The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Volume XLVIII-1/W2-2023
ISPRS Geospatial Week 2023, 2–7 September 2023, Cairo, Egypt

POTENTIAL-GUIDED UAV-FLIGHT PATH PLANNING FOR THE INSPECTION OF COMPLEX STRUCTURES

P. Debus[1] *, V. Rodehorst[1]

[1] Computer Vision in Engineering, Bauhaus-Universität Weimar, 99423 Weimar, Germany - paul.debus@uni-weimar.de

KEY WORDS: UAS, Inspection, Photogrammetry, Path Planning, Optimization.

ABSTRACT:

This works presents a method to compute a potential field for UAV flight path planning for the image-based inspection of complex structures. The potential field is based on the photogrammetric requirements of the inspection and allows to estimate the potential of a viewpoint position to improve the coverage and reconstructability of the structure, based on a model of the structure and the previous viewpoints. This potential field can be used to guide the planning of the flight path to produce globally optimized flight paths, that minimize the number of required images and the length of the flight path, while achieving complete coverage, high resolution and stable reconstructability. The estimation of the potential is implemented efficiently and continuously differentiable to allow for the use in optimization algorithms. In addition, a general framework for potential-based flight path planning is presented that allows to integrate different requirements and constraints into the planning process, outlining the requirements for the implementation of a flight path planning method using this framework. The method is evaluated and validated in a synthetic realistic scenario, where it demonstrated high performance and reliable results. It is able to identify regions of high potential even in complex environments, for example the underside of an overpass of tight corners that often pose challenges for existing methods.

1. INTRODUCTION

The application of unmanned aircraft systems (UAS) for the inspection of important infrastructure has been gaining traction recently. Using these mobile sensor platforms to deliver various equipment, such as photo cameras, to flexible locations around the structure allows efficiently inspecting those parts that are not easily accessible to humans. This allows to simplify and speed up structural inspections, which are essential for guaranteeing save operation of critical infrastructure, such as bridges, dams, or power plants, described exemplarily in (Jeong et al., 2020). Using digital images for the inspection also allows using digital methods for the evaluation of the data, supporting human experts, reducing manual efforts, and increasing reliability and objectivity of the inspection (Brilakis and Haas, 2020). Important uses of the captured images include the automatic detection of different anomalies (Benz and Rodehorst, 2022) or the generation of 3D models of the inspected structure (Morgenthal et al., 2019) using photogrammetric methods.

The success of these methods depends on the quality of the available data, the digital images that are captured. These have to fulfill different requirements regarding the complete coverage of the structure, a sufficient resolution on the surface of the structure (sometimes also called ground sampling distance, GSD), and the arrangement of the camera positions around the structure to allow for a successful 3D reconstruction. These requirements need to be considered from the beginning, especially during the on-site capturing of images. To achieve this, the flight paths for the unmanned aerial vehicle (UAV) are commonly planned in advance, as described in (Maboudi et al., 2023). Planning the flight path before the acquisition not only increases the quality of the captured data, but also allows to reduce the time and effort needed for the inspection and enables the complete automation of this process. These planned flight paths typically consist of a series of waypoints for the UAV, defined by their position and orientation that the drone flies through using a global navigation satellite system (GNSS) and captures images at the defined positions.

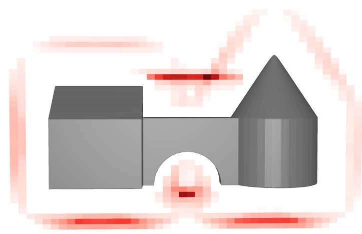

Figure 1. The potential field identifies regions, where viewpoints can improve the target qualities of an inspection. The figure shows a section through the field around a reference scene, where darker color indicates locations of higher potential to inspect the scene.

In this work, we propose a model-based offline method as the foundation of advanced flight path planning algorithms, which allows predicting the potential a given viewpoint position has to improve coverage and reconstructability of a given structure, considering the coverage already achieved by the flight path up to that point. This potential can guide the planning procedure to produce globally optimal flight paths, that minimize the number of required images and the length of the flight path, while achieving complete coverage, high resolution and stable reconstructability. The estimation of the potential is implemented efficiently and continuously differentiable to allow for the use in optimization algorithms.

In contrast to existing work, different photogrammetric requirements are explicitly encoded in the potential, such as sufficient overlap of adjacent images, and a full 3D problem is considered, where the structure can have any shape, for example with undersides or curves, and the UAV is free to move in 3D space and turn to any direction. This is required to make this method suit-

* Corresponding author

able for the inspection of complex structures such as bridges, where the UAV has to fly below the bridge deck and capture images of the underside of the deck.

After introducing and defining the problem of flight path planning in Section 2 and discussing important related work in Section 3, the proposed method is presented in Section 4, consisting of the generation of the potential field and the update of the field depending on the planned flight path, as well as the description of a general framework for potential-guided path planning. The method is evaluated and validated in Section 5 and the conclusions and future work are presented in Section 6.

