Overcoming the data challenge in medical image analysis

A new take on convolutional neural networks

aidence
human sense in artificial intelligence

Motivation

Lung cancer detection

- Lung cancer is the leading cause of cancer-related death worldwide
- *Screening* may help reduce mortality rate
- Software can assist the radiologist

Motivation

Data challenge

- CNNs methodology of choice for image analysis
- CNNs require a large amount of annotated data to learn from
- *Data-efficiency*: the ability to learn in complex domains without requiring large quantities of data
- G-CNNs reduce the sample complexity for lung nodule detection

Convolutional Neural Networks

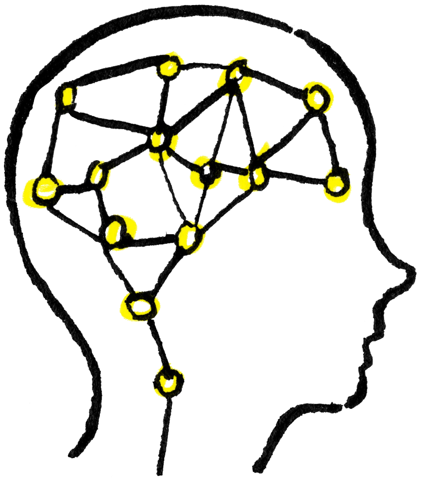

What are CNNs?

- Fully connected neural net not feasible for image analysis
- A *convolutional neural network* is a neural network with a *convolutional* layer
- Convolution: apply a set of weights (*filter* or *kernel)* at every position in the image to extract features
- *Feature maps* are then convolved with the next set of filters

Convolutional Neural Networks

Same dog, but translated.

Why are CNNs efficient?

- Same set of weights used at *every* position
- Therefore it doesn't matter where an object in an image is located
- This is called *translational equivariance*

$$T_g(f(\mathbf{x})) = f(T_g(\mathbf{x}))$$

- Leads to *translational invariance*

$$f(\mathbf{x}) = f(T_g(\mathbf{x}))$$

Group Convolutions

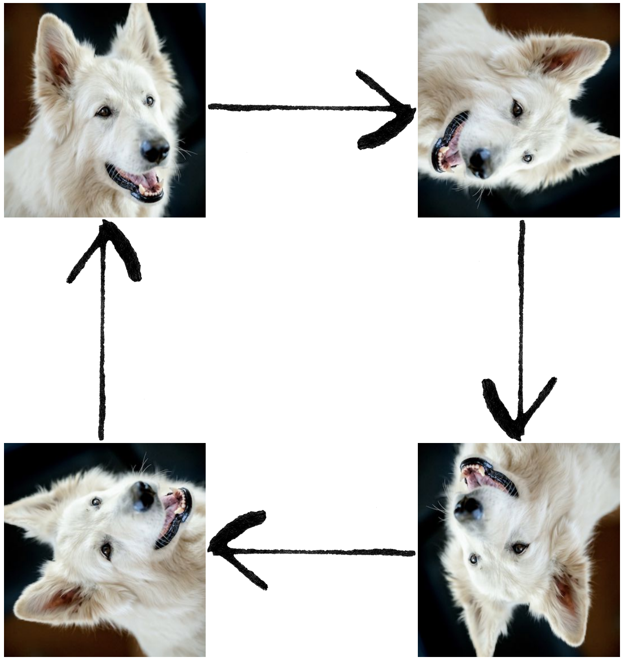

Same dog, but rotated!

Intuition

- *Translational* weight-sharing and equivariance make CNNs relatively data-efficient
- Does not currently work for *reflections* and *rotations*
- *G-CNN*: generalises equivariance to other groups of transformations

Group Convolutions

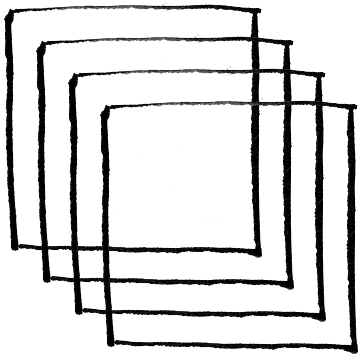

Implementation

- Create *augmented* filter bank: transform each filter by each $g \in G$
- Apply regular convolution on augmented filter bank
- Produces $|G|$ orientation channels per feature
- Equivariant to $g \in G$: $Tg(x)$ shuffles the orientation channels depending on the group structure

