

Figure 4: Ratio between the variance of the gradient and its norm as a function of the similarity between the viewing directions \mathbf{z}_i and \mathbf{z}_j .

A. Gradient of Common-Line loss at co-planar poses

In this section, we show how the presence of images with co-planar poses affect our common-lines based loss and, more specifically, its gradient.

Consider two images o_i and o_j and their predicted poses $R_i = (\mathbf{x}_i, \mathbf{y}_i, \mathbf{z}_i) \in \mathbb{R}^{3 \times 3}$ and $R_j = (\mathbf{x}_j, \mathbf{y}_j, \mathbf{z}_j) \in \mathbb{R}^{3 \times 3}$. Note we assume $R_* \in \mathbb{R}^{3 \times 3}$ rather than $R_* \in \text{SO}(3)$ since our loss depends on R_i, R_j parameterized as matrices in $\mathbb{R}^{3 \times 3}$ and the following arguments are independent of how our neural networks predict elements of $\text{SO}(3)$ (although the particular choice of parameterization can affect how this gradient is back-propagated through the model).

Note that the loss $\mathcal{L}(R_i, R_j)$ from Eq. 17 or Eq. 18 only depends on the coordinates in Eq. 16, i.e.:

$$\mathbf{x}_i = \mathbf{l}_{ij}^T \mathbf{x}_i \quad \mathbf{y}_i = \mathbf{l}_{ij}^T \mathbf{y}_i \quad \mathbf{x}_j = \mathbf{l}_{ij}^T \mathbf{x}_j \quad \mathbf{y}_j = \mathbf{l}_{ij}^T \mathbf{y}_j$$

For convenience, define $\mathbf{v}_i = (x_i, y_i)^T \in \mathbb{R}^2$ and $\mathbf{v}_j = (x_j, y_j)^T \in \mathbb{R}^2$ and write $\mathcal{L}_{ij} = \mathcal{L}(R_i, R_j)$. Then, using the chain rule:

$$\frac{\partial \mathcal{L}_{ij}}{\partial R_i} = \frac{\partial \mathbf{v}_i}{\partial R_i} \frac{\partial \mathcal{L}_{ij}}{\partial \mathbf{v}_i} + \frac{\partial \mathbf{v}_j}{\partial R_i} \frac{\partial \mathcal{L}_{ij}}{\partial \mathbf{v}_j} \quad (21)$$

Hence, in this section, we can focus only on the partial derivatives $\frac{\partial \mathbf{v}_i}{\partial R_i}$ and $\frac{\partial \mathbf{v}_j}{\partial R_i}$. We will study the partial derivatives with respect to $\mathbf{x}_i, \mathbf{y}_i$ and \mathbf{z}_i independently.

First, note that:

$$\frac{\partial x_i}{\partial \mathbf{x}_i} = \frac{\partial y_i}{\partial \mathbf{y}_i} = \mathbf{l}_{ij} \quad (22)$$

$$\frac{\partial y_i}{\partial \mathbf{x}_i} = \frac{\partial x_i}{\partial \mathbf{y}_i} = 0 \quad (23)$$

and

$$\frac{\partial x_j}{\partial \mathbf{x}_i} = \frac{\partial y_j}{\partial \mathbf{x}_i} = \frac{\partial x_j}{\partial \mathbf{y}_i} = \frac{\partial y_j}{\partial \mathbf{y}_i} = 0 \quad (24)$$

Hence, we only need to consider the following quantities: $\frac{\partial \mathbf{v}_i}{\partial \mathbf{x}_i}, \frac{\partial \mathbf{v}_i}{\partial \mathbf{z}_i} = (\frac{\partial x_i}{\partial \mathbf{z}_i}, \frac{\partial y_i}{\partial \mathbf{z}_i})$ and $\frac{\partial \mathbf{v}_j}{\partial \mathbf{z}_i} = (\frac{\partial x_j}{\partial \mathbf{z}_i}, \frac{\partial y_j}{\partial \mathbf{z}_i})$.

To understand how the similarity of \mathbf{z}_i and \mathbf{z}_j , we study the variance of the gradients when R_j is perturbed by a small amount of noise, as a function of the similarity $|\mathbf{z}_i^T \mathbf{z}_j|$.

