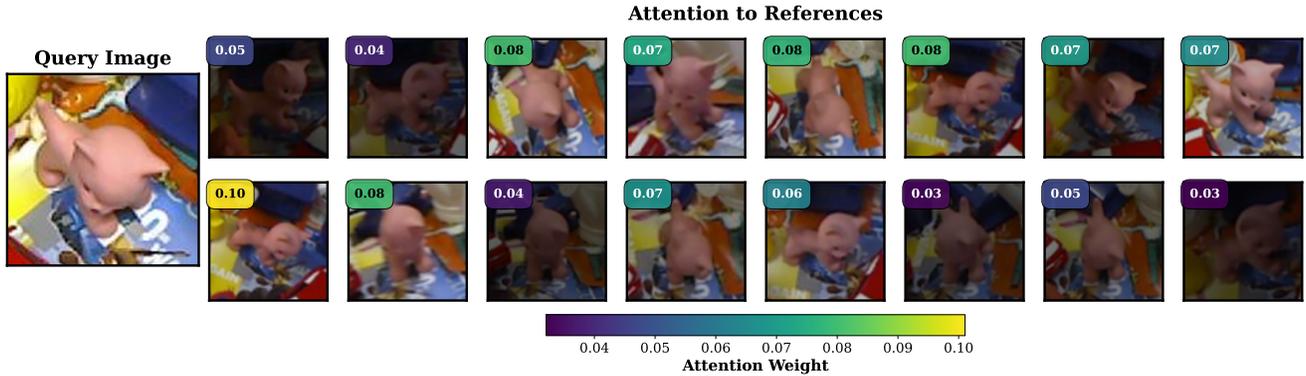


Supplementary Material

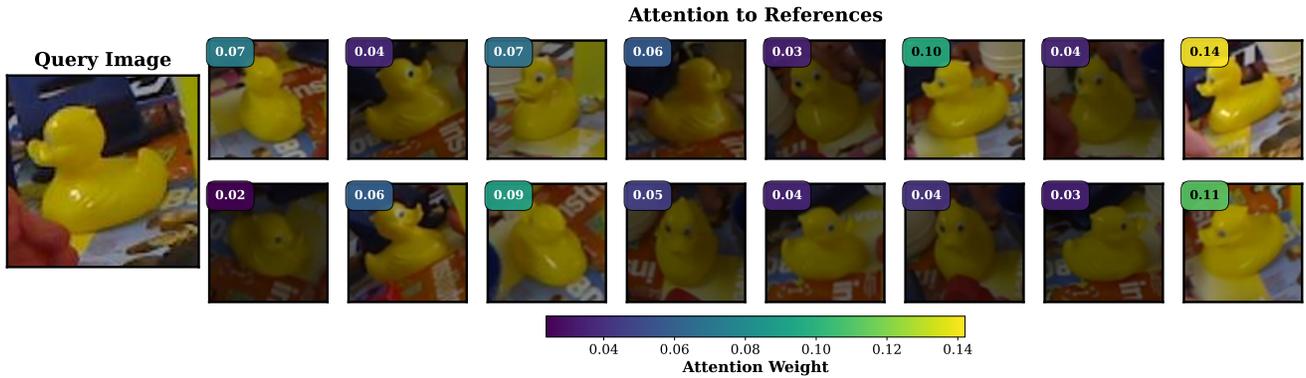
Eff-GRot: Efficient and Generalizable Rotation Estimation with Transformers

