

Uncertainty-aware hybrid paradigm of nonlinear MPC and model-based RL for offroad navigation: Exploration of transformers in the predictive model

Faraz Lotfi, Khalil Virji, Farnoosh Faraji, Lucas Berry, Andrew Holliday, David Meger, and Gregory Dudek¹

Abstract—In this paper, we investigate a hybrid scheme that combines nonlinear model predictive control (MPC) and model-based reinforcement learning (RL) for navigation planning of an autonomous model car across offroad, unstructured terrains without relying on predefined maps. Our innovative approach takes inspiration from BADGR, an LSTM-based network that primarily concentrates on environment modeling, but distinguishes itself by substituting LSTM modules with transformers to greatly elevate the performance of our model. Addressing uncertainty within the system, we train an ensemble of predictive models and estimate the mutual information between model weights and outputs, facilitating dynamic horizon planning through the introduction of variable speeds. Further enhancing our methodology, we incorporate a nonlinear MPC controller that accounts for the intricacies of the vehicle’s model and states. The model-based RL facet produces steering angles and quantifies inherent uncertainty. At the same time, the nonlinear MPC suggests optimal throttle settings, striking a balance between goal attainment speed and managing model uncertainty influenced by velocity. In the conducted studies, our approach excels over the existing baseline by consistently achieving higher metric values in predicting future events and seamlessly integrating the vehicle’s kinematic model for enhanced decision-making. The code and the evaluation data are available at (Github-repo).

Index Terms—Model-based RL, transformers, nonlinear MPC, uncertainty-aware planning, offroad navigation

I. INTRODUCTION

Researchers in the field of off-road autonomous navigation have developed robust techniques to effectively guide vehicles towards their destinations [1], [2]. A key strategy involves the use of image-based predictive models, as referenced in [3], [4], trained based on a loss function to meet specific requirements when proposing a trajectory in urban scenarios. Another approach involves optimizing reward functions in an RL scheme, where a model is trained to dynamically model the environment as it evolves when an agent takes actions within it. In this latter area, BADGR [5] stands out for its fusion of model-based and model-free RL, effectively addressing path planning on smooth terrains while minimizing collision risks. Subsequent researchers have capitalized on varied data sources and implemented architectural improvements to further advance this domain. One notable example fuses drone aerial imagery with onboard visuals, improving perceptual capabilities [6], while another introduces a trajectory planner with constrained attention mechanisms

integrated into the BADGR framework [7]. However, these models lack the inclusion of uncertainty and focus solely on the robot’s environment, neglecting its kinematics and dynamics. Addressing this issue would involve integrating motion constraints into planning and considering factors such as road friction, slope, and vehicle states, which are difficult to estimate from images alone.



Fig. 1: A model scale vehicle as used in our data collection.

The previously mentioned predictive models are related to model based control as used in control theory. Deep learning-based controllers excel at processing complex, high-dimensional input data, whereas MPC traditionally relies on a mathematical model to evaluate actions [8]. MPC has seen extensive use in planning, particularly in navigation. An influential method in this domain is the Model-Predictive Path Integral (MPPI) [9] approach, which introduces a stochastic optimal control framework. Numerous studies have been conducted using this technique, aimed at enhancing various facets of its application. For example, [10] incorporates Conditional Value-at-Risk (CVaR) to generate optimal control actions for safety-critical robotic applications, resulting in fewer collisions compared to the baseline MPPI. Additionally, [11] proposes a variable horizon planning MPPI, a topic further explored in the subsequent literature section on uncertainty incorporation. However, these approaches predominantly assume a static world when planning and neglect to model the dynamics of the environment based on the given set of actions. In other words, they tackle the planning problem by predominantly focusing on the vehicle’s dynamics while inadequately representing the environment’s dynamics.

