# Deep Declarative Networks: A New Hope 

A/Prof. Stephen Gould

Research School of Computer Science


The Australian National University
2019


## What did we gain?

$\checkmark$ Better-than-human performance on closed-world classification tasks
$\checkmark$ Very fast inference (with the help of GPU acceleration)
$\checkmark$ versus very slow iterative optimization procedures
$\checkmark$ Common tools and software frameworks for sharing research code
$\checkmark$ Robustness to variations in realworld data if training set is sufficiently large and diverse

## What did we lose?

x Clear mathematical models; separation between algorithm and objective (loss function)
$\times$ Theoretical performance guarantees
x Interpretability and robustness to adversarial attacks
$x$ Ability to enforce hard constraints
$x$ Intuition guided by physical models
$\times$ Parsimony - capacity consumed learning what we already know

What if we could have the best of both worlds?

## Deep learning models



- Linear transforms (i.e., convolutions)
- Elementwise non-linear transforms
- Spatial/global pooling



## Deep learning layer

$x$ : input
$y$ : output
$\theta$ : local parameters
$f$ : forward function


## End-to-end computation graph



## End-to-end learning

- Learning is about finding parameters that maximize performance,


## $\operatorname{argmax}_{\theta} \quad$ performance $(\operatorname{model}(\theta))$

- To do so we need to understand how the model output changes as a function of its input and parameters
- (Local based learning) incrementally updates parameters based on a signal back-propagated from the output of the network
- This requires calculation of gradients

$$
\frac{d J}{d x}=\frac{d J}{d y} \frac{d y}{d x} \text { and } \frac{d J}{d \theta}=\frac{d J}{d y} \frac{d y}{d \theta}
$$



## Example: Back-propagation through a node

Consider the following implementation of a node

```
fwd_fcn(x)
    y0}=\frac{1}{2}
    for t = 1, ..., T do
        \mp@subsup{y}{t}{}\leftarrow\frac{1}{2}(\mp@subsup{y}{t-1}{}+\frac{x}{\mp@subsup{y}{t-1}{}})
    return }\mp@subsup{y}{T}{
```

We can back-propagate gradients as

$$
\begin{gathered}
\frac{\partial y_{t}}{\partial y_{t-1}}=\frac{1}{2}\left(1-\frac{x}{y_{t-1}^{2}}\right) \\
\frac{\partial y_{t}}{\partial x}=\frac{1}{2 y_{t-1}}+\frac{\partial y_{t}}{\partial y_{t-1}} \frac{\partial y_{t-1}}{\partial x}
\end{gathered}
$$

It turns out that this node implements the Babylonian algorithm, which computes

$$
y=\sqrt{x}
$$

As such we can compute its derivative directly as

$$
\begin{aligned}
\frac{\partial y}{\partial x} & =\frac{1}{2 \sqrt{x}} \\
& =\frac{1}{2 y}
\end{aligned}
$$

```
bck_fcn(x, y)
    return }\frac{1}{2y
```

Chain rule gives $\frac{\partial J}{\partial x}$ from $\frac{\partial J}{\partial y}$ (input) and $\frac{\partial y}{\partial x}$ (computed)

## Separate of forward and backward operations



## Deep declarative networks (DDNs)



In an imperative node the implementation of the forward processing function $\tilde{f}$ is explicitly defined. The output is then

$$
y=\tilde{f}(x ; \theta)
$$

where $x$ is the input and $\theta$ are the parameters of the node.

In a declarative node the inputoutput relationship is specified as the solution to an optimization problem

$$
y \in \operatorname{argmin}_{u \in C} f(x, u ; \theta)
$$ where $f$ is the objective and $C$ are the constraints.



