AI-Powered Disinfection Force: Harnessing 5G for Intelligent Distributed Object Recognition and Sanitation Fleet

NIKOLAOS K. PAPADAKIS
Department of Machines, Intelligent and Distributed Systems
Infili Technologies S.A.
60, Kousidi str, Zografou 15772, Athens
GREECE

Abstract: With the increasing prevalence of embedded cameras in devices such as phones, robots, and vehicles, there's a growing need to enhance their perception capabilities through advanced deep learning techniques. These systems, however, often struggle under suboptimal conditions like environmental noise. The paper addresses these challenges by proposing a 5G-enabled solution that uses a fleet of unmanned ground vehicles (UGVs) for coordinated disinfection tasks in hospital environments. Utilizing high-bandwidth, low-latency 5G networks, these UGVs can share high-quality images for distributed object recognition tasks and execute precise disinfection routines using onboard UV lamps. The proposed system includes an optimal formation algorithm for UGVs to maintain effective positioning and coordination. By leveraging 5G connectivity, the fleet can efficiently exchange sensor data and perform real-time, innovative object recognition and disinfection, presenting a significant advancement in intelligent sanitation technology.

Key-Words: AI; Computer Vision; Intelligent Fleet; Distributed Object Recognition; Decentralized Architectures; Edge Computing; 5G

Received: April 6, 2024. Revised: September 11, 2024. Accepted: October 2, 2024. Published: October 29, 2024.

1 Problem significance and motivation

Nowadays, cameras are widely embedded into many devices like phones, wearable devices, robots, and cars and they are used in a large variety of applications providing efficient tools for improving the perception capabilities of these devices. In order to achieve this, deep-learning approaches are introduced to improve performance quality of traditional computer-vision tasks. However, these approaches are well known to perform well only under optimal conditions, i.e., with input images/videos without noise, with high quality, etc. However, such conditions are obviously not always met, especially when these cameras are operating under real-world environmental conditions (also called environmental noise), including rain, dust, fog-haze, fire-smog, light reflections, limited illumination conditions, etc.; under well-known hardware camera limitations (i.e., sensor noise, non-linearity, white balance, image artifacts); or under security breaches (e.g., adversarial attacks).

Despite the aforementioned limitations, deep learning is taking over as the single most important and powerful method for object-recognition and classification tasks. Deep Neural Networks (DNNs), despite the diversity in their architectures, all share a common nature. Their performance is as good as the data they are trained with. In this context noise and distortions in the training data become important issues since high-quality annotated data are not easy to find for every specific training task. Current systems mainly rely on three general approaches in addressing noise in training data:

- Incorporating noise in the recognition task by training standard DNNs for classification/recognition [1], [2], [3], [4], [5], [6], [7] (of the suitable kind for each cue) with a mixture of noisy and clear inputs, so the network will learn how to correctly map noisy inputs as well. In this scenario, however, a huge amount of *high-quality annotated* noisy data is required, leading to a *chicken - egg* dilemma very common in deep-learning tasks today.
- Generating noisy training data of high annotation quality by applying *extended augmentation techniques* [8], [9], [10], meaning producing synthetic images with various kinds of noise and distractions, different views, angles etc., similar to the ones expected in real-life scenarios. Even though this is a valid approach, there is no obvious way for one to assess how *realistic* these synthetic data would be or how its *divergence from reality* affects the learning accuracy and performance in real-life scenarios.
- Applying *denoising* [15],[14],[16],[17] as an explicit preprocessing step to the inputs of the

recognition/classification task. The advantage of this approach lies in decoupling noise removal from the recognition task in a way that the former can be learned explicitly. Deep learning can thus be used for denoising as well, but with modified architectures that favor this specific task. The challenges of this approach lies primarily in choosing the *correct* denoising architecture depending on the task, as a *one-size-fits-all* denoising architecture does not exist.

No matter the approach, the need of extensive and high-quality training is thus the major limitation. This is because manifestations of the same object can take arbitrary forms in the wild, e.g., images taken from different angles, with occlusions, lighting conditions and with varying noise effects as was discussed above. It is thus impossible for one to train with images that cover all possible cases. If only a single image is presented for recognition the probability of misclassification cannot be measured under these varying conditions. However, if more images of the same object from different angles and under different conditions are presented simultaneously for classification with different algorithms, the probability of a correct classification increases without the need of additional training.