Regarding the image-based predictive model, a conventional approach involves using a long short-term memory (LSTM) network for temporal feature extraction; however, transformers have demonstrated superior performance across a broad spectrum of applications [12], [13]. Initially designed to tackle bottlenecks in natural language processing [14], they have since been found useful in many other domains, including vision [15]. Transformers have been shown to be better than LSTMs at identifying and exploiting long-

¹Authors are with the Mobile Robotics Lab (MRL), Faculty of Computer Science, McGill University, Montreal, Canada F.Lotfi@CIM.McGill.ca

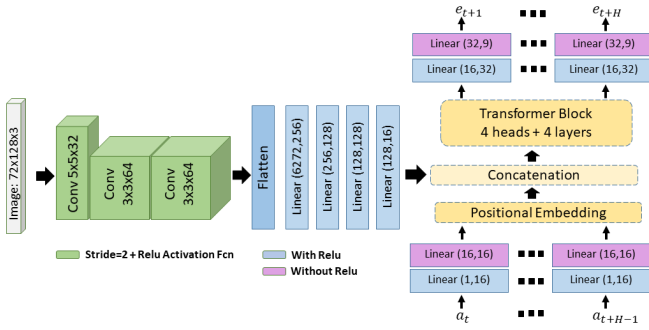


Fig. 2: Our proposed transformer-based architecture.

range dependencies in sequential data. Motivated by this, we explore the application of transformers in constructing a predictive model for more accurate anticipation of future events.

An essential consideration in the design of a planner involves the integration of uncertainty, which serves as an indicator of distribution shifts and the model’s prediction confidence. The utilization of uncertainty as a concept has been extensively explored in various fields [16], [17]. Within the context of navigation, efforts have been made to incorporate uncertainty-aware exploration [18]. In the realm of balancing uncertainty and speed, a comparable strategy to ours can be found in RAMP [11]. This method utilizes LiDAR data and ground point inflation to map uncharted areas. It uses an MPPI-based planner to increase speed in low-uncertainty zones and decrease speed in high-uncertainty areas. Unlike RAMP, which prioritizes reducing speed in uncertainty, our approach optimizes the robot’s trajectory to minimize uncertainty and it incorporates the world’s dynamic into the planning.

In a nutshell, we address mapless navigation challenges using a hybrid approach. We improve an environment predictive model with transformers and integrate uncertainty and nonlinear MPC for a hybrid planner. This approach optimizes robot navigation for speed, collision avoidance, and reduced uncertainty. The practical validation of our approach is conducted using real-world data gathered by an RC car, depicted in Fig. 1. Subsequent sections delve into the methodology and present the results, accompanied by insightful observations. Ultimately, the conclusion section provides the culmination of our paper’s findings.

II. METHODOLOGY

Our method is comprised of two key modules. The first is the image-based predictive model, which models the environment and leverages an RL approach suggesting a set of steering angles at each planning step. The second module is the MPC planner, which utilizes the vehicle’s kinematic/dynamics model as well as the uncertainty in the environment’s model to recommend throttle actions. In the subsequent sections, we provide an in-depth exploration of both modules. Besides, to train the image-based planner, we gathered a dataset and performed manual annotations to

Parameter	Unit	Description	Initial Value
C_1	-	Geometrical	0.5
C_2	m^{-1}	Geometrical	1.69
C_{m_1}	m/s^2	Motor Parameter	12
C_{m_2}	$1/s$	Motor Parameter	2.5
C_{r_2}	$1/m$	2nd-order Friction Parameter	0.15
C_{r_0}	m/s^2	0th-order Friction Parameter	0.7
g	m/s^2	Gravity	9.81

TABLE I: The kinematic model parameters.

ensure its quality and have better evaluation studies.

