

Perceive, Attend, and Drive: Learning Spatial Attention for Safe Self-Driving

Bob Wei^{*,1,2}, Mengye Ren^{*,1,3}, Wenyuan Zeng^{1,3}, Ming Liang¹, Bin Yang^{1,3}, Raquel Urtasun^{1,2}

Abstract—In this paper, we propose an end-to-end self-driving network featuring a sparse attention module that learns to automatically attend to important regions of the input. The attention module specifically targets motion planning, whereas prior literature only applied attention in perception tasks. Learning an attention mask directly targeted for motion planning significantly improves the planner safety by performing more focused computation. Visualizing the attention also improves interpretability of end-to-end self-driving.

I. INTRODUCTION

Self-driving is one of today’s most impactful technological challenges, one that promises to bring safety and convenience to transportation. Tremendous improvements have been made in self-driving perception systems, thanks to the success of deep learning. This has enabled accurate detection and localization of obstacles, providing a holistic understanding of the surrounding world, which is then sent to the motion planner to decide subsequent driving actions.

Despite the eminent success of these perception systems, their detection objective is *mis-aligned* with the self-driving vehicle’s overall goal, which is to drive safely to the destination. We typically train the perception systems to detect all objects in the sensor’s range, and assign each object an equal weight even if some objects are not important and will never interact with the self-driving vehicle; for example, they could be far away or parked on the other side of the road as in Figure 1. As a result, a vast amount of computation and model capacity is wasted in recognizing very difficult instances that matter only for common metrics such as average precision (AP), but not so much for driving. This is in striking contrast with how humans drive: we focus our visual attention in areas that directly impact safe planning. Inspired by the use of visual attention in our brain, which helps us to efficiently and selectively process complex scenes, we aim to bring attention to self-driving systems.

Numerous studies in the past have explored adding sparse attention in deep neural networks to improve computation efficiency in classification [1], [2] and object detection [3], [4], [5]. In order to perform well on the metrics employed in common benchmarks, the attention mask in [3] still needs to cover all actors in the scene, slowing the network when the scene has many vehicles.

In this paper, our aim is to address these inconsistencies such that each computation matters for the end goal of driving. Specifically our contributions are as follows:




Fig. 1. *Left*: a toy example. *Right*: our learned spatial attention in red with ego vehicle in blue and others in green. Not all actors impact our safe driving and so we should prioritize accordingly.

- We learn an attention mask directly towards the motion planning objective for safe self-driving.
- We use the attention mask to reweight object detection and motion forecasting losses in our joint end-to-end training, focusing the model capacity on objects that matter most. Different from manually prioritizing instances [6], here the weighting is entirely data-driven.
- Our attention-based model achieves significantly reduced collision rate and improved planning performance, while at a much lower computation cost.
- Attention mask visualization improves interpretability of end-to-end deep learning models in self-driving.

II. RELATED WORK

Attention mechanism in deep learning: Human and other primate visual perception systems feature visual attention to reduce the complexity of the scene and speed-up inference [7], [8]. Earlier studies in visual saliency aim to predict human gaze with no particular task in mind [9]. Attention mechanisms nowadays are built in as part of end-to-end models to optimize towards specific tasks. The attention modules are typically implemented as multiplicative gates to select features. This schema has shown to improve performance and interpretability on downstream tasks such as object recognition [10], [11], [1], instance segmentation [5], image captioning [12], question answering [13], [14], as well as other natural language processing applications [15], [16], [17]. The visualization of the end-to-end learned attention suggests that deep attention-based models have an intelligent understanding of the inputs by focusing on the most informative parts of the input.

Sparse activation in neural networks: Sparse coding models [18] use an overcomplete dictionary to achieve sparse activation in the feature space. In modern CNNs, sparsity is typically brought by the widespread use of ReLU activation functions, but these are rather unstructured, and speed-up has only been shown on specially designed hardware [19], [20]. Structured spatial sparsity, on the other hand, can be made efficient by using a sparse convolution operator [3], [21], which in turn allows the network to shift its focus on more

^{*}Equal contribution ¹Uber ATG ²University of Waterloo ³University of Toronto. Correspondence: q25wei@uwaterloo.ca, [mren3, wenyuan, ming.liang, byang1, urtasun]@uber.com




Fig. 2. The full neural motion planner (NMP) [35] with backbone network and header networks.

difficult parts of the inputs [2], [22], [3]. In self-driving, [6] proposes a ranking function to prioritize computations that would have the most impact on motion planning. Weight pruning [23], [24] is another popular way to achieve sparsity in the parameter space, which is an orthogonal direction to our proposed method.

Attention and loss weighting in multi-task learning:

Our end-to-end self-driving network is an instance of multi-task learning as all three tasks—perception, prediction and motion planning—are simultaneously solved by individual output branches with shared features. It is common to use a summation of all the loss functions, but sometimes there are conflicting objectives among the tasks. Prior literature in multi-task learning has studied dynamic weighting towards different loss components, by using training signals such as uncertainty [25], gradient norm [26], difficulty level [27], or entirely data-driven objectives [28], [29]. In [28], [29], task and example weights are learned by optimizing the performance of the main task. The attention mechanism has also been used in multi-task learning: in [30], a network applies task-specific attention masks on shared features to encourage the outputs to be more selective. Similar to dynamic loss weighting models [28], we exploit the learned attention towards weighting instance detection losses. Instead of using multiple attentions, as was done in [30], we use one single attention mask to optimize our main task: driving.

