On the Minimal Recognizable Image Patch

Presented By: Mark Fonaryov

Related Work – Human Recognition

[8] - Atoms of recognition in human and computer vision.

Discovered MIRCs:

- MIRCs were of different sizes and positions in each image.
- Each image was covered by multiple MIRCs (15.1 ± 7.6).

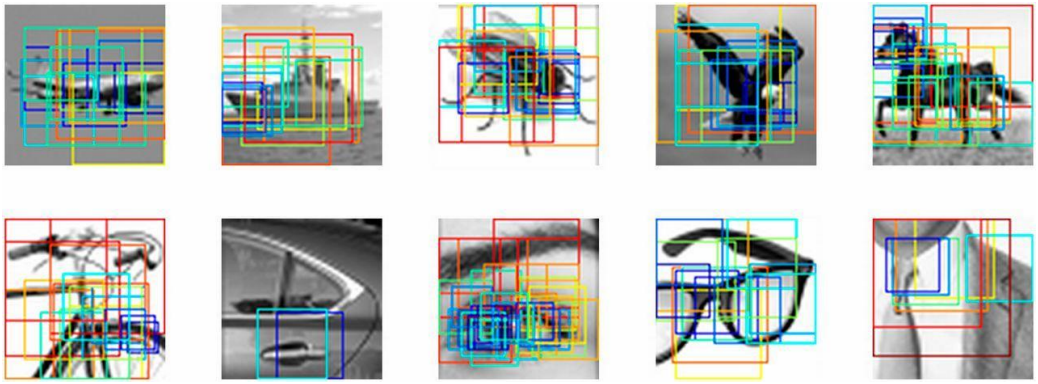

[8] - Shimon Ullman, Liav Assif, Ethan Fetaya, and Daniel Harari. **Atoms of recognition in human and computer vision**. Proceedings of the National Academy of Sciences, 2016.

Related Work – Human Recognition

[8] - Atoms of recognition in human and computer vision.

Discovered MIRCs:

- MIRCs were associated with sharp reduction in accuracy - 0.71 ± 0.05.
- Recognition algorithms tested on MIRCs – did not produce sharp drops - 0.14 ± 0.24.

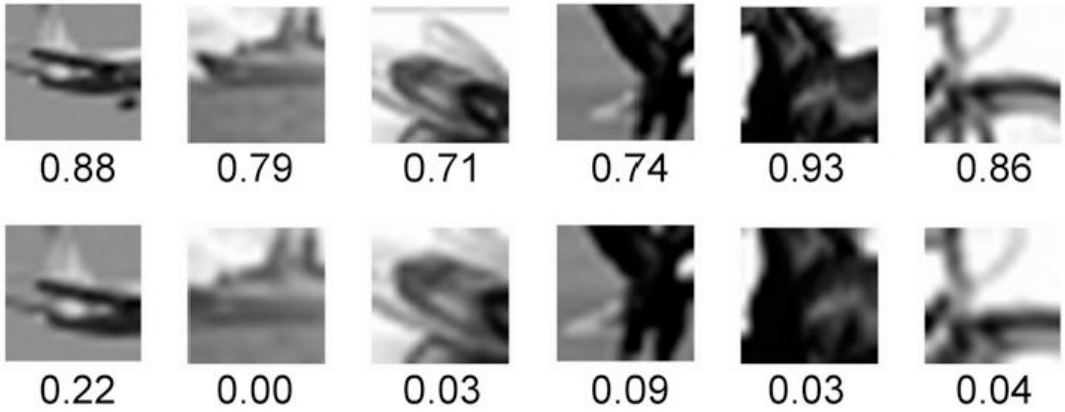

[8] - Shimon Ullman, Liav Assif, Ethan Fetaya, and Daniel Harari. **Atoms of recognition in human and computer vision**. Proceedings of the National Academy of Sciences, 2016.

