

# Vision-Based Uncertainty-Aware Motion Planning based on Probabilistic Semantic Segmentation

Ralf Römer\*, Armin Lederer\*, Samuel Tesfazgi, and Sandra Hirche

**Abstract**—For safe operation, a robot must be able to avoid collisions in uncertain environments. Existing approaches for motion planning under uncertainties often assume parametric obstacle representations and Gaussian uncertainty, which can be inaccurate. While visual perception can deliver a more accurate representation of the environment, its use for safe motion planning is limited by the inherent miscalibration of neural networks and the challenge of obtaining adequate datasets. To address these limitations, we propose to employ ensembles of deep semantic segmentation networks trained with massively augmented datasets to ensure reliable probabilistic occupancy information. To avoid conservatism during motion planning, we directly employ the probabilistic perception in a scenario-based path planning approach. A velocity scheduling scheme is applied to the path to ensure a safe motion despite tracking inaccuracies. We demonstrate the effectiveness of the massive data augmentation in combination with deep ensembles and the proposed scenario-based planning approach in comparisons to state-of-the-art methods and validate our framework in an experiment with a human hand as an obstacle.

**Index Terms**—Planning under Uncertainty, Object Detection, Segmentation and Categorization, Deep Learning for Visual Perception.

## I. INTRODUCTION

WHEN robots operate in the real world, safety requires the avoidance of unintended collisions. In order to detect any form of obstacles, visual perception using deep learning (DL) has gained growing attention in recent years. While DL-based perception systems have achieved impressive results in various tasks, several issues prevent their applicability in many safety-critical systems. Firstly, DL models are typically trained using large datasets [1], [2], which are often not available for custom robotics tasks. However, the creation and annotation of task specific datasets is generally

costly, such that small datasets are desirable. Moreover, DL models are prone to causing prediction errors in previously unseen situations [3], such that a reliable quantification of their predictive uncertainty is important to enable a cautious behavior of the robotic system. Finally, even when uncertainty information is available, it is often not in the parametric form required for uncertainty-aware planning algorithms, such that additional simplifications are necessary [4].

### A. Related Work

A common approach in existing work is to address visual perception by semantic segmentation [5]. In recent years, the focus of segmentation approaches has been on DL architectures due to the impressive accuracy achieved with fully convolutional networks [2]. However, DL is inherently miscalibrated and therefore often produces over-confident predictions [3], such that research has focused on improving uncertainty quantification of DL. While Bayesian neural networks with dropout have attracted interest early on [6], deep ensembles [7] have gained increasing attention due to their often demonstrated superior performance, reasonable computational cost for inference and parallelizability [8]. These advantages have also been shown when applying them to semantic segmentation [9], but they typically require large training datasets and do not extend to low data regimes, in which ensembling alone is generally insufficient to obtain well-calibrated uncertainty estimates [10]. In order to mitigate this limitation, numerous methods for data augmentation have been proposed [11], [12] with the goal of improving the generalizability and robustness. The main focus of these augmentation schemes is usually the avoidance of overfitting, such that the dataset size is commonly increased by merely a low factor, often two, and by applying not more than three augmentation methods [13]. However, this is insufficient to achieve robust calibration for small task-specific datasets which are often not representative in terms of diversity.

The results of semantic segmentation can directly be used for robotic motion planning, but such approaches ignore uncertainties in the perception [14]. Previous works on motion planning with probabilistic environment representation mostly assume known obstacle geometry and Gaussian distributed object position [4], [15], [16]. While these assumptions allow for the derivation of analytical chance constraints, a coarse over-approximation of obstacles can lead to excessive conservatism of planned paths in complex environments. A finer obstacle parameterization or the individual representation of each obstacle can become computationally challenging, especially when the number of obstacles is large, and representing a variation of

Manuscript received: April, 29, 2023; Revised July, 26, 2023; Accepted September, 26, 2023.

This paper was recommended for publication by Editor J. Kober upon evaluation of the Associate Editor and Reviewers' comments. This work was supported by the European Research Council (ERC) Consolidator Grant "Safe data-driven control for human-centric systems (CO-MAN)" under grant agreement number 864686, by the Horizon 2020 research and innovation programme of the European Union under grant agreement number 871767 of the project ReHyb, and by TUM AGENDA 2030, funded by the Federal Ministry of Education and Research (BMBF) and the Free State of Bavaria under the Excellence Strategy of the Federal Government and the Länder as well as by the Hightech Agenda Bavaria.

The authors are with the TUM School of Computation, Information and Technology, Technical University of Munich, 80333 Munich, Germany. Ralf Römer is with the Learning Systems and Robotics Lab (LSY). Armin Lederer, Samuel Tesfazgi, and Sandra Hirche are with the Chair of Information-Oriented Control (ITR). {ralf.roemer; armin.lederer; samuel.tesfazgi; hirche}@tum.de

\* Both authors contributed equally.

Digital Object Identifier (DOI): see top of this page.

the uncertainty in the vicinity of the obstacles is difficult. Safe perception-based navigation has also been addressed using Hamilton-Jacobi reachability [17], but the computational cost increases significantly for irregular obstacle surfaces. Other approaches rely on deterministic error bounds for positions estimated using visual perception, such that robust planning methods can be employed [18]. Since these error bounds require dense training data in practice, large datasets must generally be available for such approaches. Hence, there exists no uncertainty-aware and flexible approach for safe motion planning based on semantic segmentation when merely a small, non-representative dataset is available.

### B. Contribution

We propose a general framework for motion planning with uncertainty based on visual perception via probabilistic semantic segmentation. In the perception module, we enable the training of well-calibrated semantic segmentation models for small, non-representative datasets by combining deep ensembles with massive data augmentation. By modifying training images with eight different methods, we increase the dataset size by a factor of 20, which allows us to obtain reliable probabilistic occupancy information. In the safe motion planning module, we avoid conservatism and computational complexity due to parametric obstacle representations by formulating the path planning problem in an uncertain environment as a scenario optimization problem. This allows us to directly determine collision probabilities using the results from probabilistic semantic segmentation, which we exploit in a scenario chance constrained version of the popular RRT\* algorithm [19]. For the resulting path, we propose safe velocity scheduling to ensure safe robot motion despite tracking inaccuracies. We demonstrate the effectiveness of the proposed data augmentation approach for visual perception and the scenario-based path planning formulation in a comparison to state-of-the-art methods. Moreover, we evaluate the safety and robustness of our framework in an experiment with a robotic manipulator.

