

Newton Methods for Neural Networks: Gauss Newton Matrix-vector Product

Chih-Jen Lin
National Taiwan University

Outline

- 1 Backward setting
 - Jacobian evaluation
 - Gauss-Newton Matrix-vector products
- 2 Forward + backward settings
 - R operator
 - Gauss-Newton matrix-vector product
- 3 Complexity analysis



Outline

- 1 Backward setting
 - Jacobian evaluation
 - Gauss-Newton Matrix-vector products
- 2 Forward + backward settings
 - R operator
 - Gauss-Newton matrix-vector product
- 3 Complexity analysis



Outline

- 1 Backward setting
 - Jacobian evaluation
 - Gauss-Newton Matrix-vector products
- 2 Forward + backward settings
 - R operator
 - Gauss-Newton matrix-vector product
- 3 Complexity analysis



Jacobian Evaluation: Convolutional Layer I

- For an instance i the Jacobian can be partitioned into L blocks according to layers

$$J^i = [J^{1,i} \quad J^{2,i} \quad \dots \quad J^{L,i}], \quad m = 1, \dots, L, \quad (1)$$

where

$$J^{m,i} = \begin{bmatrix} \frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(W^m)^T} & \frac{\partial \mathbf{z}^{L+1,i}}{\partial (\mathbf{b}^m)^T} \end{bmatrix}.$$

- The calculation seems to be very similar to that for the gradient.



Jacobian Evaluation: Convolutional Layer II

- For the convolutional layers, recall for gradient we have

$$\frac{\partial f}{\partial W^m} = \frac{1}{C} W^m + \frac{1}{I} \sum_{i=1}^I \frac{\partial \xi_i}{\partial W^m}$$

and

$$\frac{\partial \xi_i}{\partial \text{vec}(W^m)^T} = \text{vec} \left(\frac{\partial \xi_i}{\partial S^{m,i}} \phi(\text{pad}(Z^{m,i}))^T \right)^T$$



Jacobian Evaluation: Convolutional Layer III

- Now we have

$$\frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(W^m)^T} = \begin{bmatrix} \frac{\partial z_1^{L+1,i}}{\partial \text{vec}(W^m)^T} \\ \vdots \\ \frac{\partial z_{n_{L+1}}^{L+1,i}}{\partial \text{vec}(W^m)^T} \end{bmatrix}$$

$$= \begin{bmatrix} \text{vec}\left(\frac{\partial z_1^{L+1,i}}{\partial S^{m,i}} \phi(\text{pad}(Z^{m,i}))^T\right)^T \\ \vdots \\ \text{vec}\left(\frac{\partial z_{n_{L+1}}^{L+1,i}}{\partial S^{m,i}} \phi(\text{pad}(Z^{m,i}))^T\right)^T \end{bmatrix}$$



Jacobian Evaluation: Convolutional Layer IV

- If \mathbf{b}^m is considered, the result is

$$\begin{aligned}
 & \begin{bmatrix} \frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(\mathbf{W}^m)^T} & \frac{\partial \mathbf{z}^{L+1,i}}{\partial (\mathbf{b}^m)^T} \end{bmatrix} \\
 = & \begin{bmatrix} \text{vec} \left(\frac{\partial z_1^{L+1,i}}{\partial S^{m,i}} \left[\phi(\text{pad}(Z^{m,i}))^T \mathbf{1}_{a_{\text{conv}}^m b_{\text{conv}}^m} \right] \right)^T \\ \vdots \\ \text{vec} \left(\frac{\partial z_{n_{L+1}}^{L+1,i}}{\partial S^{m,i}} \left[\phi(\text{pad}(Z^{m,i}))^T \mathbf{1}_{a_{\text{conv}}^m b_{\text{conv}}^m} \right] \right)^T \end{bmatrix}^T \cdot
 \end{aligned}$$



Jacobian Evaluation: Convolutional Layer

V

- We can see that it's more complicated than gradient.
- Gradient is a vector but Jacobian is a **matrix**



Jacobian Evaluation: Backward Process I

- For gradient, earlier we need a backward process to calculate

$$\frac{\partial \xi_i}{\partial S^{m,i}}$$

- Now what we need are

$$\frac{\partial z_1^{L+1,i}}{\partial S^{m,i}}, \dots, \frac{\partial z_{n_{L+1}}^{L+1,i}}{\partial S^{m,i}}$$