2. PROBLEM STATEMENT

The planning of UAV flight paths for the inspection of complex structures has been considered in different formulations and with different constraints. Two fundamentally different problems concern the planning of the flight path online, during acquisition, or offline in advance.

Online methods, that process sensory input from the UAV need to simultaneously explore the scene and capture images with the required quality. Their main advantage is the quick applicability, as no model of the structure is required and the inspection can start without preparation, and the possibility to work with dynamic scenes. However, they are limited in their ability to optimize the flight path globally, as they can only consider the current state of the inspection and initially cannot predict the complete geometry of the scene.

Offline methods use an existing geometric model of the structure to guide the planning, having all information about the structure accessible. While they can only consider static scenes and need initial setup and acquisition of this model, the resulting path can be optimized and reduce the effort during the acquisition. This is especially important for the inspection of complex structures, where the acquisition can be time-consuming and the inspection process is repeated regularly.

In this work, flight path planning is formulated as an offline problem with a geo-referenced surface model of the structure available and a clearly defined inspection task. The goal of the inspection is to produce images that are suitable for two main purposes of analyses, image-based assessment of the surface condition and photogrammetric 3D reconstruction of the scene. These purposes produce a set of requirements on the images to be captured that need to be considered during acquisition and therefore during flight path planning. First, it is required that the structure is completely captured by the images. Second, a specific resolution on the surface (GSD) is required, to allow the detection of anomalies in the images, for example surface damages like cracks or corrosion. This inspection parameter is defined from the specific inspection task. The third requirement stems from the photogrammetric 3D reconstruction, which is commonly performed using the Structure from Motion (SfM) method (Schönberger and Frahm, 2016). This requires a suitable arrangement of the positions of the UAV around the structure while taking the images with regard to the overlap of adjacent images and the angles between them to allow for the reliable triangulation of object points. The third requirement is also called reconstructability in literature (Liu et al., 2022b, Maboudi et al., 2023).

In addition to these quality requirements, which need to be fulfilled for a flight path to be admissible, the number of images and the length of the flight path need to be minimized to reduce the effort of the inspection. Additional images increase the complexity of all further evaluations and computations while not adding value or quality to the results. In conclusion, the path planning problem can be formulated as the search for a flight path that minimizes the number of images and the length of the flight path, while fulfilling the quality requirements of complete coverage, sufficient resolution and stable reconstructability.

A flight path C consists of an ordered series of viewpoints $c_i \in C$ that are visited by the UAV and at which images are captured. A viewpoint is defined by its position in space and its viewing direction, as the exterior orientation of the camera while capturing the corresponding image. In the scope of this work, a viewpoint c_i has five degrees of freedom, three for the position (X, Y, Z) and two for the orientation (θ, φ), which correspond to the Euler angles around the y and z axis. The viewing direction of the UAV is along the positive x axis and rotation around that axis is not considered here. Both position and orientation of a viewpoint are defined in global coordinates, the position with regard to a reference system, for example a local coordinate system suitable for the specific structure, and the rotation as a rotation around the y axis, followed by a rotation around the unrotated z axis. To use such a flight path for image acquisition using a UAV, the coordinates have to be converted into the coordinate system used by the UAS, commonly WGS84.

3. RELATED WORK

The problem of UAS flight path planning has been part of active research and development in recent years, as shown in the review compiled in (Maboudi et al., 2023). Many different approaches exist that use different fundamental ideas and formulations to compute suitable flight paths. In (Liu et al., 2022a), the authors fit four different geometric primitives to the structures and provide optimized flight paths for each shape. By placing viewpoints around a structure and optimizing the viewing directions for coverage and reconstructability, (Koch et al., 2019) provide a method for the integration of semantic information, such as restricted airspace and regions of interest, into the flight path planning. The authors of (Tan et al., 2021) use geometric information from BIM models to compute raster flight paths along building facades for close range inspection. Similarly, the approach proposed in (Wang et al., 2022) uses BIM information, but extends the solution to freeform surfaces and complex 3D shapes with special consideration for edges and corners. By dividing the structure into topological regions, the authors of (Shang and Shen, 2022) solve multiple smaller problems, by iteratively selecting the next best view.

Recently, the reconstructability of the scene as the target quantity for flight path planning has come into focus and multiple approaches have been published. In (Zhou et al., 2023), the authors find a set of initial candidate viewpoints that are iteratively selected based on their estimated improvement in reconstructability of the scene. In (Li et al., 2023), the authors select a minimal set of viewpoints from a large sampled viewpoint candidates to minimize redundancy while maximizing the reconstructability, by iteratively removing those candidates that provide large redundancy and only small reconstructability improvements. To better guide the flight path planning, the authors of (Liu et al., 2022b) train a neural network on previously computed 3D reconstructions and their quality to predict the reconstructability given a rough model of the scene and a set of viewpoint positions. The network inputs are the relative orientation of a surface point and a viewpoint candidate, the output is the predicted improvement in reconstructability.