## II. PROBLEM STATEMENT

We consider widely-used camera-based perception based on a DL model, which requires labeled data for training. However, due to the lack of suitable training data for many specific robotic applications and the inherent miscalibration of neural networks [3], it is generally difficult to obtain an accurate representation of the environment with reliable information on uncertainty. Therefore, we investigate this problem in this letter as formalized in the following.

*Problem 1 (Uncertainty in visual perception):* Assume a small dataset with low variety is given, which contains RGB images  $\mathcal{C} \in \mathbb{R}^{H \times W \times 3}$  with height/width  $H/W$  of marked obstacles with arbitrary pose, geometry and size. Based on the dataset, we consider the problem of training a DL model  $\mathbf{f} : \mathbb{R}^{H \times W \times 3} \times \Theta \rightarrow [0, 1]^{H \times W}$  with parameters  $\theta \in \Theta$  which outputs for each point  $(i, j)$  in an image  $\mathcal{C}$  the probability of being occupied by an obstacle  $\mathcal{O}$ , i.e.,  $P((i, j) \in \mathcal{O}) = f_{ij}(\mathcal{C}, \theta)$ .

Based on the probabilities of obstacles in the image space, the robotic system must be capable of planning a safe trajectory of poses  $\{\mathbf{p} | \mathbf{p} \in \mathcal{T}\}$  in the task space  $\mathcal{T}$ . Since each pose




Fig. 1: We aim to plan a trajectory  $\pi : [0, \tau] \rightarrow \mathcal{T}$  in the task space such that during the executed motion  $\hat{\pi}(\cdot)$ , any point  $x \in \mathcal{R}(\hat{\pi}(t))$  occupied by the robot  $\mathcal{R}(\hat{\pi}(t))$  lies inside the free workspace  $\mathcal{W}_{\text{free}}$  with a probability of at least  $1 - \delta$  for all  $t \in [0, \tau]$ .

$\mathbf{p}$  implies that the robot occupies some region  $\mathcal{R}(\mathbf{p}) \subset \mathcal{W}$  of the physical workspace  $\mathcal{W} \subseteq \mathbb{R}^3$ , a pose  $\mathbf{p}$  is only collision-free if the set  $\mathcal{R}(\mathbf{p})$  lies completely in the obstacle-free subset  $\mathcal{W}_{\text{free}} \subset \mathcal{W}$ . However, this cannot be ensured deterministically in general as merely an uncertain estimate of the obstacles and consequently  $\mathcal{W}_{\text{free}}$  is available from the visual perception module. Therefore, we aim to satisfy the constraints imposed by  $\mathcal{W}_{\text{free}}$  probabilistically via individual chance constraints with a prescribed probability threshold  $\delta \in (0, 1)$ . This leads to the following formal definition of a  $\delta$ -safe motion, which is illustrated in Fig. 1.

*Definition 1 ( $\delta$ -safe motion):* The motion  $\hat{\pi} : [0, \tau] \rightarrow \mathcal{T}$  executed within the time interval  $[0, \tau]$  is called  $\delta$ -safe if it satisfies  $P(x \in \mathcal{W}_{\text{free}}) \geq 1 - \delta, \forall x \in \mathcal{R}(\hat{\pi}(t)), \forall t \in [0, \tau]$ .

In this definition, safety is introduced using an individual condition for each time instance, which is a commonly considered requirement for planning in uncertain environments [4], [15]. However, these conditions are posed on the motion  $\hat{\pi}(\cdot)$  realized by the robotic system, which is a significantly stronger notion of safety than merely requiring their satisfaction for the planned trajectory  $\pi : [0, \tau] \rightarrow \mathcal{T}$ , which is also illustrated in Fig. 1. Therefore, this safety notion clearly cannot be ensured without additional information about the dynamics of the robotic system, such that we assume the availability of a control law with guaranteed tracking error bounds as formalized in the following.

*Assumption 1 (Velocity-dependent tracking error bound):* The tracking error of the robotic system is bounded by a non-decreasing function  $\gamma : \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+$  of its reference velocity, i.e.,  $\|\hat{\pi}(t) - \pi(t)\|_2 \leq \gamma(\|\dot{\hat{\pi}}(t)\|)$ .

This assumption reflects the fact that the tracking error of robotic systems often grows with increasing reference velocity due to unmodeled effects in control laws, e.g., friction or imprecise inertia parameters. Note that it does not require exact tracking for zero velocity, such that it can also be employed for underactuated robots, which can exhibit relatively large tracking errors at low velocities. Reducing the conservatism for such systems using more sophisticated controllers, e.g., based on barrier functions [18], is subject to future work. Tracking error bounds can be obtained, e.g., statistically from experiments [20]. However, generating test trajectories that cover a sufficiently wide range of operating conditions can be challenging for higher-dimensional task spaces.

In addition to safety, other criteria often have to be considered during planning. While they can be represented using general cost functions in principle, we restrict ourselves

**IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.**

to commonly found path integrals over immediate costs  $c : \mathcal{T} \rightarrow R_{0,+}$ , such that their velocity-independence can be exploited to employ computationally more efficient solution approaches [21]. Therefore, we address the following safe vision-based planning problem.

*Problem 2 (Safe vision-based motion planning):* Given the uncertain estimate of the obstacles obtained from the visual perception system and the velocity-dependent tracking error bound in Assumption 1, we consider the problem of finding a trajectory  $\pi(\cdot)$ , which minimizes the cost  $c$  along the path defined through  $\pi(\cdot)$  and ensures  $\delta$ -safety of the resulting executed motion  $\hat{\pi}(\cdot)$ , i.e.,

$$\min_{\pi(\cdot)} \int_0^T c(\pi(t)) \|\dot{\pi}(t)\|_2 dt \quad \text{such that } \hat{\pi}(\cdot) \text{ is } \delta\text{-safe.} \quad (1)$$

### III. VISUAL PERCEPTION WITH UNCERTAINTY REPRESENTATION

In order to address Problem 1, it is necessary to probabilistically solve a classification problem for each pixel, which is commonly referred to as semantic segmentation. Therefore, we briefly introduce the fundamentals of training DL models for semantic segmentation in Section III-A, before we show the extension to an ensemble of DL models to obtain occupancy probabilities for each pixel in Section III-B. Finally, we develop a data augmentation approach to achieve more robust uncertainty estimation in Section III-C.

#### A. Deep Learning for Semantic Segmentation

For classifying each pixel in an image using semantic segmentation, we employ the commonly used approach of fully convolutional networks (FCNs) with an encoder-decoder structure [2]. The encoder applies convolutional and pooling layers, capturing contextual information in a feature vector with downsampled spatial dimensions. In the decoder, the spatial dimensions are upsampled back to the input image size  $H \times W$ , allowing for pixel-wise dense predictions. We apply atrous convolutional layers with different rates in parallel (Atrous Spatial Pyramid Pooling), which allows to capture objects at multiple scales [22]. In the final layer of the model, we employ the softmax activation function for the output channel of every image pixel  $(i, j)$ . Denoting the parameters of the DL model by  $\hat{\theta}$ , this yields a function  $\hat{f}_{ij}(\cdot, \hat{\theta})$  for every pixel, which outputs a probability-like value for pixel  $(i, j)$  not being occupied by an obstacle, i.e.,  $\hat{f}_{ij}(C, \hat{\theta}) \in [0, 1]$ . For the training process, we employ the commonly used cross-entropy loss  $\mathcal{L}'(y_{ij}, \hat{\theta}) = -y_{ij} \log(\hat{f}_{ij}(C, \hat{\theta})) - (1 - y_{ij}) \log(1 - \hat{f}_{ij}(C, \hat{\theta}))$ , where  $y_{ij} = 0$  if  $(i, j) \in \mathcal{O}$ , i.e., pixel  $(i, j)$  lies in an obstacle, and  $y_{ij} = 1$  otherwise. It is computed pixel-wise and summed over the spatial dimensions of the final layer, resulting in the total loss  $\mathcal{L}(\hat{\theta}) = \sum_{ij} \mathcal{L}'(y_{ij}, \hat{\theta})$ .

#### B. Probabilistic Segmentation using Deep Ensembles

Even though DL models for semantic segmentation yield probability-like outputs, these values are generally not well suited as a measure of uncertainty due to the inherent miscalibration and overconfidence of NNs [3]. We address this issue




Fig. 2: We employ eight methods for data augmentation to systematically add missing variety to the dataset. This increases robustness of the trained segmentation model to differences in perceived scenes and reduces the epistemic uncertainty typically caused by insufficient coverage of the input space by the training samples.

by employing an ensemble of models which is known to produce well calibrated uncertainty predictions with suitable training data [8]. An ensemble consists of multiple distinct models, called ensemble members. For training the ensemble, we initialize each ensemble member with random model parameters and randomly shuffle the training data before each epoch [7]. The models are trained independently, such that they capture different features within the data. For inference, the individual member predictions are combined to the final prediction by considering the ensemble as a uniformly-weighted mixture of  $M$  models, i.e.,  $\mathbf{f}(C, \theta) = \sum_{m=1}^M \mathbf{f}(C, \hat{\theta}_m) / M$ , where  $\hat{\theta}_m$  denotes the parameters of the  $m$ -th member concatenated into the overall parameter vector  $\theta$ . Since each function  $f_{ij}(\cdot, \theta)$  also yields values in the range  $[0, 1]$ , but generally exhibits a better calibration, we use it to determine the occupancy probabilities of pixels in the following.

#### C. Data Augmentation for Dataset Diversification

While deep ensembles can also be used for learning from small datasets in principle, it has been demonstrated that the resulting performance in semantic segmentation strongly depends on the size of the training set [1]. Moreover, due to the low variation in the training dataset considered in Problem 1, the estimated epistemic uncertainty may not be well-calibrated when actually employing the deep ensemble in applications with more diverse images [10]. Therefore, directly training a deep ensemble using a small dataset with low variation would be unreliable in safety-critical applications.

In order to mitigate this effect, we propose to massively augment the training data to artificially add the missing variation to the data. The underlying idea of this augmentation is to apply label-preserving transformations to the annotated images, as well as to the corresponding segmentation masks [23]. This idea can be exploited, e.g., when merely few annotated images of the different operating environments are available, to artificially modify their backgrounds as proposed in [24]. Moreover, the robustness of the ensemble against partial object occlusions can be increased by randomly erasing parts of the

**IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.**

labeled object [25]. In total, we identify eight augmentation methods for straightforwardly adding missing information to segmentation datasets, which are illustrated in Fig. 2. To create as diverse training samples as possible, we combine the augmentation methods by successively applying them to the same image. The key idea of our augmentation scheme is to randomly combine seven augmentation methods by applying each with a probability of 0.5, after replacing the background. In that way, we not only apply multiple augmentation methods to the same image, but also vary the set of methods used. As a result, our augmentation scheme introduces much more variety than if certain methods were always applied in the exact same way and order. This allows us to substantially improve the segmentation performance and the quality of the uncertainty quantification using deep ensembles when only a small non-representative dataset is available.

*Remark 1:* While our proposed approach can improve uncertainty quantification, it does not provide calibration guarantees. In order to obtain them, re-calibration techniques can be applied, e.g., by re-scaling the output of the deep ensemble such that obstacle pixels are guaranteed to be correctly classified with the desired probability or higher [26]. Suitable scaling factors can be obtained from test images, e.g., using data-driven optimization approaches [27]. Alternatively, the empirical accuracy of the deep ensemble can be determined using test images, such that generalization guarantees such as [28, Theorem 1] can be exploited to certify a reduced accuracy. Note that the results from both of these approaches are best when the calibration is already high, which underlines the importance of our proposed approach.

#### IV. SAFE VISION-BASED MOTION PLANNING

Based on the probabilistic semantic segmentation results, the goal is to plan and execute a  $\delta$ -safe motion as introduced in Problem 2. To this end, we split the planning problem into path planning and velocity scheduling. For solving the former, we present a scenario optimization approach in Section IV-A and discuss its integration into the RRT\* algorithm in Section IV-B. Subsequently, we determine the maximum velocity profile along the obtained path still admitting a  $\delta$ -safe motion in Section IV-C.