To do so, we sample 300 random pairs of rotations $R_i, R_j \in \text{SO}(3)$ and compute their similarity $|\mathbf{z}_i^T \mathbf{z}_j|$. Then, for each pair, we generate 50 variations of R_j by perturbing it with a small Gaussian noise with standard deviation $\sigma = 0.04$ in the quaternion space, and compute each gradient $\frac{\partial x_i}{\partial \mathbf{x}_i}$, $\frac{\partial v_i}{\partial \mathbf{z}_i}$ and $\frac{\partial v_j}{\partial \mathbf{z}_i}$.

For each pair, we compute the average norm of the gradients (Frobenious norm for the Jacobians) and the standard deviation (over the 50 samples) of each partial derivative, which we average to obtain a single number. In Fig. 4, we plot the ratio between the standard deviation and the average norm for each pair, as a function of the similarity $|\mathbf{z}_i^T \mathbf{z}_j|$.

Whenever \mathbf{z}_i is close to $\pm \mathbf{z}_j$, the variance of the gradient is very close to its average norm (the ratio approaches 1); this is particularly true for the gradients of \mathbf{z}_i , see Fig. 4a but less severe for \mathbf{z}_i and \mathbf{y}_i . That result suggests that the training process can be particularly unstable in this setting, especially since the gradient on \mathbf{z}_i is necessary to leave this situation but it is also the most affected by that.

B. Common-Line Loss Landscape

In this section, we provide a simple study of the common-line loss landscape. To do so, we compute the common line loss between two random images o_i and o_j , respectively at the poses $R_i R_{\theta_1}$ and $R_j R_{\theta_2}$.

$R_i = (\mathbf{x}_i, \mathbf{y}_i, \mathbf{z}_i)$ and $R_j = (\mathbf{x}_j, \mathbf{y}_j, \mathbf{z}_j)$ are the ground-truth poses of o_i and o_j .

$R_\theta = \begin{bmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix}$ is a 2D rotation by $\theta \in [0, 2\pi)$.

Hence, $R_i R_\theta$ is the pose of o_i rotated by θ around its projection axis \mathbf{z}_i and simply corresponds to rotating the common line $\mathbf{v}_i = (\mathbf{x}_i^T \mathbf{l}_{ij}, \mathbf{y}_i^T \mathbf{l}_{ij})^T \in \mathbb{R}^2$ in the image o_i by θ . Note that if $\theta_1 = \theta_2 = \pi$, the loss is unchanged since the common line predicted is the same, only reflected.

In Fig. 5, we plot the common-line loss in Eq. 18 as a function of θ_1 and θ_2 , for different random pairs (i, j) . Note that any possible predicted pair of common lines corresponds to a point in the figure (multiple choices of $R_i, R_j \in \text{SO}(3)$ lead to the same common lines). This enables us to study the complete loss landscape for the simple case of $N = 2$ images. The right column of Fig. 5 highlights the global minima of the loss.

First, we note the expected periodicity of the loss by $\theta_1 = \theta_2 = \pi$ in all images.

In the first pair (first row), we also observe spurious global minima at $\theta_1 = 0$ and $\theta_2 = \pi$ (and the opposite), which corresponds to a reflection of the correct common line in only one of the two images. This is likely related to the spurious planar symmetry described in Levy et al. (2022a) and in Sec. 4, which motivated the use of a "symmetrized loss" in Levy et al. (2022a).

We also note that the landscape can vary a lot over different pairs. While the first pair has a smooth landscape with two clear global minima at expected locations $(0, 0)$ and (π, π) , other pairs show multiple global optima. In some cases, like the last row, the two locations $(0, 0)$ and (π, π) are close but not exactly global optima.

We also compare the original formulation of the loss in Eq. 17 with the modified version in Eq. 18 which we use in our experiments. Fig. 6 shows similar plots obtained using Eq. 17 (pairs are randomly sampled and don't necessary match those in Fig. 5). When using the original loss in Eq. 17, the global optima often do not include the ground-truth $(0, 0)$ and (π, π) .

Finally, we emphasise that this study is limited to the case $N = 2$. However, during training, the loss is averaged over multiple pairs.