It is this idea that will be pursued in the use-case scenario presented herein. Based on the above-mentioned observations a new technology is proposed that will enable a fleet of unmanned ground vehicles (UGVs) to perform object recognition and coordinated fleet-formation tasks for the purpose of disinfecting indoor and outdoor hospital areas using onboard UV lamps in a 5G-enabled setting. The availability of high bandwidth low latency 5G networks makes it possible to coordinate the exchange of high-quality images among a fleet of UGVs that collectively perform a recognition task by running different classification algorithms in a suitable distributed manner and apply predefined disinfection actions on the recognized object. We propose to demonstrate how a fleet of UGVs carrying sensors, cameras and adjustable UV lamps in a 5G infrastructure setting, can run different onboard recognition systems and perform disinfection actions. By arranging specific formations between them, each UGV can collect images from multiple angles of the same object with the purpose to optimize recognition and disinfection actions. 5G connectivity enables not only the exchange of high-resolution images and video between single UGVs, but also a distributed UGV formation algorithm which can run in real time. By presenting many simultaneous views of the same object at the querying phase one facilitates the classification algorithms to generalize better on the unseen object and decide the correct disinfection routine. At the same time 5G connectivity enables a mixture of the above techniques to be unified under a boosting framework [13] by utilizing the fleet hardware collectively in real time.

Additionally, for enhanced coordination and system reliability, a hierarchical structure is proposed, where the swarm is composed of leader UGVs at the top level, regional UGVs at the mid-level, and individual UGVs at the bottom layer. This structured approach ensures efficient coordination and task execution, Also it enables effective emergency situation handling through a novel weighted consensus mechanism that is proposed through this work based on the Byzantine Fault Tolerance algorithm [21].

2 5G use-case scenario and its significance

Using 5G connectivity we propose a new, distributed visual-surveillance technology for extracting actionable information and perform specific predefined tasks, more specifically, disinfection indoors/outdoor hospital areas using UV light lamps installed on UGV fleets. Focusing on visual (and, possibly, infrared) imaging sensors the system should offer identification, classification and analysis on visual data from UGVs as well as stationary visual surveillance sources and enable real-time, onboard decisions and system-wide planning regarding route, speed, and disinfection tasks. The suggested scenario is a small fleet of 10-15 UGVs deployed at a 5G-enabled hospital installation with the mission to perform visual recognition tasks and UV-light disinfection on demand, from various points around one building sharing corridors and across uncontrolled lanes or through busy parking lots to another building on the installation with speeds ranging from 3-25mph. The UGVs are aided by stationary networked cameras that cover the area, a priori visual background data (enabling background subtraction) and knowledge of own location. The demand for stationary cameras are not restrictive since any UGV can play this role and in fact stationary UGVs will model a fixed coordinate system from optimally chosen positions through which other UGVs will tag their position by intra-fleet 5G communications. The roles between stationary and moving UGVs can change depending on the conditions.

The system should offer the following capabilities:

1. Obstacle detection and Obstacle Avoidance (ODOA).

2. Correct positioning and speed regulation with respect to moving and stationary objects.

3. Co-ordinated and optimized system-wide responses across the fleet.
4. Data collection and communications.
5. Extraction of actionable information from the sensor stream.
6. Disinfection in rooms, corridors and individual objects using coordinated fleet movements and the use of onboard UV lamps.

A high-bandwidth and low-latency 5G environment enables powerful distributed algorithms for fleet coordination and actionable information extraction. In current systems, units performing the recognition and sensor carriers themselves, such as UGVs, have only limited availability (bandwidth and delay constrained) of the data coming from stationary and other UGV nodes while they typically have full availability of the data flow in their own sensors. This constraint is lifted in 5G-connectivity scenarios. In the proposed system, therefore, each UGV analyses data and extracts data significant for object identification for the dual purpose of: a) to locally provide ODOA, b) to use the 5G bandwidth to share with other UGVs high-resolution images and videos of its own angle of view. These two purposes are highly intertwined because other UGVs receiving extracted data become rather effective in extracting their own information for further sharing.

