Detecting Critical Infrastructures in Disaster Images by Combining PSPNet and Genetic Algorithm-driven Hyperparameter Optimization

Iyed Dhahri[1], Mahmoud Golabi[1], Karim Hammoudi [1] and Lhassane Idoumghar[1]

***Abstract*— Natural disasters have intensified in recent years, directly affecting populations, especially in urban and semi-urban areas where infrastructure is damaged and causes casualties. Rapid damage assessment is essential for decision-making and developing relief measures for victims of such disasters. In the aftermath of such disasters, the application of deep learning to aerial imagery has become increasingly crucial for damage assessment. By categorizing damage levels to buildings and other environmental elements, such as roads, these methods can significantly support informed decision-making processes. RescueNet is proposed as a database for training these intelligent damage estimation models. In this study, we introduce a GA-driven hyperparameter optimization for adapting PSPNet to the RescueNet hurricane imagery dataset. Using a genetic algorithm to explore key design choices (e.g, learning rate, backbone depth, pooling bin sizes, dropout probability, ...) we automatically evolve high-performance configurations over successive generations. Our best model boosts mean Intersection over Union on undamaged buildings from 95.16% to 96.57% and on vehicles from 85.97% to 89.76%, outperforming the default PSPNet settings. By sharply improving the model's ability to distinguish intact versus damaged regions, we deliver faster, more reliable disaster-impact assessments. In ongoing work, we are extending this metaheuristic approach to optimize additional critical classes such as roads to further enhance end-to-end operational readiness.**

Keywords: Image Processing, Optimization, Damage Assessment, Deep Learning.

I. INTRODUCTION AND MOTIVATION

In recent years, numerous natural disasters have resulted in significant human and material losses worldwide. Despite substantial efforts to deliver timely assistance, rapid assessment and prioritization of the most affected areas remain challenging due to the limited availability of data in the immediate aftermath of a disaster. In such scenarios, the ability to promptly evaluate the extent of the damage is crucial for minimizing response delays and formulating an effective emergency response strategy.

One of the most effective ways to gather information about affected areas is to capture aerial images or use satellite platforms directly after the disaster to visualize damaged areas. Therefore, it is useful to leverage recent advances in deep learning to automatically categorize and assess damage from captured images, especially for buildings, as this reflects the scale of the disaster and is also the most suitable place to locate victims. With this in mind, many datasets have been proposed to the literature for building damage assessment, such as RescueNet [1], which consists of drone images captured after a hurricane in Florida, USA. This dataset provides semantic segmentation labels for ten distinct classes, including four standardized levels of building damage (e.g., no damage to totally destroyed), as well as other features relevant to post-disaster assessment, such as road conditions (e.g., blocked or clear) and vehicles. Using this dataset, it is possible to train a damage assessment model applicable to newly acquired images following future disaster events, such as the one that occurred in Mayotte, France [2]. The original study introduced a set of models trained on the RescueNet dataset, including PSPNet [3], which demonstrated strong performance in image segmentation and damage level classification [1]. In this work, we aim to further enhance its capacity to accurately detect critical damage classes.

Tuning hyperparameters is essential for optimizing deep learning models, as it directly influences their generalization ability and sensitivity to critical classes in imbalanced datasets. Conventional methods such as grid search or random search are typically inefficient or intractable when faced with the high-dimensional and non-convex search spaces characteristic of deep architectures. In contrast, metaheuristic algorithms offer a more scalable and flexible alternative by effectively balancing exploration and exploitation [4]. This approach has shown promising results in domains such as medical image segmentation [5][6], where algorithms like Genetic Algorithms (GAs) have been successfully employed to optimize non-differentiable architectural parameters (e.g., backbone depth, pooling strategies, dilation rates), thereby enhancing segmentation quality.