## Imperative vs. declarative node example: global average pooling

$$
\left\{x_{i} \in \mathbb{R}^{m} \mid i=1, \ldots, n\right\} \rightarrow \mathbb{R}^{m}
$$

Imperative specification:

## Declarative specification:

$$
y=\frac{1}{n} \sum_{i=1}^{n} x_{i}
$$

$$
y=\underset{u \in \mathbb{R}^{m}}{\operatorname{argmin}} \sum_{i=1}^{n}\left\|u-x_{i}\right\|^{2}
$$

"the vector $u$ that is the minimum distance to all input vectors $x_{i}{ }^{\prime \prime}$

## Deep declarative nodes: special cases



| Unconstrained <br> (e.g., robust pooling) | $y(x) \in \operatorname{argmin}_{u \in \mathbb{R}^{m} f(x, u)}$ |
| :---: | :---: |
| Equality Constrained (e.g., projection onto $L_{p}$-sphere) | $y(x) \in\left\{\begin{array}{c}\operatorname{argmin}_{u \in \mathbb{R}^{m}} f(x, u) \\ \text { subject to } h(x, u)=0\end{array}\right\}$ |
| Inequality Constrained <br> (e.g., projection onto $L_{p}$-ball) | $y(x) \in\left\{\begin{array}{c}\operatorname{argmin}_{u \in \mathbb{R}^{m}} f(x, u) \\ \operatorname{subject~to~} h(x, u) \leq 0\end{array}\right\}$ |

## Imperative and declarative nodes can co-exist



## Learning as bi-level optimization

learning problem

| minimize (over $\boldsymbol{x})$ | objective $(\boldsymbol{x})$ |
| :--- | :--- |
| subject to | constraints $(\boldsymbol{x})$ |

bi-level learning problem
minimize (over $\boldsymbol{x}$ ) objective $(\boldsymbol{x}, \boldsymbol{y})$
subject to constraints $(\boldsymbol{x})$
declarative node problem
minimize (over $\boldsymbol{y}$ ) objective $(\boldsymbol{x}, \boldsymbol{y})$
subject to constraints $(\boldsymbol{y})$

## A game theoretic perspective

- Consider two players, a leader and a follower
- The market dictates the price its willing to pay for some goods based on supply, i.e., quantity produced by both players, $P\left(q_{1}+q_{2}\right)$
- Each player has a cost structure associated with producing goods, $C_{i}\left(q_{i}\right)$ and wants to maximize profits, $q_{i} P\left(q_{1}+q_{2}\right)-C_{i}\left(q_{i}\right)$
- The leader picks a quantity of goods to produce knowing that the follower will respond optimally. In other words, the leader solves

$$
\begin{array}{cc}
\operatorname{maximize}_{q_{1}} & q_{1} P\left(q_{1}+q_{2}\right)-C_{1}\left(q_{1}\right) \\
\text { subject to } & q_{2} \in \operatorname{argmax}_{q} q P\left(q_{1}+q\right)-C_{2}(q)
\end{array}
$$



## Solving bi-level optimization problems

$$
\begin{array}{cc}
\operatorname{minimize}_{x} & J(x, y) \\
\text { subject to } & y \in \operatorname{argmin}_{u} f(x, u)
\end{array}
$$

- Closed-form lower-level problem: substitute for $\boldsymbol{y}$ in upper problem


## $\operatorname{minimize}_{x} J(x, y(x))$

- May result in a difficult (single-level) optimization problem


## Solving bi-level optimization problems

$$
\begin{array}{cc}
\operatorname{minimize}_{x} & J(x, y) \\
\text { subject to } & y \in \operatorname{argmin}_{u} f(x, u)
\end{array}
$$

- Convex lower-level problem: replace lower problem with sufficient conditions (e.g., KKT conditions)
$\begin{array}{cc}\text { minimize }_{x, y} & J(x, y) \\ \text { subject to } & h(y)=0\end{array}$
- May result in non-convex problem if KKT conditions are not convex


## Solving bi-level optimization problems

$$
\begin{array}{cc}
\operatorname{minimize}_{x} & J(x, y) \\
\text { subject to } & y \in \operatorname{argmin}_{u} f(x, u)
\end{array}
$$

- Gradient descent: compute gradient with respect to $\boldsymbol{x}$

$$
x \leftarrow x-\eta\left(\frac{\partial J(x, y)}{\partial x}+\frac{\partial J(x, y)}{\partial y} \frac{d y}{d x}\right)
$$

- But this requires computing the gradient of $\boldsymbol{y}$ (itself a function of $\boldsymbol{x}$ )


## Algorithm for solving bi-level optimization

## SolveBiLevelOptimization:

initialize $x$
repeat until convergence:
solve $y \in \operatorname{argmin}_{u} f(x, u)$
compute $J(x, y)$
compute $\frac{d J}{d x}=\frac{\partial J(x, y)}{\partial x}+\frac{\partial J(x, y)}{\partial y} \frac{d y}{d x}$
update $x \leftarrow x-\eta \frac{d J}{d x}$
return $x$


How do we compute $\frac{d}{d x} \operatorname{argmin}_{u \in C} f(x, u)$ ?

