DAA*: Deep Angular A Star for Image-based Path Planning
Supplementary Material

Zhiwei Xu
The University of Melbourne, Australia
danny.xu@unimelb.edu.au

A. Learned Weights

To highlight the importance and validation of the path angular freedom (PAF), we show the learned α in the main paper. Here, we provide all the learned weights in Table 6.

Dataset	Method	α	λ	κ
MPD	DAA*-min	100.0	19.9	29.6
	DAA*-max	0.0	52.5	36.2
	DAA*-mix	33.4	66.0	75.3
TMPD	DAA*-min	100.0	17.2	65.9
	DAA*-max	0.0	32.1	0.1
	DAA*-mix	76.9	48.1	83.8
CSM	DAA*-min	100.0	47.3	100.0
	DAA*-max	0.0	32.1	0.1
	DAA*-mix	68.7	51.7	76.8
Aug-TMPD	DAA*	96.4	63.3	97.8
	DAA*-path	58.7	55.2	56.3
	DAA*-mask	69.6	50.5	77.4
	DAA*-weight	62.2	49.9	70.3
Warcraft	DAA*-min	100.0	75.4	16.7
	DAA*-max	0.0	78.8	43.9
	DAA*-mix	27.7	80.5	74.4
Pokémon	DAA*-min	100.0	78.5	22.8
	DAA*-max	0.0	83.4	44.0
	DAA*-mix	32.3	83.2	70.2
SDD-intra	DAA*-min	100.0	53.5	0.2
	DAA*-max	0.0	73.4	68.4
	DAA*-mix	9.5	77.9	91.4
SDD-inter	DAA*-min	100.0	29.6	3.0
	DAA*-max	0.0	69.0	57.7
	DAA*-mix	17.4	73.0	84.1

Table 6. Learned weights, scaled by $\times 100$ for high readability. We set $\alpha = 1$ for DAA*-min and $\alpha = 0$ for DAA*-max. We report the mean values where multiple training seeds are applied.

B. Relation of α and Predicted Paths

In Eq. (5), a large α aims to minimize the path angle to often achieve linear paths, while a small one maximizes it for smooth paths. The learned $\alpha = 0.69$ in CSM encourages small path angles for short paths, such as linear path segments. For example, in Fig. 2(a), the path consists essentially of 3 main linear segments by DAA*. This inspection also explains Fig. 2(b) for the learned $\alpha = 0.62$ by DAA*.

Meanwhile, unlike these binary maps in Figs. 2(a)–(b) where obstacles and accessible areas are provided in the datasets, video-game maps and SDD have more complex scenarios and have to learn cost maps in Figs. 2(c)–(f). These learned costs are often continuous on the edges of objects such as rocks, lakes, and roundabouts. Thus, smooth paths can generally be achieved by increasing the path angle. This aligns with the small α, 0.28 and 0.32 on video-game datasets and particularly 0.1 and 0.17 on SDD where realistic road scenes require smaller α to obtain smoother paths considering driving safety.

C. Training Procedure

In addition to the pathfinding of DAA* in Alg. 1, we also detail the training procedure with path loss and the combination of path loss and PPM loss in Alg. 2.

D. Ethical and Societal Impact

This work aims to automate effective vision-based path planning through end-to-end learning of path shortening and smoothing. While our DAA* promises significant improvements in imitating human and machine demonstrations, it also raises potential ethical and societal concerns such as auto-driving safety in complex environments, model robustness under attacks, and energy consumption in model training. For instance, under certain attacks in a system such as ADAS, α can be maliciously changed to 1, where predicted paths are forced to be linear crossing inaccessible areas say roundabouts or unsafe zigzags causing traffic chaos.

However, according to the trade-off between searching efficiency and path optimality, the search cost 0.5%–6.4% is minor. We also reveal the necessity of only binary path

[1] The discretized activation disables gradient accumulation by $(\cdot)_{\text{detach}}$, see Line 74 of A*. This is equivalent to computing the gradients from p_k.

