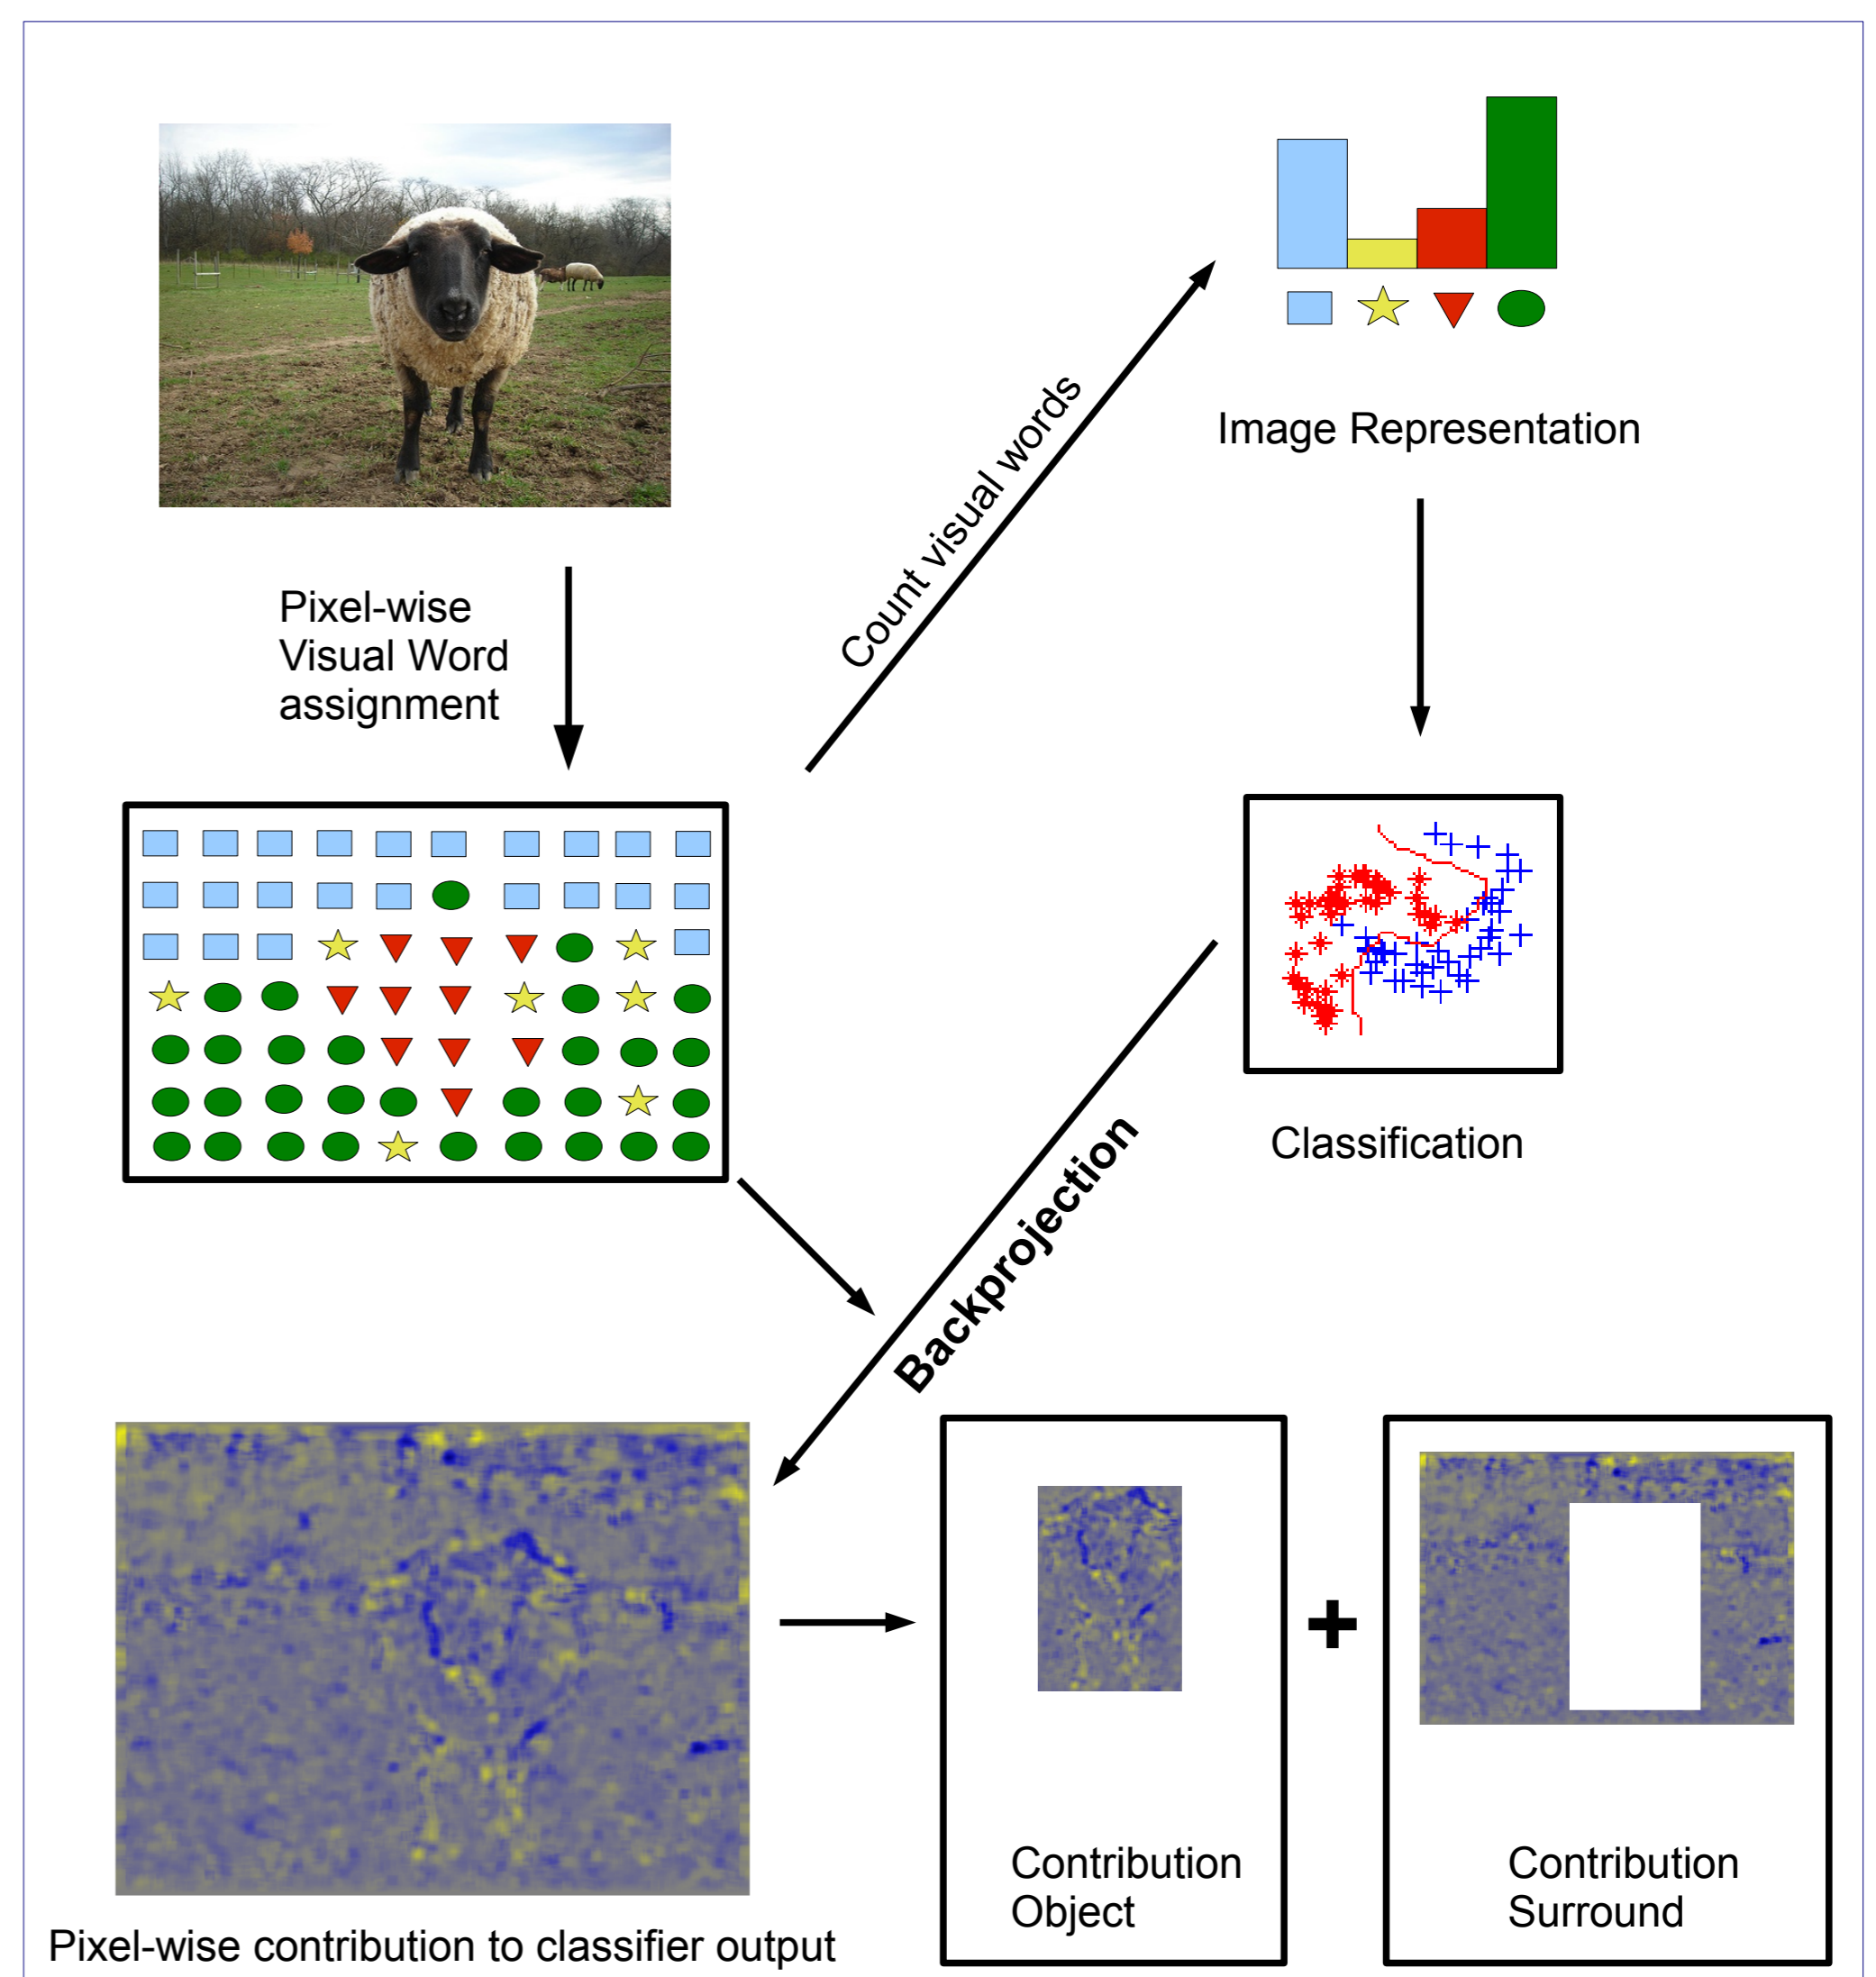


Abstract

- The Highly Successful Bag-of-Words technique forms the basis of all top image classification systems
- In this demo we use a Bag-of-Words implementation that stood at the basis of our highly competitive Pascal VOC 2008-2011 entries.
- To make its workings more insightful, our IJCV paper measures for each pixel how much it contributes to the classification of an object.
- We divide the contribution into Object and Surround.
- We visualise the per-pixel classification contribution.
- Our interactive demo enables exploration of the Pascal VOC 2010 validation dataset based on the Object/Surround contributions