##### A. Path Planning as Scenario Optimization Problem

In order to admit a planning in the task space  $\mathcal{T}$  of a robotic system, it is generally necessary to augment the probabilistic semantic segmentation result to 3D by using depth information, e.g., from a LiDAR scanner, a stereo camera system or available knowledge about the scene. Potential errors in the depth measurements can be considered, e.g., by enlarging the obstacle accordingly in the direction of the camera. This allows us to compute the occupancy probability  $P(\mathbf{x} \in \mathcal{W}_{\text{free}})$  for each point  $\mathbf{x}$  in the 3D workspace  $\mathcal{W}$ . Moreover, we can determine the possibly occupied region  $\mathcal{R}(\boldsymbol{\pi}(t)) \oplus \mathcal{B}(\eta(\gamma(\|\dot{\boldsymbol{\pi}}(t)\|)))$  in the workspace for each pose along a trajectory  $\boldsymbol{\pi}(t) \in \mathcal{T}$ , where  $\mathcal{B}(r)$  is a sphere with radius  $r$ ,  $\gamma(\|\dot{\boldsymbol{\pi}}(t)\|)$  represents the tracking error bound defined in Assumption 1, and  $\eta(\cdot)$  maps the task space error to the

maximum resulting workspace error. The function  $\eta$  can be computed based on the Jacobian mapping task space velocities to the work space. The computation of  $\eta$  simplifies and the conservatism reduces if the tracking error in the task space orientation can be neglected. The condition for a pose  $\boldsymbol{\pi}(t) \in \mathcal{T}$  being  $\delta$ -safe can be expressed as

$$P(\mathbf{x} \in \mathcal{W}_{\text{free}}) \geq 1 - \delta, \forall \mathbf{x} \in \mathcal{R}(\boldsymbol{\pi}(t)) \oplus \mathcal{B}(\tilde{\gamma}(\|\dot{\boldsymbol{\pi}}(t)\|)), \quad (2)$$

where  $\tilde{\gamma}(\cdot) := \eta(\gamma(\cdot))$ . As the dependence on the reference velocity  $\dot{\boldsymbol{\pi}}(\cdot)$  is not suitable for standard path planning algorithms, we substitute it with a desired minimum velocity  $\underline{v} \in \mathbb{R}_+$ . This allows us to split the safe motion planning problem into a simple path planning problem followed by a velocity scheduling problem.

For path planning with uncertainty, strong assumptions are usually made about the shape of obstacles and their probability distribution to derive analytic expressions for the constraint in Problem 2 [4], [15], [16]. To avoid the associated conservatism and directly use the uncertainty estimates obtained from the solution of Problem 1 for path planning, we reformulate Problem 2 as a scenario problem [27]. The path is discretized into  $K \in \mathbb{N}$  poses  $(\mathbf{p}_1, \dots, \mathbf{p}_K)$  and a fixed set  $\mathcal{R}_0 := \mathcal{R}(\mathbf{0})$  is defined, such that for all poses  $\mathbf{p} \in \mathcal{T}$ , the set  $\mathcal{R}(\mathbf{p})$  can be described as a rigid motion  $T^{\mathbf{p}}(\cdot)$  of  $\mathcal{R}_0$ , i.e.,  $\mathcal{R}(\mathbf{p}) = T^{\mathbf{p}}(\mathcal{R}_0)$ . This allows us to approximate the safety constraint using random samples  $\mathbf{x}^{(n,k)} = T^{\mathbf{p}_k}(\mathbf{x}_0^{(n)})$ , where  $N_x \in \mathbb{N}$  vectors  $\mathbf{x}_0^{(n)}$  are drawn from a uniform distribution  $\mathcal{U}(\mathcal{R}_0 \oplus \mathcal{B}(\tilde{\gamma}(\underline{v})))$ . This is a common approach for reformulating robust into scenario constraints [27]. It leads to the scenario optimization problem

$$\min_{(\mathbf{p}_1, \dots, \mathbf{p}_K)} \sum_{j=1}^K c(\mathbf{p}_j) \quad (3a)$$

$$\text{s.t. } P(\mathbf{x}^{(n,k)} \in \mathcal{W}_{\text{free}}) \geq 1 - \delta, \quad (3b)$$

$$\forall n = 1, \dots, N_x, k = 1, \dots, K,$$

with  $\mathbf{x}^{(n,k)} = T^{\mathbf{p}_k}(\mathbf{x}_0^{(n)})$ ,  $\mathbf{x}_0^{(n)} \sim \mathcal{U}(\mathcal{R}_0 \oplus \mathcal{B}(\tilde{\gamma}(\underline{v})))$ . While the planned path is directly affected by the choice of  $\delta$  through (3b), reducing  $\delta$  can have different impact around the obstacle, as the extent of the uncertain area usually varies along the obstacle boundary. In problem (3), the safety constraint is only evaluated at discrete points  $\mathbf{p}_1, \dots, \mathbf{p}_K$  on the path. Still,  $\delta$ -safety of the continuous path can be achieved by enlarging the set  $\mathcal{R}_0 \oplus \mathcal{B}(\tilde{\gamma}(\underline{v}))$ , accordingly. For simple calculation, enlarging its enclosing sphere by  $\mathcal{B}(l/2)$ , where  $l$  is the distance between the positions occupied at the ends of the line segment, can be performed. Based on the solution  $(\mathbf{p}_1^*, \dots, \mathbf{p}_K^*)$  of problem (3), we define the continuous path  $\boldsymbol{\pi}^* : [0, 1] \mapsto \mathcal{T}$  by linearly interpolating between all  $\mathbf{p}_k^*$  such that  $\boldsymbol{\pi}^*(k\Delta s) = \mathbf{p}_{k+1}^*$  for  $\Delta s = 1/(K-1)$ . Since the reliability of the scenario approximation grows with the number of random samples  $N_x$  [27], this approach provides a well-suited obstacle representation for the proposed uncertainty-aware visual perception.