A. Image-based Predictive Model

The image-based planner operates by analyzing the current image in conjunction with a predefined set of actions, spanning “p” future steps as inputs. This process generates predictions for the forthcoming events. This work is based on the model introduced in [5], which includes the estimation of the heading angle alongside events, facilitating combined classification and regression training. Utilizing these predicted events, a reward function is then formulated, enabling an online model-based RL approach. This involves the use of a global optimizer, such as the cross-entropy method [19] or covariance matrix adaptation evolution strategy (CMA-ES) [20], [21], to maximize the expected reward signal across time. Notably, re-planning occurs at each step, implying that only the initial action from the proposed set is executed. Existing models mainly use LSTM networks, but transformers have shown better performance across various applications. This study replaces LSTM modules with transformers, creating a new architecture as shown in Fig. 2.

In order to estimate epistemic uncertainty we trained an ensemble of 5 image-based predictive models. This enables one to estimate the mutual information (MI), $I(\cdot, \cdot)$ between model outputs, Z , and weights, W , an established measure of epistemic uncertainty within the community [16], [22]. Our model adopts a multitask objective, assigning equal weight to the losses from both the image class prediction (classification task) and heading angle estimation (regression task). For the classification task, we model an ensemble of categorical distributions, making it possible to derive the closed-form solution for $I(Z, W)$. However, when estimating epistemic uncertainty for the regression task, we create an ensemble of Gaussian distributions and $I(Z, W)$ no longer has a closed-form solution [23]. To address this challenge, we rely on the pairwise-distance estimators (PaiDEs) [24]. PaiDEs have been shown to efficiently estimate epistemic uncertainty for regression tasks for probabilistic neural networks by leveraging closed-form distributional distance between ensemble components [25]. We utilize both Kullback-Liebler (KL) divergence and the Bhattacharyya distance between ensemble components for our PaiDEs.

B. MPC

In our MPC module, we’ve implemented the state-space model outlined in [26], incorporating both the slope effect and the uncertainty state, σ , given by the vision-based planner. This integration yields the following equation:

Label	Tree	Other Obstacles	Human	Waterhole	Mud	Jump	Traversable Grass	Smooth Road	Wet Leaves
Samples	586	1631	517	66	267	164	6421	10632	698

TABLE II: Our datasets label distribution.

Model	Macro-Precision	Macro-Recall	Macro-F1-Score	Accuracy
ResNet18	0.440	0.741	0.499	0.654
ResNet50	0.627	0.838	0.701	0.808
ResNet101	0.644	0.846	0.719	0.808
VGG11	0.791	0.818	0.793	0.875
VGG16	0.719	0.839	0.766	0.851
ViT-B/16	0.685	0.844	0.749	0.854
ViT-L/16	0.760	0.857	0.802	0.878

TABLE III: Results of the classification study on our datasets test set.

$$\dot{X} = V \cos(\psi + C_1 \delta) \quad (1)$$

$$\dot{Y} = V \sin(\psi + C_1 \delta) \quad (2)$$

$$\dot{\psi} = V \delta C_2 \quad (3)$$

$$\dot{V} = (C_{m_1} - C_{m_2} V) D - C_{r_2} V^2 - C_{r_0} - (V \delta)^2 C_1 C_2^2 - Mg \sin(\phi) \quad (4)$$

$$\dot{\phi} = 0 \quad (5)$$

$$\dot{\sigma} = 0 \quad (6)$$

The state vector $[X, Y, \psi, V, \phi, \sigma] \in \mathbb{R}^6$ encapsulates the vehicle’s pose in GPS coordinates, including its (X, Y) position and heading angle, along with its linear velocity and the road slope. Meanwhile, the input vector $[\delta, D]$ encompasses both the steering angle and throttle value. Detailed parameters of the kinematic model are provided in Table I. To enhance the accuracy of this model during MPC iterations that assess various action sets, a key focus is placed on state and parameter estimation, aiming to minimize the model uncertainty. To achieve this, we employ a moving horizon estimator (MHE) [27], which leverages a measurement window to estimate states and parameters. The measurements encompass GPS data, which includes the vehicle’s position and orientation, alongside accelerometer data and camera data. While the standard model in the mentioned equation includes an additional uncertainty state, the dynamic nature of this uncertainty is encapsulated by the image-based module. Consequently, within the optimization process, MPC accommodates uncertainty, preventing undue speed increments irrespective of the world model’s uncertainty. The flowchart of the proposed planner is presented in Fig. 4. For MPC action set formulation and practical viability, we harness global optimization techniques as mentioned earlier, which eliminates the need for gradient flow in optimization, streamlining a potentially challenging and time-intensive aspect.

C. Dataset

Our dataset encompasses trials of hands-on robot navigation, deliberately exposing the robot to intricate scenarios to enhance the dataset’s diversity. This effort resulted

Ground Truth	Tree	Other Obstacles	Human	Waterhole	Mud	Jump	Traversable Grass	Smooth Road	Wet Leaves
Tree	52	2	0	0	0	0	4	1	1
Other Obstacles	2	149	0	0	0	1	5	6	1
Human	0	0	51	0	0	0	1	1	0
Waterhole	0	0	0	5	1	0	0	2	0
Mud	0	1	0	0	21	0	4	2	0
Jump	0	0	0	0	0	17	0	0	0
Traversable Grass	4	27	4	0	1	2	573	31	1
Smooth Road	4	36	18	4	6	2	54	922	18
Wet Leaves	0	5	1	0	0	0	0	5	60

Fig. 3: The ViT-L/16 model’s confusion matrix on our datasets test set.

in a carefully synchronized repository of 20,982 samples collected at a frequency of 5 Hz to minimize potential sample correlations. Unlike previous approaches, we opted for manual annotations to minimize potential labeling noise. Our annotations encompass nine distinct events, including “tree”, “other-obstacles”, “human”, “waterhole”, “mud”, “jump”, “traversable-grass”, “smooth-road”, and “wet-leaves”. A detailed summary of the annotation distribution across our dataset is provided in Table II. To evaluate the complexity of our dataset on a classification task, we trained and evaluated three Residual Network (ResNet) [28] models, two Vision Transformer (ViT) [29] models, and two Visual Geometry Group (VGG) [30] models on our dataset. We trained each model using transfer learning and assessed their classification performance using established metrics such as macro-F1-score, macro-recall, and macro-precision.

Table III presents the outcomes of our classification study, revealing that while the classification of our dataset proves challenging for ResNet18, it avoids excessive complexity that would necessitate a substantial alteration in results by augmenting the model’s layers from ResNet50 to ResNet101. The confusion matrix depicted in Fig. 3 demonstrates the Vision Transformer ViT-L/16 model’s ability to effectively differentiate among the predefined categories. These results solidify the dataset’s optimal level of intricacy, lending confidence to the feasibility of the classification task at hand. Lastly, preceding research restricted the scope of analyzed events, notably omitting labels for dynamic entities like humans. In contrast, our dataset meticulously annotates nine distinct events, opening avenues for comprehensive exploration across diverse applications.




Fig. 4: Flowchart of the proposed planner.

III. RESULTS

A. Image-based predictive model

In this section, we describe the studies we conducted to compare the LSTM and transformer-based techniques. Initially, we trained the pre-existing baseline model known as BADGR [5] on our dataset. To delve deeper into the ramifications of the planning horizon, we train three models of each type with three different planning horizons: 10, 20, and 40 timesteps, and compare their performance. Please note that a horizon of 40 can be considered the ultimate planning limit, as it covers the next 8 seconds of planning. This duration is sufficient, given our maximum speed of $3m/s$. In every study, we conduct training for a total of 700 epochs, utilizing a batch size of 64, and employing an NVIDIA GeForce RTX 3090 GPU. We evaluate each model on a test set including 4000 random samples, and the same random trajectories are used for all studies.

It's worth noting that this facet of inquiry has been largely overlooked in the existing literature, which has predominantly concentrated on the realm of RL, focusing on metrics such as expected returns.