Patch Based Classification

Determining patch recognizability: the Patch-Based Classification (PBC) model:

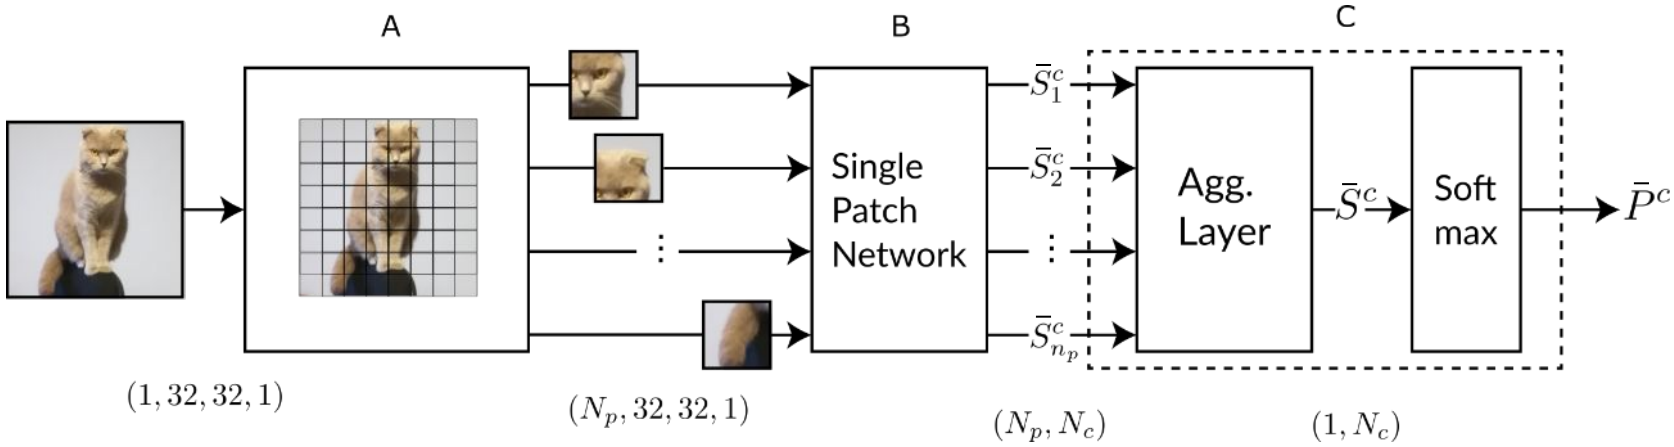

(A) Input image is split into N_p patches, each resized to 32×32.

(B) Each patch passes thru the single-patch-network (SPN).

(C) Aggregation – patch-level scores -> image-level scores -> image-level probabilities.

Patch Score Aggregation

Aggregation influences the confidence associated with the different categories.

confidence –> prediction loss -> training process

Category-independent max - Maximum score is evaluated separately for each category.

$$S^c_{max-ind} = \max_p \{S^c_p\}$$

Winner-directed max – Scores are taken from a single patch with overall maximum score.

$$S^c_{max-dir} = S^c_{p*}$$
$$p^* = argmax_p(\max_c S^c_p)$$

Single-Image Recognizability

Histograms of maximal confidence drops:

- Most images include a sharp and significant sharp drop.
- Category-independent max – average maximal drop is 0.624.
- Winner-directed max – average maximal drop is 0.72.