One notable approach presented in (Ivić et al., 2023) uses a potential field to guide the flight path planning. The potential field is computed from the distance to the surface, depending on the required GSD. By using the Helmholtz partial differential equation, they can use the gradients of the potential field to find a good direction for the next viewpoint. They choose the viewing direction to the closest surface point and update the potential in a region around the viewpoint. While this approach is able to produce plausible viewpoints, the missing integration of the viewing direction into the computation of the potential leads to missing coverage in some regions of the structure, as the coverage model considers the space around the structure and not the surface, which is of interest for the inspection. This approach has however provided the motivation for this work, as it shows the capability of potential fields for flight path planning.

4. PROPOSED SOLUTION

This work proposes a method to compute the localized potential for flight path planning. The potential is computed in the 5D parameter space in which viewpoints are defined and encodes a prediction how good a viewpoint placed at that specific position with the given orientation would be. The proposed method extends the ideas presented by (Ivić et al., 2023) by also considering the orientation of the UAV at the viewpoints for the computation of the potential and including the reconstructability of the scene. By fusing these concepts into one novel approach, the potential field gains expressivity and allows to predict the improvement a possible next viewpoint would provide, considering the already planned flight path and the bundle of the camera positions and object geometry. This enables a computational model to relate surface coverage and improvements to reconstructability to the space of possible viewpoints around the structure. This model can be implemented efficiently and fully differentiable, making it suitable for the use in various optimization methods.

Based on the potential field, a general framework for the planning of UAV flight paths is presented. This framework describes the integration of the potential field into flight path planning methods and outlines the future work to build on the ideas presented here.

4.1 Potential Field Generation

The evaluation of flight paths is an important task for the comparison of different planning methods and the analysis, if a flight path is suitable for a given inspection task. In the literature, different methods are used, in many cases, for example in (Li et al., 2023, Zhou et al., 2023, Shang and Shen, 2022), the completeness and accuracy of the reconstructed scene are computed by simulating the reconstruction and comparing the result to the ground truth, the original model of the scene. In a previous work (Debus and Rodehorst, 2021), a unified evaluation metric for the assessment of flight paths was presented that considers coverage, resolution on the surface, and the accuracy of the 3D reconstruction, without simulating the acquisition and computing a reconstruction. The metric provides localized information about which regions are covered as well as a global score for the entire flight path.

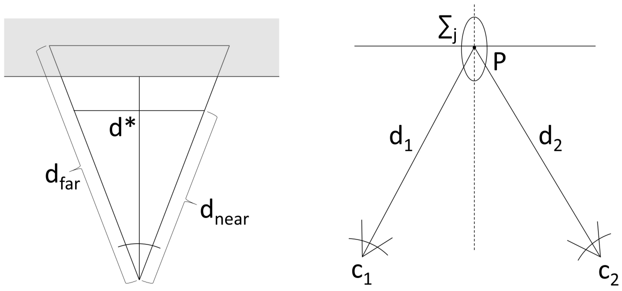

(a) Distances d^*, d_{near}, and d_{far} for the computation of the resolution measure.

(b) Scheme of the accuracy measure

Figure 2. Schematic relations for the accuracy and resolution criteria as defined in (Debus and Rodehorst, 2021).

The resolution metric of this evaluation describes, whether the surface of the object is covered with images from a distance that results in the required GSD on the surface. For this, as depicted in Figure 2a, the minimum and the maximum admissible distance d_{near} and d_{far} are determined around the target distance d^*, which corresponds to the distance from which the used camera takes pictures with precisely the required resolution on the surface. If the distance between a point on the surface and a viewpoint looking at it is between those distances, this part of the structure is considered as covered, and the resolution measure is fulfilled.

The accuracy measure, schematically shown in Figure 2b, is derived from the triangulation error of a surface point. Computing the triangulation from all viewpoints that see the point, the accuracy of the resulting 3D position is computed using error propagation. If the error is below a threshold, which needs to be defined for the specific inspection task, the accuracy measure is fulfilled.

These two measures together are an extension of the commonly used coverage measure, as they also consider the resolution on the surface and use error propagation to compute the accuracy of the 3D reconstruction, instead of simulating the entire process and evaluating the performance of the rendering and reconstruction algorithms.

The improvement of these scores after adding an additional viewpoint can serve as an insightful value for the potential field, as it directly describes the quantities of interest for a successful inspection. Even with the mentioned improvements, the evaluation of an entire flight path is still computationally expensive and therefore unsuitable to be used in an optimization loop, where it would be evaluated many times.

To remedy this, we propose a proxy formulation for this potential, which uses simpler heuristics to reduce computational effort, similar to the computations in (Koch et al., 2019). This function $P : \mathbb{R}^5 \to \mathbb{R}$ estimates the improvement a viewpoint V provides to an existing flight path C. Here it is distinguished between viewpoints $c_i \in C$ that are already planned, and the viewpoint candidate V, which is not part of the flight path. To reduce the amount of computation required, the structure of interested is represented by a set of M points on the surface of the structure $S = \{s_1, s_2, ..., s_m\}$ and the corresponding normal vectors $N = \{n_1, n_2, ..., n_m\}$ at those points. The potential of a viewpoint V is then defined as

$$P(V) = \sum_{s_i \in S} W_i * P_{s_i}(V) \tag{1}$$

where W_i is a weight of the point s_i and $P_{s_i}(V)$ is the potential of the viewpoint V for the point s_i. The weight W_i encodes the already achieved coverage of point s_i and decreases with improved coverage. The potential $P_{s_i}(V)$ is used to scale the weight according to the observation quality of the point and is computed as a function of the distance d between the viewpoint and the point s_i, the angle α between the viewing direction of V and the viewing ray r, and the orientation β of the viewpoint and the surface normal n_i. These measures are illustrated in Figure 3a.

