

Image Retrieval by Cross-Media Relevance Fusion

Jianfeng Dong¹, Xirong Li², Shuai Liao²
Jieping Xu², Duanqing Xu¹, Xiaoyong Du²

¹Zhejiang University



²Renmin University of China



Task

Produce a relevance score for a given **image-query** pair.

Image




Query

new year
baby clip art

Key question: How to compute cross-media relevance?

Cross-media relevance computation

Image and query have to be represented in a common space

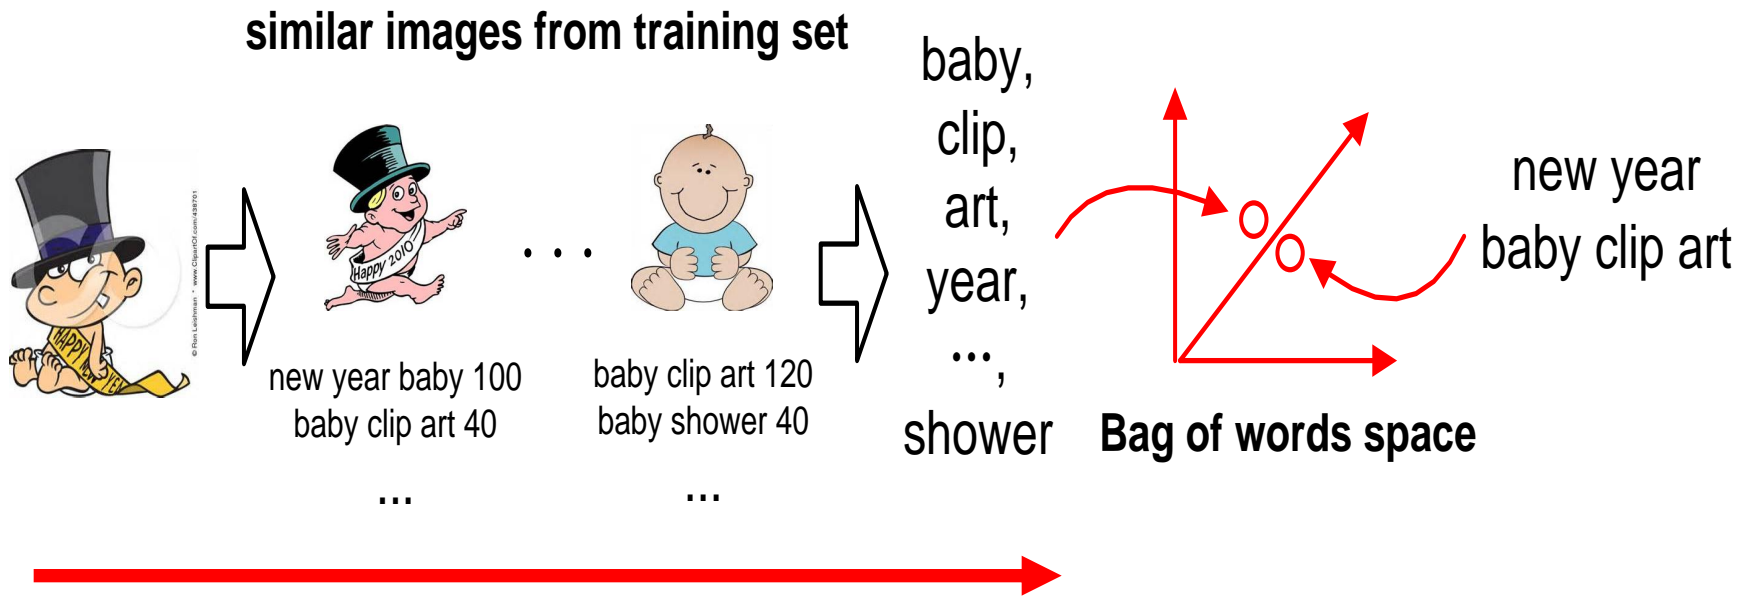


Choices of the common space:

1. Bag of words space
2. Visual feature space
3. Learned space

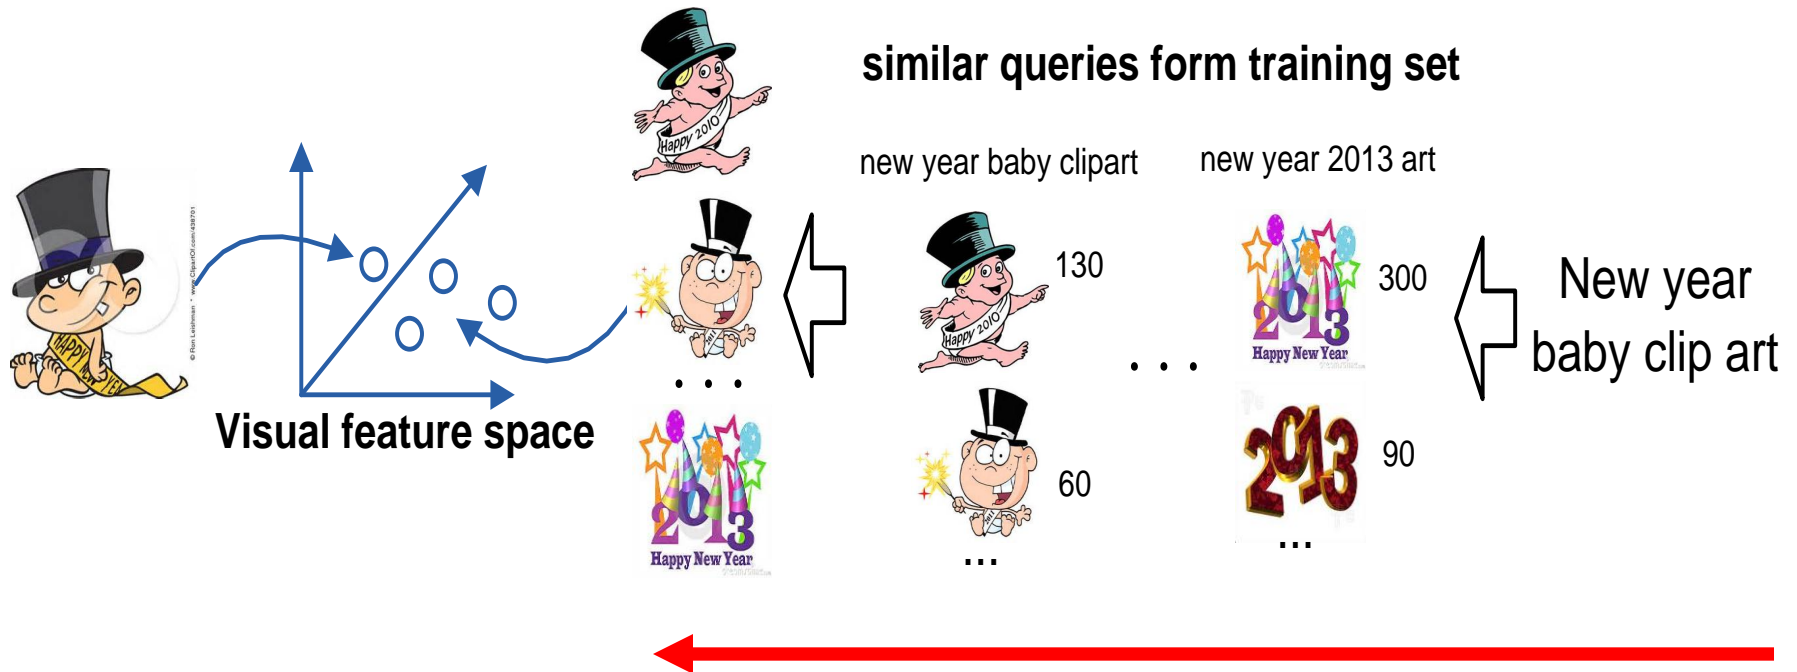
Cross-media 1: Image2Text

Compute relevance in a bag of words space.