## Implicit differentiation

Let $f: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ be a twice differentiable function and let

$$
y(x)=\operatorname{argmin}_{u} f(x, u)
$$

The derivative of $f$ vanishes at $(x, y)$. By Dini's implicit function theorem (1878)

$$
\frac{d y(x)}{d x}=-\left(\frac{\partial^{2} f}{\partial y^{2}}\right)^{-1} \frac{\partial^{2} f}{\partial x \partial y}
$$

The result extends to vector-valued functions, vector-argument functions and (equality) constrained problems. See [Gould et al., 2019].

## Proof sketch


$y \in \operatorname{argmin}_{u} f(x, u) \Rightarrow \frac{\partial f(x, y)}{\partial y}=0$

LHS: $\quad \frac{d}{d x} \frac{\partial f(x, y)}{\partial y}=\frac{\partial^{2} f(x, y)}{\partial x \partial y}+\frac{\partial^{2} f(x, y)}{\partial y^{2}} \frac{d y}{d x}$

RHS: $\quad \frac{d}{d x} 0=0$

Rearranging gives $\frac{d y}{d x}=-\left(\frac{\partial^{2} f}{\partial y^{2}}\right)^{-1} \frac{\partial^{2} f}{\partial x \partial y}$.

## Deep declarative nodes: what do we need?



- Forward pass
- A method to solve the optimization problem
- Backward pass
- Specification of objective and constraints
- (And cached result from the forward pass)
- Do not need to know how the problem was solved



## examples

## Global average pooling

$$
\left\{x_{i} \in \mathbb{R}^{m} \mid i=1, \ldots, n\right\} \rightarrow \mathbb{R}^{m}
$$



## Robust penalty functions, $\phi$

| Quadratic | Pseudo-Huber | Huber | Welsch | Truncated Quad. |
| :---: | :---: | :---: | :---: | :---: |
| $\frac{1}{2} Z^{2}$ | $\sqrt{1+\left(\frac{z}{\alpha}\right)^{2}}-1$ | $\left\{\begin{array}{c} \frac{1}{2} z^{2} \text { for }\|z\| \leq \alpha \\ \text { else } \alpha\left(\|z\|-\frac{1}{2} \alpha\right) \end{array}\right.$ | $1-\exp \left(\frac{-z^{2}}{2 \alpha^{2}}\right)$ | $\begin{cases}\frac{1}{2} z^{2} & \text { for }\|z\| \leq \alpha \\ \frac{1}{2} \alpha^{2} & \text { otherwise }\end{cases}$ |
|  |  |  |  |  |
| closed-form, convex, smooth, unique solution | convex, smooth, unique solution | convex, non-smooth, non-isolated solutions | non-convex, smooth, isolated solutions | non-convex, non-smooth, isolated solutions |

## Example: robust pooling



| minimize (over $x)$ | $J(x, y) \triangleq \frac{1}{2}\\|y\\|^{2}$ |
| :---: | :---: |
| subject to | $y \in \operatorname{argmin}_{u} \sum_{i=1}^{n} \phi\left(u-x_{i} ; \alpha\right)$ |



## Example: Euclidean projection


closed-form, smooth, unique solution*


## Example: quadratic programs

$$
\begin{array}{cc}
\operatorname{argmin}_{u \in \mathbb{R}^{m}} & \frac{1}{2} u^{T} P u+q^{T} u+r \\
\text { subject to } & A u=b \\
G u \leq h
\end{array}
$$

Can be differentiated with respect to its parameters:

$$
P \in \mathbb{R}^{m \times m}, \quad q \in \mathbb{R}^{m}, \quad A \in \mathbb{R}^{n \times m}, \quad b \in \mathbb{R}^{n}
$$