Visualisation Framework



Bag-of-Words details

- Pixel-wise Sampling. Using dense acceleration (Uijlings 2010)
- 5 Colour SIFT (Lowe 2004, Sande 2010)
- Random Forest as vocabulary (Moosman 2006, Uijlings 2010)
- Normalize Codebook Vectors using square root (Peronnin 2010, Uijlings 2010)
- Support Vector Machine (libsvm)
- Histogram Intersection Kernel
- Fast Additive Classification (Maji 2009)
- Good Pascal 2007 performance: 0.585 MAP (e.g. Chatfield 2011, Perronnin 2010 report 0.593-0.617 MAP)

Screenshot



References

- K. Chatfield, V. Lempitsky, A. Vedaldi, A. Zisserman. *The devil is in the details: an evaluation of recent feature encoding methods*. BMVC, 2011.
- D. Lowe, *Distinctive image features from scale-invariant keypoints*. IJCV, 2004.
- S. Maji, A. Berg, J. Malik. *Classification using intersection kernel support vector machines is efficient*. CVPR, 2008.
- F. Moosmann, B. Triggs, and F. Jurie. *Fast discriminative visual codebooks using randomized clustering forests*. NIPS, 2006.
- F. Perronnin, J. Sanchez, and T. Mensink. *Improving the Fisher Kernel for Large-Scale Image Classification*. ECCV, 2010.
- J. Uijlings, A. Smeulders, R. Scha. *Real-Time Visual Concept Classification*, TMM, 2010.
- K. vd Sande, T. Gevers, C. Snoek. *Evaluating Color Descriptors for Object and Scene Recognition*, PAMI, 2010

This demo showcases: