

Variational End-to-End Navigation and Localization

Alexander Amini¹, Guy Rosman², Sertac Karaman³ and Daniela Rus¹

Abstract—Deep learning has revolutionized the ability to learn “end-to-end” autonomous vehicle control directly from raw sensory data. While there have been recent advances on extensions to handle forms of navigation instruction, these works are unable to capture the full distribution of possible actions that could be taken and to reason about localization of the robot within the environment. In this paper, we extend end-to-end driving networks with the ability to understand maps. We define a novel variational network capable of learning from raw camera data of the environment as well as higher level roadmaps to predict (1) a full probability distribution over the possible control commands; and (2) a deterministic control command capable of navigating on the route specified within the map. Additionally, we formulate how our model can be used to localize the robot according to correspondences between the map and the observed visual road topology, inspired by the rough localization that human drivers can perform. We evaluate our algorithms on real-world driving data, and reason about the robustness of the inferred steering commands under various types of rich driving scenarios. In addition, we evaluate our localization algorithm over a new set of roads and intersections which the model has never driven through and demonstrate rough localization in situations without any GPS prior.

I. INTRODUCTION

Human drivers have an innate ability to reason about the high-level structure of their driving environment even under severely limited observation. They use this ability to relate high-level driving instructions to concrete control commands, as well as to better localize themselves even without concrete localization information. Inspired by these abilities, we develop a learning engine that enables robot vehicle to learn how to use maps within an end-to-end autonomous driving system.

Coarse grained maps afford us a higher-level of understanding of the environment, both because of their expanded scope, but also due to their distilled nature. This allows for reasoning about the low-level control within a hierarchical framework with long-term goals [1], as well as localizing, preventing drift, and performing loop-closure when possible. We note that unlike much of the work in place recognition and loop closure, we are looking at a higher level of matching, where the vehicle is matching intersection and

Support for this work was given by the National Science Foundation (NSF) and Toyota Research Institute (TRI). However, note that this article solely reflects the opinions and conclusions of its authors and not TRI or any other Toyota entity. We gratefully acknowledge the support of NVIDIA Corporation with the donation of the V100 GPU and Drive PX2 used for this research.

¹ Computer Science and Artificial Intelligence Lab, Massachusetts Institute of Technology {amini,rus}@mit.edu

² Toyota Research Institute {guy.rosman}@tri.global

³ Laboratory for Information and Decision Systems, Massachusetts Institute of Technology {sertac}@mit.edu

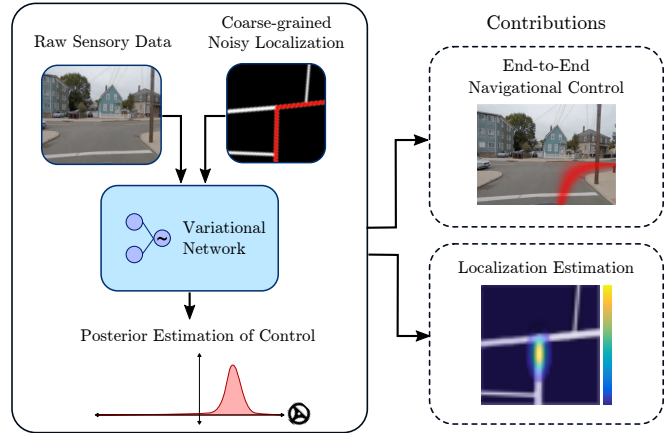


Fig. 1. **Variational end-to-end model.** Our model learns from raw sensory data as well as coarse grained topology maps to navigate and localize within complex environments.

road patterns to the coarse scale geometry found in the map, allowing handling of different appearances and small variations in the scene structure, or even unknown fine-scale geometry, as long as the overall road network structure matches the expected structures.

While end-to-end driving [2] holds promise due to its easily scalable and adaptable nature, it has a limited capability to handle long-term plans, relating to the nature of imitation learning [3], [4]. Some recent methods incorporate maps as inputs [5], [6] to capture longer term action structure, yet they ignore the uncertainty maps inherently allow us to address – uncertainty about the location, and uncertainty about the longer-term plan.

In this paper, we address these limitations by developing a novel model for integrating navigational information with raw sensory data into a single end-to-end variational network, and do so in a way that preserves reasoning about uncertainty. This allows the system to not only learn to navigate complex environments entirely from human perception and navigation data, but also understand when localization or mapping is incorrect, and thus correct for the pose (cf. Fig. 1).

Our model processes coarse grained, unrouted roadmaps, along with forward facing camera images to produce a probabilistic estimate of the different possible low-level steering commands which the robot can execute at that instant. In addition, if a routed version of the same map is also provided as input, our model has the ability to output a deterministic steering control signal to navigate along that given route.