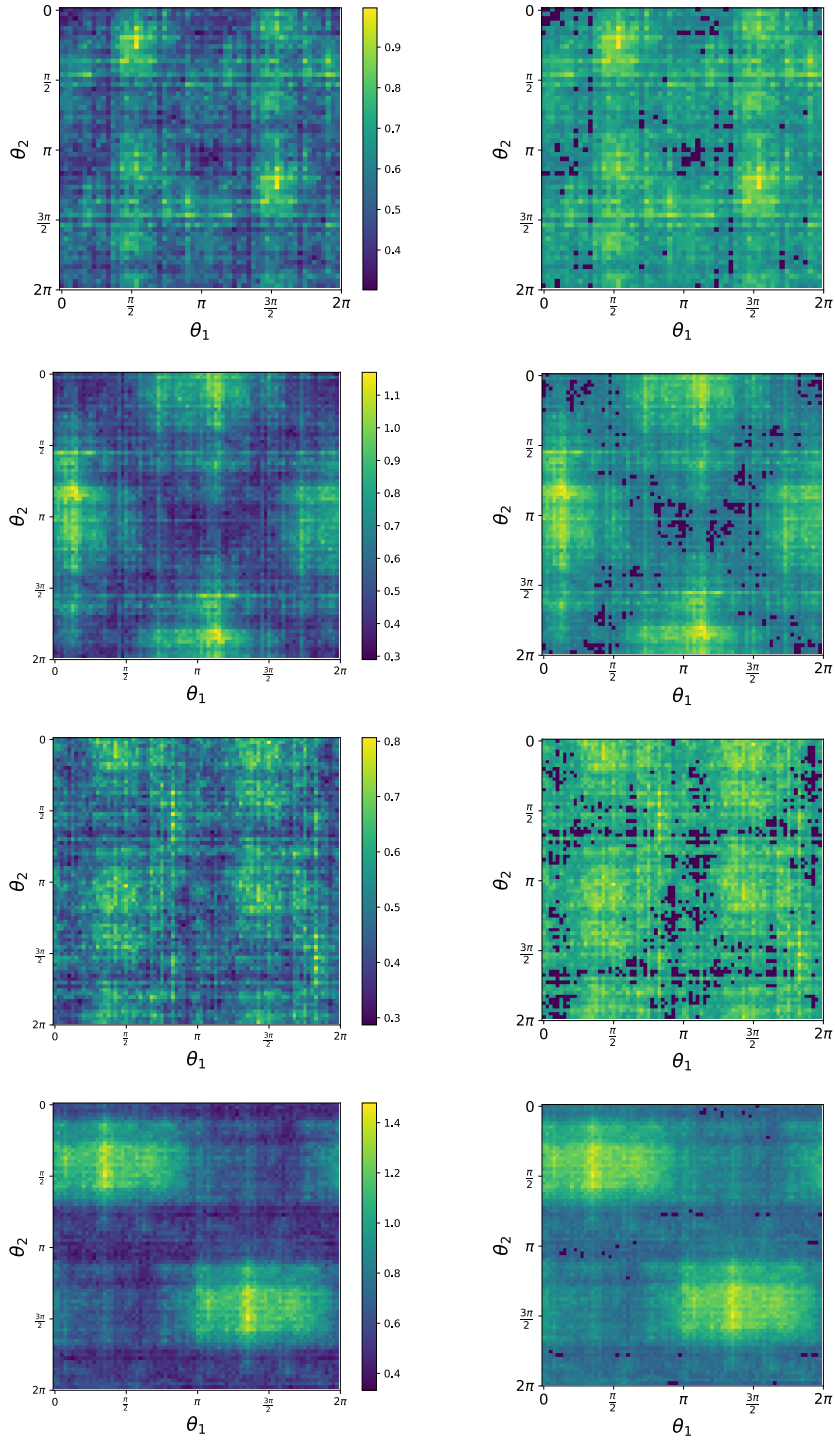


Figure 5: Common line loss using Eq. 18 between two random images o_i and o_j using the poses $R_i R_{\theta_1}$ and $R_j R_{\theta_2}$, with $\theta_1, \theta_2 \in [0, 2\pi)$. Each row is a different random pair (i, j) . In the right column, areas where $\mathcal{L} < 0.4$ are highlighted with a darker color.

