Medical Image Retrieval from 3D Lesion Content

Blaine Rister
Department of Electrical Engineering

Content-Based Image Retrieval

- **Image retrieval:** Rank images in a database in terms of similarity to a given query image
- **Content-based image retrieval (CBIR):** rank by image content, rather than metadata
- Can supplement metadata-based retrieval
- Useful when metadata is unavailable or unreliable

Query

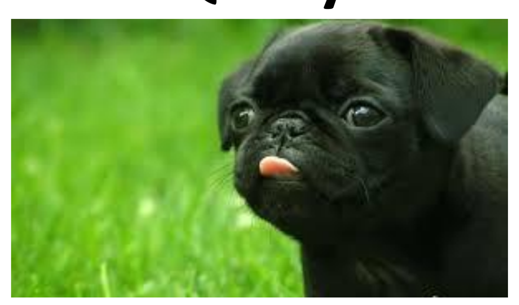

Results

Relevance

Liver Lesions

- Abnormal spots in liver tissue
- Usually indicate tumors, benign or malignant
- Data comes from abdominal computed tomography (CT) scans
  - 3D x-ray
- Lesions annotated by radiologists
  - Demarcation of boundaries
  - Relevance matrix gives learning supervision
    - $R_{ij} \in \{1,2,3\}$, 3 most relevant

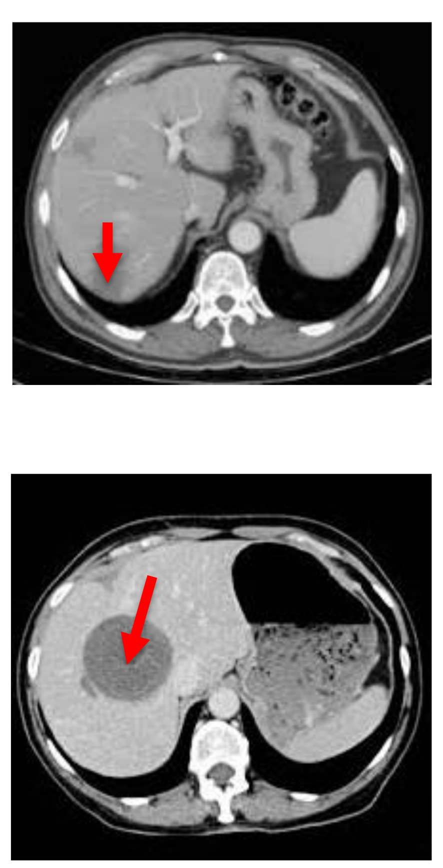

Examples of liver lesions in CT scans

3D Scale-Invariant Feature Transform

- Robust descriptor of local image content
  - Quantifies lesion shape
- Array of gradient histograms
  - Bin gradient vectors by orientation
  - Tessellate with regular icosahedron for isotropy
  - Interpolate between nearest bins by Barycentric coordinates
- Open-source implementation