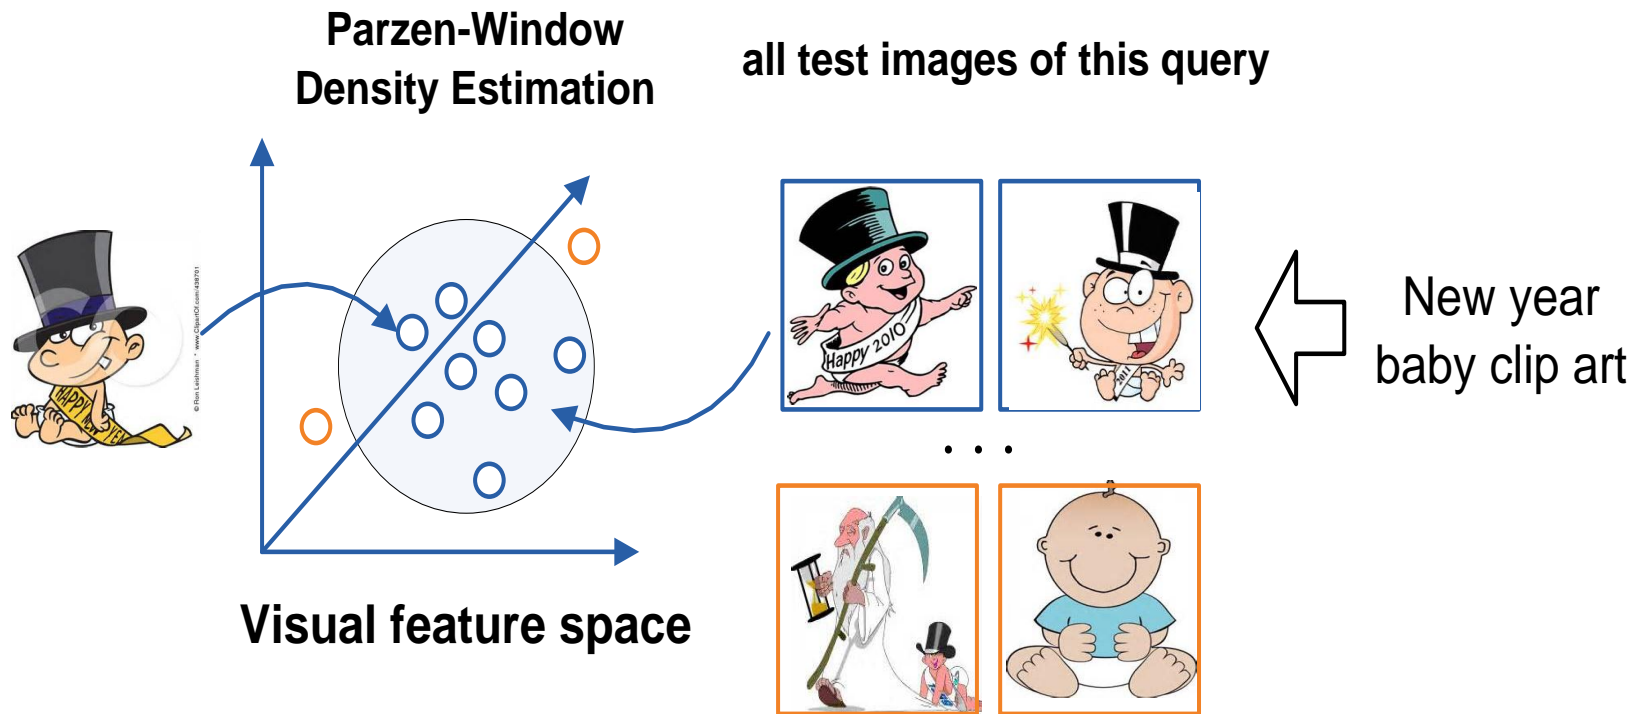
Cross-media 2: Text2Image

Compute relevance in a visual feature space.