## Example: convex programs

$$
\begin{array}{cc}
\operatorname{argmin}_{u \in \mathbb{R}^{m}} & c^{T} u \\
\text { subject to } & b-A u \in K
\end{array}
$$

Can be differentiated with respect to its parameters:

$$
A \in \mathbb{R}^{n \times m}, \quad b \in \mathbb{R}^{n}, \quad c \in \mathbb{R}^{m}
$$

## Implementing deep declarative nodes

- Need: objective and constraint functions, solver to obtain $y$
- Gradient by automatic differentiation

$$
\frac{d y(x)}{d x}=-\left(\frac{\partial^{2} f}{\partial y^{2}}\right)^{-1} \frac{\partial^{2} f}{\partial x \partial y}
$$

```
import autograd.numpy as np
from autograd import grad, jacobian
def gradient(x, Y, f)
    fY = grad(f, 1)
    fYY = jacobian(fY, 1)
    fXY = jacobian(fY, 0)
    return -1.0 * np.linalg.solve(fYY(x,y), fXY(x,y))
```


## cvxpylayers

- Disciplined convex optimization
- Subset of optimization problems

- Write problem using cvx
- Solver and gradient computed automatically!

```
x = cp. Parameter(n)
y = cp. Variable(n)
obj = cp. Minimize(cp.sum_squares(x - y ))
cons = [ y >= 0]
prob = cp. Problem(obj, cons)
layer = CvxpyLayer(prob, parameters=[x], variables=[y])
```


## applications

## Robust point cloud classification


[Gould et al., 2019]

## Robust point cloud classification







| O | Top-1 Accuracy \% |  |  |  |  |  | Mean Average Precision $\times 100$ |  |  |  |  |  |
| ---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\%$ | $[34]$ | Q | PH | H | W | TQ | $[34]$ | Q | PH | H | W | TQ |
| 0 | 88.4 | 84.7 | 84.7 | 86.3 | 86.1 | 85.4 | 95.6 | 93.8 | 95.0 | 95.4 | 95.0 | 93.8 |
| 10 | 79.4 | 84.3 | 85.6 | 85.5 | 86.6 | 85.5 | 89.4 | 94.3 | 94.6 | 95.1 | 94.6 | 94.7 |
| 20 | 76.2 | 84.8 | 84.8 | 85.2 | 86.3 | 85.5 | 87.8 | 94.8 | 95.0 | 95.0 | 94.8 | 95.0 |
| 50 | 72.0 | 84.0 | 83.1 | 83.9 | 84.3 | 83.9 | 83.3 | 93.8 | 93.5 | 94.3 | 94.8 | 94.8 |
| 90 | 29.7 | 61.7 | 63.4 | 63.1 | 65.3 | 61.8 | 38.9 | 76.8 | 78.7 | 78.5 | 79.1 | 76.6 |


| O | Top-1 Accuracy \% |  |  |  |  | Mean Average Precision $\times 100$ |  |  |  |  |  |  |
| ---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\%$ | $[34]$ | Q | PH | H | W | TQ | $[34]$ | Q | PH | H | W | TQ |
| 0 | 88.4 | 84.7 | 84.7 | 86.3 | 86.1 | 85.4 | 95.6 | 93.8 | 95.0 | 95.4 | 95.0 | 93.8 |
| 1 | 32.6 | 84.9 | 84.7 | 86.4 | 86.2 | 85.3 | 48.6 | 93.8 | 95.1 | 95.3 | 95.1 | 93.0 |
| 10 | 6.47 | 83.9 | 84.6 | 85.3 | 86.0 | 85.9 | 8.20 | 93.4 | 94.8 | 94.4 | 94.9 | 93.9 |
| 20 | 5.95 | 79.6 | 82.8 | 81.1 | 84.7 | 84.9 | 7.73 | 91.9 | 93.4 | 92.7 | 94.2 | 94.6 |
| 30 | 5.55 | 70.9 | 74.2 | 72.2 | 77.6 | 83.2 | 6.00 | 87.8 | 89.5 | 85.1 | 90.9 | 92.8 |
| 40 | 5.35 | 55.3 | 59.1 | 55.4 | 63.1 | 75.6 | 6.41 | 77.6 | 80.2 | 72.7 | 83.2 | 90.6 |
| 50 | 4.86 | 32.9 | 36.0 | 34.6 | 44.1 | 57.9 | 5.68 | 62.3 | 60.2 | 60.1 | 66.4 | 85.3 |
| 60 | 4.42 | 14.5 | 16.2 | 18.1 | 27.1 | 30.6 | 5.08 | 39.1 | 36.3 | 38.5 | 42.7 | 68.5 |
| 70 | 4.25 | 5.03 | 6.33 | 7.95 | $\mathbf{1 4 . 1}$ | 11.9 | 4.66 | 22.5 | 19.3 | 18.4 | 25.7 | 47.9 |
| 80 | 3.11 | 4.10 | 4.51 | 5.64 | 8.88 | 5.11 | 4.21 | 10.8 | 8.91 | 8.98 | 14.9 | 26.7 |
| 90 | 3.72 | 4.06 | 4.06 | 4.30 | 5.68 | 4.22 | 4.49 | 8.20 | 5.98 | 5.80 | 8.37 | 9.78 |