Motivated by these advantages, we employ a GA-based hyperparameter tuning procedure to optimize the PSPNet architecture as used in the RescueNet framework. By carefully defining the search space and tailoring the GA operators to the specific characteristics of our segmentation task, we observed notable improvements in the detection of critical semantic classes, particularly undamaged buildings and vehicles. Accurate identification of undamaged buildings is crucial for post-disaster assessment, as it enables a clearer delineation between affected and unaffected zones. Furthermore, when integrated with cadastral maps containing geospatial building footprints [7], this information can significantly enhance emergency response planning by facilitating precise geolocation of intact structures and prioritization of rescue resources. These improvements also support the formulation of more effective disaster preparedness and mitigation strategies.

[1] Université de Haute-Alsace, IRIMAS, Mulhouse, France, iyed.dhahri@uha.fr, mahmoud.golabi@uha.fr, karim.hammoudi@uha.fr, lhassane.idoumghar@uha.fr

II. PROPOSED APPROACH

Semantic segmentation involves pixel-wise classification, wherein each pixel is assigned a specific class label [8]. The task becomes increasingly complex when the dataset includes a large number of classes (10 in our case) [1], and even more so when these classes exhibit subtle inter-class variations such as the four levels of building damage or varying object scales. To address these challenges, several deep learning architectures have been proposed. One notable example is PSPNet, which enhances feature representation by aggregating multi-scale contextual information through its Pyramid Pooling Module (PPM) [2]. This module performs pooling operations at multiple scales using different kernel sizes to capture both global and local context, which are then unified into a single feature representation.

Although PSPNet has demonstrated strong segmentation performance, its effectiveness is highly sensitive to hyperparameter configurations. Suboptimal settings can lead to poor segmentation results, particularly for specific classes. Additionally, the fixed pyramid pooling module in PSPNet may not generalize well across diverse disaster imagery, limiting its adaptability. To address these limitations, we propose integrating a genetic algorithm into the hyperparameter tuning process. Genetic algorithms, inspired by the principles of natural selection, are well-suited for exploring high-dimensional and complex parameter spaces. Our GA-based optimization evolves a population of candidate configurations over successive generations through crossover and mutation operations. In each iteration, the population is evaluated based on a validation metric, specifically MIoU, to guide the selection of the most promising configurations. The parameters tuned include dropout rates, backbone architecture options, and pooling strategies. By combining the state-of-the-art PSPNet with a GA-based hyperparameter optimization strategy, our method enhances the robustness of the segmentation model, particularly improving the detection of critical yet challenging classes such as undamaged buildings. Figure 1 illustrates the overall architecture of the proposed approach.

III. EXPERIMENTAL RESULTS AND EVALUATION

A. RescueNet Dataset

RescueNet is a dataset designed for post-disaster damage assessment using UAV (Unmanned Aerial Vehicle) imagery, specifically captured after Hurricane Michael in Florida, USA. It comprises 4,494 high-resolution images (3000 × 4000 pixels), taken at altitudes ranging from 30 to 100 meters. Each pixel in the images is annotated with one of ten semantic classes, including four levels of building damage ranging from no damage to total destruction. An example from RescueNet is shown in Figure 2, illustrating damaged buildings, vehicles, and road infrastructure. Although RescueNet comprises only hurricane imagery, its representations generalize across disaster types. In 2, a PSPNet model pre-trained on RescueNet was fine-tuned on FloodNet (a flood-damage dataset) and significantly outperformed a model trained from scratch, demonstrating the effectiveness of cross-disaster transfer learning.

In our experiments, we adopted the same preprocessing and data augmentation procedures as described in the reference study, in order to ensure a fair comparison and isolate the impact of our proposed methodology. Specifically, we investigated the effect of employing a genetic algorithm to optimize the hyperparameters of the PSPNet segmentation model. From the ten available semantic classes, we retained only those deemed critical for post-disaster response, particularly those directly related to human needs and safety: the four building damage levels (as proxies for shelter conditions), along with vehicle and road classes (reflecting transportation accessibility). We excluded categories of marginal relevance to emergency response such as pools and trees from our analysis, since they contribute little to assessing disaster impact. Although these classes can offer contextual clues (for example, damaged vegetation may indicate wind severity, and empty pools can imply structural failure), we limited our evaluation to the target classes directly tied to human safety and rescue logistics.

B. GA-based Architecture Opimization

TABLE I
SEARCH SPACE DEFINITION FOR GENETIC ALGORITHM

Hyperparameter	Search Space / Description
Dropout Rate	Float in [0.1, 0.5]
Backbone Index	1–3 (ResNet50, ResNet101, ResNet152)
Pooling Bins Option	1–3 ((1,2,3,6), (1,3,6,8), (1,2,4,8))
Activation Function	(0 = ReLU, 1 = LeakyReLU)
Filters Scaling Factor	Float in [0.5, 2.0]
Optimizer Type	(0 = Adam, 1 = SGD, 2 = RMSprop)
Learning Rate	Float in [1e-5, 1e-2]
Batch Size Index	(0 = batch size 2, 1 = batch size 4)
Augmentation Size	Fixed to 713
Classifier Conv Layers	Integer (1 to 3)
Normalization Type	Fixed to 0 (BatchNorm2d)

TABLE II
BEST HYPERPARAMETER CONFIGURATION

Hyperparameter	Best Configuration
Dropout Rate	0.1
Backbone	ResNet152
Pooling Bins	(1, 2, 3, 6)
Activation Function	LeakyReLU
Filters Scaling Factor	2.0
Optimizer Type	RMSprop
Learning Rate	1.0e-5
Batch Size	4
Augmentation Size	713
Classifier Conv Layers	2
Normalization Type	BatchNorm2d

Defining an appropriate hyperparameter search space is crucial for improving the performance of the model. Guided by domain knowledge and supported by a series of preliminary experiments, we selected a set of hyperparameters and

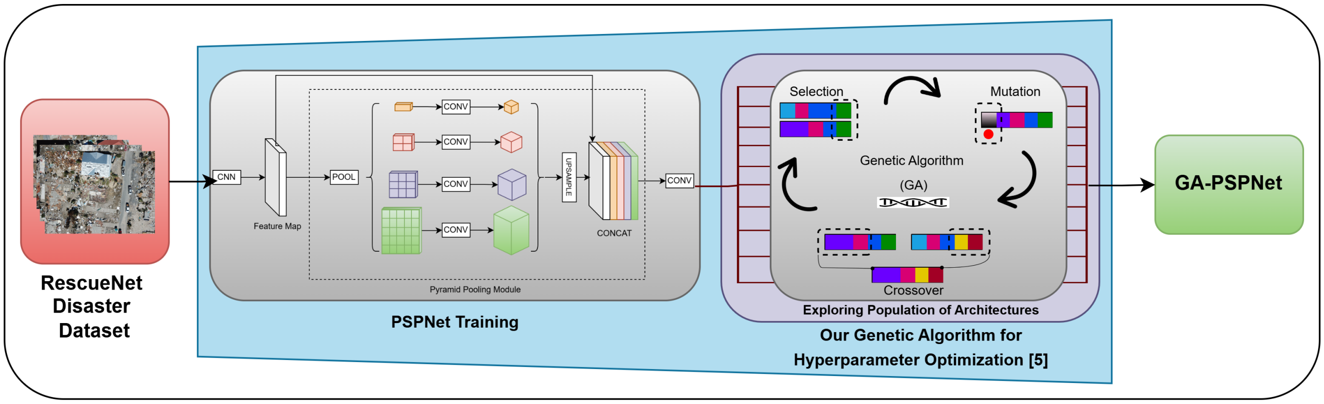

Fig. 1. Proposed GA-PSPNet Approach for Detecting Critical Infrastructures in Disaster Images

Fig. 2. Sample of Damage Scene from RescueNet. Damaged and undamaged buildings as well as road parts and vehicles have been captured by local UAV.

their corresponding ranges, as summarized in Table I. Below, we provide a rationale for the inclusion of key parameters in the search space design:

- **Dropout Rate:** An ideal dropout rate is essential for reducing overfitting without losing network ability in learning complex features.
- **Backbone Index:** Varying the depth of the backbone network enables the identification of architectures that are most effective in extracting diverse and informative features from disaster imagery.
- **Pooling Bins Option:** The pyramid pooling module's configuration affect on the way multi-scale features are combined. So finding the optimal configuration will enhance the segmentation performance of our model.
- **Filters Scaling Factor:** Adjusting the number of filters in the classifier block directly influences the model's representational capacity and overall performance.
- **Classifier Convolutional Layers:** Modifying the depth of the classifier block impacts the model's ability to learn and represent complex, high-level features relevant to post-disaster analysis.
- **Optimizer Type:** Each optimizer has its own characteristics in terms of convergence speed and stability, so the optimal choice is the one that best suits our training dynamics.

Subsequently, the GA is employed to explore the defined hyperparameter search space and identify an optimal configuration. The GA initializes a population of candidate hyperparameter sets sampled from the predefined search space, and evolves this population over several generations. Each candidate is evaluated by training a corresponding instance of the PSPNet model on a subset of the RescueNet dataset for a fixed number of epochs, using the validation MIoU as the fitness metric. To guide the evolutionary process, we adopt a tournament selection strategy, where small groups of candidates are randomly sampled, and the individual with the highest validation MIoU within each group is selected for reproduction. Reproduction is performed through a single-point crossover operator (with a probability of 0.9), which exchanges portions of the hyperparameter vectors between two parent solutions to produce new offspring. Mutation is then applied to maintain diversity in the population and prevent premature convergence, by randomly altering individual hyperparameter values within valid ranges. After several generations, the GA converges to a high-performing configuration, which is reported in Table II. This final configuration is subsequently used to train PSPNet on the full training set for 100 epochs, after which performance is evaluated on the held-out test set. We plan to deploy our trained PSPNet for rapid, large-scale damage assessment, so we benchmarked its inference speed on high-resolution drone imagery. On a single GPU, our model processed and segmented 450 full-resolution frames in under 15 minutes which is much faster than ground teams conducting manual surveys. In practice, the bottleneck is image capture and transmission, even without model compression or distributed inference, our end-to-end pipeline (from takeoff to segmentation) can easily meet the timeline requirements of emergency responders.

C. *Comparison of Disaster-focused PSPNet Models*

Table III presents the test IoU scores for each class within the shelter category, providing insight into the model's ability to discriminate between different levels of building damage. Statistically significant improvements over the baseline are highlighted in bold. Notably, our GA-optimized PSPNet

TABLE III

ACCURACY COMPARISON BETWEEN RESCUENET PSPNET AND GA-PSPNET ON SHELTER CLASSES

Critical Category	**RescueNet PSPNet**	**GA-PSPNet**
Shelter	**Accuracy (%)**	
Building - No Damage	0.9516	**0.9657**
Building - Med. Damage	0.9436	0.9381
Maj. Damage	**0.9691**	0.9341
Tot. Destruction	**0.9895**	0.9448

TABLE IV

ACCURACY COMPARISON BETWEEN RESCUENET PSPNET AND GA-PSPNET ON TRANSPORTATION-RELATED CLASSES

Critical Category	**RescueNet PSPNet**	**GA-PSPNet**
Transportation	**Accuracy (%)**	
Vehicle - Various Types	0.8597	**0.8976**
Road - Clear	**0.9807**	0.8439
Blocked	**0.9483**	0.9348

outperforms the original RescueNet PSPNet in detecting undamaged buildings, which is crucial for accurately delineating safe zones from damaged areas. Since detecting undamaged structures implicitly requires reliable identification of damaged ones, this result suggests improved overall robustness. Rapid and reliable differentiation between intact and damaged zones enables search-and-rescue teams to focus their efforts on affected buildings, avoiding time spent in safe areas and thereby optimizing overall response efficiency. For the intermediate damage levels, both models yield comparable performance, while the original PSPNet achieves better results for the most severely damaged categories.

Table IV reports the IoU scores for the transportation category. Our model demonstrates superior performance in detecting vehicles, including both cars and trucks. Accurate detection of vehicles within post-disaster imagery provides critical insights into transportation network status and resource availability. However, the baseline PSPNet shows better detection capability for roads.