In the present use case 5G low latency enables a coordinated action to minimize intra-group disruption but also to optimize coordinated responses in a deterministic fashion (subject to mission-critical, time and other constraints). The usefulness of the sensor data, collectively acquired by the fleet, is directly related to field parameters (angle, distance, relative speed, etc.), and, thus, the loss or lack of formation directly undermines the ability of combining sensor data from different UGVs towards credible recognition. Images taken from the same angle do not contribute to further collective analysis. It becomes then evident that certain formations may greatly facilitate corresponding fusing algorithms by providing diversified sensor data. This way, significant gains in computational time and accuracy in recognition tasks can be achieved. A controlled formation can also diversify the response among the UGVs and provide the means to bind responses in a meaningful manner following deterministic mission-critical specifications. For instance, UGVs might coordinate in sharing the disinfection task in a room.

2.0.1 The Optimal Formation Algorithm

The formation algorithm runs on each of the UGVs and determines the movement of each of the UGVs for the disinfection task. The UGVs try to keep a certain formation by solving a point-correspondence problem. Each UGV knows the positions of the other UGVs, the position of the obstacles and other objects in the environment in space-time, and the formation to be kept during disinfection. Therefore, the movement of the whole fleet to the next position can be modeled as a geometric point-correspondence problem where each UGV must move from its current position to the desired position that will bring it in the desired formation, covering the minimum distance to the ultimate target for disinfection. If a UGV has knowledge of the full formation then it can move simultaneously to cover a new position – a benefit of keeping a formation while moving. If the desired position cannot be occupied by a specific UGV, due to an obscuring object or another constraint, the formation becomes approximate. However, the UGVs can learn the new approximation. The UGVs that are out of formation enter a following mode by solving a different correspondence problem, namely that of staying out of the way of the UGVs that are in formation, but following them while waiting for a chance to get in formation again and opportunistically exchanging positions with other UGVs already in the formation. Note that all UGVs solve the same correspondence problem of moving in formation to new positions in a distributed manner knowing their unique ID but also what all the other UGVs will be attempting, given a predefined disinfection algorithm.

The proposed distributed algorithm is modeled by solving a shortest-path problem on a graph in space time. The space is quantized as a 3D bin space. Each bin is represented by its tag or the coordinates of its center and represents a vertex of the graph. Each bin represents a location where objects are allowed to exist. Transitions from one bin to another are allowed explicitly by inserting edges connecting respective bins. As objects come in and go out of the scene (e.g., other UGVs), bins can be either *occupied* or *free*; in the former case all edges to this bin are erased; in the latter case edges connecting the bin to other neighboring *unoccupied* bins are inserted. This is a local to the bin operation and can efficiently be performed in a completely distributed fashion (bits in a binary mask representing the graph). Each UGV runs the same algorithm, and at any given time each UGV has the following information, common to all UGVs (5G guarantees common real-time information).

1. The space bin graph (defines occupied bins in space, including other UGVs).
2. A certain formation the UGVs have to follow.
3. A protocol of synchronous movement related to

the current formation.

4. A target object to be disinfected and a predefined algorithm of movement to do so.

Step 3 above requires a synchronous phase. Alternatively, it can be implemented by circulating a token between UGVs in an order dictated by the specific formation, this approach can be implemented seamlessly on top of 5G networking protocols necessary for the communications between UGVs. This unification of movement, recognition and disinfection task in the context of networking protocols is the main innovation of the proposed use case system particularly important in view of the mutual optimization choices and seamless 5G integration.

The UGV currently holding the token solves the shortest path algorithm considering as occupied bins the ones occupied by all the existing objects in the environment together with their (immediate) future positions (extrapolated by speed and direction information derived from the predefined disinfection algorithm of movement) and the UGVs that already had the token and already moved to their new positions. In other words the order in which the token is exchanged (a topological parameter of the current formation) governs the UGVs' order of movement and is the temporal part of the proposed algorithm. For a certain UGV there is a preference for the bins closest to it that keep it in the task and in formation. This preference is modeled by introducing weights related to the edges connecting the bins. If the UGV can move to the desired bin that keeps it in formation the algorithm continues as above and the collective movement remains normal and in formation.