Fig. 5 unveils the outcomes, highlighting a compelling ascending trend for the LSTM-based model as we progress through the prediction steps. We think this pattern might arise from the LSTM networks' recurrent nature. As we progress in prediction, they extract increasingly valuable temporal features, enhancing predictive performance until

reaching a saturation point. Moving on to the second study, we substituted the LSTM component, as delineated in the methodology section, with a transformer-based model, subsequently reiterating the study.

Fig. 5 underscores a remarkable enhancement in the model's performance when using transformers. Notably, this improvement is most pronounced in its consistent prediction of future events. This trend continues in the domain of heading angle estimation, as depicted in Fig. 8, a pivotal component for guiding the vehicle toward the final goal. The LSTM network has poor accuracy for the short-term predictions, but its accuracy increases for predictions farther in the future. In contrast, the transformer-based model demonstrates impressive performance right from the start, although it experiences a slight decrease in predictive accuracy at later timesteps.

The superiority of the transformer-based model over the LSTM lies in its capacity to effectively filter essential information through self-attention mechanisms. In contrast, LSTM networks, due to their recurrent nature, tend to exhibit biases towards future events, making them reliant on prior action inputs for accurate future event predictions. This key distinction underscores the advantage of the transformer-based model. To further probe the performance of the LSTM-based model, we conducted an additional study, increasing the number of time-steps between predictions to minimize correlations among future events. However, this adjustment did not alter the model's performance; the weakness in early predictions persisted.

This discrepancy holds great significance when the robot re-plans after each action. In this context, precise early-stage predictions become crucial, making the transformer-based architecture a preferable choice. To further assess the initial performance of LSTM, we present confusion matrices for the initial prediction steps in Fig. 6. These matrices unveil the LSTM model's bias towards predicting "smooth-road" or "traversable-grass" in its early predictions.

B. Uncertainty-aware hybrid planner

In this section, we execute the complete pipeline, which incorporates the uncertainty module as well as the MPC with a specific objective in mind: to minimize both uncertainty and the time taken to achieve the goal at the same time. We employ a strategy involving simulated higher speeds, achieved by augmenting the time steps in the planning process. This enhances input action diversity, improving the training and exploration of the uncertainty-aware planner. To elaborate, we designate a data sampling time of 0.2 seconds as the baseline speed, while 0.4 seconds equate to 2 times that speed (2X), and 0.6 seconds correspond to 3 times (3X) the baseline. Consequently, we adapt the training of the image-based component of the planner to align with these variations.

Within this framework, the throttle range $[0,1]$ is mapped to a speed range of $[0, 3X \text{ baseline}]$. Following this, we train the predictive model with random speeds; we then compare the average uncertainty in each prediction step for both the




Fig. 5: When assessing three distinct metrics, it becomes evident that the transformer-based model excels in terms of sustaining consistent performance throughout the planning horizon. Note that each model’s evaluation is limited to the horizon it was originally trained for.




Fig. 6: This figure depicts confusion matrices from a trained LSTM-based model with planning horizon of 40 at prediction steps 0, 3, and 6. Initially, it mostly predicts "smooth-road" or "traversable-grass," but as the prediction horizon advances, the performance improves. Obstacles start to be recognized, and although some errors persist, the obstacles and smooth road become more distinguishable. Furthermore, we have 4,000 test samples, in contrast to the 2,000 samples shown in Fig. 3.




Fig. 7: MPC effectively mitigates epistemic uncertainty while also improving performance on favored terrains. Notably, the average of expected returns are 64.6 and 59.7 for the integrated and ignored uncertainty models, respectively.

LSTM and transformer based models. Fig. 9 presents the results, revealing that the uncertainty of both the LSTM and transformer models increases as the planning horizon increases for each task. This observation aligns with previous findings [17], and is consistent with our intuition that there should be greater certainty for steps closer to the current state. Note that for the classification task, the $I(Z, W)$ is calculated directly. In contrast, for the regression task, we employed two pairwise-distance estimators, namely KL

divergence and Bhattacharyaa distance (Bhatt), as outlined in [25].