The key contributions of this paper are as follows:

- Design of a novel variational end-to-end control net-

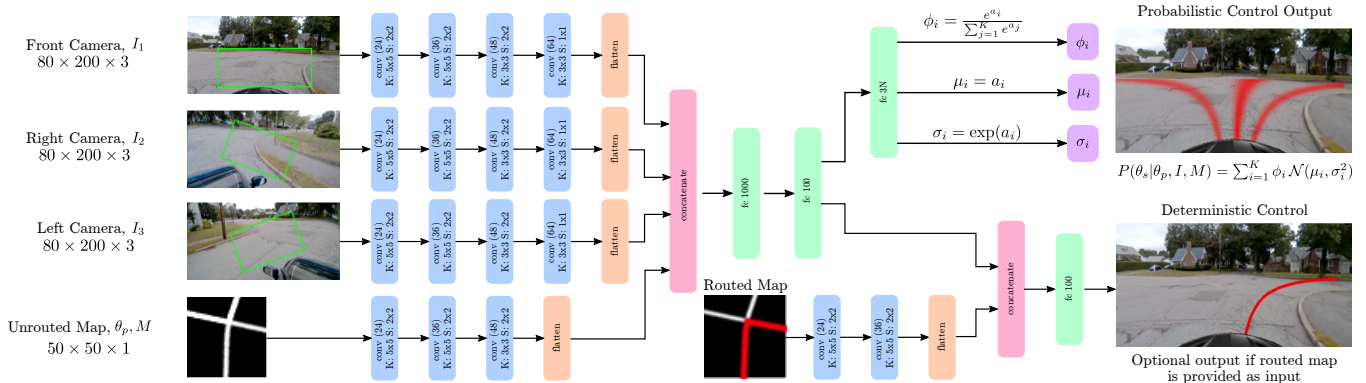


Fig. 2. **Model architecture.** An overview of our end-to-end learning system where raw camera images and noisy roadmap are fed to parallel convolutional pipelines, then merged into fully-connected layers to learn a full parametric Gaussian Mixture Model (GMM) over control. If a routed map is also available, it is merged at the penultimate layer to learn a deterministic control signal to navigate along the provided route. Green rectangles in road imagery in our figures denote the image region actually fed to the network.

work which integrates raw sensory data with routed and unrouted maps, enabling navigation and localization in complex driving environments; and

- Formulation of a localization algorithm using our trained end-to-end network to reason about a given road topology and infer the robot’s pose by drawing correspondences between the map and the visual road appearance; and
- Evaluation of our algorithms on a challenging real-world dataset, demonstrating navigation with steering control as well as improved pose localization even in situations with severely limited GPS information.

The remainder of the paper is structured as follows: we summarize the related work in Sec. II, formulate the model and algorithm for posterior pose estimation in Sec. III, describe our experimental setup, dataset, and results in Sec. IV, and provide concluding remarks in Sec. V.

II. RELATED WORK

Our work ties in to several related efforts in both control and localization. As opposed to traditional methods for autonomous driving which typically rely on distinct algorithms for localization and mapping [7], [8], [9], planning [10], [11], [12], and control [13], [14], end-to-end algorithms attempt to collapse the problem (directly from raw sensory data to output control commands) into a single learned model. The ALVINN system [15] originally proposed the use of multilayer perceptron to learn the direction a vehicle should steer in 1989. Recent advancements in convolutional neural networks (CNNs) have revolutionized the ability to learn driving control (i.e. steering wheel angle or road curvature) from raw imagery [2]. Followup works have incorporated conditioning on additional cues [3], including mapped information [5], [6]. However, these works are still quite limited in that they do not capture the uncertainty of multiple outcomes, nor present the ability to reason about discrepancy between their input modalities.

A recent line of work has tied end-to-end driving networks to variational inference [16], allowing us to look at cases

where multiple directions are possible, as well as reason about robustness, atypical data, and dataset normalization. Our work extends this line in that it allows us to use the same outlook and reason about maps as an additional conditioning factor.

Our work also relates to several research efforts in reinforcement learning in subfields such as bridging different levels of planning hierarchies [1], [17], and relating to maps as agents plan and act [18], [19]. This work relates to a vast literature in localization and mapping [7], [8], such as visual SLAM [9], [20] and place recognition [21], [22]. However, our notion of visual matching is much more high-level, more akin to semantic visual localization and SLAM [23], [24], where the semantic-level features are driving affordances.

III. MODEL

In this section, we describe the model used in our approach. We use a variational neural network, which takes raw camera images, I , and an image of a noisy, unrouted roadmap, M_U , as input. At the output we attempt to learn a full, parametric probability distribution over inverse curvature to navigate that instant. We use a Gaussian Mixture Model (GMM) with $K = 3$ modes to describe the possible steering control command, and penalize the $L_{1/2}$ norm of the weights to discourage extra components. Additionally, the model will optionally output a deterministic control command if a routed version of the map is also provided as input. The overall network can be written separately as two functions, representing the stochastic (unrouted) and deterministic (routed) parts respectively:

$$\{(\phi_i, \mu_i, \sigma_i^2)\}_{i=1}^K = f_S(I, M_U, \theta_p), \quad (1)$$

$$\hat{\theta}_s = f_D(I, M_R, \theta_p),$$

where $\theta_p = [p_x, p_y, p_\alpha]$ is the current pose in the map (position and heading), and $f_S(I, M_U, \theta_p)$, $f_D(I, M_R, \theta_p)$ are network outputs computed by cropping a square region of the relevant map according to θ_p , and feeding it, along with the forward facing images, I , to the network. M_U denotes the *unrouted* map, with only the traversible areas marked,




Fig. 3. **Control output under new, rich road environments.** Demonstration of our system which takes as input (A) image perception (green box denotes patch fed to the model); and (B) coarse unrouted roadmap (and routed, if available). The output (C) is a full continuous probability distribution for unrouted maps, and a deterministic command for navigating on a routed map. We demonstrate the control output on five scenarios, of roughly increasing complexity (left to right), ranging from straight road driving, to intersections, and even a roundabout.

while M_R denotes the *routed* map, containing the desired route highlighted. The deterministic control command is denoted as θ_s . In this paper, we refer to steering command interchangeably as the road curvature: the actual steering angle requires reasoning about road slip and control plant parameters that change between vehicles, making it less suitable for our purpose. Finally, the parameters (i.e. weight, mean, and variance) of the GMM’s i -th component are denoted by $(\phi_i, \mu_i, \sigma_i^2)$, which represents the steering control in the absence of a given route.

The overall network structure is given in Fig. 2. Each camera image is processed by a separate convolutional pipeline similar to the one used in [2]. Similarly, the cropped, non-routed, map patch is fed to a set of convolutional layers, before concatenation to the image processing outputs. However, here we use fewer layers for two main reasons: a) The map images contain significantly fewer features and thus don’t require a complex feature extraction pipeline and b) We wish to avoid translational invariance effects often associated with convolutional layers and subsampling, as we are interested in the pose on the map. The output of the convolutional layers is flattened and fed to a set of fully connected layers to produce the parameters of a probability distribution of steering commands (forming f_S). As a second task, we the previous layer output along with a convolutional module processing the routed map, M_R , to output a single deterministic steering command, forming f_D . This network structure allows us to handle both routed and non-routed maps, and later affords localization and driver intent, as well as driving according to high level navigation (i.e. turn-by-turn instruction).

We learn the weights of our model using backpropagation with the cost defined as:

$$\mathbb{E} \left\{ \begin{aligned} &\mathcal{L}(f_S(I, M, \theta_p), \theta_s) + \|\phi\|_p + \\ &\sum_i \psi_S(\sigma_i) + (f_D(I, M, \theta_p) - \theta_s)^2 \end{aligned} \right\} \quad (2)$$

where ψ_S is a per-component penalty on the standard deviation σ_i . We chose a quadratic term in $\log\text{-}\sigma$ as the

regularization,

$$\psi_S(\sigma) = \|\log \sigma - c\|^2. \quad (3)$$

$\mathcal{L}(f_S(I, M, \theta_p), \theta_s)$ is the negative log-likelihood of the steering command according to a GMM with parameters $\{(\phi_i, \mu_i, \sigma_i)\}_{i=0}^N$ and

$$P(\theta_s | \theta_p, I, M) = \sum \phi_i \mathcal{N}(\mu_i, \sigma_i^2). \quad (4)$$

