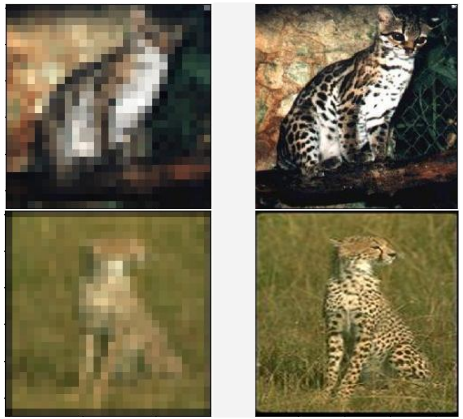
Very Low-Res Image Classification

Adam Vest

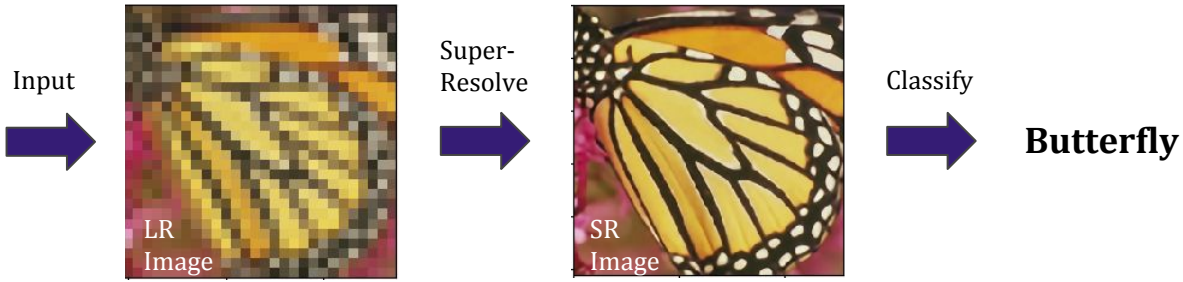
Problem Overview

- Low-res (LR) images contain limited information making recognition tasks such as classification more difficult
  - Evidenced by high classification error rates on LR datasets with a non-trivial number of classes
- Real-world classification systems need to be flexible and handle any LR images they may encounter without suffering serious performance degradation



Proposed Solution

- Use super-resolution to upscale images into the high-res domain before performing classification
- **Hypotheses:**
  - Through the use of super-resolution (SR), we may improve LR image classification performance
  - The performance of SR-based classification systems should improve as the quality of SR images improves, so deep ConvNet and GAN approaches should outperform BC
- **Goal:** to develop a resolution-agnostic image classification system that utilizes super-resolution to improve LR image classification performance

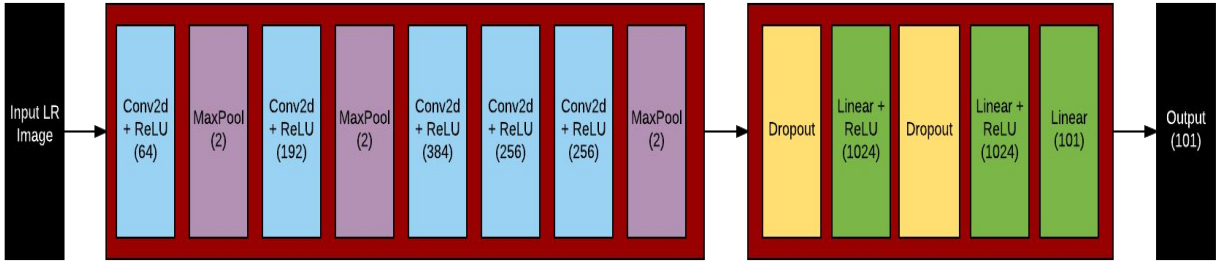


Experiment Setup

- Used Caltech101 dataset (50%, 25%, 25% train, validation, testing split)
- Preprocessed images by scaling to 256, center-cropping to 224x224, then downscaling to 56x56
- All models trained using early stopping on validation accuracy
- **Performed Experiments:**
  - Simple AlexNet-based LR image classifier
  - HR image classification using (1) AlexNet as a feature extractor and (2) fine-tuning AlexNet
  - Classified LR images super-resolved by Bicubic Interpolation using (1) pretrained HR feature extractor AlexNet and (2) pretrained HR fine-tuned AlexNet
  - Classified LR images super-resolved by SRResNet using (1) pretrained HR feature extractor AlexNet and (2) pretrained HR fine-tuned AlexNet

Classifier Details

- All models are based on the AlexNet architecture
- HR image classifiers were initialized with weights for AlexNet pretrained on ImageNet
- Pretrained layers were given a reduced learning rate while fine-tuning
- LR Classifier Architecture (Modified AlexNet):



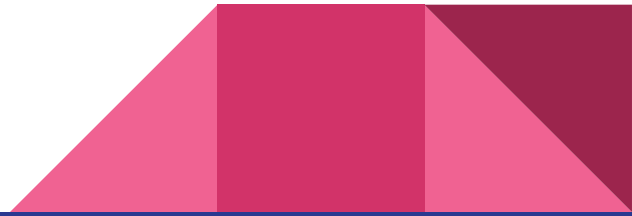
Results and Discussion

Model	Top-1 Accuracy	Top-5 Accuracy
Simple LR Classifier	50.07%	70.80%
HR Images w/ Feature Extractor AlexNet	79.91%	93.55%
HR Images w/ Fine-tuned AlexNet	85.68%	96.57%
BC SR Images w/ Pretrained HR Feature Extractor Model	54.14%	80.87%
BC SR Images w/ Pretrained Fine-tuned HR Model	67.60%	88.24%
SRResNet SR Images w/ Pretrained HR Feature Extractor Model	70.62%	88.38%
SRResNet SR Images w/ Pretrained Fine-tuned HR Model	78.63%	92.68%

- General Model Performance Rankings:
  - HR Classifiers
  - SR Classifiers
  - LR Classifiers
- Models utilizing SRResNet to super-resolve LR images outperform models using Bicubic Interpolation by a large margin
- Fine-tuning gives significant performance boosts

Current/Future Work

- Work on Super-Resolution for Classification, i.e., a method of fine-tuning the SRResNet to produce better classification results
  - Standard fine-tuning based on classification loss
  - LWF
  - Knowledge distillation
  - Adversarial examples inspired techniques
- Test higher upscaling factors (8x)



Thank you! Questions?