##### B. Path Planning with the SCC-RRT\* Algorithm

For solving the scenario optimization problem (3), we exemplarily employ a modified version of the popular

**IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.**

RRT\* path planning algorithm [19], which we refer to as scenario chance-constrained RRT\* (SCC-RRT\*). The collision checking represents the main difference between our SCC-RRT\* algorithm and previous uncertainty-aware RRT\* variants [29]. In order to evaluate whether a line segment  $[p_1, p_2] \subset \mathcal{T}$  is eligible, the safety constraint (3b) is evaluated for discrete poses  $p^{(i)} = p_1 + i\Delta_p p_2 - p_1 / \|p_2 - p_1\|_2$ ,  $i = 0, 1, \dots, \lceil \|p_2 - p_1\| / \Delta_p \rceil$ , where  $\Delta_p \in \mathbb{R}_+$ . For each  $p^{(i)}$ , we draw  $N_x$  samples  $x_0^{(n)}$ ,  $n = 1, \dots, N_x$ , uniformly from the set  $(\mathcal{R}_0 \oplus \mathcal{B}(\tilde{\gamma}(\underline{v})))$  and apply the rigid motion  $T^{p^{(i)}}(\cdot)$  to each  $x_0^{(n)}$ . If all points  $T^{p^{(i)}}(x_0^{(n)}) \in \mathcal{W}_{\text{free}}$  lie inside the free workspace  $\mathcal{W}_{\text{free}}$  with probability at least  $1 - \delta$ , the line segment  $[p_1, p_2]$  is considered  $\delta$ -safe. The quality of a path is evaluated using a line cost function. By employing the common RRT\* rewiring procedure, our SCC-RRT\* algorithm asymptotically converges to the minimum-cost path if the cost function is monotonous and bounded [19]. As our planning algorithm is global, the risk of getting stuck in local minima due to the unstructured nature of the occupancy estimates obtained from solving Problem 1 is avoided.

### C. Safe Velocity Scheduling

While the path obtained using (3) admits a trajectory executed at the specified minimum velocity  $\underline{v}$ , a faster execution is possible when the obstacles are sufficiently far away. To find an upper bound  $v^*(s)$  on the velocity for  $\delta$ -safety at some point  $\pi^*(s)$ , we need to determine the distance between the robot and the closest point which is not  $\delta$ -safe. This distance can be compactly expressed as

$$d_o(\pi(s)) = \min_{x_r, x_o} \|x_r - x_o\|_2 \quad (4a)$$

$$\text{s.t. } x_r \in \mathcal{R}(\pi(s)) \quad (4b)$$

$$P(x_o \in \mathcal{W}_{\text{free}}) < 1 - \delta, \quad (4c)$$

which can be effectively approximated by sampling multiple positions  $x_o$  and merely optimizing with respect to  $x_r$ . Using the distance  $d_o(\pi(s))$ , we can easily determine  $v^*(s)$  since  $\tilde{\gamma}(\dot{\pi}(s)) \leq d_o(\pi(s))$  must hold for  $\delta$ -safety. Since the chain rule  $\dot{\pi}(s) = \frac{d\pi(s)}{ds} \dot{s} = \pi'(s)\dot{s}$  admits a parameterization of the velocity in terms of  $\dot{s}$ , we can obtain  $v^*(s)$  by solving

$$v^*(s) = \max_{a \geq 0} \pi'(s)a \quad (5a)$$

$$\text{s.t. } \tilde{\gamma}(\|\pi'(s)\|a) \leq d_o(\pi(s)) \quad (5b)$$

$$\|\pi'(s)\|a \leq \bar{v}, \quad (5c)$$

where  $\bar{v}$  denotes the maximum executable velocity of the robot. Unlike path planning, velocity scheduling is not directly affected by  $\delta$ , as (4) is solved only for poses on the path already determined based on  $\delta$ . We discretize the range of  $s \in [0, 1]$  into  $l$  steps and use a line-search to efficiently compute (5) at the discretization points. Finally, numerical integration of the obtained velocity profile yields a smooth  $\delta$ -safe trajectory  $\pi^*(t)$  [21].

## V. EVALUATION

We evaluate the visual perception and motion planning modules separately and deploy the proposed framework for

Training set	Methods	PA $\pm 1$ std [%]	mIoU $\pm 1$ std [%]
200 images	Cutout [11]	88.2 $\pm$ 0.9	78.7 $\pm$ 1.4
	Mixup [12]	89.4 $\pm$ 1.4	78.3 $\pm$ 2.0
	Flip. + 90° Rot. [32]	88.5 $\pm$ 1.0	79.4 $\pm$ 1.5
	Flip. + Rot. + Crop. [13]	86.9 $\pm$ 1.3	76.7 $\pm$ 1.9
	Our scheme	<b>95.4 <math>\pm</math> 0.7</b>	<b>89.1 <math>\pm</math> 1.4</b>
1000 images	Cutout [11]	91.7 $\pm$ 0.6	83.6 $\pm$ 0.8
	Mixup [12]	88.1 $\pm$ 0.8	77.1 $\pm$ 1.1
	Flip. + 90° Rot. [32]	88.2 $\pm$ 1.0	78.1 $\pm$ 1.4
	Flip. + Rot. + Crop. [13]	89.2 $\pm$ 1.1	79.7 $\pm$ 1.7
	Our scheme	<b>96.2 <math>\pm</math> 0.8</b>	<b>91.0 <math>\pm</math> 1.6</b>

TABLE I: Our augmentation scheme significantly outperforms other methods for creating datasets with small (200) and large (1000) size.

Training set	Methods	NLL $\times 10^2$	BS $\times 10^3$
200 images	Cutout [11]	88.1	41.8
	Mixup [12]	68.3	25.0
	Flip. + 90° Rot. [32]	91.7	47.5
	Flip. + Rot. + Crop. [13]	97.4	41.2
	Our scheme	<b>25.4</b>	<b>10.5</b>
1000 images	Cutout [11]	63.1	32.3
	Mixup [12]	87.1	36.9
	Flip. + 90° Rot. [32]	84.1	47.0
	Flip. + Rot. + Crop. [13]	80.2	39.4
	Our scheme	<b>23.4</b>	<b>11.4</b>

TABLE II: Data augmentation has a significant impact on uncertainty quantification. The impact of the size of the training set depends on the used approach for data augmentation.

a real-world task. First, we examine the effect of data augmentation and ensembling in Section V-A. As a segmentation model, we employ the DeepLabv3 architecture [30] with the ResNet50 CNN [31] as the backbone. In Section V-B, we show the high flexibility of the proposed scenario formulation for planning in uncertain environments. Finally, in Section V-C, the effectiveness of the safe vision-based motion planning framework is demonstrated in an experiment with a 7 DOF robotic manipulator and a hand as obstacle.