## Video activity recognition

## Stand Up



## Sit Down



[Fernando and Gould, 2016]


## Video clip classification pipeline



## Temporal pooling

- Max/avg/robust pooling summarizes an unstructured set of objects

$$
\left\{x_{i} \mid i=1, \ldots, n\right\} \rightarrow \mathbb{R}^{m}
$$

- Rank pooling summarizes a structured sequence of objects

$$
\left\langle x_{i} \mid i=1, \ldots, n\right\rangle \rightarrow \mathbb{R}^{m}
$$

## Rank Pooling

- Find a ranking function $r: \mathbb{R}^{n} \rightarrow \mathbb{R}$ such that $r\left(x_{t}\right)<r\left(x_{s}\right)$ for $t<s$
- In our case we assume that $r: x \mapsto u^{T} x$ is a linear function
- Use $u$ as the representation



## Experimental results

| Method | Accuracy (\%) |
| :--- | :---: |
| Max-Pool + SVM | 66 |
| Avg-Pool + SVM | 67 |
| Rank-Pool + SVM | 66 |
| Max-Pool-CNN (end-to-end) | 71 |
| Avg-Pool-CNN (end-to-end) | 70 |
| Rank-Pool-CNN (end-to-end) | 87 |
| Improved trajectory features + <br> fisher vectors + rank-pooling | $\mathbf{8 7}$ |

## Visual attribute ranking

1. Order a collection of images according to a given attribute
2. Recover the original image from shuffled image patches


## Birkhoff polytope

- Permutation matrices form discrete points in Euclidean space which imposes difficulties for gradient based optimizers
- The Birkhoff polytope is the convex hull for the
 set of $n \times n$ permutation matrices
- This coincides exactly with the set of $n \times n$ doubly stochastic matrices
- We relax our visual permutation learning problem over permutation matrices to a problem over doubly stochastic matrices

$$
\left\{x_{1}, \ldots, x_{n}\right\} \rightarrow B^{n}
$$

## End-to-end visual permutation learning



## Sinkhorn normalization or projection onto $B^{n}$

$$
\begin{aligned}
& \text { sinkhorn_fcn (A) } \\
& \begin{array}{l}
Q=A \\
\text { for } t=1, \ldots, T \text { do } \\
Q_{i, j}
\end{array} \frac{Q_{i, j}}{\sum_{k} Q_{i, k}} \\
& Q_{i, j} \leftarrow \frac{Q_{i, j}}{\sum_{k} Q_{k, j}} \\
& \text { return } Q
\end{aligned}
$$

Alternatively, define a deep declarative module

$$
\begin{array}{ll}
\underset{Q \in \mathbb{R}_{+}^{n \times n}}{\operatorname{minimize}} & \|Q-A\| \\
\text { subject to } & Q \mathbf{1}=\mathbf{1} \\
& Q^{T} \mathbf{1}=\mathbf{1}
\end{array}
$$

## Visual attribute learning results



## Blind perspective-n-point



## Blind perspective-n-point



## Blind perspective-n-point



Australian National University

code and tutorials at http://deepdeclarativenetworks.com CVPR 2020 Workshop (http://cvpr2020.deepdeclarativenetworks.com)