Transformation groups

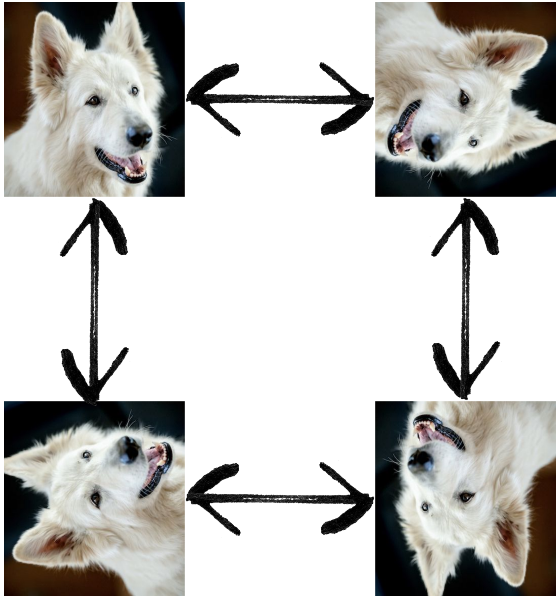

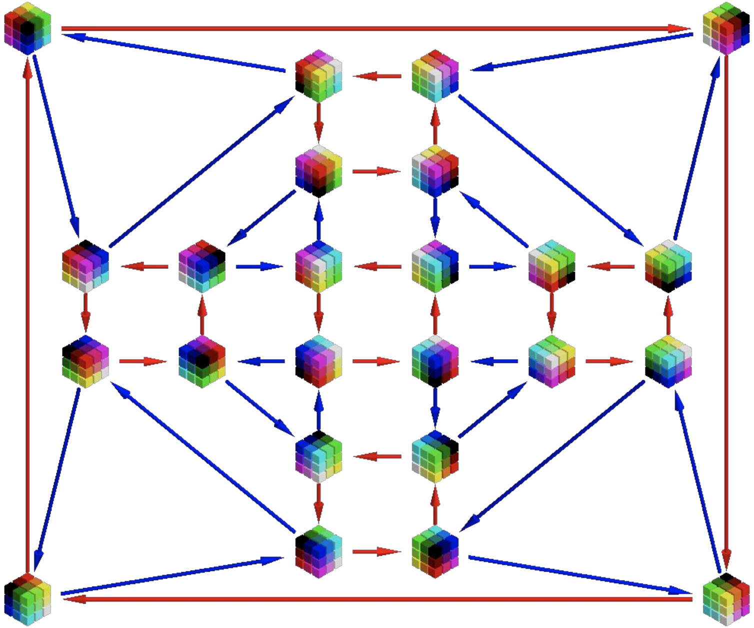

Results

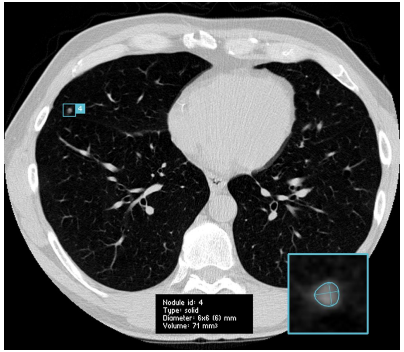

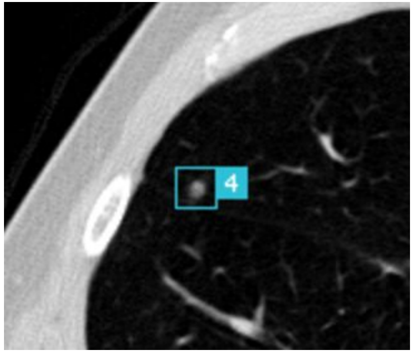

Lung nodule detection

- *Lung nodules*: suspect lesions in the lung that may be malignant
- Two-stage pipeline:
  - Candidate generation
  - False positive reduction
- Large training set available from NLST
- High-quality testset available from LIDC/IDRI