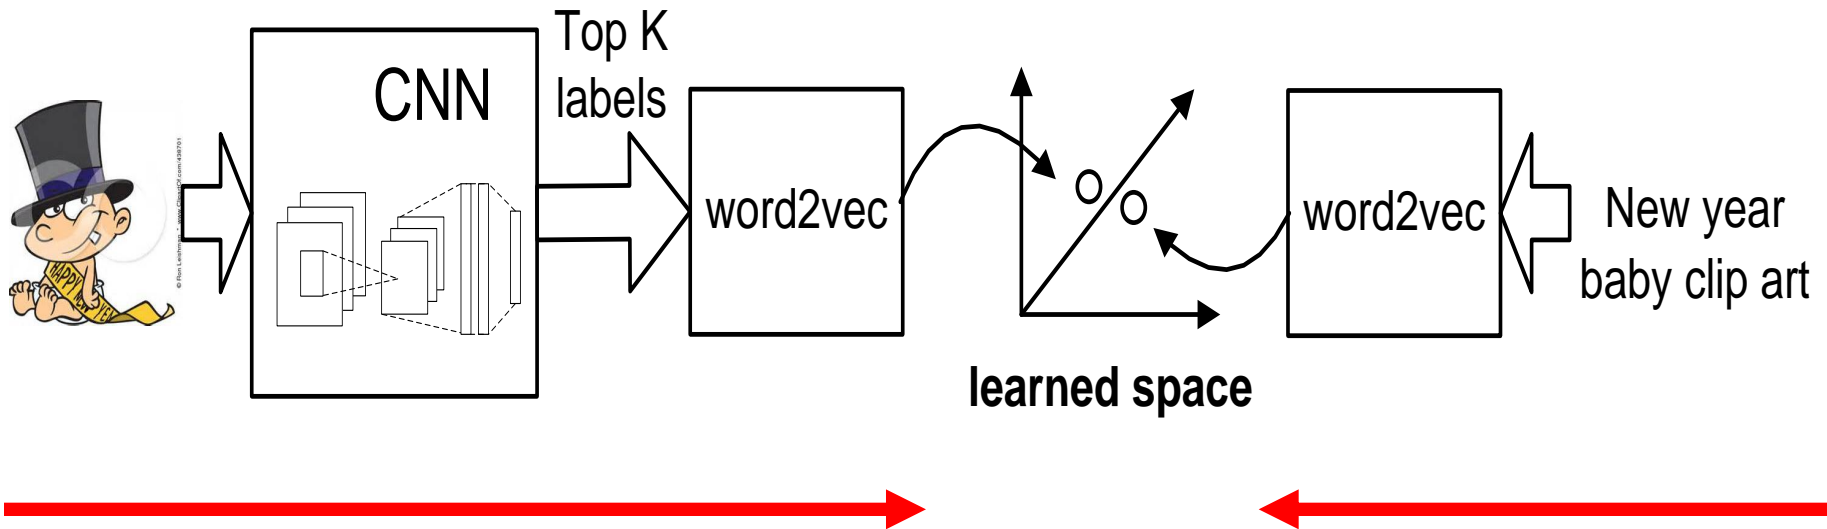
Cross-media 3: Parzen window

Hypothesis: relevant images are visually similar



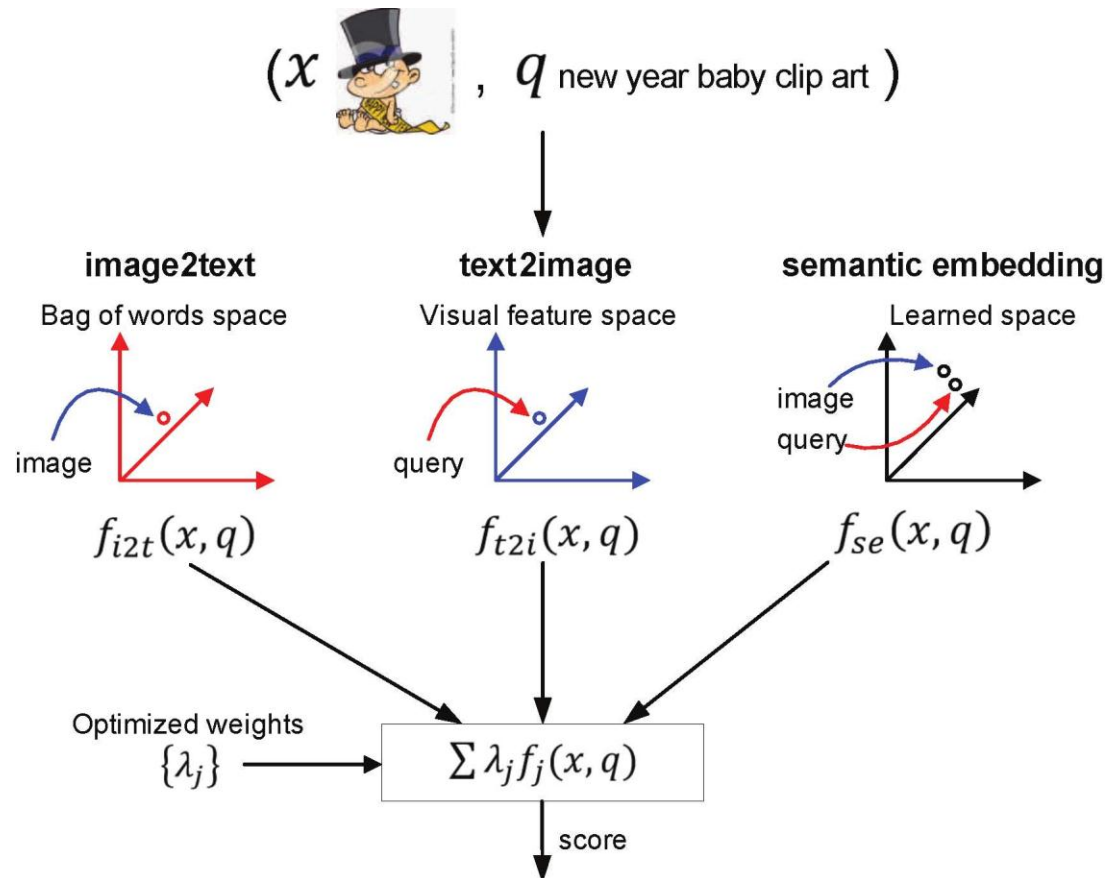
Cross-media 4: Semantic Embedding

Compute relevance in a learned semantic space.




Cross-media relevance fusion Framework

Fusion for better performance.



Official evaluation

Our runs are ranked at the top.



Take-home messages

- Text2Image (Parzen window) is a winning component.
- Relevance fusion gives the best performance.



<https://github.com/danieljf24/cmrf>