Algorithm 2: Training Procedure of Deep Angular A*

Input: A dataset $\mathcal{D} = \{\mathcal{I}_i, s_i, t_i, \hat{\mathcal{P}}_i, \hat{\boldsymbol{\theta}}_i\}$ with a map image $\mathcal{I}_i$ containing N pixels, a source node s_i, a target node t_i, a reference path $\hat{\mathcal{P}}_i$, and a reference PPM $\hat{\boldsymbol{\theta}}_i$ (only required by PPM loss) for all $i \in \{1, ..., M\}$ given M samples, a mask threshold $\mathcal{T} = 0.5$, learning rate l, and loss weights w_1 and w_2 for path loss and PPM loss, respectively.
Parameter: A cost-map encoder network $f(\mathcal{I}, s, t)$, containing weights $\mathbf{w}$ and bias $\mathbf{b}$, and term weights $\{\alpha, \lambda, \kappa\}$.
Note: On *Aug-TMPD*, we set $w_1 = w_2 = 1$ for DAA*, $w_1 = 1$ and $w_2 = 0$ for DAA*-path, $w_1 = w_2 = 1$ for DAA*-mask with a cost-map mask, and $w_1 = 10$ and $w_2 = 1$ for DAA*-weight. *On the other datasets*, we set $w_1 = 1$ and $w_2 = 0$ since only the path loss is required.
Procedure:

1: Initialize $\alpha = 0.5$, $\lambda = 0.5$, and $\kappa = 1$.
2: **for** $i \in \{1, ..., M\}$ **do**
3: Compute a cost map $\boldsymbol{\theta}_i = f(\mathcal{I}_i, s_i, t_i)$.
4: Compute a confidence list $\boldsymbol{p} = \{p_k\}$ in Eq. (9) for all $k \in \mathcal{M}_c$ by Step 1, Alg. 1 given inputs $\{\boldsymbol{\theta}_i, s_i, t_i\}$.
5: Form a reference path map $\hat{\mathcal{Y}}_i \in \mathcal{B}^N$ with 1 for all nodes in $\hat{\mathcal{P}}_i$ and 0 otherwise.
6: Form a search history map $\mathcal{Y} \in \mathcal{B}^N$ with $(1 - p_k)_{\text{detach}} + p_k$ for all $k \in \mathcal{M}_c$ using discretized activation[1] and 0 otherwise.
7: Compute path loss $\mathcal{L}_i^p = ||\mathcal{Y} - \hat{\mathcal{Y}}_i||_1$.
8: For DAA*-mask, compute a cost-map mask $\boldsymbol{m} = \mathbb{1}[\boldsymbol{\theta}_i \geq \mathcal{T}]$, then update $\boldsymbol{\theta}_i \leftarrow \boldsymbol{\theta}_i \odot \boldsymbol{m}$ and $\hat{\boldsymbol{\theta}}_i \leftarrow \hat{\boldsymbol{\theta}}_i \odot \boldsymbol{m}$ using Hadamard product.
9: Compute PPM loss $\mathcal{L}_i^c = ||\boldsymbol{\theta}_i - (1 - \hat{\boldsymbol{\theta}}_i)||_2$.
10: Compute a weighted loss $\mathcal{L}_i = w_1 \mathcal{L}_i^p + w_2 \mathcal{L}_i^c$.
11: **end for**
12: Compute the average loss $\mathcal{L} = \frac{1}{M} \sum_{i=1}^{M} \mathcal{L}_i$.
13: Compute the gradients of learnable parameters, $\nabla_{\boldsymbol{w}} \mathcal{L}$, $\nabla_{\boldsymbol{b}} \mathcal{L}$, $\nabla_{\alpha} \mathcal{L}$, $\nabla_{\lambda} \mathcal{L}$, and $\nabla_{\kappa} \mathcal{L}$, by backpropagting $\mathcal{L}$.
14: Update parameters $\boldsymbol{w} \leftarrow \boldsymbol{w} - l \nabla_{\boldsymbol{w}} \mathcal{L}$, $\boldsymbol{b} \leftarrow \boldsymbol{b} - l \nabla_{\boldsymbol{b}} \mathcal{L}$, $\alpha \leftarrow \alpha - l \nabla_{\alpha} \mathcal{L}$, $\lambda \leftarrow \lambda - l \nabla_{\lambda} \mathcal{L}$, and $\kappa \leftarrow \kappa - l \nabla_{\kappa} \mathcal{L}$.
15: Repeat Lines 2-14 till the training loss converges.

labelling for supervised learning with more accessible path reference rather than PPM reference. We aim to further mitigate these issues by developing fair, transferable, robust, and more efficient algorithms in future work. Our commitment is to ensure that autonomous path planning is socially responsible and environmentally sustainable.