A. Localization via End-to-End Networks

The conditional structure of the network affords updating a posterior belief about the vehicle’s pose, based on the relation between the map and the road topology seen from the vehicle. For example, if the network is provided visual input, I , which appears to be taken at a 4 way intersection we aim to compute $P(\theta_p | I, M)$ over different poses on the map to reason about where this input could have been taken. Note that our network only computes $P(\theta_s | \theta_p, I, M)$, but due to the conditioning structure of our model, we are able to estimate our pose given the visual input through double marginalization over σ_s and θ_p . Given a prior belief about the pose, $P(\theta_p)$, we can write the posterior belief after seeing an image I as:

$$\begin{aligned} P(\theta_p | I, M) &= \mathbb{E}_{\theta_s} P(\theta_p | \theta_s, I, M) \\ &= \mathbb{E}_{\theta_s} \left[\frac{P(\theta_p, \theta_s | I, M)}{P(\theta_s | I, M)} \right] \\ &= \mathbb{E}_{\theta_s} \left[\frac{P(\theta_p, \theta_s | I, M)}{\mathbb{E}_{\theta_{p'}} P(\theta_s | \theta_{p'}, I, M)} \right] \\ &= \mathbb{E}_{\theta_s} \left[\frac{P(\theta_s | \theta_p, I, M)}{\mathbb{E}_{\theta_{p'}} P(\theta_s | \theta_{p'}, I, M)} P(\theta_p) \right], \end{aligned} \quad (5)$$

where the equalities are due to full probability theorem and Bayes theorem. The posterior belief can be therefore computed via marginalization over θ_p, θ_s . While marginalization over two random variables is traditionally inconvenient, in two cases of interest, marginalizing over θ_p becomes easily tractable: a) when the pose is highly localized due to previous observations, as in the case of online localization and b)

where the pose is sampled over a discrete road network, as is done in mapmatching algorithms. The algorithm to update the posterior belief is shown in Algorithm 1. Intuitively, the algorithm computes, over all steering angle samples, the probability that a specific pose and images/map explain that steering angle, with the additional loop required to estimate the partition function and normalize the distribution. We note the same algorithm can be used with small modifications within the map-matching framework [25], [26].

Algorithm 1 Posterior Pose Estimate from Driving Direction

Input: $I, M, p(\theta_p)$
Output: $P(\theta_p|I, M)$

```

for  $i = 1 \dots N_s$ : do
  Sample  $\theta_s$ 
  Compute  $P(\theta_s|\theta_p, I, M)$ 
  for  $j = 1 \dots N_p$ : do
    Compute  $P(\theta_s|\theta'_p, I, M)$ 
    Aggregate  $\mathbb{E}_{\theta'_p} P(\theta_s|\theta'_p, I, M)$ 
  end for
  Aggregate  $\mathbb{E}_{\theta_s} \left[ \frac{P(\theta_s|\theta_p, I, M)}{\mathbb{E}_{\theta'_p} P(\theta_s|\theta'_p, I, M)} P(\theta_p) \right]$ 
end for
Output  $P(\theta_p|I, M)$  according to Equation 5.

```

IV. RESULTS

In this section, we demonstrate results obtained using our method on a real-world train and test dataset of rich driving environments. We start by describing our system setup and dataset and then proceed to demonstrate driving using an image of the roadmap and steering angle estimation with both routed and unrouted maps. Finally, we demonstrate how our approach allows us to reduce pose uncertainty based on the agreement between the map and the camera feeds.

A. System Setup

We evaluate our system on a 2015 Toyota Prius V outfitted with autonomous drive-by-wire capabilities [27]. Additionally, we made several advancements to the sensor and compute platform specifically for this work. Three Leopard Imaging LI-AR0231-GMSL cameras [28], capable of capturing 1080p RGB images at approximately 30Hz, are used as the vision data source for this study. We mount the three cameras on the front of the vehicle at various yaw angles: one forward facing and the remaining two rotated on the left/right of the vehicle to capture a larger FOV. Coarse grained global localization is captured using the OXTS RT3000 GPS [29] while an Xsense MTi 100-series IMU [30] is used to capture acceleration, rotation, and orientation data from the vehicle. Specifically, at any time, t , we use the yaw rate $\gamma^{(t)}$ (*rad/sec*), and the speed of the vehicle, $v^{(t)}$ (*m/sec*), to compute the curvature (or inverse radius) of the path which the human executed as $\theta_s^{(t)} = \frac{\gamma^{(t)}}{v^{(t)}}$. Finally, all of the sensor processing was done onboard an NVIDIA Drive PX2 [31].

In order to build the road network, we gather edge information from Open Street Maps (OSM) throughout the traversed region. We are given a directed topological graph, $G(V, E)$, where $v_i \in V$ represents an intersection on the road network and $e_i \in E$ represents a single directed road between two intersections. The weight of every edge, $w(e_i)$, is defined according to its great circle length, but with slight modifications could also capture more complexity by incorporating the type of road or even real-time traffic delays. The problem of offline map-matching involves going from a noisy set of ordered poses, $\{\theta_p^{(t)}\}_{t=1}^N$, to corresponding set of traversed road segments $\{e_i^{(t)}\}_{t=1}^N$ (i.e. the route taken). We implemented our map matching algorithm as an offline pre-processing step as in [26].