### A. Uncertainty Quantification in Visual Perception

In accordance with Problem 1, we aim to evaluate our uncertainty-aware visual perception approach for a small dataset with low variety containing images of complex-shaped objects. To this end, we consider highly accurate semantic segmentation of human hands and create our own dataset for this task<sup>1</sup>. Since the labelling procedure is time consuming, only 100 images are created and divided equally into a training and a testing set as depicted in Figures 3a and 3b. It can clearly be seen that the training set is chosen such that it exhibits only limited diversity, e.g., all images have monotonous and very similar backgrounds. In contrast, the test set includes highly diverse images with hands from multiple people in various environments. Therefore, the training set is not representative of the test images, such that the resulting segmentation task is a good instance for Problem 1, and thus well-suited to demonstrate the effectiveness of the proposed massive augmentation procedure in adding missing variability to a dataset.

We evaluate the combination of the data augmentation methods within our augmentation scheme. For this purpose, we apply our scheme and the methods proposed in [11],

<sup>1</sup>The images in the dataset show only hands of co-authors, who consented to the usage and publication.



Fig. 3: We split our available annotated samples into a training and a test set. The former contains only little variety and is thus not representative for the test set, making the training task a good example for Problem 1.

M	1	2	3	4	5	6	7	8	9	10
$BS \times 10^3$	30.8	29.7	28.1	26.9	25.1	24.3	23.9	23.3	23.6	23.4
$NLL \times 10^2$	18.0	16.1	13.9	13.0	12.3	12.1	11.9	11.7	11.4	11.4

TABLE III: Predictive uncertainty improves with increasing ensemble size  $M$ , which is shown by decreasing Brier score (BS) and negative log-likelihood (NLL).




Fig. 4: The reliability curves illustrate the good calibration of ensembles trained with massive data augmentation, justifying the interpretation of the semantic segmentation output as probabilities.

[12], [13], [32] to create two augmented datasets from the 50 training images: A small one containing 200 images and a large one containing 1000 images. In Table I, the results for our scheme in terms of pixel accuracy (PA) and mean intersection-over-union (mIoU) [2] are presented together with the results of the comparison methods. It can be seen that our scheme provides significantly stronger improvements in model performance for both degrees of dataset inflation. The best performance is achieved when the dataset is enlarged from 50 to 1000 images using our scheme, i.e., by a factor of 20, which is ten times the factor of two used in [13]. We also examine the impact of our data augmentation scheme on calibration and compare it to the approaches from [11], [12], [13], [32] for an inflated training set of 200 and 1000 images, respectively. To this end, we consider two common metrics, the Brier score (BS) and negative log-likelihood (NLL) [7]. As shown in Table II, our augmentation scheme yields much better predictive uncertainty estimates than the comparison methods, including mixup, which is a popular method to improve calibration [12]. Our results show the importance of addressing the lacking variety within the training set through massive data augmentation combining different augmentation methods.

Our safe planning framework builds on the assumption that for each pixel, the segmentation output can be used

as an occupancy probability. To justify this assumption, we first examine calibration via the reliability diagram. For this, the confidence interval  $[0.5, 1)$  is partitioned into ten equally sized bins. For each bin, the pixel accuracy is calculated and plotted against the average confidence value of the pixels within the bin. We evaluate three differently sized ensembles,  $M \in \{2, 5, 10\}$ , and the individual members of the medium-sized ensemble, which are trained with our augmented training dataset. The results depicted in Figure 4 show that calibration improves with increasing ensemble size. Moreover, as shown in Table III, BS and NLL significantly decrease with increasing ensemble size, which is consistent with the results reported in [7]. The reliability curve as well as BS and NLL indicate only small improvements for  $M > 5$ . Therefore, we conclude an ensemble size of  $M = 5$  to maintain a good balance between computational demand and calibration suitable for the considered visual perception task.

### B. Uncertainty Representations for Collision Avoidance

We compare the performance of path planning with our perception-based collision avoidance with three popular methods [4], [15], [16] that are based on the common assumption of uncertain position  $\mathbf{x}_o \sim \mathcal{N}(\hat{\mathbf{x}}_o, \sigma^2 \mathbf{I}_3)$ , and known geometry and orientation of the obstacles. We consider a 4D task space  $\mathcal{T}$  composed of the robot position  $\mathbf{x} \in \mathcal{W} \subset \mathbb{R}^3$  and its orientation  $\phi$  around the vertical  $z$ -axis. The set  $\mathcal{R}(p)$  is described as an ellipse parallel to the  $x$ - $y$ -plane, we set  $\underline{v} = 0.01$  m/s,  $\bar{v} = 0.2$  m/s,  $\tilde{\gamma}(v) = v/\bar{v} \cdot 0.01$  m, and consider spheres and cuboids as obstacles. For applying our approach, we consider that the occupancy probability decreases linearly with the distance  $d$  to the obstacle surface<sup>2</sup>, becoming zero at  $d = d_{\text{stop}}$ . Since we employ the chance constraint with probability  $\delta = 0.05$ , this leads to  $\tilde{d} = 0.95d_{\text{stop}}$  representing the extended boundaries of the uncertain object. Due to the Gaussian distribution of the object positions, these boundaries are equivalently parameterized using  $2\sigma$  in the existing approaches [4], [15], [16]. For comparing the different methods, we create a simple scene containing three large obstacles and a more cluttered scenario with eight small obstacles. For solving (3), we employ the SSC-RRT\* algorithm proposed in Section IV-B. The number of search tree iterations is set to

<sup>2</sup>We could just as well assume other profiles for the probabilistic segmentation around the object boundaries. However, many predictions we obtain when perceiving real objects show a roughly linear decrease.