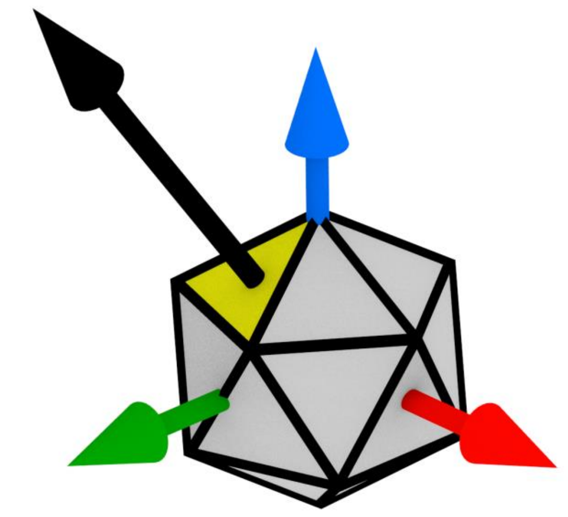

Gradient vector, shown in black, intersecting a histogram tile, show in yellow

Haralick Texture Features

- Quantifies image texture
- Originally developed for satellite image classification
  - Rock formations
  - City vs. farmland
- Compute gray-level co-occurrence matrix P
  - Quantize image into N intensity levels
  - Estimate P_{ij}, the probability of a pixel of intensity i neighboring one with probability j.
- Texture features are statistics computed from P
  - Contrast: $\sum_{ij} p_{ij}|i-j|^2$
  - Correlation: $\sum_{ij} p_{ij} \frac{(i-\mu_i)(j-\mu_j)}{\sigma_i \sigma_j}$
  - Energy: $\sum_{ij} {p_{ij}}^2$
  - Homogeneity: $\sum_{ij} \frac{p_{ij}}{1+|i-j|}$

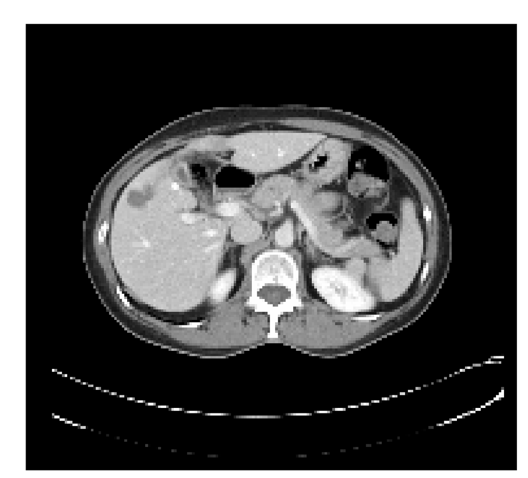

Contrast	Correlation	Energy	Homogeneity
.36	.92	.46	.97

Sample image (above) and texture features

Supervised learning

- Given query image features vector f_q, training image feature vectors $\{f^{(i)}\}$, predict relevance scores $\{r^{(i)}\}$
  - $r^{(i)}$ relevance of i^{th} training image to query
  - Known as the **pointwise** approach to ranking
- Regularized linear regression
  - Minimize $\sum_i \left\| a^T f^{(i)} - r^{(i)} \right\|_2^2 - \lambda \|a\|_2$
  - Convex optimization problem
  - Cross-validate parameter λ

Results and future work

- Dataset: 30 annotated lesions with 30×30 relevance matrix
- Search results quantified with discounted cumulative gain (DCG)
  - $DCG = \sum_i \frac{2^{r^{(i)}}-1}{log_2(i+1)}$
- Normalized DCG: $\text{NDCG} = \frac{DCG_{predicted}}{DCG_{truth}}$