Networking communications are also optimal (maximizing the benefit of 5G infrastructure in terms of bandwidth requirements for the ad hoc network between UGVs and stationary nodes) and sensor data are also optimally exchanged. If a certain UGV is not able to move to the desired bin (according to formation position) it drops to a following mode, by moving at the end of the queue in the next token run, or changes disinfection task by moving into another area. This, in practice, means that at the next run it will move last, trying to occupy some other valid position. Its token order will move upwards progressively as other UGVs lose formation and become last in the token queue. When a 'following' node succeeds in entering a valid formation bin, it moves upwards in the token queue, in front of all the UGVs that remain in following mode. Some important observations:

- Each UGV solves an "all shortest paths" algorithm according to Ref. [11] which is suitable for real time application. It can thus predict movements of other UGVs and thus optimize its own movements during disinfection.
- The concept of "formation" is used only as a "topological" concept, i.e., it dictates the order of exchanging a token for movement and possibly the networking protocol for communications that is optimized by the certain formation and 5G infrastructure. Specific to the certain topological formation metric parameters (distance between the UGVs, closest distance permitted for other objects, etc.), can be dynamically adapted by modifying bin sizes in the proposed graph representation in 3D space.
- Bin sizes, e.g., quantization/granularity of time and space resolution can be dynamically adjusted according to terrain/mission constraints; speed regulation can also be modeled this way.
- The direction of onboard UGV cameras and UV lamps can be dynamically directed to certain bins resolving ambiguous graph-connectivity issues and optimizing bin resolution parameters.

3 The progressive object detection algorithm

According to the terrain modeling as a 3D bin grid, the object-detection procedure initially will mark certain bins as 'occupied' when an object appears in the range of detection. The contribution of the stationary cameras (or 'stat UGVs') is important towards this direction. Each stationary camera observes a certain range of bins. A simple background removal can be implemented with simple hardware at the stationary camera by superimposing the current captured image on a clean plate containing a static background. A binary quantization of the resulting image can produce a binary mask marking the coordinates occupied by the object. All such masks are transmitted to UGVs whenever a significant change happens in the scene.

At the UGV side, binary masks are collected from stationary camera broadcasts. A pair of orthogonal views of the same scene (same bins), results in extending the bin mask as a 3D representation. A combination of several masks from various stationary views of the same scene further refines the 3D bin mask showing the occupied bins. This procedure of assembling the 2D stationary masks to 3D bin masks and the respective updates of the occupied bins happens at the UGV side using the stationary data. It is the object-detection phase where the interest is restricted to the specific bins occupied by the objects. A UGV can now use its onboard camera to focus on its respective bins and take a high-definition image/video of the respective object. These high-resolution (hi-res) images containing the object of interest can now be used locally in combination with the masks for recognizing the object. The proposed use-case innovation is due to 5G infrastructure that allows transmitting hi-res images to the fleet (moving and stationary nodes) for coordinated disinfection response based on distributed object recognition.

A benefit of formation at this stage lies in the ability of the UGVs to share the observed space by the onboard cameras in a distributed manner in a way that will maximize the collectively covered area (without a given formation this problem of arranging the onboard cameras of different UVGs to cover a wide area scene with non-overlapping views would be highly complex and in some cases unsolvable). At the same time, recognition algorithms based on DNNs perform better in the presence of multiple views of the same object.

The ability to share high-resolution image/videos, taken by onboard UGV cameras and stationary nodes, among all the members of the group with the purpose of performing distributed object recognition from multiple angles and combining results for the appropriate disinfection algorithms to be applied, is what makes this use-case scenario specific to 5G infrastructures. Without 5G high-bandwidth, low-latency and low-error characteristics sharing this volume of data between the fleet would be impossible and, in fact, is one of the major limitations in state-of-the-art object-recognition systems.