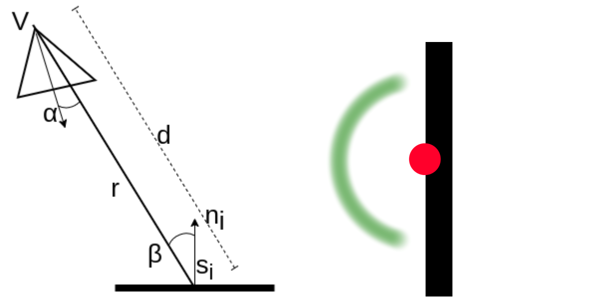

(a) Relation between the viewpoint and the surface point with the different measures considered for the scaling of the visibility.

(b) 2D schematic presentation of potential for different positions to view one point on the surface. The intensity of the color indicates the potential value.

Figure 3. Illustrations of the different concepts and relations involved in the computation of the potential.

As shown in Figure 2a, distance d needs to be inside the admissible distance range between d_{near} and d_{far} to achieve the required resolution on the surface, as described in (Debus and Rodehorst, 2021). The angle α between viewing ray r and the viewing direction of the camera gives an indication, if the point is visible in the image from the viewpoint, by comparing it to the field of view (FOV) of the used camera hardware. Points outside the FOV do not increase the potential of a viewpoint. The third criterion that is considered for $P_{s_i}(V)$ concerns the incidence angle β of the viewing ray r on the surface of the structure, expressed as the angle between the ray and the normal n_i. If this angle becomes too large, the distortions of the surface in the images become very significant and coverage decreases. A common threshold for the angle β, for example in (Koch et al., 2019), is $75°$.

To reward viewpoints close to the target quantities for distance d and the angles α and β, the comparison to those targets is implemented via scaled sigmoid functions. The final potential value $P_{s_i}(V)$ is the product of the three scaled measures, which then is used to scale the weight W_i.

The potential of a viewpoint V is the sum of the potentials for all surface points S, as defined in Equation 1. The reasoning behind this formulation is to detect regions with high potential for good viewpoints. Figure 3b shows a symbolic representation of the potential for capturing a single point on the surface. With a suitable rotation towards that point, the areas of high potential are located on an arc around the surface point, whose middle axis is the target distance d^* away from the point, while the width of the arc corresponds to the admissible distance range.

(a) Viewing direction perpendicular to the surface.

(b) Viewing direction rotated with regard to the surface.

Figure 4. 2D visualizations of the potential field for two points on the surface and different viewing directions.

Figure 4a shows the potential field for two surface points, where the viewing direction is fixed perpendicular to the surface. Two arcs show the locations with potential to observe the object points, while the overlap of these arcs is able to observe both points and accordingly has higher potential. Figure 4b shows the same setup with the viewing direction rotated with regard to the surface. The arcs are rotated accordingly, but do not overlap any more, due to limited FOV and incidence angles.

4.2 Potential Field Update

To consider the already achieved coverage of a given flight path, the potential field needs to be updated, for example after adding a new viewpoint during planning. Areas that have already been covered well do not provide potential for new viewpoints. This update is done by adjusting the weights W_i of the points s_i according to the coverage achieved by the flight path. This coverage can be computed using the computationally expensive metric defined in (Debus and Rodehorst, 2021). For each point $s_i \in S$, the accuracy score Φ_i is computed, as the requirements of the resolution score are already considered in it. The accuracy score is then used to compute the updated weight W_i for each point s_i using a scaled sigmoid function by comparing it to the target accuracy Φ^*, which is defined as part of the inspection task. The weights are initialized with $W_i = 1$ for all points, as no coverage is available, and then updated as follows, where A is a scaling factor and in this work is set to 100:

$$W_i = \frac{1}{1 + \exp((\Phi^* - \Phi_i) * A)} \tag{2}$$

The effect of this update is visualized in Figure 5. The left image shows the potential field before the update, while the images to the right show the updated potential field after adding one and two new viewpoints. The areas close to the placed viewpoints have less potential and the covered object points at the bottom provide less potential to those areas. Note that bad viewpoint constellations can lead to worse scores then before, so adding viewpoints does not always reduce the potential values.

By using the expressive formulation of the accuracy and resolution for updating the weights, the fast implementation of the potential field can be informed with high quality information, providing a good balance between performance and quality.