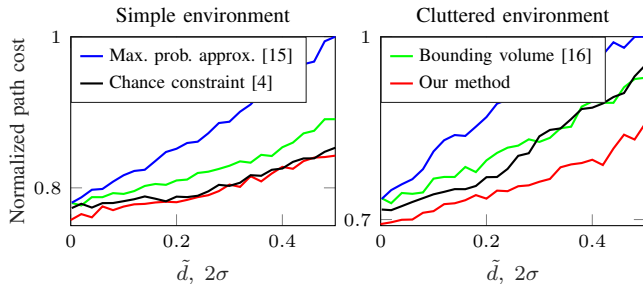


Fig. 5: We compare different uncertainty-aware collision avoidance strategies for path planning within the CC-RRT\* framework. Our perception-based approach leads to less conservative paths, with the advantage becoming stronger for more cluttered environments.

$N_{\text{iter}} = 2000$ . We aim for a path that is short in Cartesian space and also avoids unnecessary rotations. To this end, we define the line cost function for a line segment  $[\mathbf{p}, \mathbf{p}'] \subset \mathcal{T}$  between two poses  $\mathbf{p} = [\mathbf{x}^\top, \phi]^\top$ ,  $\mathbf{p}' = [\mathbf{x}'^\top, \phi']^\top$  in the task space as

$$c_1(\mathbf{p}, \mathbf{p}') = \|\mathbf{x} - \mathbf{x}'\|^2 + r(\phi - \phi')^2, \quad (6)$$

where  $r > 0$  determines the penalty for orientation changes.

We run the SCC-RRT\* algorithm using this cost function and  $N_x = 100$  for 100 times per method and scene over a range of different values of  $\tilde{d}$  and  $2\sigma$ , which represent different uncertainty levels in the perception. The resulting normalized average costs are illustrated in Fig. 5. While the behavior of the cost for the simple scenario is almost identical for our method and the approach proposed in [4], the cost achieved using the existing methods [4], [15], [16] increases significantly faster in the cluttered environment. The proposed SCC-RRT\* algorithm achieves this slower deterioration of the planning performance through the sampling-based obstacle representation, which ensures that our approach does not directly depend on the number of obstacles. This is in contrast to parametric obstacle representations in existing methods [4], [15], [16], where the conservatism of the uncertainty approximations for the individual obstacles accumulates, such that the path performance crucially suffers from growing uncertainties as measured by  $\sigma$ .

Additionally, we compare the computational efficiency of the proposed scenario approach with existing parametric approaches. For this, we create scenes with different numbers of randomly placed spherical obstacles. The average computation times for running the RRT\* algorithm 100 times with the collision checking methods [4], [15], [16] are recorded and shown in Fig. 6. For comparison, we also include the average computation time of our sampling-based approach for different values of  $N_x$ . The parametric methods exhibit a significant increase in computation time with the number of obstacles, which must all be checked for collision. The SCC-RRT\* algorithm does not suffer from this issue as it exploits a joint environment representation that is only sampled at test points instead of individual object representations. Thus, the proposed scenario approach simultaneously achieves flexibility and efficiency and is applicable to highly cluttered environments containing many obstacles.




Fig. 6: The average computation times for the RRT\* algorithm with parametric collision checking increase significantly with the number of obstacles. In contrast, for our vision-based obstacle representation, the computation time only depends on the number of samples  $N_x$ , which is especially beneficial in cluttered environments.




Fig. 7: We experimentally validate our approach with a robotic manipulator.

### C. Experimental Evaluation with a KUKA iiwa robot

In order to show the real-world applicability of our approach, we conduct an experiment with an impedance controlled KUKA iiwa robotic manipulator with seven degrees of freedom and a hand as obstacle<sup>3</sup>. A Logitech C270 USB webcam positioned 0.5 m above the working area takes images of the workspace. We apply the deep ensemble trained as discussed in Section V-A and augment the 2D probabilistic semantic segmentation result to 3D by assuming fixed hand height. We use the same expression for  $\underline{v}$ ,  $\bar{v}$  and  $\mathcal{R}(\mathbf{p})$  as in Section V-B and plan a  $\delta$ -safe path with the SCC-RRT\* algorithm with  $N_{\text{iter}} = 2000$ . From tracking experiments with the manipulator, we obtain an approximate tracking error bound  $\tilde{\gamma}(v) = v/\bar{v} \cdot 0.01$  m. This procedure requires test trajectories that cover a sufficiently wide range of operating conditions, which can be challenging for higher-dimensional task spaces. To obtain a motion in close proximity to the fingers, we replace the first term in the cost function (6) with a term that penalizes the area between the path and the table. The motion is executed with a maximum velocity  $\bar{v} = 0.2$  m/s, see Figure 7. As shown in Figure 8 and in our supplementary video, the experiment is successfully repeated with different individuals, varying lighting conditions and backgrounds, which demonstrates the practicability and robustness of our safe perception-based planning framework. In some situations, parts of the hand are detected by only a subset of the semantic segmentation models, showing the particularly strong impact of ensembling on safety.

<sup>3</sup>Approval for this type of experiments involving the close interaction between a Kuka iiwa manipulator and a human has been obtained by the ethics committee of the medical faculty of the Technical University of Munich.



Fig. 8: The experiment is conducted with different participants, lighting conditions and backgrounds, which are not represented by our training data.

## VI. CONCLUSION

In this letter, we present a framework for vision-based motion planning with uncertainty using semantic image segmentation. We show that combining massive data augmentation and deep ensembles yields good uncertainty quantification for semantic segmentation even for highly specific tasks lacking representative training data. This allows us to interpret the output of the semantic segmentation probabilistically and use it for motion planning with uncertainty. We avoid the conservatism of existing uncertainty-aware path planning approaches by employing a sampling-based method for collision checking that is based on scenario optimization. As a result, our planning method makes no assumptions about the obstacle geometry and can be applied in highly cluttered environments. Our framework is evaluated in simulation and experiment with a robotic manipulator. Directions for future work include using smaller semantic segmentation models, extending our method to estimate obstacle dynamics and enabling adaptive online planning based on a scenario MPC approach, so that it can be safely employed in dynamic environments.