In the proposed use-case scenario the recognition is enhanced due to the existence of 5G infrastructure that permits sharing high-bandwidth, real-time critical data for distributed object detection. An object occupying certain bin(s) can be seen from multiple views performed from all UGV(s) (some with better view than others). Intermediate recognition steps depending on the specific algorithm used can also be transmitted back and forth among the members of the fleet to better guide the formation for acquiring more and better images. At the same time, different UGVs can focus and recognize different objects by transmitting data and findings to each other, and this way maximizing the speed of disinfection response if multiple objects exist around the fleet.

To achieve high-speed response in formation change for object recognition and subsequent disinfection, a progressive algorithm based on shape analysis from 2D views is proposed here as the first step in a distributed pipeline of recognition tasks that will end in DNN recognition from multiple sources based on boosting techniques. From the 3D bin mask a UGV can perform a contour extraction from various views and perform credible shape-based recognition from a locally kept database of objects. At the same time, high-resolution images of the same object taken by onboard fleet UVG cameras can be used for DNN recognition at each UGV, with different NN architectures. The combination of multiple methods in the context of boosting [13], by means of distributed fleet processing can lead to enhanced object recognition due to 5G infrastructure. This, together with the ability to exchange/transmit high-resolution/bandwidth real-time data back and forth among the fleet during the recognition and disinfection tasks, casts the proposed solution as a significant use-case scenario of 5g networking technologies.

Regarding the shape-based recognition task, in particular, it can initially guide the formation. To this end, the GLS/VAR method [11] is proposed because its benchmark matching scores are, to the best of our knowledge, the best one available today for real-time applications. GLS/VAR is a descriptor of planar shapes computed on the closed boundary. It uses a distance map between all pairs of points to represent shape and the VAR descriptor to define point correspondences between different shapes. The descriptor for each shape is extracted from its boundary together with its landmark points. The *landmark points* are used for cross-shape correspondences at the comparison phase. At the time of writing this proposal, it is, to the best of our knowledge, the fastest method that can achieve above 70% bulls-eye agnostic score on the MPEG 1400 benchmark dataset of shapes, thus it is better suited for large scale search and indexing applications in the wild. Another important property of the GLS/VAR descriptor is its global nature and the resulting resistance to boundary noise[12].

3.1 Proposed Weighted BFT algorithm

A weighted Byzantine Fault Tolerance mechanism is proposed, considering the hierarchical structure of the UGV fleet. This paves the way to implement a more robust decentralized system that would effectively handle the situation in the case of UGV malfunction or coordination failures of the fleet that would lead to non-proper disinfection of an area. This approach significantly improves the quality of the decision process by taking into account inputs coming from various types of UGVs in a way that, based on their roles and importance in their area of responsibility, would be appropriate. As the influence of such critical nodes will be balanced by the localized insights of individual UGVs, this weighted approach could lead to a much more well-rounded decision.

The hierarchical structure that we propose can be analyzed as follows: Firstly, the leader UGVs (L-UGVs at the top of the hierarchy, is the first-in-line coordinator and holds the highest decision-making weight through strategic overviewing and command capabilities. Therefore, it is responsible for superordinate mission planning overall resource allocation, and critical decision-making to ensure that the fleet's activities are executed according to the set objectives. Secondly, regional coordinator UGVs (RC-UGVs) at the mid-level, are intermediate, moderate-weight UGVs responsible for managing small, specific subgroups of individual UGVs within defined regions. They aggregate information from their assigned UGVs, process it into more specific forms, make decisions relevant to the small areas they cover, and send critical information back to the lead UGV. At the base of this hierarchy, individual UGVs have the last word in execution and relate to tasks such as capturing real-time data, disinfection, and detecting problems immediately. UGVs bear the least weight in the decision process but are, nevertheless, front-line devices for localizing information and grounding all fleet operations in environmental actualities. In combination, these three classes of UGVs thus form a complementary and well-balanced system capable of strongly decentralized decision-making and robust execution of assigned missions.

From a mathematical perspective, several key points should be analyzed to identify the response as formulated by the majority. Firstly, each UGV type has a specific weight w_i, where $i = l, r, u$ and $w_l > w_r > w_u$, for leader UGVs, regional UGVs, and individual UGVs respectively. The total weight is calculated by the sum:

$$\mathcal{W} = \sum_{i=l,r,u} (n_i \times w_i),$$

where n_i is the number of specific role UGVs.