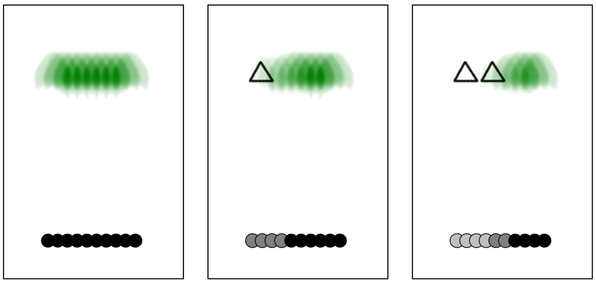

Figure 5. Potential field update after adding a new viewpoint, represented by the black triangles. The object points are shown at the bottom, where lighter color means better coverage and less potential. The intensity of the green color indicates the value of the potential field.

4.3 General Framework for Potential-guided Path Planning

As demonstrated in (Ivić et al., 2023), using potential fields to guide path planning can be very effective. The application of different planning strategies is enabled by having an expressive measure how well a potential viewpoint improves the current flight path. We propose a general framework for potential-guided path planning, which can serve as a basis for different approaches. The fundamental idea corresponds to the next best view (NBV) planning, that is regularly used in literature. The general workflow consists of the following steps:

1. Initialize the potential field and pick a starting viewpoint.
2. Update the potential field with the first viewpoint.
3. Select a suitable viewpoint V using the potential field and add it to the flight path.
4. Update the potential field with the new viewpoint V.
5. Check if the flight path is complete, e.g. all object points are covered with sufficient accuracy. If not, go to step 3.

Specific implementations of this framework need to define at least two parts, the termination criterion and the method to select the next viewpoint. The obvious and very strong termination criterion is to require the potential to be 0 everywhere, so no new viewpoint can improve the flightpath, which can also be expressed as $\sum W_i = 0$. This is however not practical, as small inaccuracies in the rough model of the scene can lead to parts of the surface that are not visible from the outside. A more reasonable termination criterion can be to require the mean weight to be below a certain threshold, for example 0.1, which should provide good results while being practically achievable.

The more complex and more distinguishing part of an implementation is the method to select a next viewpoint, which generally is the focus of research in this area. As demonstrated in (Liu et al., 2022b), a new method to identify promising viewpoint candidates can efficiently be integrated with existing flight path planning methods, such as the planner proposed by (Smith et al., 2018). The proposed potential field can also be integrated with those, as the proposed framework is fundamentally compatible with those other methods.

The specific implementation of the potential field presented here is composed only of continuously differentiable functions and through the use of automatic forward mode differentiation as implemented by (Revels et al., 2016) allows for the automatic computation of gradients of the potential with regard to the viewpoint. This property enables the use of gradient-based methods to select a next viewpoint, which can be very efficient. Through smart choices of the step size, gradient descent can be used to find a local maximum of the potential field, which corresponds to the viewpoint that improves the flight path the most in a local area. To explicitly encode the requirements from photogrammetry – overlap of adjacent images, limited rotation between them and a minimal translation between viewpoints to avoid pure rotations – a set of admissible moves can be generated. Using the potential at each of the resulting positions, the most promising viewpoint can be selected.

Apart from these NBV based planning methods, other methods building on the potential field are possible, for example placing a viewpoint in the maximum of the field, until convergence is reached. These are not considered as part of this framework, as the explicit consideration of photogrammetric requirements is important for successful data acquisition.

5. EVALUATION AND VALIDATION

To evaluate the proposed potential field, different scenarios are considered, where the results show the plausibility and suitability of the potential field. Figure 6 shows two constellations, that commonly are challenging for flight path planning algorithms. Corners require special attention, as the rotations between adjacent images quickly become very large, which can lead to difficulties in the SfM process. As Figure 6a shows, the potential field can provide multiple regions of high potential for such a corner. Depending on the orientation of the viewpoints, expressed by the color in that figure, different regions have high potential and together provide good coverage for the corner. Also, inner corners require special viewpoint placement, as observing them from multiple viewpoints at the required distance is not trivial. Here the potential field demonstrates the ability to identify those suitable regions, which can guide the planning algorithm to place viewpoints in those regions.

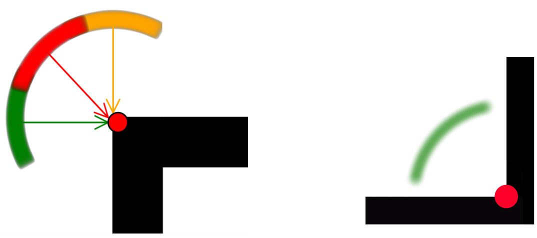

(a) The potential for a point on a corner, color coded according to the viewing direction. Green means looking to the right, orange looking down, and red looking to the lower right

(b) Schematic 2D visualizations of the potential for looking at a point in an inner corner.

Figure 6. 2D visualization of the potential field at corners, which often are challenges for path planning algorithms.