## REFERENCES

- [1] C. Sun, A. Shrivastava, S. Singh, and A. Gupta, "Revisiting unreasonable effectiveness of data in deep learning era," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 843–852.
- [2] E. Shelhamer, J. Long, and T. Darrell, "Fully convolutional networks for semantic segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, no. 4, pp. 640–651, 2017.
- [3] A. Kendall and Y. Gal, "What uncertainties do we need in Bayesian deep learning for computer vision?" in *Advances in Neural Information Processing Systems*, 2017, pp. 5580–5590.
- [4] H. Zhu and J. Alonso-Mora, "Chance-constrained collision avoidance for MAVs in dynamic environments," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 776–783, 2019.
- [5] A. Milioto and C. Stachniss, "Bonnet: An open-source training and deployment framework for semantic segmentation in robotics using CNNs," in *Proceedings of the International Conference on Robotics and Automation*, 2019, pp. 7094–7100.
- [6] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian approximation: Representing model uncertainty in deep learning," in *Proceedings of the International Conference on Machine Learning*, 2016, pp. 1050–1059.
- [7] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," in *Advances in Neural Information Processing Systems*, 2017, pp. 6405–6416.
- [8] Y. Ovadia, E. Fertig, J. Ren, Z. Nado, D. Sculley, S. Nowozin, J. Dillon, B. Lakshminarayanan, and J. Snoek, "Can you trust your model's uncertainty? Evaluating predictive uncertainty under dataset shift," in *Advances in Neural Information Processing Systems*, 2019, pp. 13991–14002.
- [9] A. Mehrtash, W. M. Wells, C. M. Tempny, P. Abolmaesumi, and T. Kapur, "Confidence calibration and predictive uncertainty estimation for deep medical image segmentation," *IEEE Transactions on Medical Imaging*, vol. 39, no. 12, pp. 3868–3878, 2020.
- [10] R. Rahaman and A. H. Thiery, "Uncertainty quantification and deep ensembles," *Advances in Neural Information Processing Systems*, pp. 20063–20075, 2021.
- [11] T. Devries and G. W. Taylor, "Improved regularization of convolutional neural networks with cutout," 2017. [Online]. Available: <https://arxiv.org/pdf/1708.04552.pdf>
- [12] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond empirical risk minimization," in *Proceedings of the International Conference on Learning Representations*, 2018.
- [13] D. Uzun and D. Yi, "Deep ensembles for semantic segmentation on road detection," in *Proceedings of the International Conference on Automation and Computing*, 2021, pp. 1–6.
- [14] L. Bartolomei, L. Teixeira, and M. Chli, "Perception-aware path planning for UAVs using semantic segmentation," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020, pp. 5808–5815.
- [15] C. Park, J. S. Park, and D. Manocha, "Fast and bounded probabilistic collision detection for high-DOF trajectory planning in dynamic environments," *IEEE Transactions on Automation Science and Engineering*, vol. 15, no. 3, pp. 980–991, 2018.
- [16] M. Kamel, J. Alonso-Mora, R. Siegwart, and J. Nieto, "Robust collision avoidance for multiple micro aerial vehicles using nonlinear model predictive control," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2017, pp. 236–243.
- [17] A. Bajcsy, S. Bansal, E. Bronstein, V. Tolani, and C. J. Tomlin, "An efficient reachability-based framework for provably safe autonomous navigation in unknown environments," in *Proceedings of the Conference on Decision and Control*, 2019, pp. 1758–1765.
- [18] S. Dean, A. Taylor, R. Cosner, B. Recht, and A. Ames, "Guaranteeing safety of learned perception modules via measurement-robust control barrier functions," in *Proceedings of the Conference on Robot Learning*, 2020, pp. 654–670.
- [19] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *The International Journal of Robotics Research*, vol. 30, no. 7, pp. 846–894, 2011.
- [20] J. Nubert, J. Köhler, V. Berenz, F. Allgöwer, and S. Trimpe, "Safe and fast tracking on a robot manipulator: Robust mpc and neural network control," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3050–3057, 2020.
- [21] J. E. Bobrow, S. Dubowsky, and J. S. Gibson, "Time-optimal control of robotic manipulators along specified paths," *The International Journal of Robotics Research*, vol. 4, no. 3, pp. 3–17, 1985.
- [22] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 4, pp. 834–848, 2018.
- [23] C. Shorten and T. M. Khoshgoftaar, "A survey on image data augmentation for deep learning," *Journal of Big Data*, vol. 6, no. 60, 2019.
- [24] G. Ghiasi, Y. Cui, A. Srinivas, R. Qian, T.-Y. Lin, E. D. Cubuk, Q. V. Le, and B. Zoph, "Simple copy-paste is a strong data augmentation method for instance segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 2918–2928.
- [25] Z. Zhong, L. Zheng, G. Kang, S. Li, and Y. Yang, "Random erasing data augmentation," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 34, no. 7, 2020, pp. 13001–13008.
- [26] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On calibration of modern neural networks," in *Proceedings of the International Conference on Machine Learning*. PMLR, 2017, pp. 1321–1330.
- [27] M. C. Campi, S. Garatti, and M. Prandini, "The scenario approach for systems and control design," *Annual Reviews in Control*, vol. 33, no. 2, pp. 149–157, 2009.
- [28] E. Bradford, L. Imsland, D. Zhang, and E. A. del Rio Chanona, "Stochastic data-driven model predictive control using gaussian processes," *Computers & Chemical Engineering*, vol. 139, p. 106844, 2020.
- [29] B. D. Luders, S. Karaman, and J. P. How, "Robust sampling-based motion planning with asymptotic optimality guarantees," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, 2013, pp. 1–25.
- [30] L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam, "Rethinking atrous convolution for semantic image segmentation," 2017. [Online]. Available: <https://arxiv.org/pdf/1706.05587.pdf>
- [31] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.
- [32] L. Nanni, D. Cuza, A. Lumini, A. Loreggia, and S. Brahmam, "Deep ensembles in bioimage segmentation," 2021. [Online]. Available: <https://arxiv.org/abs/2112.12955>