# Automated Surf Reports from Image Data

Stuart Cornuelle | Stanford University
stuartcc@stanford.edu

## Introduction

- Surfers rely on wave forecasting and reporting services to know where and when to find the best waves
- These services still rely heavily on a human expert to interpret and translate the relevant swell and weather data
- We explore a multitask learning approach to extract the two most salient features of current surfing conditions from surf cam image data alone; inputs are still frame images, outputs are (1) predicted wave height and (2) a predicted "conditions" label that encapsulates elements such as wave quality, wave shape and oceanic surface texture

## Methods

**Approach**

- Single-task CNN models are implemented for both Task A (wave height) and Task B (conditions) as benchmarks
- Multitask architecture uses hard parameter sharing in convolutional layers and separate branching networks of fully connected layers specific to Tasks A & B
- Multitask loss is non-weighted mean of softmax cross entropy loss for each task

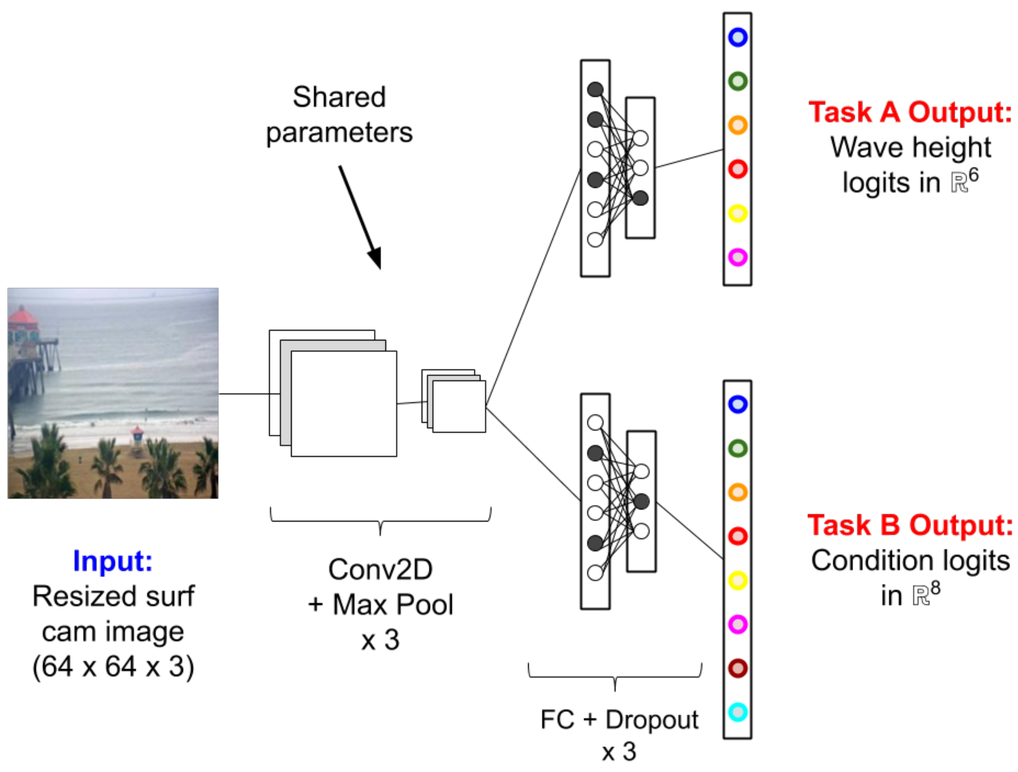


**(a)** Multitask learning architecture with hard parameter sharing.

## Data + Features

- 13,548 surf cam still frame images
- 365 days' surf reports $\times$ 3 spots
- Wave height (Task A) $\in$ 0-25 ft. binned to $\{0, 1, \ldots, 5\}$
- Conditions (Task B) e.g. "good" mapped to $\{0, 1, \ldots, 7\}$

## Performance

| Model | Test Acc. (%) |
|---|---|
| Task A (Single) | 71.7 |
| Task A (Multi) | 66.1 |
| Task B (Single) | 49.8 |
| Task B (Multi) | 58.7 |
| Task A & B (Single) | 60.1 |
| Task A & B (Multi) | 62.3 |

## Select References

[1] Sebastian Ruder. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*, 2017.