To demonstrate the performance in a synthetic scenario and validate the suitability for complex scenes, the potential field is evaluated on a reference model and the corresponding inspection task definition from the Bauhaus Pathplanning Challenge, as presented in (Debus and Rodehorst, 2021). The chosen scene is an artificial house, shown in Figure 7a, which provides challenging features in the form of an overpass, a round segment, and a pitched roof. The potential field is evaluated in the surrounding of the model. As the potential field has five degrees of freedom and only three dimensions can be visualized, the two rotation dimensions are reduced by computing the sum over them. This results in the visualized potential depending only on the position in space to highlight the performance to adapt to the shape of the structure and the requirements of the inspection task.

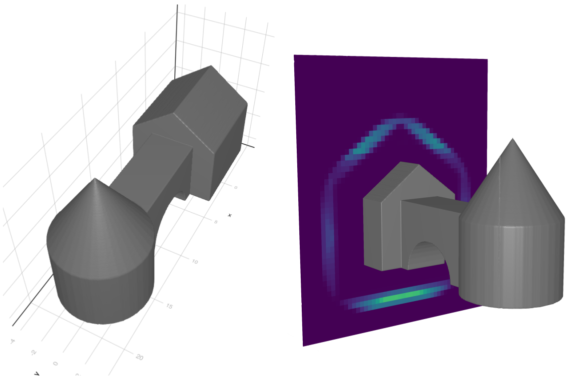

(a) The reference model of the artificial house model, created using a CAD application to provide a challenging reference scene.

(b) A slice through the potential field perpendicular to the x axis. The potential is color coded, with blue being the lowest and yellow the highest potential.

Figure 7. The reference model of the artificial house and a section through the potential field for this scene.

Figure 7b shows one slice through the potential field perpendicular to the x axis of the model. The lighter regions indicate higher potential values. This shows, that the potential is high at the required distance d^* to the surface of the structure and reproduces the shape of the structure. As the ground level is not explicitly encoded in the model, there are also areas of high potential below the structure, which have to be considered by the planning algorithm.

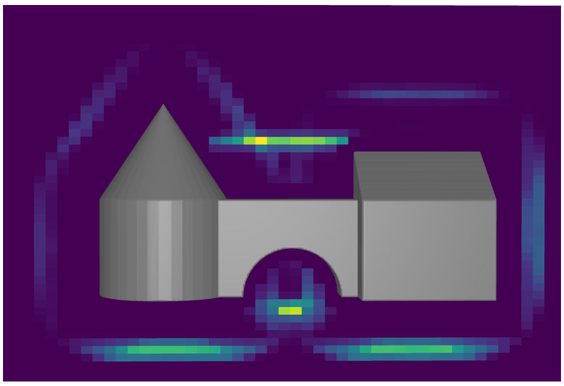

Figure 8. A slice through the potential field for the artificial house model, perpendicular to the y axis.

A section of the potential field perpendicular to the y axis is shown in Figure 8. This shows that the potential field identifies the part below the overpass as important and with high potential, because multiple viewpoints with different orientations are required to capture this part. The top part of the overpass is flat and can efficiently be captured from above, resulting in high potential for that area. This effect of high potential below the overpass is also displayed in Figure 9, which shows a section of the potential field perpendicular to the z axis.

The potential field is not only able to encode the potential with regard to the position, but also to the orientation of the view-

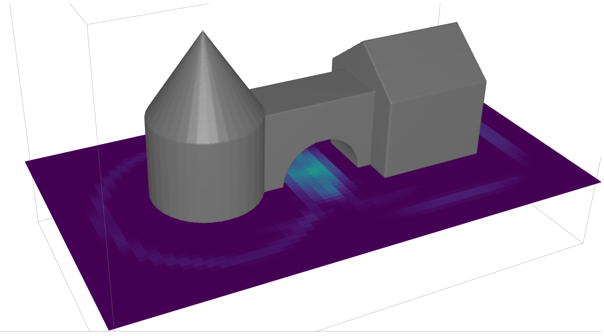

realistic

Figure 9. A slice through the potential field for the artificial house model, perpendicular to the z axis.

points. Figure 10 and Figure 11 show the 10% and 30% of the viewpoint candidates with the highest potential. The arrows begin at the position of the viewpoint and are pointing along the viewing direction of the viewpoint.

The viewpoint markers are arranged around the structure in the target distance, forming a hull, comparable to the distance field computed in (Ivić et al., 2023). The key difference is that multiple viewpoints with high potential can have similar positions in space, but different orientations. This allows good coverage also at corners or tight spaces, for example under the overpass or at the transition between the segments of the structure.

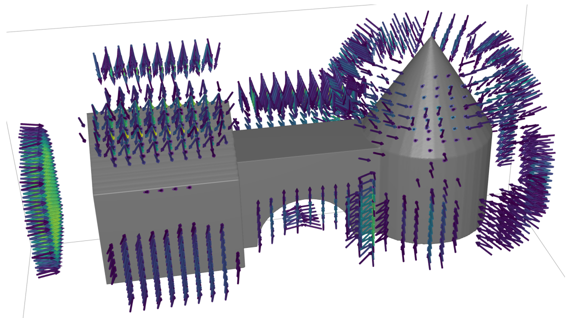

Figure 10. Visualization of 10% of the viewpoints with the highest potential before placing any viewpoint, encoded in the color of the arrows. The arrows are oriented according to the viewing direction of the viewpoint.