Results

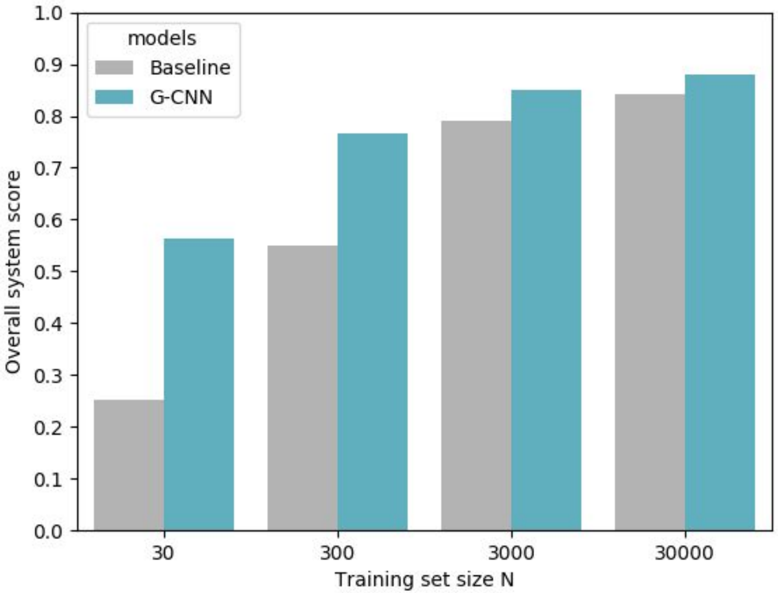

Experiment outline

- Baseline CNN vs. G-CNN
- Four training set sizes: 30, 300, 3.000 and 30.000 samples
- Evaluation with FROC analysis score: sensitivity vs. average false positives per scan

Conclusion

G-CNN roughly as good as the regular CNN trained on 10x the amount of data!

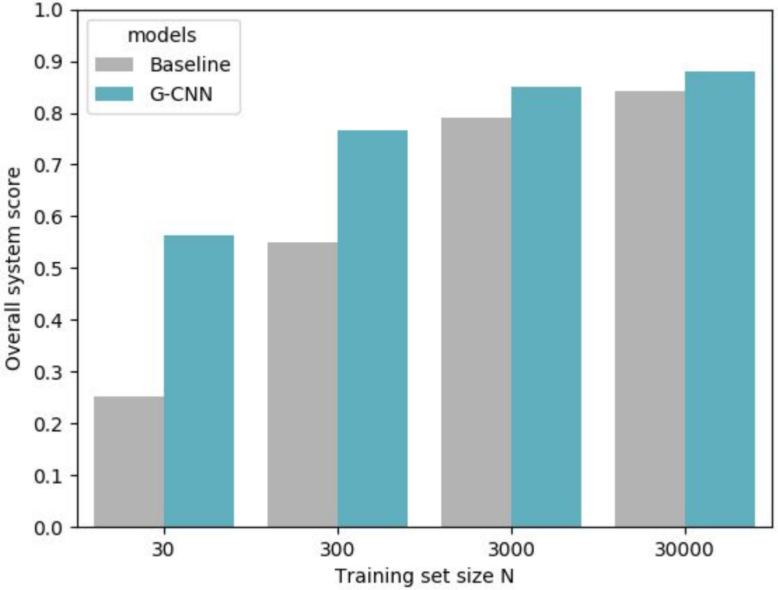

Additional observations

- G-CNNs way more sensitive to *malignant* nodules
- G-CNNs converge faster, reducing training time
- Applicable to any 3D volume, including CT and MRI
- Easy to use: simply replace your **convolution** with a **g-convolution**, and keep everything else the same!

Hints

- Data scarcity is an issue in the medical domain
- G-CNNs extend the weight-sharing properties of CNNs
- 3D G-CNNs proved beneficial for false positive reduction in lung nodule detection
- G-convolutions are easy to use

We're hiring!