Safety-driven learnable motion planning: One of the primary motivations of introducing attention into an end-to-end motion planning network is to improve safety. Traditionally, safety for self-driving models is done in terms of formal model checking and validation [31], [32], [33], [34]. More recently, with the widely available driving data, imitation learning has been introduced in self-driving [35], [36], [37] to learn from cautious human driving. Safety has also been considered in terms of explicitly learning a risk sensitive measure from human demonstration [38], [39]. In our work, although safety is not explicitly encoded in our loss function, we have experimentally verified that the sparse attention models are significantly better at avoiding collisions.

III. PERCEIVE, ATTEND AND DRIVE

In this section, we present our framework for using learned, motion-planning aware attention. We first describe the end-to-end neural motion planner that serves as the starting point of our work, and then introduce our proposed attention module and attention-driven loss function, which enable our novel neural motion planner to focus its computation in areas that matter for the end task of driving.




Fig. 3. Our sparse attention neural motion planner (SA-NMP), which takes in LiDAR and HDMaps data, and outputs perception, prediction, and planning. Attention is generated from a U-Net and applied on the input branches of residual blocks. Our attention is learned towards the planning task, which is a direct model output.

A. A review on Neural Motion Planner (NMP)

Our proposed model extends upon the neural motion planner (NMP), which jointly solves the perception, prediction and planning problems for self-driving. In this section we briefly review NMP depicted in Figure 2, and refer the reader to [35] for more details.

Input and backbone: NMP voxelizes LiDAR point clouds to a birds-eye-view (BEV) feature map and fuses them with M channels of rasterized HDMap features to produce an input representation of size $(ZT' + M) \times H \times W$, where Z, H, W are the height and spatial dimensions and $T' = 10$ is the number of input LiDAR sweeps. The backbone consists of 5 blocks, with the first 4 producing multi-scale features that are concatenated and fed to the final block. Overall the backbone downsamples the spatial dimension by 4.

Multi-task headers: Given the features computed by the backbone, $X \in \mathbb{R}^{128 \times \frac{H}{4} \times \frac{W}{4}}$, NMP uses two separate headers for perception & prediction, and motion planning. The perception & prediction header consists of separate branches for classification and regression. The classification branch outputs a score for each anchor box at each spatial location over the feature map X , while the regression branch outputs regression targets for each anchor box, including targets for localization offset, size, and heading angle. The planning header consists of convolution and deconvolution layers to produce a cost volume $C \in \mathbb{R}^{T \times \frac{H}{4} \times \frac{W}{4}}$ representing the cost for the self-driving vehicle to be at every possible location and time, with T the fixed future planning horizon.

Planning inference: At inference time, NMP samples N trajectories that are physically realizable, and chooses the lowest cost trajectory for the ego-car:

$$\tau^*(X) = \arg \min_{\tau_{1..N}} c(\tau_i, X). \quad (1)$$

where the cost of a trajectory $\tau = (x_t, y_t)_{t=1}^T$ is the sum of all its waypoints in the cost volume:

$$c(\tau, X) = \sum_{t=1}^T C_{t, x_t, y_t}(X), \quad (2)$$

We sample trajectories using a mixture of Clothoid, circle, and straight curves [40]. We refer the readers to [35] for more details on the sampling procedure.

B. Sparse Attention Neural Motion Planner (SA-NMP)

In this section, we propose our sparse spatial attention module for self-driving, shown in Figure 3, which learns to

save computation while performing well on the end task of driving safely to the goal.

Input and backbone: We exploit the same input representation as NMP and use the same perception, prediction, and planning headers. We replace NMP backbone with the state-of-the-art backbone network of PnPNet [41], which uses cross-scale blocks throughout to fuse BEV sensory input. Each cross-scale block consists of 3 parallel branches at different resolutions that downsample the feature map, perform bulk computations, and then upsample back to the backbone resolution, before finally fusing cross-scale features across all branches. There is an additional residual connection across each cross-scale block. The final output feature from the backbone consists of 128 channels at 4x downsampled resolution, which is forwarded to the planning and detection headers. In addition to the improved performance, PnPNet can be easily scaled for different computational budgets by varying the depth and width of the cross-scale blocks.

Existing attention-driven approaches [3] tackle only the perception task and use either a road mask obtained from HDMaps or a vehicle mask produced by a different perception module. As a consequence, they waste computation on areas that will not affect the self-driving car. Inspired by human driving, we instead propose a novel approach that is end-to-end trainable and performs computation selectively for planning a safe maneuver. As shown in Fig. 3, the learned attention mask then gates the backbone network, limiting computation to areas where attention is active. By using binary attention, we can leverage sparse convolution to improve the computational efficiency.

Attention generator: We utilize a network to predict a scalar score for each spatial location. For efficiency and simplicity, we use a small U-Net [42] with skip connections and two downsample/upsample stages. The baselines in Sec. IV-A replace this with static attention maps.