To ensure that only a supermajority agreement can make a decision, the total threshold weight for consensus is set to be equivalent to two-thirds of the total weight. Promote stability rather than a quick and rash decision that may be unbalanced, given the state at a particular moment. This level of support will improve the decision, as every action taken is most likely well-thought-out and supported by an array of inputs, while preventing domination at any level or type of UGV, leading to balanced participation of different kinds of vehicles. Therefore the threshold $\mathcal{T}$ is calculated by $\mathcal{T} = 2/3 \times \mathcal{W}$, where $\mathcal{W}$ is the total weight.

For a specific action $\mathcal{A}$, the total weight sum of UGVs supporting such an action within the context provides the total weighted vote $V_{\mathcal{W}}(\mathcal{A})$. Here, the total weighted vote for a specific action ensures that influence from each UGV is correctly accounted for according to its hierarchical weight. The weight indicates the role and importance of UGVs to the fleet. This uses an indicator function based on UGVs responses R_i, which is denoted by $\delta(\mathcal{A}, R_i)$ to sum only the weights of UGVs that support the given action, where it equals one if a UGV supports the given action and otherwise is equal to 0. This adds weighted support to any particular stance that underpins the consequential decision in a fully transparent manner and balances the strategic oversight, where higher-weighted UGVs are, with the operational insights, where lower-weighted UGVs are, to drive rounded and robust decisions that reflect the overall fleet's input. The mathematical equation for this is the following:

$$V_{\mathcal{W}}(\mathcal{A}) = \sum_{i=l,r,u} (n_i \times w_i \times \delta(\mathcal{A}, R_i))$$

In this way, while incorporating the critical real-time data from the operational UGVs, strategic decisions get to leverage the much broader perspective provided by these higher-weighted nodes.

In order to determine the majority response M_R, the algorithm will compare the weighted vote for each action against the threshold weight. It will iterate through all possible actions and tests if the weighted votes of the considered action are meeting or exceeding the threshold weight; once this condition is met, a choice based on this result will be made as the majority response to execute. Thus, only those actions with substantial support from the fleet, as seen by the weighted consensus, get executed, hence upholding stability and avoiding the making of unbalanced or premature decisions:

$$M_R = \mathcal{A}, \quad \text{if} \quad V_{\mathcal{W}}(\mathcal{A}) \geq \mathcal{T}.$$

The proposed consensus mechanism is illustrated in the figure below (Figure 1), while the pseudocode for this approach is presented by Algorithm 1.

Algorithm 1 Weighted Consensus Mechanism for Emergency Situations in UGV Fleet

```
Initialization:
Initialize weights for each type of UGV:
Define consensus threshold
Initialize the list of UGVs and their roles
while True do
    for each UGV in UGVs do
        if UGV detects an emergency then
            Set emergency ← Active
            Get emergency data from UGV
            Broadcast the emergency data to all UGVs
        end if
    end for
    if emergency is Active then
        emergency responses ← ∅
        for each UGV in UGVs do
            Each UGV creates a response based on the emergency
            Add the response to the list of emergency responses
        end for
        Calculate the weighted votes based on the emergency responses
        for each response in emergency responses do
            UGV_id, action ← response
            weight ← get_node_weight(UGV_id)
            if action ∉ weighted_votes then
                weighted_votes[action] ← 0
            end if
            weighted_votes[action]+ = weight
        end for
        total_weight ← sum(weighted_votes.values())
        threshold_weight ← total_weight × 2/3
        majority_response ← None
        for each action in weighted_votes do
            if weighted_votes[action] ≥ threshold_weight then
                majority_response ← action
                break
            end if
        end for
        if the action in the majority response is "reroute" then
            for each UGV in UGVs do
                Update the route based on the majority response
            end for
        else if the action in the majority response is "assist" then
            for each UGV in UGVs do
                if UGV's location is in the list of assist locations from the majority response then
                    UGV provides assistance
                end if
            end for
        else if the action in the majority response is "redistribute" then
            for each UGV in UGVs do
                Update task assignment based on the new tasks from the majority response
            end for
        end if
        Set emergency ← False
    end if
end while
```