Figure 10 and Figure 11 only visualize those regions with the highest potential, which is mostly at the flat sides of the structure, as there many points are visible from one viewpoint. The potential however exists also around the corners to provide adequate coverage in these areas.

To get an impression of the performance improvement of the proposed formulation for the potential field over the complete evaluation presented in (Debus and Rodehorst, 2021), the processing times for both can be compared. Both were computed on the same machine for the same structure with the same set of 200 000 points sampled on the surface. The computation of the full evaluation depends on both the number of points and the number of viewpoints in the flight path, while the computation of the potential field only depends on the number of points on the surface. For a flightpath with 157 cameras, the computation of the full evaluation took around 60 seconds, for a flightpath 1300 viewpoints, the computation took 300 seconds. In both scenarios, the computation of the potential field took less than 5ms.

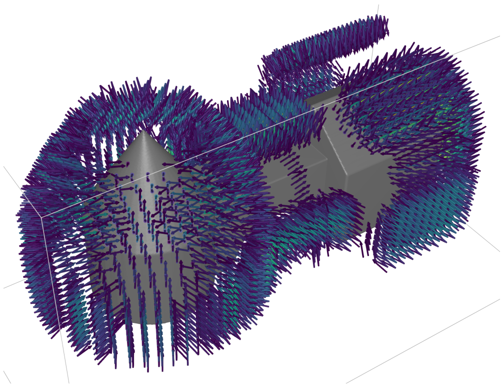

Figure 11. Visualization of 30% of the viewpoints with the highest potential before placing any viewpoint, encoded in the color of the arrows. The arrows are oriented according to the viewing direction of the viewpoint.

In a combined method, where the full evaluation is used to determine the weights for the surface points, the use of the potential field to find promising next viewpoints can accelerate the process by several orders of magnitude, while using the detailed information from the evaluation procedure to compute the achieved coverage.

A simpler updating procedure for the potential field was evaluated, where instead of computing the full evaluation after adding a new viewpoint to the flight path, the weights for those surface points, that are visible from that viewpoint, were reduced by a constant decay factor $\lambda < 1$. A sensible choice is $\lambda = 0.5$, so that a surface point visible from three viewpoints has a weight of 0.125, close to the possible termination criterion discussed in Section 4.3. While not as accurate and explainable as the full evaluation, this approach is much faster, as the visibility can be computed using the same approach as for the potential field, taking only a few milliseconds.

6. CONCLUSIONS AND FUTURE WORK

This work proposes a novel formulation of a potential field for the UAS flight path planning problem, which encodes the predicted improvement a viewpoint placed at a given position provides to the coverage of a structure and the fulfilment of the requirements defined in the specific inspection task. The proposed formulation not only considers the geometry of the structure of interest and the properties of the used camera, but also the coverage achieved by previous viewpoints. Due to a very efficient implementation, the potential field can be used to guide the flight path planning in optimization procedures, where the potential can be evaluated many times.

To update the potential after adding a new viewpoint, an evaluation as described in (Debus and Rodehorst, 2021) can be performed. This allows informing the potential field with precise information about the already achieved coverage.

The method proposed in this work combines multiple partial approaches to generate the potential field. As an extension of similar ideas, a viewpoint is described by five parameters, three for the position of the UAV and two for the orientation. By also including the orientation of the UAV and the visibility of the surface, the coverage of the surface of the structure can be related to positions in space and provide useful information for flight path planning. In addition to the coverage of the structure, the reconstructability of the scene is considered, which has been identified as an important factor for the suitability of a flight path in recent research. This ensures that a 3D reconstruction can be computed from the images captured by the UAV by including photogrammetric requirements into the planning, such as overlap of adjacent images or sufficient baseline between stereo images. In contrast to other approaches to flight path planning, this work puts additional focus on the resolution of the images on the surface, the GSD. Some applications require very high resolutions, that directly influence the flight path and need to be considered from the beginning.

The evaluations performed with the proposed method show, that the update of the potential field after adding viewpoints or otherwise changing the flight path allows the potential field to encode the achieved coverage and the reconstructability of the scene, while being sensitive to changes in the flight path. They also highlighted the ability of the potential field to reliably identify regions of high potential, especially in challenging scenarios, such as bridges, overpasses or sharp geometric features. This makes the approach especially suitable for flight path planning, as existing algorithms often struggle in these scenarios.

The results of this work form a strong basis for advanced flight path planning methods, building further on this method. Using the proposed framework, the integration of the potential field into existing planners, as well as the development of new approaches are possible. The flexible formulation of the potential field can be used as the foundation for flight path planning that aims at achieving global optimality with regard to the number of images, while also achieving complete coverage and high reconstructability.

Building on the results of this work, the following extensions are planned for future research: To provide even more information in the potential, the direction of the predicted errors in the reconstruction can be encoded in the potential field to further give weight to those viewpoints, that are able to reduce the error in the reconstruction.