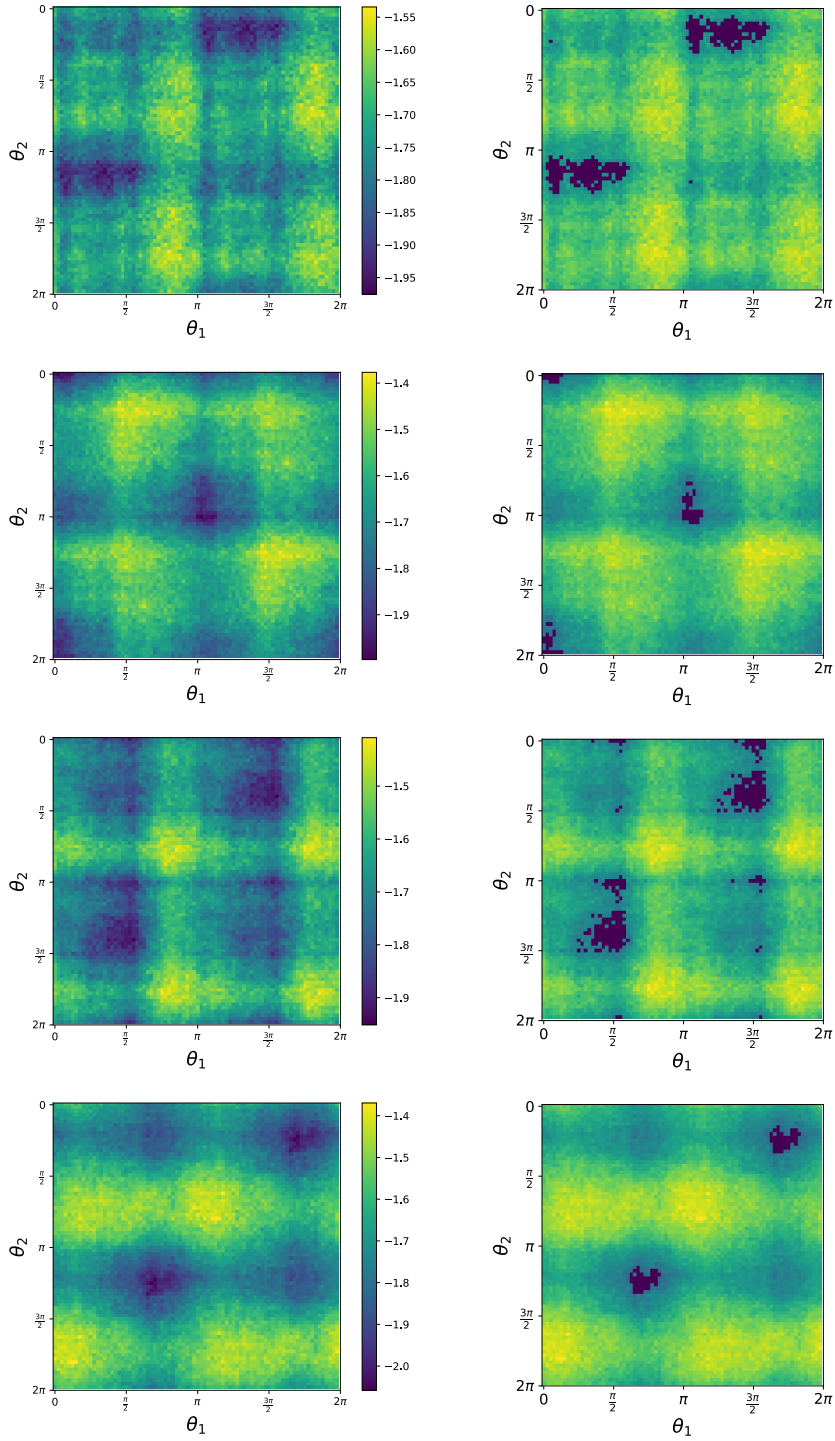


Figure 6: Common line loss using Eq. 17 between two random images o_i and o_j using the poses $R_i R_{\theta_1}$ and $R_j R_{\theta_2}$, with $\theta_1, \theta_2 \in [0, 2\pi)$. Each row is a different random pair (i, j) . In the right column, the points closer to global optima are highlighted with a darker color.