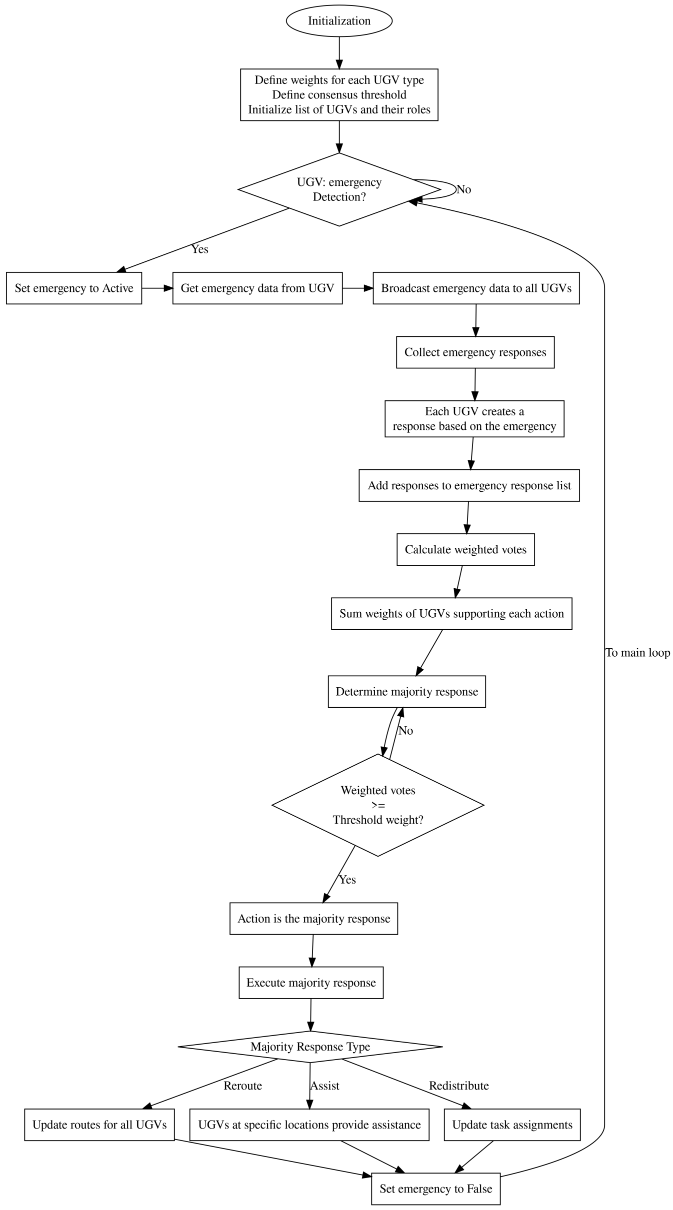

Figure 1: A flowchart diagram for the proposed BFT consensus mechanism

4 Conclusion

Leveraging its extensive expertise in developing specialized hardware for unmanned vehicle platforms (like the fully autonomous UUVs Synoris: [18], [19], [20]), along with its advanced research in computer vision, shape analysis, and machine learning, Infili Technologies S.A. proposes to create a use-case scenario utilizing cutting-edge 5G-enabled recognition and formation technology. The new system can perform recognition and subsequent disinfection tasks in hospitals from mainly visual sensor data by appropriate and controlled formation positioning of a fleet of UGVs by using their onboard adjustable UV lamps, in a distributed and optimal manner, empowered by 5G connectivity between the system's nodes.

Acknowledgment:
It is an optional section where the authors may write a short text on what should be acknowledged regarding their manuscript.

References:

[1] Joseph Redmon and Ali Farhadi, YOLOv3: An Incremental Improvement, arXiv, 2018

[2] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI), pages 234–241. Springer, 2015.

[3] S. Yu, B. Park, and J. Jeong. Deep iterative down-up CNN for image denoising. In CVPR Workshops, 2019.