Finally, we consider the transformer based architecture with the uncertainty incorporated, coupled with the MPC controller. This combination enables strategic planning for diverse scenarios using an RL approach, akin to the methodology outlined in previous works such as [5], [6], [7]. For the image-based module, we adhere to the same reward function

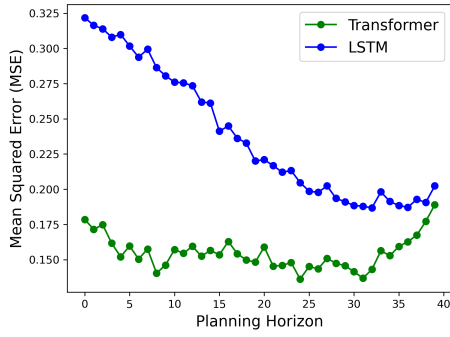


Fig. 8: The results from the comparative study on heading angle estimation errors further affirm the superior performance of the transformer-based model.

definition as:

$$R(\hat{e}_{t:t+H}^{0:K}) = -\sum_{t'=t}^{t+H-1} \hat{e}_{t'}^{coll} + \alpha^{POS} R^{POS}(\hat{e}_{t'}^{0:K}) + \alpha^{BUM} R^{BUM}(\hat{e}_{t'}^{0:K}) \quad (7)$$

Where $R^{POS} = (1 - \hat{e}_{t'}^{coll}) \frac{1}{\pi} \angle(\hat{e}_{t'}^{POS}, p^{GOAL}) + \hat{e}_{t'}^{coll}$ and $R^{BUM} = (1 - \hat{e}_{t'}^{coll}) \hat{e}_{t'}^{BUM} + \hat{e}_{t'}^{coll}$, with a total of K events within a given horizon of H , and predicted probabilities for each event represented as \hat{e} e.g. $\hat{e}_{t'}^{coll}$ stands for the probability of having collision at time t' . These values are instrumental in directing the vehicle towards preferred terrains while ensuring collision avoidance. This function transforms the model's event predictions into rewards, and the RL agent then seeks to maximize the cumulative discounted reward over the defined horizon. In the MPC component, we tackle uncertainty by introducing a distinctive reward function, detailed as follows:

$$R^{MPC} = \frac{\beta_{\sigma}}{\sigma^2} + \beta_V V^2 \quad (8)$$

with β_{σ} and β_V as tuning parameters, and σ representing the overall uncertainty. This helps to incorporate the imperative of uncertainty minimization. With this framework in place, we proceed to evaluate the effectiveness of our planner in real-world scenarios. The reason we can effectively balance the trade-off between speed and uncertainty using Eq. (8) is due to the incorporation of the MPC algorithm. MPC proposes a set of throttle commands that inherently considers the kinematic model of the vehicle, allowing us to obtain critical vehicle's states information such as speed for each prediction step. Extracting this information directly from images is notably more challenging, although it is feasible, it could result in a complex and resource-intensive architecture.

Fig. 7 summarizes our final study, which involved multiple trials. We used $\beta_{\sigma} = 10$ and $\beta_V = 1$ in our experiments. Our reward function optimizes throttle settings to minimize uncertainty, addressing various factors like input data distribution shifts. This doesn't always mean just using lower speeds during high uncertainty periods. Given the transformer-based predictive model's consistent and promising performance, lower uncertainty suggests improved results