Learning binary attention mask: Computational efficiency has been shown to be one of the most prominent advantages of using the attention mechanism. For soft attention masks the computation is still dense across the entire activation map, and therefore no computation savings can be achieved. In SBNet [3], the authors showed that their sparse convolution operator can achieve significant speed-ups with a given discrete binary attention mask. Here, we would also like to exploit the computational benefit of sparse convolution by using discrete attention outputs. To achieve this, we make use of the Gumbel softmax technique [43], [44] since the step function is not differentiable, and using the standard sigmoid function suffers from a more severe bias-variance trade-off [43]. Let i, j denote spatial coordinates, and $z_{i,j}$ the scalar output from the attention U-Net. We first add the Gumbel noise on the logits as follows:

$$\pi_{i,j} = \text{sigmoid}(z_{i,j}) \quad (3)$$

$$\alpha_{i,j}^{(0)} = \log \pi_{i,j} + g_{i,j}^{(0)} \quad (4)$$

$$\alpha_{i,j}^{(1)} = \log(1 - \pi_{i,j}) + g_{i,j}^{(1)}, \quad (5)$$

where $g_{i,j} = -\log(-\log u)$, and u is sampled from Uniform[0, 1]. At inference time, hard attention $A_{i,j}$ can be obtained by comparing the logits,

$$A_{i,j} = \begin{cases} 1 & \text{if } \alpha_{i,j}^{(0)} \geq \alpha_{i,j}^{(1)} \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

During training, however, we would like to approximate the gradient by using the straight-through estimator [43], [45]. Hence, in the backward pass, the step function is replaced with a softmax function with a temperature constant K (where \tilde{A} is the underlying soft attention):

$$\tilde{A}_{i,j} = \frac{\exp(\alpha_{i,j}^{(0)}/K)}{\exp(\alpha_{i,j}^{(0)}/K) + \exp(\alpha_{i,j}^{(1)}/K)}. \quad (7)$$

Applying attention: Given a generated hard attention map, we would like to apply it back to the BEV features in the model backbone so as to sparsify the spatial information, allowing computation to be focused on the important regions only. We choose to do so in a residual manner [1], [3] to avoid deteriorating the features throughout the backbone. Let $x + F(x)$ denote the normal residual block. Our attention mask is multiplied with the input to the residual block:

$$\text{ResAttend}(x) = x + F(x \odot A), \quad (8)$$

where \odot denotes elementwise multiplication. See Fig. 3 for an illustration of our architecture.

C. Learning

We train our sparse neural motion planner (including the attention) end-to-end with a joint multi-task learning objective that combines perception, motion forecasting (L_{class} , L_{reg}) and planning (L_{plan}):

$$L = \lambda_{\text{class}} L_{\text{class}} + \lambda_{\text{reg}} L_{\text{reg}} + \lambda_{\text{plan}} L_{\text{plan}} + \lambda_A L_A + \lambda \|w\|_2^2, \quad (9)$$

where L_A is an L1 loss, defined in Equation 17, that controls the sparsity of the attention mask, and $\|w\|_2^2$ is the standard weight decay term. From [35], [41], we fix $\lambda_{\text{class}} = 1.0$, $\lambda_{\text{reg}} = 0.5$, $\lambda_{\text{plan}} = 0.001$. More optimization details are provided in Section IV-A

Perception & prediction (PnP) loss: This loss follows the standard classification and regression objectives for object detection. The classification part uses binary cross-entropy:

$$L_{\text{class},i,j} = \sum_k -\hat{y}_{i,j,k} \log(y_{i,j,k}) - (1 - \hat{y}_{i,j,k}) \log(1 - y_{i,j,k}), \quad (10)$$

where \hat{y} is the predicted classification score between 0 and 1, and y is the binary groundtruth. For each detected instance, the model outputs a bounding box, and a pair of coordinates and angles for each future step. We reparameterize the shift of a bounding box (x, y, w, h, θ) from an anchor bounding box $(x_a, y_a, w_a, h_a, \theta_a)$ in a 6-dimensional vector δ :

$$\delta_t = \left(\frac{x_a - x}{w}, \frac{y_a - y}{h}, \log \frac{w_a}{w}, \log \frac{h_a}{h}, \sin(\theta_a - \theta), \cos(\theta_a - \theta) \right). \quad (11)$$

A regression loss is then applied for the trajectory of the instance up to time T . For each spatial coordinate (i, j) , we sum up the losses of all bounding boxes b where the ground truth belongs to this location:

$$L_{\text{reg},i,j} = \sum_{b \in (i,j)} \sum_{t=0}^T \text{SmoothL1}(\hat{\delta}_{b,t}, \delta_{b,t}), \quad (12)$$

where $\hat{\delta}$ is the model predicted shifts and δ is the groundtruth shifts.