[4] B. Park, S. Yu, and J. Jeong. Densely connected hierarchical network for image denoising. In CVPR Workshops, 2019.

[5] [15] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. In CVPR, pages 4700–4708, 2017.

[6] [14] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In CVPR, pages 770–778, 2016.

[7] [43] Y. Zhang, Y. Tian, Y. Kong, B. Zhong, and Y. Fu. Residual dense network for image super-resolution. In CVPR, pages 2472–2481, 2018.

[8] Y. Zhang, Y. Tian, Y. Kong, B. Zhong, and Y. Fu. Residual dense network for image restoration. arXiv preprint arXiv:1812.10477, 2019.

[9] Y. Zhang, K. Li, K. Li, L. Wang, B. Zhong, and Y. Fu. Image super-resolution using very deep residual channel attention networks. In ECCV, pages 286–301, 2018.

[10] M. Haris, G. Shakhnarovich, and N. Ukita. Deep backprojection networks for super-resolution. In CVPR, June 2018.

[11] Konstantinos A. Raftopoulos and Stefanos D. Kollias. The Global–Local transformation for noise resistant shape representation. Computer Vision and Image Understanding,volume = 115, number = 8, pages = 1170-1186, year = 2011

[12] Konstantinos A. Raftopoulos, Stefanos D. Kollias, Dionysios D. Sourlas, and Marin Ferecatu. 2018. On the Beneficial Effect of Noise in Vertex Localization. Int. J. Comput. Vision 126, 1 (January 2018), 111–139. https://doi.org/10.1007/s11263-017-1034-6

[13] Chen, Tianqi and Guestrin, Carlos. XGBoost: A Scalable Tree Boosting System. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pages=785–794, year=2016, organization=ACM, doi=10.1145/2939672.2939785

[14] A. Artusi and K. A. Raftopoulos, "A Framework for Objective Evaluation of Single Image De-Hazing Techniques," in IEEE Access, vol. 9, pp. 76564-76575, 2021, doi: 10.1109/ACCESS.2021.3082207.

[15] Chen, Chang Wen et al. "Deep Boosting for Image Denoising." European Conference on Computer Vision (2018).

[16] Xu, Jun et al. "External Prior Guided Internal Prior Learning for Real-World Noisy Image Denoising." IEEE Transactions on Image Processing 27 (2017): 2996-3010.

[17] Zhang, K., Zuo, W., Zhang, L. (2018).FDNet: Toward a fast and flexible solution for CNN-based image denoising. IEEE Transactions on Image Processing, 27(9), 4608–4622. DOI: [10.1109/TIP.2018.2839891] (https://doi.org/10.1109/TIP.2018.2839891)

[18] Nikolaos Papadakis, Dimitris Paraschos, "Synoris: an Unmanned Underwater Platform Based on Hydrophone Arrays for Detection and Tracking From Sound Signatures," International Journal of Circuits, Systems and Signal Processing, vol. 16, pp. 780-786, 2022, DOI: 10.46300/9106.2022.16.96.

[19] Dimitrios Paraschos, Nikolaos Papadakis, Evangelos Mantas, 12th NMIOTC Annual Conference 1-2 June 2021, SOUDA BAY GREECE, SYNORIS – The ”Chariot” of Modern Maritime Warfare, NATO Matitime Interdiction Operations journal Issue 22 (1st Issue 2021) ISSN 2242-441X.

[20] Paraschos D, Papadakis N. K. Autonomous underwater vehicle challenge: design and construction of a medium-sized, AI-enabled low-cost prototype. The Journal of Defense Modeling and Simulation. 2024;21(3):269-281. doi:10.1177/15485129211027236.

[21] Castro, Miguel, and Barbara Liskov. ”Practical byzantine fault tolerance.” OsDI. Vol. 99. No. 1999. 1999.

Contribution of Individual Authors to the Creation of a Scientific Article (Ghostwriting Policy)

The authors equally contributed in the present research, at all stages from the formulation of the problem to the final findings and solution.

Sources of Funding for Research Presented in a Scientific Article or Scientific Article Itself

No funding was received for conducting this study.

Conflicts of Interest

The authors have no conflicts of interest to declare that are relevant to the content of this article.