, T. D.; Berezkin, A.; and Kirichek, R. 2023. Graph neural networks for traffic classification in satellite communication channels: A comparative analysis. *Proceedings of Telecommunication Universities* 9:14–27.

Huang, G.; Liu, Z.; van der Maaten, L.; and Weinberger, K. Q. 2018. Densely connected convolutional networks.

Kim, D.; Heo, B.; and Han, D. 2024. 2. densenets reloaded: Paradigm shift beyond resnets and vits. *arXiv.org*.

Mahum, R.; Munir, H.; Mughal, Z.-U.-N.; Awais, M.; Khan, F.; Saqlain, M.; Mahamad, S.; and Tlili, I. 2022. A novel framework for potato leaf disease detection using an efficient deep learning model. *Human and Ecological Risk Assessment: An International Journal* 29:1–24.

Matijašević, N.; Samcovic, A.; and ogatović, M. 2022. Application of convolutional neural networks for the classification of human emotions. *SINTEZA*.

Mikhaylov, I. S.; Aung, Y. T.; Win, M. H.; and Aung, Z. 2023. Neural networks application for the classification problem solving with domain constraints. In *2023 3rd International Conference on Technology Enhanced Learning in Higher Education (TELE)*, 220–224.

Morinaga, M.; Mori, J.; and Yamamoto, I. 2023. Aircraft model identification using convolutional neural network trained by those noises in a wide area around an airfield. *Acoustical Science and Technology*.

Pajolla, M. 2022. Do tráfico ao trabalho escravo: rede de atividades ilegais impulsiona devastação da amazônia. *Brasil de Fato*.

Perić, S. L.; Vukić, N.; Antić, D.; Milojković, M.; and orević, A. 2023. Application of convolutional neural networks for road type classification. *Facta Universitatis*.

Potter, H. 2022. As pistas da destruição. *Itercept Brasil*.

Potter, H. 2023. Tráfico e garimpo ilegal compartilham aviões e pilotos para lavar dinheiro na amazônia. *Repórter Brasil*.

Rosenfeld, A.; Zemel, R. S.; and Tsotsos, J. K. 2019. 6. high-level perceptual similarity is enabled by learning diverse tasks. *arXiv: Computer Vision and Pattern Recognition*.

Valente, R. 2023. Controle ‘efetivo’ do espaço aéreo é fundamental para terra yanomami, diz diretor da pf. *Pública*.

Vincent-Boulay, N., and Marsden, C. 2024. Aircraft categorization approach using machine learning to analyze aircraft behavior. *Journal of Air Transportation* 32(4):218–229.

Zang, H., and Zhu, J. 2022. Deep learning architecture for flight flow spatiotemporal prediction in airport network. *Electronics*.