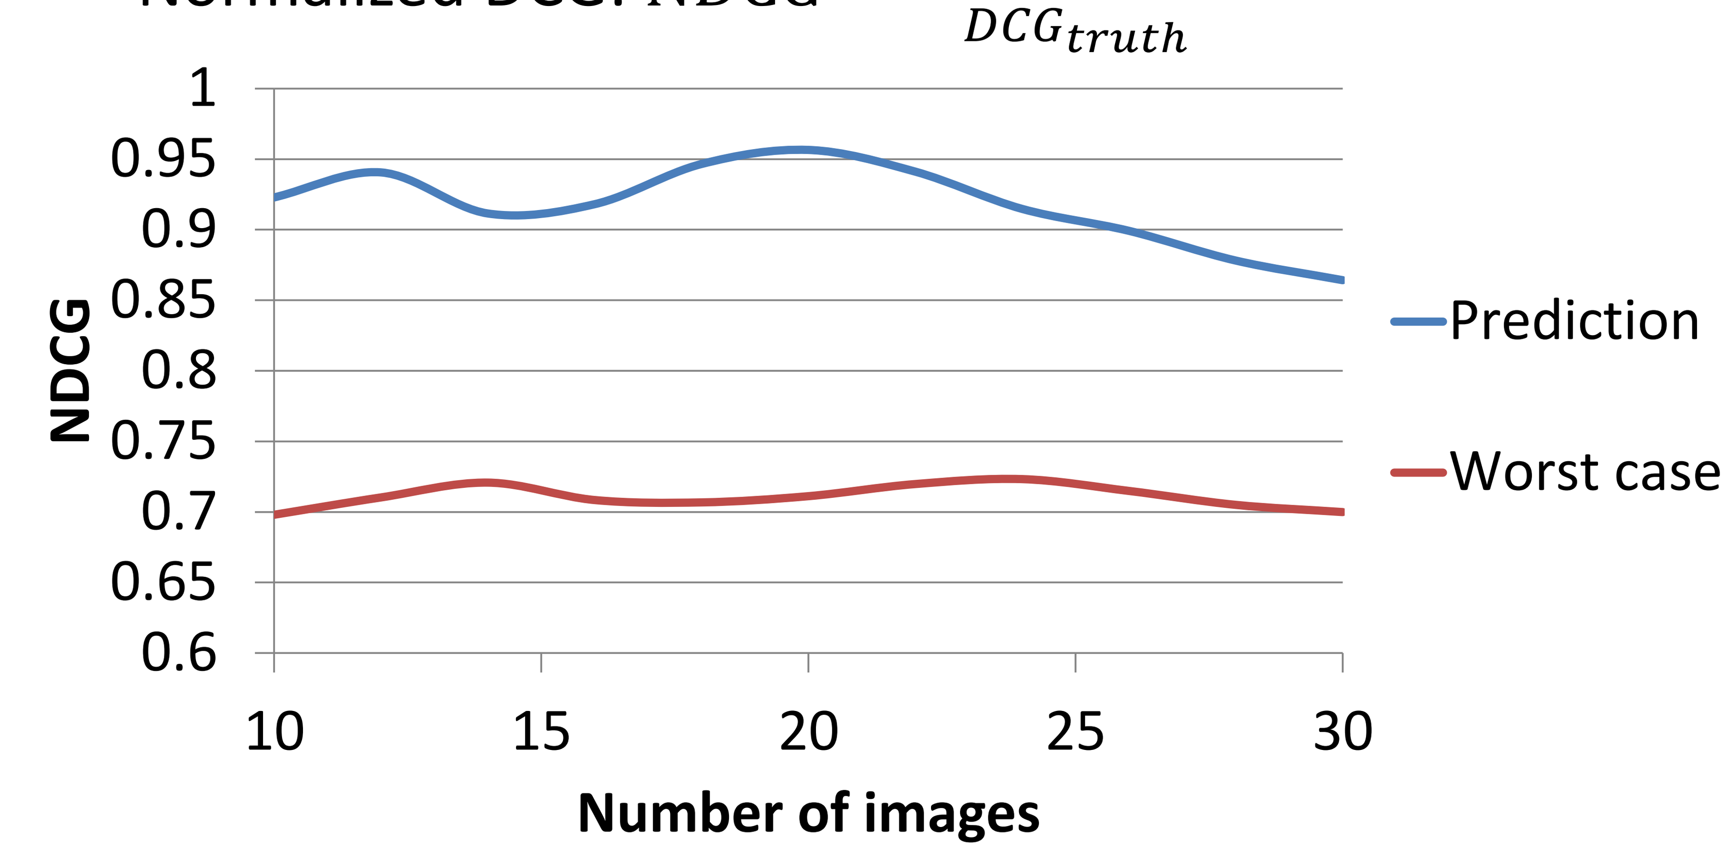

Cross-validated NDCG for various training set sizes. Predicted NDCG is computed from the predicted relevance scores. Worst-case NDCG is the worst possible on that data. The best possible NCDG is always 1.

- Results encouraging, leaving room for improvement
  - Rotation-invariant features
  - New regression algorithm for ranking