Auxiliary loss masking: Our overall objective is to achieve good performance in motion planning, so PnP is an auxiliary task. Since the majority of our computation happens within the attended area, intuitively the model should not be penalized as severely for mis-detecting objects not in the attended area. We therefore propose to use our spatial attention mask A to re-weight the PnP losses as follows:

$$L_{\text{class}} = \gamma_1 \sum_{i,j} A_{i,j} L_{\text{class},i,j} + \gamma_0 \sum_{i,j} L_{\text{class},i,j} \quad (13)$$

$$L_{\text{reg}} = \gamma_1 \sum_{i,j} A_{i,j} L_{\text{reg},i,j} + \gamma_0 \sum_{i,j} L_{\text{reg},i,j} \quad (14)$$

where γ_1 weights attended instances, and γ_0 weights all instances. We fix $\gamma_0 = 0.1$ and $\gamma_1 = 0.9$.

Motion planning loss: The motion planning loss utilizes the max-margin objective, where the ground-truth driving trajectory (performed by a human) should be of lower cost than other trajectories sampled by the model. Let (x_t, y_t) be the groundtruth trajectory and let c_t be the cost volume output by the model. We sample N lowest cost trajectories based on the cost volume: $\{x_t^{(i)}, y_t^{(i)}\}_{i=1}^N$, and penalize the maximum margin violation between the groundtruth and model samples:

$$L_{\text{plan}} = \max_{i=1 \dots N} \sum_{t=1}^T \max\{0, c_t - c_t^{(i)} + \Delta_t^{(i)}\}, \quad (15)$$

where $\Delta_t^{(i)}$ is the task loss capturing spatial differences, and traffic violations v_t encoded in binary:

$$\Delta_t^{(i)} = \left\| (x_t, y_t) - (x_t^{(i)}, y_t^{(i)}) \right\|_2 + v_t^{(i)}. \quad (16)$$

Attention sparsity loss: To encourage focused attention and high sparsity, we use an L1 regularizer on the attention mask as follows. We control sparsity with λ_A in Equation 9.

$$L_A = \sum_{i,j} A_{i,j}. \quad (17)$$

IV. EXPERIMENTAL EVALUATION

We evaluate on a real-world driving dataset (Drive4D), training on over 1 million frames from 5,000 scenarios and validating on 5,000 frames from 500 scenarios, using both LiDAR and HD-maps. We also evaluate on nuScenes [46], a large-scale public dataset, with a train-set of over 200,000 frames and a test-set of 5,000 frames. Due to inaccurate localization provided with the benchmark, we omit HDMaps and only use LiDAR.

A. Implementation Details and Metrics

Training: To jointly train SA-NMP with attention, we use pretrained weights for the backbone and headers from training a SA-NMP without attention (dense) for two epochs. We train all our models with batch size 5 across 16 GPUs in parallel using Adam [47] optimizer. We use an initial learning rate of 0.0001 and a decay of 0.1 at 1.0 and 1.6 epoch(s), for a total of another 2.0 epochs.

Evaluation: To evaluate driving and safety performance, we focus on the following planning metrics which are accumulated over all 6 future timesteps (3s): *Planning L2* is the L2 distance between waypoints of the predicted future ego trajectory and those of the ground-truth trajectory (characterized by human driving). *Collision rate* is the frequency of collisions between the planned ego trajectory and the ground truth trajectories of other actors in the scene. *Lane violation rate* measures the number of lane boundary violations by the planned ego trajectory. Due to lack of accurate HDMaps [48], we don't evaluate this on nuScenes v1.0.

Baselines: We compare our learned attention to baselines that are end-to-end trained using static attention masks obtained from priors. *Road Mask* covers the entire road as provided from the map data. *Vehicle Mask* strictly covers all detections in the input space, obtained from a PSPNet [49] trained for segmentation. *Proximity Mask* is a circular radius around the ego vehicle. *Dense* is not using sparse attention.

B. Results

Quantitative results: With a sparse attention mask learned towards motion planning, we can leverage the sparsity in the network backbone to greatly reduce computational costs, while not only maintaining but improving model performance. In our experimental results, we use theoretical FLOPs to show the efficiency of our network, but this also translates to real time gains as SBNet [3] has been shown to leverage sparsity to achieve real speed-ups. The increase in efficiency from leveraging sparsity is shown in Table I, where our learned attention model uses $\sim 80\%$ fewer FLOPs than *Dense SA-NMP* thanks to its 95% sparse attention mask. Also, even with an identical SA-NMP backbone as the baselines (except *NMP*), our model with learned attention performs better in all motion planning metrics, which indicates that focused backbone computation is greatly advantageous to the overall goal of safe planning. *NMP+Road* performs slightly better in *Lane Violation* due to the road mask attention focusing on all road/lane markings. However, this baseline uses more than double the FLOPs since its attention mask looks at the entire road surface at only 68.9% sparsity. From Fig. 4, our learned attention model clearly outperforms other baselines in collision rate and planning L2, across all computational budgets.