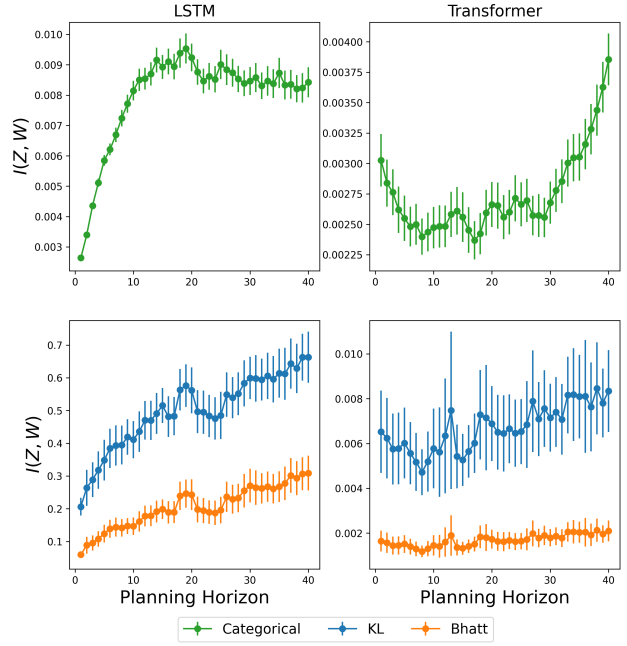


Fig. 9: An exploration of uncertainty in LSTM and transformer-based models highlights that uncertainty grows along the planning horizon, results displayed with 95% confidence interval.

when the input closely aligns with the training data. The prediction uncertainty curve shows a U-shaped pattern over prediction time, indicating increased confidence as more actions are considered. Initial uncertainty is high due to uncertainty about the robot's path, and it rises again at a prediction horizon threshold when current data is insufficient for accurate prediction. Finally, employing MPC in conjunction with the reward signal described in Eq. (8) leads to lower average speeds and increased speed variance. This highlights how MPC aims to balance speed and uncertainty, making the overall planner more conservative in its use of maximum speeds.

IV. CONCLUSION

This paper proposes a hybrid planner that combines non-linear MPC and RL-based approaches, effectively addressing uncertainty. The RL algorithm employs an image-based predictive model to optimize steering angles based on reward maximization, while non-linear MPC recommends throttle settings that trade off goal-reaching time and uncertainty due to various reasons. Furthermore, we introduce a transformer-based architecture to model the environment. This model surpasses the current baseline by achieving consistent and applicable prediction accuracy across various prediction steps. Our dataset covering nine diverse events ensures rigorous evaluation, producing a balance between complexity and feasibility. Our results underscore the superiority of our proposed planner over the existing baseline, seamlessly integrating vehicle and environment models to generate viable, uncertainty-aware trajectories.