One concern during map rendering process is handling *ambiguous* parts of the map. Ambiguous parts are defined as parts where the the driven route cannot be described as a single simple curve. In order to handle large scale driving sequences efficiently, and avoid map patches that are ambiguous, we break the route into non-ambiguous subroutes, and generate the routed map for each of the subroutes, forming a set of charts for the map. For example, in situations where the vehicle travels over the same intersection multiple times but in different directions, we should split the route into distinct segments such that there are no self-crossings in the rendered route. Finally, we render the unrouted map by drawing all edges on a black canvas. The routed map is rendered by adding a red channel to the canvas and by drawing the traversed edges, $\{e_i^{(t)}\}_{t=1}^N$.

Using the rendered maps and raw perceptual image data from the three cameras, we trained our model (implemented in TensorFlow [32]) using 25 km of driving data taken in a suburban area with various different types of turns, intersections, roundabouts, as well as other dynamic obstacles (vehicles and pedestrians). We test and evaluate our results on an entirely different set of road segments and intersections which the network was never trained on.

B. Driving with Navigational Inputs

We demonstrate the ability of the network to compute both the continuous probability distribution over steering control for a given unrouted map as well as the deterministic control to navigate on a routed map. In order to drive with a navigational input, we feed both the unrouted and the routed maps into the network and compute $f_D(I, M, \theta_p)$. In Fig. 3 we show the inputs and parametric distributions of steering angles of our system. The roads in both the routed and unrouted maps are shown in white. In the routed map, the desired path is painted in red. In order to generate the trajectories shown in the figure, we project the instantaneous control curvature as an arc, into the image frame. In this sense, we visualize the projected path of the vehicle if it was to execute the given steering command. Green rectangles in camera imagery denote the region of interest (ROI) which is actually actually fed to the network. We crop according to these ROIs so the network is not provided extra non-essential




Fig. 4. **Model evaluation.** (A) The fraction of the instances when the true (human) steering command was within a certain z-score of our probabilistic network. (B) The probability density of the true steering command as a function of spatial location in the test roadmap. As expected the density decreases before intersections as the overall measure is divided between multiple paths. Points with gross GPS failures were omitted from visualization.

data which should not contribute to its control decision (e.g., the pixels above the horizon line do not impact steering).

We visualize the model output under various different driving scenarios ranging (left to right) from simple lane following to richer situations such as turns and intersections (cf. Fig. 3). In the case of lane following (left), the network is able to identify that the GMM only requires a single Gaussian mode for control, whereas multiple Gaussian modes start to appear for forks, turns, and intersections. In the case where a routed path is also provided, the network is able to disambiguate from the multiple modes and select a correct control command to navigate towards the route. We also demonstrate generalization on a richer type of intersection such as the roundabout (right) which was never included as part of the training data.

To quantitatively evaluate our network we compute the mixture estimation accuracy over our entire test set (cf. Fig 4). Specifically, for a range of z-scores over the steering control distribution we compute the number of samples within the test set where the true (human) control output was within the predicted range. To provide a more qualitative understanding of the spatial accuracy of our variational model we also visualized a heatmap of GPS points over the test set (cf. Fig. 4B), where the color represents the probability density of the predicted distribution evaluated at the true control value. We observed that the density decreases before intersections as the modes of the GMM naturally spread out to cover the greater number of possible paths.

C. Reducing Localization Uncertainty

We demonstrate how our model can be used to localize the vehicle based on the observed driving directions using Algorithm 1. We investigate in our experiments the reduction of pose uncertainty, and visualize areas which offer better types of pose localization.

For this experiment, we began with the pose obtained from the GPS and assumed an initial error in this pose with some uncertainty (Gaussian over the spatial position,

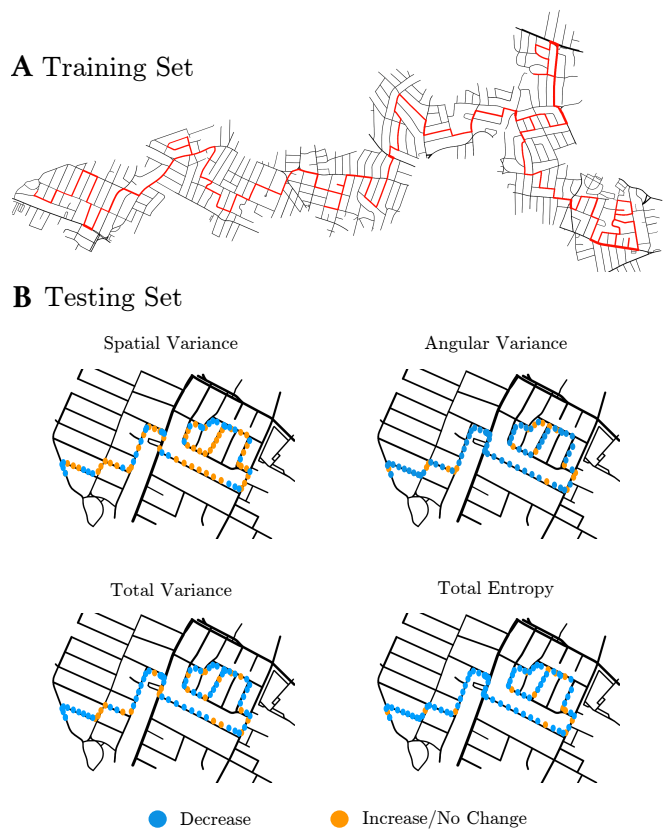


Fig. 5. **Evaluation of posterior uncertainty improvement.** (A) A roadmap of the data used for training with the route driven in red (total distance of 25km). (B) A heatmap of how our approach increases/decreases four different types of variance throughout test set route. Points represent individual GPS readings, while the color (orange/blue) denotes the absolute impact (increase/decrease) our algorithm had on its respective variance. Decreasing variance (i.e. increasing confidence) is the desired impact of our algorithm.

heading, or both). We compute the posterior probability of the pose as given by Alg. 1, and look at the individual uncertainty measures or total entropy of the prior and posterior distributions. If the uncertainty in the posterior distribution is lower than that of the prior distribution we can conclude that our learned model was able to increase its localization confidence after seeing the visual inputs provided (i.e. the camera images). In Fig. 5 we show the training (A) and testing (B) datasets. Note that the roads and intersections in both of these datasets were entirely disjoint; the model was never trained on roads/intersections from the test set.

For the test set, we overlaid the individual GPS points on the map and colored each point according to whether our algorithm increased (blue) or decreased (orange) posterior uncertainty. When looking at uncertainty reduction, it is important to note which degrees of freedom (i.e. spatial vs angular heading) localize better at different areas in the road network. For this reason, we visualize the uncertainty reduction heatmaps four times individually across (1) spatial variance, (2) angular variance, (3) overall pose variance, and (4) overall entropy reduction (cf. Fig. 5).

While heading angle is corrected easily at both straight




Fig. 6. **Pose uncertainty reduction at intersections.** The reduction of uncertainty in our estimated posterior across varying levels of added prior uncertainty. We demonstrate improvement in A) spatial $\sigma^2(p_x) + \sigma^2(p_y)$, B) angular: $\sigma^2(p_\alpha)$, C) sum of variance over p_x, p_y, p_α , and D) entropy in p_x, p_y, p_α , Gaussian approximation. Note that we observe a “positive” improvement over all levels of prior uncertainty (averaged over all samples in regions preceding intersections).

driving and more complex areas (turns and intersections), spatial degrees of freedom are corrected best at rich map areas, and poorly at linear road segments. This is expected and is similar to the aperture problem in computer vision [33] – the information in a linear road geometry is not enough to establish 3DOF localization.

If we focus on areas preceding intersections (approx 20 meters before), we typically see that the spatial uncertainty (prior uncertainty of $2m$) is reduced right before the intersection, which makes sense since after we pass through our forward facing visual inputs are not able to capture the intersection behind the vehicle. Looking in the vicinity of intersections, we achieved average reduction of 0.31 nats. For the angular uncertainty, with initial uncertainty of $\sigma = 0.8$ radians (45 degs), we achieved a reduction in the standard deviation of 0.2 radians (11 degs).

We quantify the degree of posterior uncertainty reduction around intersections in Fig. 6. Specifically, for each of the degrees of uncertainty (spatial, angular, etc) in Fig. 5 we present the corresponding numerical uncertainty reduction as a function of the prior uncertainty in Fig. 6. Note that we obtain reduction of both heading and spatial uncertainty for a variety of prior uncertainty values. Additionally, the averaged improvement over intersection regions is always positive for all prior uncertainty values indicating that localization does not worsen (on average) after using our algorithm.

D. Coarse Grained Localization

We next evaluate our model’s ability to distinguish between significantly different locations without any prior on pose. For example, imagine that you are in a location without GPS but still want to perform rough localization given your

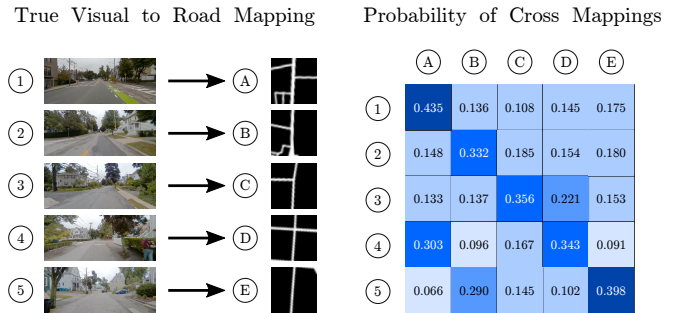


Fig. 7. **Coarse Localization from Perception.** Five example locations from the test set (image, roadmap pairs). Given images from location i , we compute the network’s probability conditioned on map patch from location j in the confusion matrix. Thus, we demonstrate how our system can establish correspondences between its camera and map input and even determine when its map pose has a gross error.

visual surroundings (similar to the kidnapped robot problem). We seek to establish correspondences between the map and the visual road area for coarse grained place recognition.

In Fig. 7 we demonstrate how we can identify and disambiguate a small set of locations, based on the the map and the camera images’ interpreted steering direction. Our results show that we can easily distinguish between places of different road topology or road geometry, in a way that should be invariant to the appearance of the region or environmental conditions. Additionally, the cases where the network struggles to disambiguate various poses is understandable. For example, when trying to determine which map the image from environment 4 was taken, the network selects maps A and D where both have upcoming left and right turns. Likewise, when trying to determine the location of environment 5, maps B and E achieve the highest probabilities. Even though the road does not contain any immediate turns, it contains a large driveway on the lefthand side which resembles a possible left turn (thus, justifying the choice of map B). However, the network is able to correctly localize each of these five cases to the correct map location (i.e. noted by the strong diagonal of the confusion matrix).

V. CONCLUSION

In this paper, we developed a novel variational model for incorporating coarse grained roadmaps with raw perceptual data to directly learn steering control of autonomous vehicle. We demonstrate deterministic prediction of control according to a routed map, estimation of the likelihood of different possible control commands, as well as localization correction and place recognition based on the map. We formulate a concrete pose estimation algorithm using our learned network to reason about the localization of the robot within the environment and demonstrate reduced uncertainty (greater confidence) in our resulting pose.

In the future, we intend to integrate our localization algorithm in the *online* setting of discrete road map-matching onboard our full-scale autonomous vehicle and additionally provide a more robust evaluation of the localization compared to that of a human driver.