[2] Rich Caruana. Multitask learning. *Machine learning*, 28(1):41–75, 1997.

[3] Olivier Moindrot and Guillaume Genthial. Hand signs recognition with tensorflow., 2018.

## Results

Single-task model outperforms for Task A; multitask outperforms for B and the aggregated performance (avg.) of A & B:

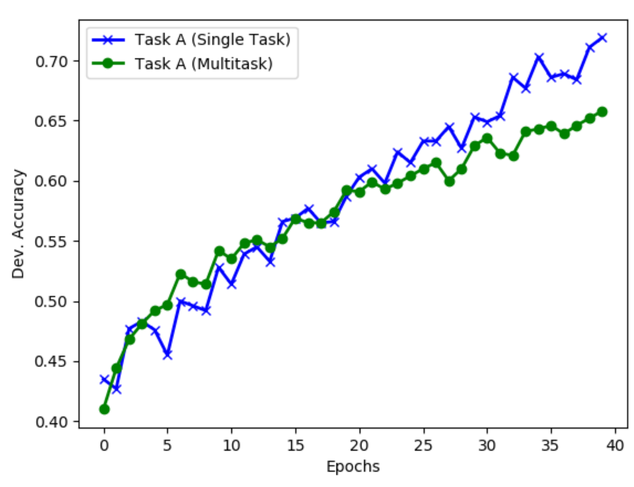


**(b)** Task A

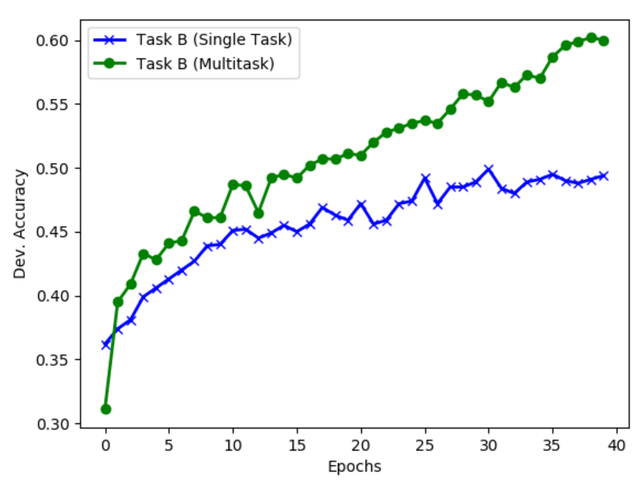


**(c)** Task B

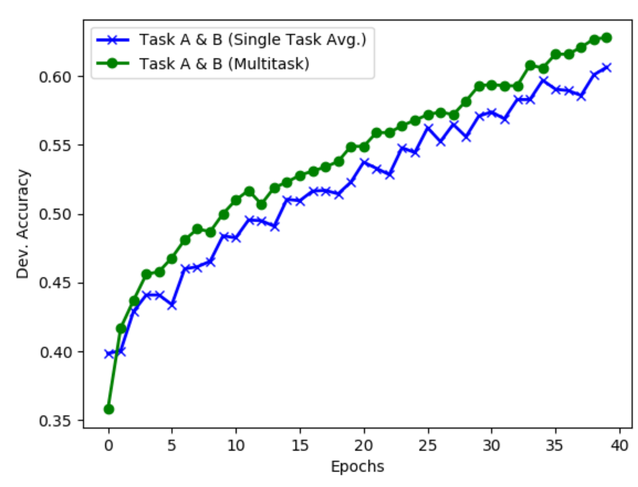


**(d)** Tasks A & B

## Conclusions

Tasks A & B are related (i.e. rely on similar features of image inputs), and so we expect and find that the multitask regime improves generalization over two single-task models in isolation. Per [1]; this is likely because helps learning focus on truly salient features, implicitly regularizes by limiting overfitting to the noise of a single task; and allows one task to benefit from ("eavesdrop on") the learning achieved by the other task in their shared layers.