For the planning of flight paths using the potential field, different approaches seem promising. By combining the gradient-based planning or the selection of the next best move from a set of admissible moves, as described in Section 4.3, with a backtracking algorithm, local minima in the planning can be avoided and exploration and exploitation can be balanced more easily.

Another promising approach is to model the flight path planning as a Markov decision process, allowing even more sophisticated algorithms to be used. Depending on the chosen design for the state space, the same admissible actions, as described above, can be used as the action space and the potential value at the next position can serve as the reward. This allows using reinforcement learning to learn a policy for flight path planning. Especially for large and continuous action and state spaces, this approach can be very promising, as classical methods often do not provide the necessary complexity to model the problem correctly.

These advanced planning methods, together with the proposed potential field, promise to advance the field of UAS flight path planning, enabling even more applications of this methodology, improving the inspection of critical infrastructure, and ultimately increasing safety for their users.

The code developed in this work is available at `https://gropius.medien.uni-weimar.de/debus/adaptivepotentialfield.jl`.

REFERENCES

Benz, C., Rodehorst, V., 2022. Image-Based Detection of Structural Defects Using Hierarchical Multi-scale Attention. B. Andres, F. Bernard, D. Cremers, S. Frintrop, B. Goldlücke, I. Ihrke (eds), *Pattern Recognition*, 13485, Springer International Publishing, Cham, 337–353.

Brilakis, I., Haas, C., 2020. *Infrastructure Computer Vision*. Butterworth-Heinemann, Oxford.

Debus, P., Rodehorst, V., 2021. Evaluation Of 3D UAS Flight Path Planning Algorithms. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLIII-B1-2021, 157–164.

Ivić, S., Crnković, B., Grbčić, L., Matleković, L., 2023. Multi-UAV trajectory planning for 3D visual inspection of complex structures. *Automation in Construction*, 147. doi.org/10.1016%2Fj.autcon.2022.104709.

Jeong, E., Seo, J., Wacker, J., 2020. Literature Review and Technical Survey on Bridge Inspection Using Unmanned Aerial Vehicles. *Journal of Performance of Constructed Facilities*, 34(6), 04020113.

Koch, T., Körner, M., Fraundorfer, F., 2019. Automatic and Semantically-Aware 3D UAV Flight Planning for Image-Based 3D Reconstruction. *Remote Sensing*, 11(13), 1550.

Li, Q., Huang, H., Yu, W., Jiang, S., 2023. Optimized Views Photogrammetry: Precision Analysis and a Large-Scale Case Study in Qingdao. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 1–17.

Liu, X., Ji, Z., Zhou, H., Zhang, Z., Tao, P., Xi, K., Chen, L., Marcato Junior, J., 2022a. An Object-Oriented UAV 3D Path Planning Method Applied In Cultural Heritage Documentation. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, V-1-2022, 33–40.

Liu, Y., Lin, L., Hu, Y., Xie, K., Fu, C.-W., Zhang, H., Huang, H., 2022b. Learning Reconstructability for Drone Aerial Path Planning. *ACM Transactions on Graphics (TOG)*, 41(6), 1–17.

Maboudi, M., Homaei, M., Song, S., Malihi, S., Saadatseresht, M., Gerke, M., 2023. A Review on Viewpoints and Path Planning for UAV-Based 3-D Reconstruction. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 16, 5026–5048.

Morgenthal, G., Hallermann, N., Kersten, J., Taraben, J., Debus, P., Helmrich, M., Rodehorst, V., 2019. Framework for Automated UAS-based Structural Condition Assessment of Bridges. *Automation in Construction*, 97, 77–95. doi.org/10.1016/j.autcon.2018.10.006.

Revels, J., Lubin, M., Papamarkou, T., 2016. Forward-Mode Automatic Differentiation in Julia. *arXiv:1607.07892 [cs.MS]*. doi.org/10.48550/arXiv.1607.07892.

Schönberger, J. L., Frahm, J.-M., 2016. Structure-from-motion revisited. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4104–4113.

Shang, Z., Shen, Z., 2022. Topology-Based UAV Path Planning for Multi-View Stereo 3D Reconstruction of Complex Structures. *Complex & Intelligent Systems*.

Smith, N., Moehrle, N., Goesele, M., Heidrich, W., 2018. Aerial Path Planning for Urban Scene Reconstruction: A Continuous Optimization Method and Benchmark. *ACM Transactions on Graphics*, 37(6), 183:1–183:15.

Tan, Y., Li, S., Liu, H., Chen, P., Zhou, Z., 2021. Automatic Inspection Data Collection of Building Surface Based on BIM and UAV. *Automation in Construction*, 131, 103881.

Wang, F., Zou, Y., del Rey Castillo, E., Ding, Y., Xu, Z., Zhao, H., Lim, J. B., 2022. Automated UAV Path-Planning for High-Quality Photogrammetric 3D Bridge Reconstruction. *Structure and Infrastructure Engineering*, 1–20.

Zhou, H., Liu, Y., Ji, Z., 2023. A UAV Photographic Path Planning Method For High-Quality Reconstruction Of Cultural Heritage. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII-1-W1-2023, 579–586.