REFERENCES

- [1] D. C. Guastella and G. Muscato, "Learning-based methods of perception and navigation for ground vehicles in unstructured environments: A review," *Sensors*, vol. 21, no. 1, p. 73, 2020.
- [2] F. Islam, M. Nabi, and J. E. Ball, "Off-road detection analysis for autonomous ground vehicles: a review," *Sensors*, vol. 22, no. 21, p. 8463, 2022.
- [3] S. Casas, A. Sadat, and R. Urtasun, "Mp3: A unified model to map, perceive, predict and plan," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14403–14412, 2021.
- [4] Y. Hu, J. Yang, L. Chen, K. Li, C. Sima, X. Zhu, S. Chai, S. Du, T. Lin, W. Wang, *et al.*, "Planning-oriented autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 17853–17862, 2023.
- [5] G. Kahn, P. Abbeel, and S. Levine, "Badgr: An autonomous self-supervised learning-based navigation system," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1312–1319, 2021.
- [6] T. Manderson, S. Wapnick, D. Meger, and G. Dudek, "Learning to drive off road on smooth terrain in unstructured environments using an on-board camera and sparse aerial images," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1263–1269, IEEE, 2020.
- [7] S. Wapnick, T. Manderson, D. Meger, and G. Dudek, "Trajectory-constrained deep latent visual attention for improved local planning in presence of heterogeneous terrain," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 460–467, IEEE, 2021.
- [8] C. E. Garcia, D. M. Prett, and M. Morari, "Model predictive control: Theory and practice—a survey," *Automatica*, vol. 25, no. 3, pp. 335–348, 1989.
- [9] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Aggressive driving with model predictive path integral control," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1433–1440, IEEE, 2016.
- [10] J. Yin, Z. Zhang, and P. Tsiotras, "Risk-aware model predictive path integral control using conditional value-at-risk," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 7937–7943, IEEE, 2023.
- [11] L. Sharma, M. Everett, D. Lee, X. Cai, P. Osteen, and J. P. How, "Ramp: A risk-aware mapping and planning pipeline for fast off-road ground robot navigation," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 5730–5736, IEEE, 2023.
- [12] A. Zeyer, P. Bahar, K. Irie, R. Schlüter, and H. Ney, "A comparison of transformer and lstm encoder decoder models for asr," in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pp. 8–15, IEEE, 2019.
- [13] S. Karita, N. Chen, T. Hayashi, T. Hori, H. Inaguma, Z. Jiang, M. Someki, N. E. Y. Soplin, R. Yamamoto, X. Wang, *et al.*, "A comparative study on transformer vs rnn in speech applications," in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pp. 449–456, IEEE, 2019.
- [14] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [15] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko, "End-to-end object detection with transformers," in *European conference on computer vision*, pp. 213–229, Springer, 2020.
- [16] N. Houlsby, F. Huszár, Z. Ghahramani, and M. Lengyel, "Bayesian active learning for classification and preference learning," *arXiv preprint arXiv:1112.5745*, 2011.
- [17] K. Chua, R. Calandra, R. McAllister, and S. Levine, "Deep reinforcement learning in a handful of trials using probabilistic dynamics models," *Advances in neural information processing systems*, vol. 31, 2018.
- [18] G. Kahn, A. Villaflor, V. Pong, P. Abbeel, and S. Levine, "Uncertainty-aware reinforcement learning for collision avoidance," *arXiv preprint arXiv:1702.01182*, 2017.
- [19] Z. I. Botev, D. P. Kroese, R. Y. Rubinstein, and P. L'Ecuyer, "The cross-entropy method for optimization," in *Handbook of statistics*, vol. 31, pp. 35–59, Elsevier, 2013.
- [20] N. Hansen, S. D. Müller, and P. Koumoutsakos, "Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (cma-es)," *Evolutionary computation*, vol. 11, no. 1, pp. 1–18, 2003.
- [21] N. Hansen, Y. Akimoto, and P. Baudis, "CMA-ES/pycma on Github." Zenodo, DOI:10.5281/zenodo.2559634, Feb. 2019.
- [22] L. Berry and D. Meger, "Normalizing flow ensembles for rich aleatoric and epistemic uncertainty modeling," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 6, pp. 6806–6814, 2023.
- [23] M. F. Huber, T. Bailey, H. Durrant-Whyte, and U. D. Hanebeck, "On entropy approximation for gaussian mixture random vectors," in *2008 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems*, pp. 181–188, IEEE, 2008.
- [24] A. Kolchinsky and B. D. Tracey, "Estimating mixture entropy with pairwise distances," *Entropy*, vol. 19, no. 7, p. 361, 2017.
- [25] L. Berry and D. Meger, "Escaping the sample trap: Fast and accurate epistemic uncertainty estimation with pairwise-distance estimators," *arXiv preprint arXiv:2308.13498*, 2023.
- [26] R. Verschuere, S. De Bruyne, M. Zanon, J. V. Frasch, and M. Diehl, "Towards time-optimal race car driving using nonlinear mpc in real-time," in *53rd IEEE conference on decision and control*, pp. 2505–2510, IEEE, 2014.
- [27] S. Lucia, A. Tătulea-Codrean, C. Schoppmeyer, and S. Engell, "Rapid development of modular and sustainable nonlinear model predictive control solutions," *Control Engineering Practice*, vol. 60, pp. 51–62, 2017.
- [28] 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*, pp. 770–778, 2016.
- [29] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," *arXiv preprint arXiv:2010.11929*, 2020.
- [30] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.