Overall, both models exhibit competitive performance, often complementing each other particularly in the shelter category. While our method excels at detecting intact buildings, the baseline performs better on severely damaged ones. Importantly, even with a basic metaheuristic optimization approach, our method improves performance in two critical classes. This suggests strong potential for further gains when employing more advanced or hybrid metaheuristic techniques for hyperparameter tuning.

IV. CONCLUSION AND FUTURE WORKS

In this paper, we proposed a metaheuristic-based approach to optimize semantic segmentation of critical infrastructure in post-disaster imagery using the RescueNet dataset. Specifically, we enhanced the original PSPNet model by integrating it with a GA for hyperparameter optimization, enabling efficient exploration of a high-dimensional and complex search space. The objective was to improve the model's ability to detect and classify critical semantic categories relevant to post-disaster response.

Experimental results demonstrate that our GA-optimized PSPNet improves the detection of undamaged buildings, a crucial component for assessing safe zones in disaster-struck areas. While the original PSPNet performs better in identifying heavily damaged buildings, our approach achieves complementary strengths, especially within the shelter category. In the transportation category, our model outperforms the baseline in detecting vehicles, whereas the original PSPNet yields better performance in identifying roads.

These findings highlight the promise of metaheuristic optimization techniques in enhancing deep learning models for damage assessment tasks. However, the cross-disaster transferability was only validated from RescueNet (hurricanes) to FloodNet (floods). Extending this capability to a truly universal damage-assessment model will require incorporating additional types of disaster datasets. Likewise, reducing our model's inference latency offers a clear path to even faster predictions and more responsive disaster relief. As future work, we plan to investigate more advanced metaheuristic algorithms, incorporate hybrid strategies, and expand the hyperparameter search space to include additional architectural and training components such as class-weighted loss functions for better balance performance across all critical damage categories. This direction aims to further improve segmentation performance across all critical classes involved in post-disaster assessment.

REFERENCES

[1] M. Rahnemoonfar, T. Chowdhury, and R. Murphy,"RescueNet: A high resolution UAV semantic segmentation dataset for natural disaster damage assessment," *Scientific Data*, vol. 10, no. 1, p. 913, 2023. doi: 10.1038/s41597-023-02799-4.

[2] H. Karim, D. Iyed, G. Mahmoud, and others, "Generalizability of a RescueNet Data-trained Damage Estimation Model for Evaluating Mayotte Disasters" *International conference on Geoinformatics and Data Analysis (ICGDA 2025)*. Springer.

[3] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, "Pyramid Scene Parsing Network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2017, pp. 6230–6239. doi: 10.1109/CVPR.2017.660.

[4] A. E. Eiben and J. E. Smith, *Introduction to Evolutionary Computing*, 2nd ed., New York, NY, USA: Springer, 2015.

[5] E. H. Houssein, G. M. Mohamed, Y. Djenouri, and others, "Nature inspired optimization algorithms for medical image segmentation: a comprehensive review," *Cluster Computing*, vol. 27, pp. 14745–14766, 2024. doi: 10.1007/s10586-024-04601-5.

[6] M. Junaid, L. Moalic, M. Essaid, and L. Idoumghar, "Training free U-Net for retinal vessel segmentation," in *Artificial Evolution 2024*, Springer; Association Evolution Artificielle, Oct. 2024, Bordeaux, France. [Online]. Available: https://hal.science/hal-04999688

[7] T. Solberg, A. Brender, and K. Stokke,"Integration of cadastral maps and remote sensing data for building footprint extraction in emergency response," *International Journal of Disaster Risk Reduction*, vol. 26, pp. 108–118, 2018, doi: 10.1016/j.ijdrr.2018.01.005.

[8] J. Long, E. Shelhamer, and T. Darrell,"Fully Convolutional Networks for Semantic Segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2015, pp. 3431–3440, doi: 10.1109/CVPR.2015.7298965.