Supplementary material

Results

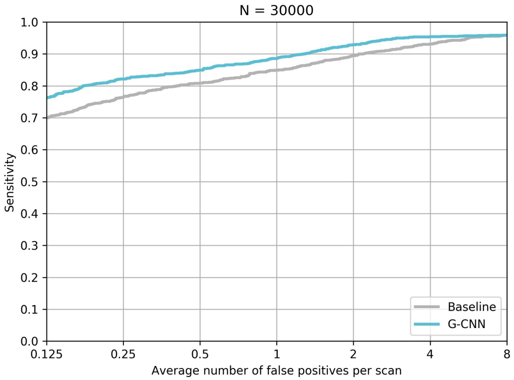

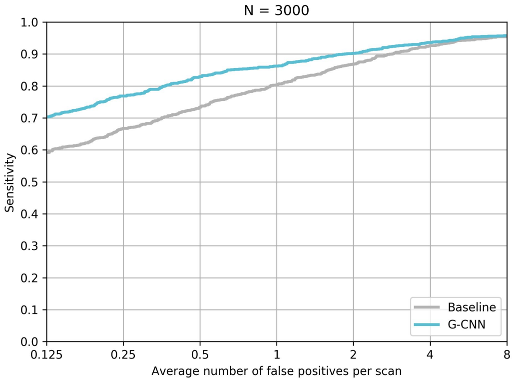

Experiment outline

- Baseline CNN vs. G-CNN
  - Competitive in LUNA16 challenge
  - Same training, validation and test data
  - Same data augmentation scheme
  - Same architecture and hyperparameters
  - Same number of parameters
- Four training set sizes: 30, 300, 3.000 and 30.000 samples
- Evaluation with FROC analysis: sensitivity vs. average false positives per scan

Results

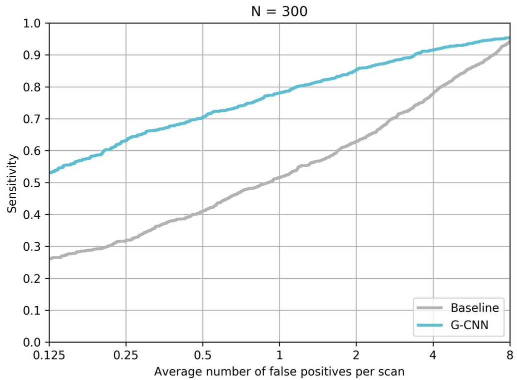

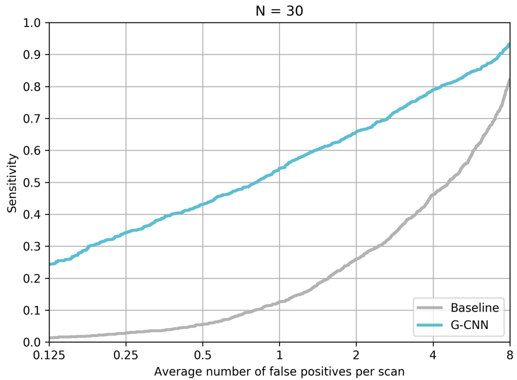

Experiment outline

- Baseline CNN vs. G-CNN
  - Competitive in LUNA16 challenge
  - Same training, validation and test data
  - Same data augmentation scheme
  - Same architecture and hyperparameters
  - Same number of parameters
- Four training set sizes: 30, 300, 3.000 and 30.000 samples
- Evaluation with FROC analysis: sensitivity vs. average false positives per scan

Group Convolutions

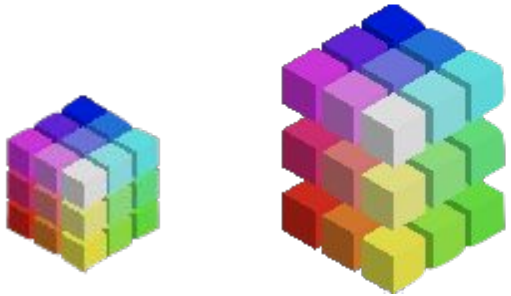

Groups in 3D

- CT and MRI are 3D, rather than 2D. Groups in 3D are highly more complex.
- Extension of the square to 3D is the *cube*
- Voxels in CT and MRI are *square prisms*
- In addition to *rotations*, we consider *reflections*
- This leads to four groups:
  - Cubic symmetry *with* and *without* reflections
  - Square prism symmetry *with* and *without* reflections