Qualitative results: In Figure 5, we show examples of our learned attention compared to baselines. As expected, our model focuses on the road and vehicles directly ahead; however, it also diverts some amount of attention to distant vehicles and road markings. This ability to dynamically distribute attention is likely why our model outperforms the

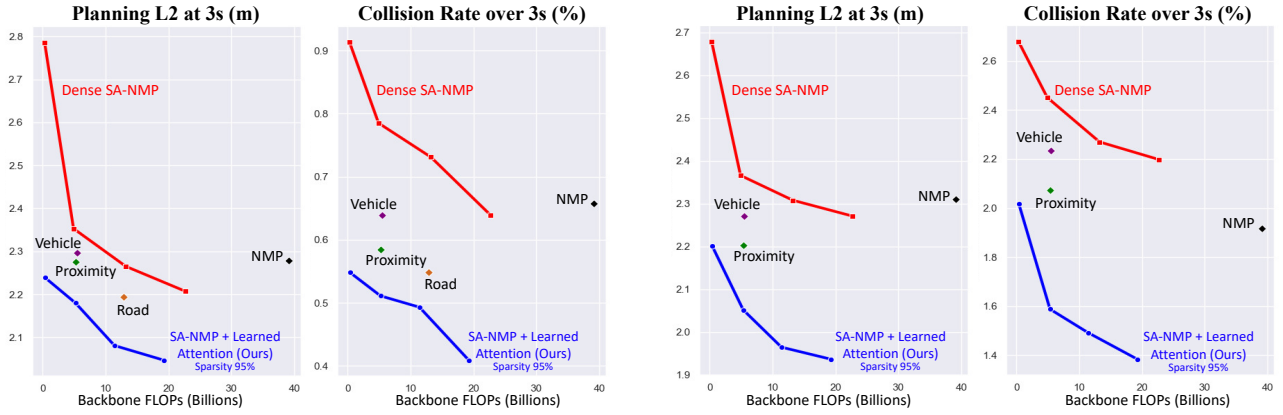


Fig. 4. Planning performance of our learned sparse attention model compared to other baselines at varying computation budgets (lower is better on both metrics). **Left:** Drive4D dataset; **Right:** nuScenes dataset. Note that all models except for *NMP* use the same SA-NMP backbone, which can be scaled by changing the depth and width, and allows us to vary the computation of *Dense SA-NMP* and *SA-NMP+Learned Attention*.

TABLE I
PERFORMANCE AND EFFICIENCY OF OUR LEARNED ATTENTION MODEL VS. DENSE AND SIMPLE ATTENTION BASELINES.

Drive4D Dataset	Backbone FLOPs	Sparsity	Planning L2 at 3s (m)	Collision Rate over 3s (%)	Lane Violation over 3s (%)
NMP [35]	39.18B	0.0%	2.279	0.657	2.780
Dense SA-NMP	22.73B	0.0%	2.207	0.639	1.350
SA-NMP+Vehicle Mask	5.43B	93.6%	2.297	0.639	1.387
SA-NMP+Proximity Mask	5.31B	94.0%	2.276	0.584	1.387
SA-NMP+Road Mask	12.85B	68.9%	2.194	0.548	1.296
SA-NMP+Learned Attn (Ours)	5.22B	95.0%	2.102	0.511	1.338

nuScenes Dataset	Backbone FLOPs	Sparsity	Planning L2 at 3s (m)	Collision Rate over 3s (%)
NMP [35]	39.18B	0.0%	2.310	1.918
Dense SA-NMP	22.73B	0.0%	2.271	2.198
SA-NMP+Vehicle Mask	5.44B	94.0%	2.263	2.234
SA-NMP+Proximity Mask	5.31B	94.0%	2.103	2.073
SA-NMP+Learned Attn (Ours)	5.34B	94.0%	2.052	1.588

simple baselines, which attend either indiscriminately (Dense and Road) or too selectively (Vehicle and Proximity). Note that our attention learns which actors to focus on and to what capacity, unlike the vehicle mask which attends fully to every actor while ignoring road features. From the visualizations, we can better understand our model’s improved collision avoidance. Since the attention is dynamic, it appears that our model is more effective at anticipating other vehicles resulting in more cautious planning. This is illustrated by Columns A, B in Figure 5, where the planned trajectory of our model avoids future collisions with others. The failure cases mostly arise from rear-end collisions, one of which is Column D where all models are hit by the trailing vehicle. Our model focuses on surrounding vehicles and not enough on the open road to its right, which would give the option of making a right turn.

Sparsity of learned attention: Table II shows the result of varying λ_A from Eq. 9, which weights the L1 regularization term from Eq. 17, with other settings held constant. We found that overall motion planning performance improves with increased sparsity, or essentially more focused computation, and peaks at 95% sparsity.

Perception and prediction (PnP) loss reweighting: Table III shows results with varying γ_1 and $\gamma_0 = 1 - \gamma_1$

TABLE II
VARYING LEARNED ATTENTION SPARSITY WITH λ_A .

Sparsity	λ_A	Planning L2 at 3s (m)	Collision Rate over 3s (%)	Lane Violation over 3s (%)
Dense 0%	-	2.207	0.639	1.350
Ours 20%	$1.0e^{-8}$	2.179	0.547	1.352
Ours 50%	$1.0e^{-7}$	2.138	0.620	1.361
Ours 75%	$5.0e^{-7}$	2.132	0.584	1.387
Ours 95%	$1.0e^{-6}$	2.102	0.511	1.338
Ours 99%	$5.0e^{-6}$	2.211	0.566	1.367

TABLE III
INFLUENCE OF LOSS REWEIGHTING RATIO γ_1 .

Model	Planning L2 at 3s (m)	Collision Rate over 3s (%)	Lane Violation over 3s (%)
$\gamma_1 = 1.00$	2.210	0.548	1.378
$\gamma_1 = 0.90$	2.102	0.511	1.338
$\gamma_1 = 0.75$	2.230	0.621	1.393
$\gamma_1 = 0.50$	2.194	0.637	1.405
$\gamma_1 = 0.00$	2.188	0.633	1.397

from Eq. 14 which control the weighting of the PnP loss computed on actors inside vs. outside the attention mask. All other variables are fixed, including sparsity at 95%. As γ_1 increases, the learned attention is less restricted by detection performance on all actors, and is able to focus on only the most important actors and parts of the road, distributing attention towards improving motion planning performance. Note that $\gamma_1 = 1.0$ is an extreme case where

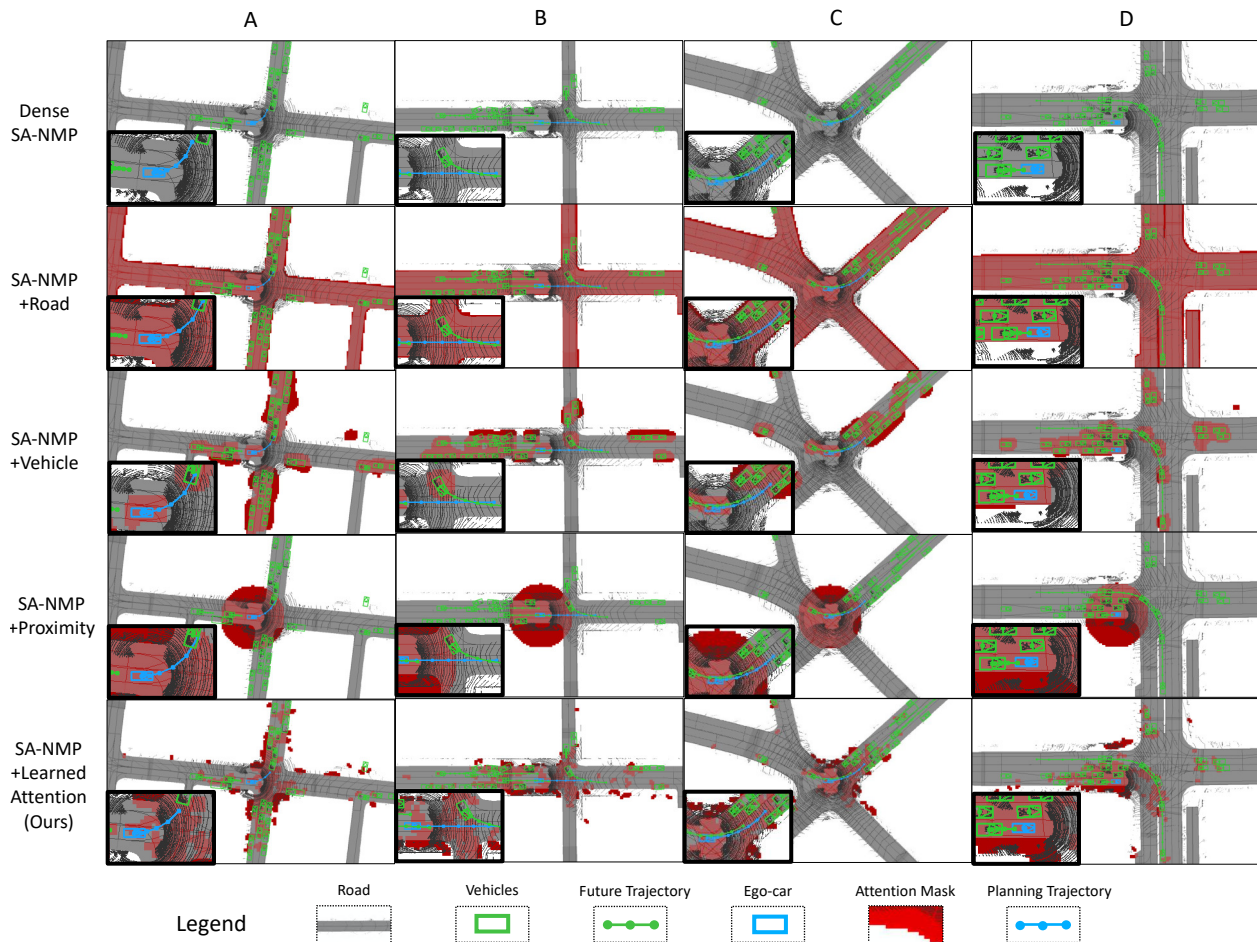


Fig. 5. Visualization of the attention masks and planned trajectory comparing dense, road mask, vehicle mask, proximity mask, and our learned attention. **Col A:** baselines turn too fast and collide with vehicle ahead. **B:** baselines collide with the future position of a left-turning vehicle. **C:** a tight left-turn where all models collide with or nearly miss parked vehicles. **D:** a rear-end collision for all models.

PnP loss is computed only on actors within attention mask: the model learns to cheat by generating attention that avoids all actors resulting in no PnP learning signal, hence the poor performance. For our main experiments, we fix $\gamma_1 = 0.9$.

Detection performance: Since the overall goal is improved motion-planning at reduced computational efficiency, focusing on accurately detecting all actors indiscriminately would be contradictory to the purpose of our learned sparse attention. We should not care as much about far away or irrelevant actors that have little effect on safe planning, and should instead focus our computation on important input regions. Table IV compares detection performance between our learned attention and the baseline dense model evaluated on different subsets of actors in the scene. *Full* includes all actors in the input, while *Attended Region* is the subset of actors that lie within the attention mask. For evaluating the dense model, we use the attention mask generated by our learned model to get the *Attended Region*, ensuring that both models are evaluated on the same actor subsets in both settings. The results shows that our 95% sparse, learned attention model is better than the dense model at detecting actors within the attention mask, meaning that its performance is better focused on actors that it believes are

TABLE IV
DETECTION PERFORMANCE ON DIFFERENT INPUT REGIONS.

Model	mAP on Full			mAP on Attended Region		
	IoU@0.3	IoU@0.5	IoU@0.7	IoU@0.3	IoU@0.5	IoU@0.7
Dense SA-NMP	97.8	94.7	80.3	94.1	93.3	87.9
Ours 95% Sparse	96.3	92.1	74.9	94.2	93.8	88.5

important. This may explain the overall improved planning performance of our attention-driven models as demonstrated in the main quantitative and qualitative results.

V. CONCLUSION

In this work we propose a biologically inspired, end-to-end learned, sparse visual attention mechanism for self-driving, where the sparse attention mask gates the feature backbone computation. As opposed to existing methods that focus on using attention for perception only, our attention masks are directly optimized for motion planning, which enables our network to output better planned trajectories while achieving more efficiency with higher sparsity. As future work, the attention module can be extended to have recurrent feedbacks from the output layers to better leverage temporal information.

REFERENCES

- [1] F. Wang, M. Jiang, C. Qian, S. Yang, C. Li, H. Zhang, X. Wang, and X. Tang, "Residual attention network for image classification," in *CVPR*, 2017.
- [2] M. Figurnov, M. D. Collins, Y. Zhu, L. Zhang, J. Huang, D. P. Vetrov, and R. Salakhutdinov, "Spatially adaptive computation time for residual networks," in *CVPR*, 2017.
- [3] M. Ren, A. Pokrovsky, B. Yang, and R. Urtasun, "Sbnet: Sparse blocks network for fast inference," in *CVPR*, 2018.
- [4] C. Guo, B. Fan, J. Gu, Q. Zhang, S. Xiang, V. Prinet, and C. Pan, "Progressive sparse local attention for video object detection," in *ICCV*, 2019.
- [5] M. Ren and R. S. Zemel, "End-to-end instance segmentation with recurrent attention," in *CVPR*, 2017.
- [6] K. S. Refaat, K. Ding, N. Ponomareva, and S. Ross, "Agent prioritization for autonomous navigation," in *IROS*, 2019.
- [7] L. Itti, G. Rees, and J. K. Tsotsos, *Neurobiology of Attention*. Burlington: Academic Press, 2005.
- [8] L. Itti, C. Koch, and E. Niebur, "A model of saliency-based visual attention for rapid scene analysis," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 20, no. 11, p. 1254–1259, 1998.
- [9] T. Judd, K. A. Ehinger, F. Durand, and A. Torralba, "Learning to predict where humans look," in *ICCV*, 2009.
- [10] J. Ba, V. Mnih, and K. Kavukcuoglu, "Multiple object recognition with visual attention," in *ICLR*, 2015.
- [11] H. Larochelle and G. E. Hinton, "Learning to combine foveal glimpses with a third-order boltzmann machine," in *NIPS*, 2010.
- [12] K. Xu, J. Ba, R. Kiros, K. Cho, A. C. Courville, R. Salakhutdinov, R. S. Zemel, and Y. Bengio, "Show, attend and tell: Neural image caption generation with visual attention," in *ICML*, 2015.
- [13] J. Lu, J. Yang, D. Batra, and D. Parikh, "Hierarchical question-image co-attention for visual question answering," in *NIPS*, 2016.
- [14] Z. Yang, X. He, J. Gao, L. Deng, and A. J. Smola, "Stacked attention networks for image question answering," in *CVPR*, 2016.
- [15] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," in *ICLR*, 2015.
- [16] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *NIPS*, 2017.
- [17] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: pre-training of deep bidirectional transformers for language understanding," in *NAACL-HLT*, 2019.
- [18] B. A. Olshausen and D. J. Field, "Sparse coding with an overcomplete basis set: A strategy employed by v1?" *Vision Research*, vol. 37, pp. 3311–3325, 1997.
- [19] P. Judd, A. D. Lascorz, S. Sharify, and A. Moshovos, "Cnvlutin2: Ineffectual-activation-and-weight-free deep neural network computing," *CoRR*, vol. abs/1705.00125, 2017.
- [20] S. Shi and X. Chu, "Speeding up convolutional neural networks by exploiting the sparsity of rectifier units," *CoRR*, vol. abs/1704.07724, 2017.
- [21] B. Graham, M. Engelcke, and L. van der Maaten, "3d semantic segmentation with submanifold sparse convolutional networks," in *CVPR*, 2018.
- [22] X. Li, Z. Liu, P. Luo, C. C. Loy, and X. Tang, "Not all pixels are equal: Difficulty-aware semantic segmentation via deep layer cascade," in *CVPR*, 2017.
- [23] B. Liu, M. Wang, H. Foroosh, M. F. Tappen, and M. Pensky, "Sparse convolutional neural networks," in *CVPR*, 2015.
- [24] Z. Liu, J. Li, Z. Shen, G. Huang, S. Yan, and C. Zhang, "Learning efficient convolutional networks through network slimming," in *ICCV*, 2017.
- [25] A. Kendall, Y. Gal, and R. Cipolla, "Multi-task learning using uncertainty to weigh losses for scene geometry and semantics," in *CVPR*, 2018.
- [26] Z. Chen, V. Badrinarayanan, C. Lee, and A. Rabinovich, "Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks," in *ICML*, 2018.
- [27] M. Guo, A. Haque, D. Huang, S. Yeung, and L. Fei-Fei, "Dynamic task prioritization for multitask learning," in *ECCV*, 2018.
- [28] X. Lin, H. S. Baweja, G. Kantor, and D. Held, "Adaptive auxiliary task weighting for reinforcement learning," in *NeurIPS*, 2019.
- [29] M. Ren, W. Zeng, B. Yang, and R. Urtasun, "Learning to reweight examples for robust deep learning," in *ICML*, 2018.
- [30] S. Liu, E. Johns, and A. J. Davison, "End-to-end multi-task learning with attention," in *CVPR*, 2019.
- [31] S. Shalev-Shwartz, S. Shammah, and A. Shashua, "On a formal model of safe and scalable self-driving cars," *CoRR*, vol. abs/1708.06374, 2017.
- [32] Ö. S. Tas, F. Hauser, and C. Stiller, "Decision-time postponing motion planning for combinatorial uncertain maneuvering," in *ITSC*, 2018.
- [33] P. F. Orzechowski, A. Meyer, and M. Lauer, "Tackling occlusions & limited sensor range with set-based safety verification," in *ITSC*, 2018.
- [34] C. Pek and M. Althoff, "Computationally efficient fail-safe trajectory planning for self-driving vehicles using convex optimization," in *ITSC*, 2018.
- [35] W. Zeng, W. Luo, S. Suo, A. Sadat, B. Yang, S. Casas, and R. Urtasun, "End-to-end interpretable neural motion planner," in *CVPR*, 2019.
- [36] H. Fan, F. Zhu, C. Liu, L. Zhang, L. Zhuang, D. Li, W. Zhu, J. Hu, H. Li, and Q. Kong, "Baidu apollo EM motion planner," *CoRR*, vol. abs/1807.08048, 2018.
- [37] A. Sadat, M. Ren, A. Pokrovsky, Y. Lin, E. Yumer, and R. Urtasun, "Jointly learnable behavior and trajectory planning for self-driving vehicles," in *IROS*, 2019.
- [38] S. Singh, J. Lacotte, A. Majumdar, and M. Pavone, "Risk-sensitive inverse reinforcement learning via semi- and non-parametric methods," *I. J. Robotics Res.*, vol. 37, no. 13-14, 2018.
- [39] J. Lacotte, M. Ghavamzadeh, Y. Chow, and M. Pavone, "Risk-sensitive generative adversarial imitation learning," in *AISTATS*, 2019.
- [40] D. H. Shin and S. Singh, "Path generation for robot vehicles using composite clothoid segments," 1990.
- [41] M. Liang, B. Yang, W. Zeng, Y. Chen, R. Hu, S. Casas, and R. Urtasun, "Object trajectory evolution for end-to-end perception and prediction," in *CVPR*, 2020.
- [42] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *MICCAI*, 2015.
- [43] E. Jang, S. Gu, and B. Poole, "Categorical reparameterization with gumbel-softmax," in *ICLR*, 2017.
- [44] C. J. Maddison, A. Mnih, and Y. W. Teh, "The concrete distribution: A continuous relaxation of discrete random variables," in *ICLR*, 2017.
- [45] Y. Bengio, N. Léonard, and A. C. Courville, "Estimating or propagating gradients through stochastic neurons for conditional computation," *CoRR*, vol. abs/1308.3432, 2013.
- [46] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nusenes: A multimodal dataset for autonomous driving," *CoRR*, vol. abs/1903.11027, 2019.
- [47] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *ICLR*, 2015.
- [48] M. Liang, B. Yang, W. Zeng, Y. Chen, R. Hu, S. Casas, and R. Urtasun, "Pnpnet: End-to-end perception and prediction with tracking in the loop," *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11 550–11 559, 2020.
- [49] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, "Pyramid scene parsing network," *CVPR*, pp. 6230–6239, 2017.