Fanis Mathioulakis Gorjan Radevski Tinne Tuytelaars
KU Leuven

fanis.mathioulakis@kuleuven.be tinne.tuytelaars@kuleuven.be



(a) Angular error = 8.6°

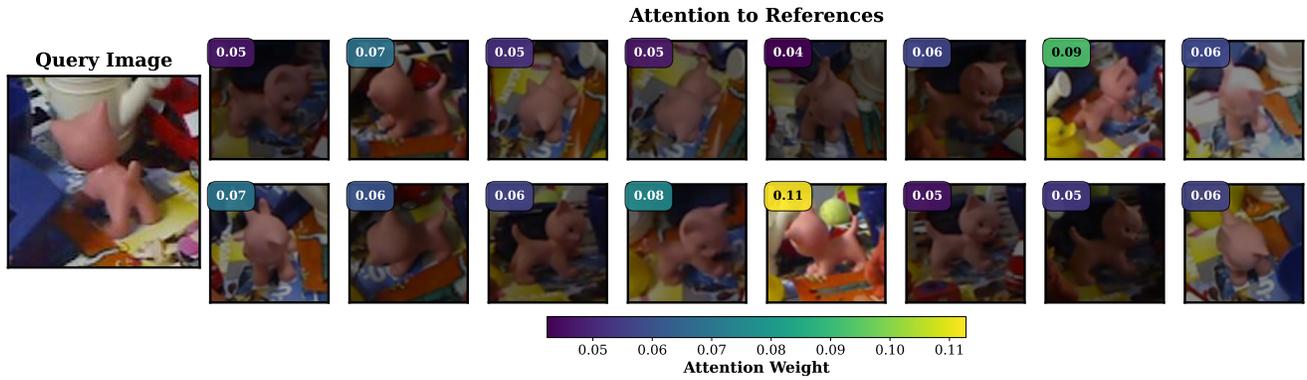


(b) Angular error = 8.3°

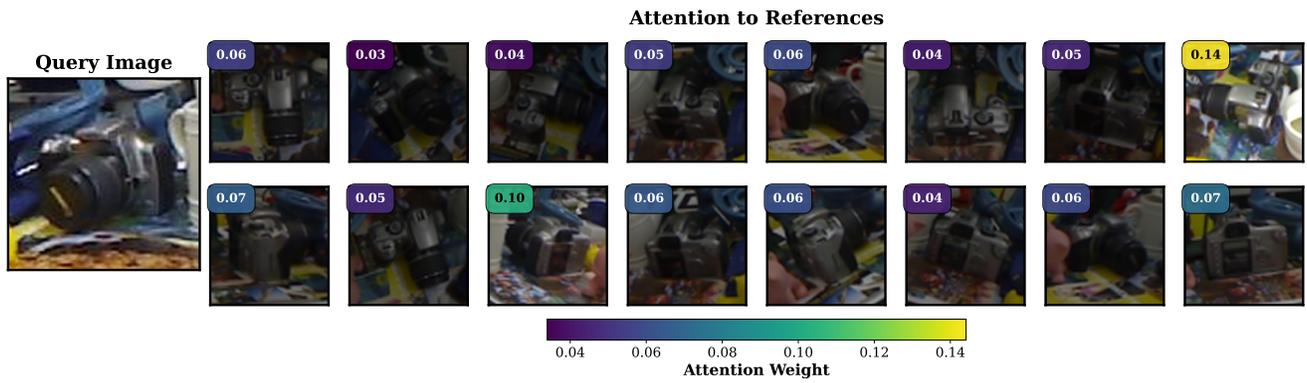
Figure 1. Transformer attention in examples with accurate rotation estimation.

Qualitative Examples of Correct and Incorrect Predictions. We provide additional examples illustrating both correct and incorrect rotation predictions. Each example shows the query image, the corresponding reference images with their transformer attention scores, and the angular error of the prediction. Figure 1 depicts cases with low angular error, where the model focuses on nearby viewpoints and makes a rotation prediction with low error. Figure 2 highlights three representative failure cases. In case (a), most reference views remain far from the query, making accu-

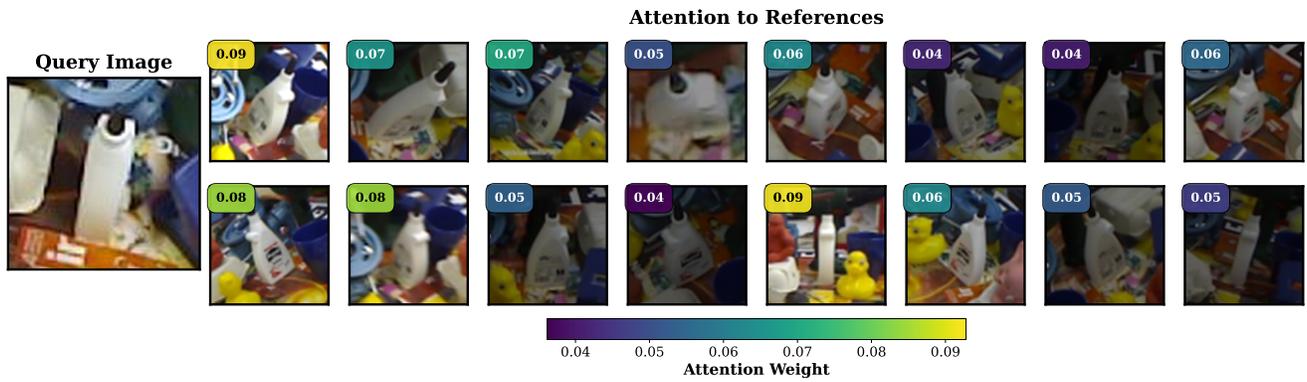
rate prediction difficult. Case (b) shows the model assigning high attention to a reference view with a large rotational difference from the query, likely because the learned feature representations make it appear similar to the query, which in turn leads to a large prediction error. Case (c) demonstrates an error due to object symmetry: the model places high attention on the fifth reference image (bottom row, from the left), which is visually similar but corresponds to the opposite side of the object (a 180° rotation), resulting in a large prediction error of 159° .



(a) Angular error = 73.1°



(b) Angular error = 70.1°



(c) Angular error = 159.0°

Figure 2. Visualization of transformer attention in cases where rotation prediction fails.