

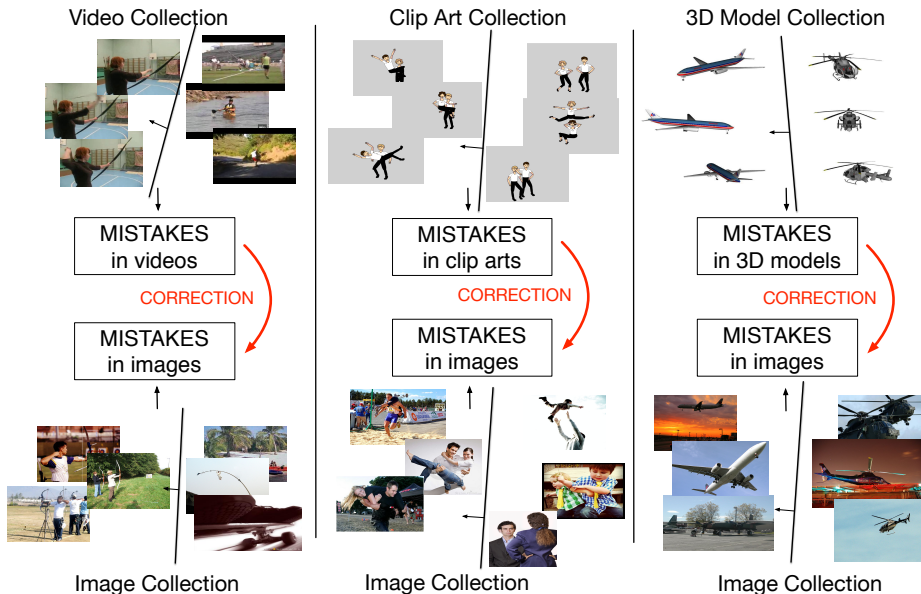


Learning from the Mistakes of Others: Matching Errors in Cross-Dataset Learning

Viktoriia Sharmanska & Novi Quadrianto



What do we propose?



How to do this?

By matching the distribution of mistakes/errors across datasets:

$$\underset{\mathbf{w} \in \mathbb{R}^d}{\text{minimize}} \quad \underbrace{\|\mathbf{w}\|^2}_{\text{regularization}} + \underbrace{C \text{ MMD}(p_{d^*}, p_d)}_{\text{loss := divergence between own mistakes and others' mistakes}}$$

mistakes: $\{1 - y_i \langle \mathbf{w}, \mathbf{x}_i \rangle\}_{i=1}^N$

Unlike other approaches:

- prior is encoded in the loss term
- we match errors (not features or outputs)

Does it work?

More at Poster 19