REFERENCES

- [1] R. S. Sutton, D. Precup, and S. Singh, "Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning," *Artificial intelligence*, vol. 112, no. 1-2, pp. 181–211, 1999.
- [2] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang, *et al.*, "End to end learning for self-driving cars," *arXiv preprint arXiv:1604.07316*, 2016.
- [3] F. Codevilla, M. Müller, A. Dosovitskiy, A. López, and V. Koltun, "End-to-end driving via conditional imitation learning," *arXiv preprint arXiv:1710.02410*, 2017.
- [4] S. Shalev-Shwartz, S. Shammah, and A. Shashua, "Safe, multi-agent, reinforcement learning for autonomous driving," *arXiv preprint arXiv:1610.03295*, 2016.
- [5] G. Wei, D. Hsu, W. S. Lee, S. Shen, and K. Subramanian, "Intention-net: Integrating planning and deep learning for goal-directed autonomous navigation," *arXiv preprint arXiv:1710.05627*, 2017.
- [6] S. Hecker, D. Dai, and L. Van Gool, "End-to-end learning of driving models with surround-view cameras and route planners," in *European Conference on Computer Vision (ECCV)*, 2018.
- [7] J. J. Leonard and H. F. Durrant-Whyte, "Simultaneous map building and localization for an autonomous mobile robot," in *Intelligent Robots and Systems '91. Intelligence for Mechanical Systems, Proceedings IROS'91. IEEE/RSJ International Workshop on*. Ieee, 1991, pp. 1442–1447.
- [8] M. Montemerlo, S. Thrun, D. Koller, B. Wegbreit, *et al.*, "FastSLAM: A factored solution to the simultaneous localization and mapping problem," *AAAI*, vol. 593598, 2002.
- [9] A. J. Davison, I. D. Reid, N. D. Molton, and O. Stasse, "MonoSLAM: Real-time single camera SLAM," *IEEE Transactions on Pattern Analysis & Machine Intelligence*, no. 6, pp. 1052–1067, 2007.
- [10] L. Kavraki, P. Svestka, and M. H. Overmars, *Probabilistic roadmaps for path planning in high-dimensional configuration spaces*. Unknown Publisher, 1994, vol. 1994.
- [11] S. M. Lavalle and J. J. Kuffner Jr, "Rapidly-exploring random trees: Progress and prospects," in *Algorithmic and Computational Robotics: New Directions*, 2000.
- [12] S. Karaman, M. R. Walter, A. Perez, E. Frazzoli, and S. Teller, "Anytime motion planning using the rrt*," Shanghai, China, May 2011, pp. 1478–1483.
- [13] W. Schwarting, J. Alonso-Mora, L. Paull, S. Karaman, and D. Rus, "Safe nonlinear trajectory generation for parallel autonomy with a dynamic vehicle model," *IEEE Transactions on Intelligent Transportation Systems*, no. 99, pp. 1–15, 2017.
- [14] P. Falcone, F. Borrelli, J. Asgari, H. E. Tseng, and D. Hrovat, "Predictive active steering control for autonomous vehicle systems," *IEEE Transactions on control systems technology*, 2007.
- [15] D. A. Pomerleau, "ALVINN: An autonomous land vehicle in a neural network," in *Advances in neural information processing systems*, 1989, pp. 305–313.
- [16] A. Amini, W. Schwarting, G. Rosman, B. Araki, S. Karaman, and D. Rus, "Variational autoencoder for end-to-end control of autonomous driving with novelty detection and training de-biasing," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018.
- [17] R. Fox, S. Krishnan, I. Stoica, and K. Goldberg, "Multi-level discovery of deep options," *arXiv preprint arXiv:1703.08294*, 2017.
- [18] A. Tamar, Y. Wu, G. Thomas, S. Levine, and P. Abbeel, "Value iteration networks," in *Advances in Neural Information Processing Systems*, 2016, pp. 2154–2162.
- [19] J. Oh, S. Singh, and H. Lee, "Value prediction network," in *Advances in Neural Information Processing Systems*, 2017, pp. 6118–6128.
- [20] J. Engel, T. Schöps, and D. Cremers, "LSD-SLAM: Large-scale direct monocular SLAM," in *European Conference on Computer Vision*. Springer, 2014, pp. 834–849.
- [21] R. Paul and P. Newman, "Fab-map 3d: Topological mapping with spatial and visual appearance," in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*. IEEE, 2010, pp. 2649–2656.
- [22] T. Sattler, B. Leibe, and L. Kobbelt, "Fast image-based localization using direct 2d-to-3d matching," in *Computer Vision (ICCV), 2011 IEEE International Conference on*. IEEE, 2011, pp. 667–674.
- [23] J. L. Schönberger, M. Pollefeys, A. Geiger, and T. Sattler, "Semantic visual localization," *ISPRS Journal of Photogrammetry and Remote Sensing (JPRS)*, 2018.
- [24] S. L. Bowman, N. Atanasov, K. Daniilidis, and G. J. Pappas, "Probabilistic data association for semantic SLAM," in *Robotics and Automation (ICRA), 2017 IEEE International Conference on*. IEEE, 2017, pp. 1722–1729.
- [25] D. Bernstein and A. Kornhauser, "An introduction to map matching for personal navigation assistants," Tech. Rep., 1998.
- [26] P. Newson and J. Krumm, "Hidden Markov map matching through noise and sparseness," in *Proceedings of the 17th ACM SIGSPATIAL international conference on advances in geographic information systems*. ACM, 2009, pp. 336–343.
- [27] F. Naser, D. Dorhout, S. Proulx, S. D. Pendleton, H. Andersen, W. Schwarting, L. Paull, J. Alonso-Mora, M. H. Ang Jr., S. Karaman, R. Tedrake, J. Leonard, and D. Rus, "A parallel autonomy research platform," in *IEEE Intelligent Vehicles Symposium (IV)*, June 2017.
- [28] "LI-AR0231-GMSL-xxxH leopard imaging inc data sheet," 2018.
- [29] "OXTS RT user manual," 2018.
- [30] A. Vydhyanathan and G. Belluscio, "The next generation xsens motion trackers for industrial applications," Xsens, white paper, Tech. Rep., 2018.
- [31] "NVIDIA Drive PX2 SDK documentation." [Online]. Available: <https://docs.nvidia.com/drive/nvlib/docs/index.html>
- [32] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, *et al.*, "Tensorflow: a system for large-scale machine learning," in *OSDI*, vol. 16, 2016, pp. 265–283.
- [33] D. Marr and S. Ullman, "Directional selectivity and its use in early visual processing," *Proc. R. Soc. Lond. B*, vol. 211, no. 1183, pp. 151–180, 1981.