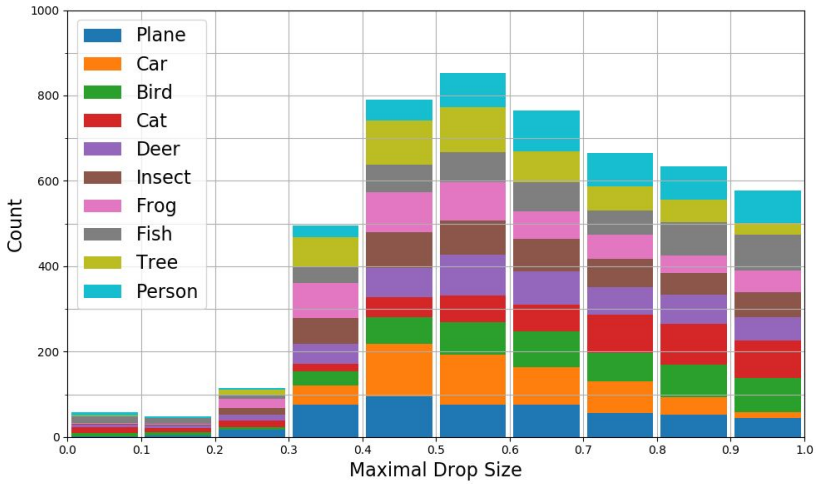

Category-independent max

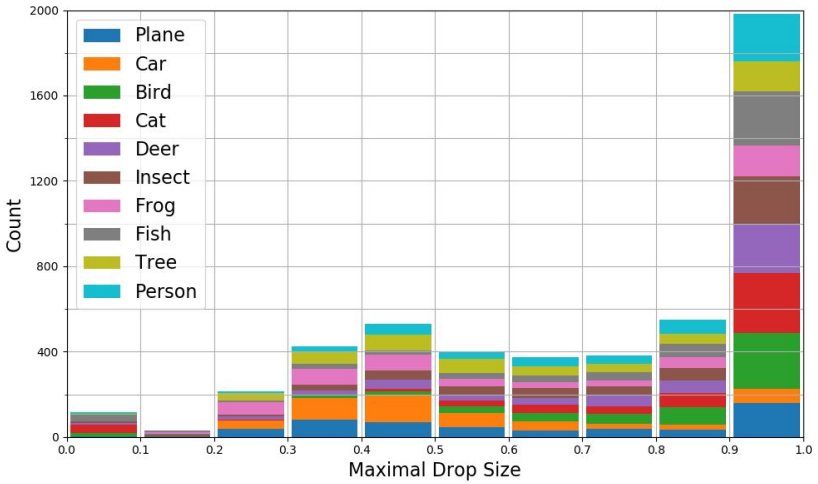

Winner-directed max

Globally Minimal Patches

MRP examples:

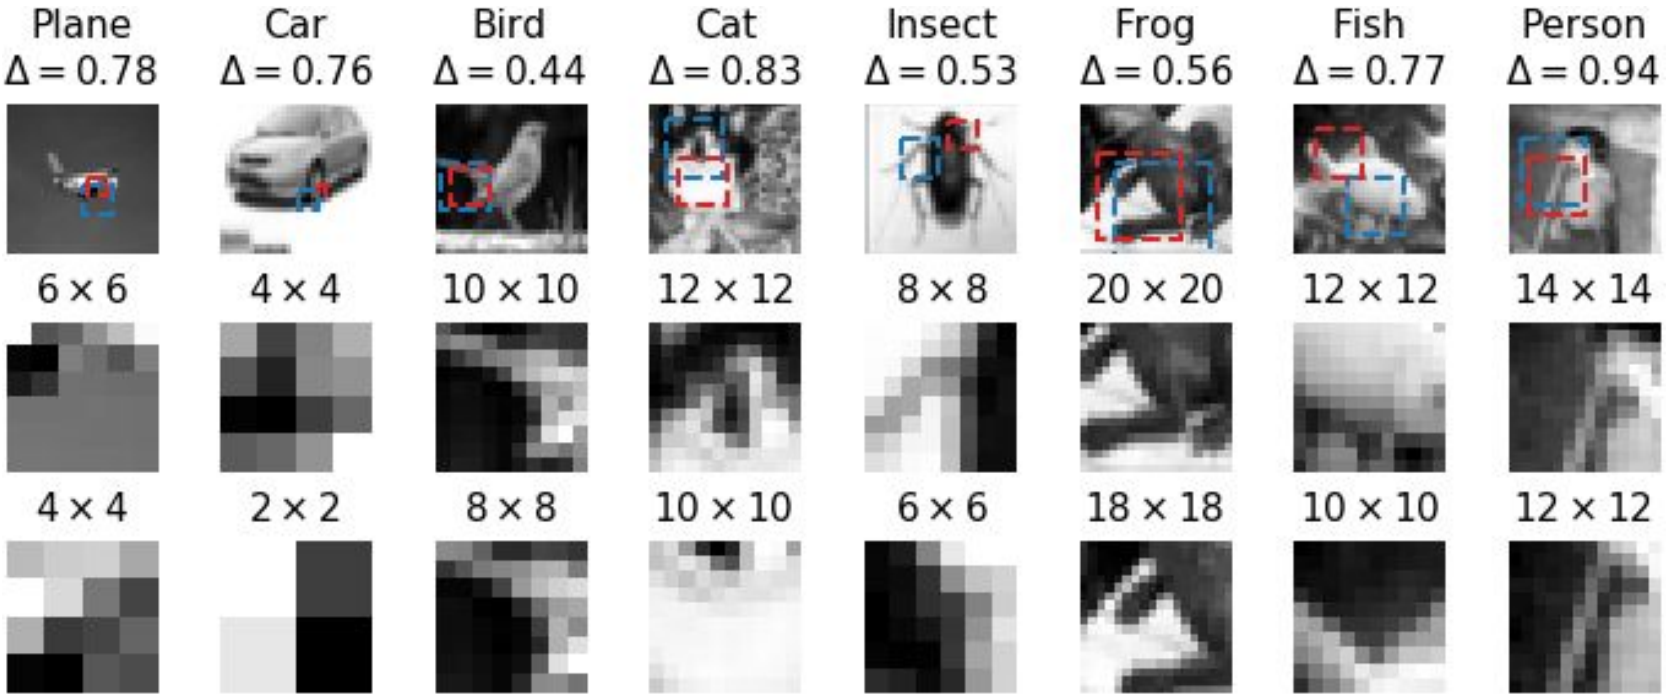

Locally Minimal Patches

- Computational *MIRC* (*cMIRC*) – a patch that is *q-locally recognizable*, while all its nine contained sub-patches are not.
- Example:

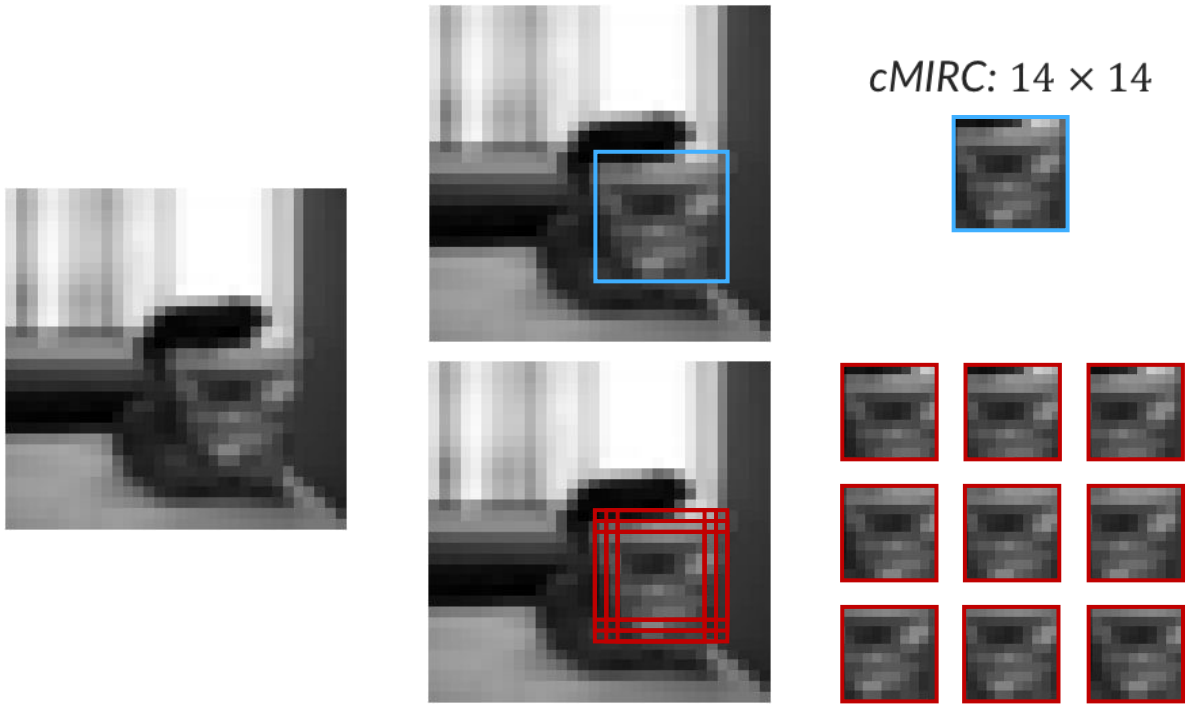

Conclusions

This work empirically characterized globally and locally minimal patches.

- Both *MRP*s and *cMIRC*s share a common property with human vision – sharp drops.
- *MRP*s were small, and usually unrecognizable by humans.
  - Likely due to closed-set setting and small number of classes.
- Further work – Estimate *MRP*s that are more consistent with human vision.
  - Using more classes.
  - Using open-set classification tools.