- The process is similar



Jacobian Evaluation: Backward Process II

- If with RELU activation function and max pooling, for gradient we had

$$\begin{aligned} & \frac{\partial \xi_i}{\partial \text{vec}(S^{m,i})^T} \\ &= \left(\frac{\partial \xi_i}{\partial \text{vec}(Z^{m+1,i})^T} \odot \text{vec}(I[Z^{m+1,i}]^T) \right) P_{\text{pool}}^{m,i}. \end{aligned}$$



Jacobian Evaluation: Backward Process III

- Assume that

$$\frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(Z^{m+1,i})}$$

are available.

$$\begin{aligned} & \frac{\partial \mathbf{z}_j^{L+1,i}}{\partial \text{vec}(S^{m,i})^T} \\ &= \left(\frac{\partial \mathbf{z}_j^{L+1,i}}{\partial \text{vec}(Z^{m+1,i})^T} \odot \text{vec}(I[Z^{m+1,i}])^T \right) P_{\text{pool}}^{m,i}, \\ & j = 1, \dots, n_{L+1}. \end{aligned}$$



Jacobian Evaluation: Backward Process IV

- These row vectors can be written together as a **matrix**

$$\begin{aligned} & \frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(S^{m,i})^T} \\ &= \left(\frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(Z^{m+1,i})^T} \odot (\mathbf{1}_{n_{L+1}} \text{vec}(I[Z^{m+1,i}]^T)^T) \right) P_{\text{pool}}^{m,i}. \end{aligned}$$

- Note that

$$\mathbf{1}_{n_{L+1}} \text{vec}(I[Z^{m+1,i}]^T)^T$$

duplicates the $\text{vec}(I[Z^{m+1,i}]^T)^T$ vector n_{L+1} times



Jacobian Evaluation: Backward Process V

- For gradient, we use

$$\frac{\partial \xi_i}{\partial S^{m,i}}$$

to have

$$\frac{\partial \xi_i}{\partial \text{vec}(Z^{m,i})^T} = \text{vec} \left((W^m)^T \frac{\partial \xi_i}{\partial S^{m,i}} \right)^T P_{\phi}^m P_{\text{pad}}^m$$

and pass it to the previous layer



Jacobian Evaluation: Backward Process VI

- Now we need to generate

$$\frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(Z^{m,i})^T}$$

and pass it to the previous layer.

- Now we have

$$\frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(Z^{m,i})^T} = \begin{bmatrix} \text{vec} \left((W^m)^T \frac{\partial z_1^{L+1,i}}{\partial S^{m,i}} \right)^T P_{\phi}^m P_{\text{pad}}^m \\ \vdots \\ \text{vec} \left((W^m)^T \frac{\partial z_{n_{L+1}}^{L+1,i}}{\partial S^{m,i}} \right)^T P_{\phi}^m P_{\text{pad}}^m \end{bmatrix} \cdot$$



Jacobian Evaluation: Fully-connected Layer l

- We do not discuss details, but list all results below

$$\frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(W^m)^T} = \begin{bmatrix} \text{vec} \left(\frac{\partial \mathbf{z}_1^{L+1,i}}{\partial \mathbf{s}^{m,i}} (\mathbf{z}^{m,i})^T \right)^T \\ \vdots \\ \text{vec} \left(\frac{\partial \mathbf{z}_{n_{L+1}}^{L+1,i}}{\partial \mathbf{s}^{m,i}} (\mathbf{z}^{m,i})^T \right)^T \end{bmatrix}$$



Jacobian Evaluation: Fully-connected Layer II

$$\frac{\partial \mathbf{z}^{L+1,i}}{\partial (\mathbf{b}^m)^T} = \frac{\partial \mathbf{z}^{L+1,i}}{\partial (\mathbf{s}^{m,i})^T},$$

$$\frac{\partial \mathbf{z}^{L+1,i}}{\partial (\mathbf{s}^{m,i})^T} = \frac{\partial \mathbf{z}^{L+1,i}}{\partial (\mathbf{z}^{m+1,i})^T} \odot (\mathbf{1}_{n_{L+1}} / [\mathbf{z}^{m+1,i}]^T)$$

$$\frac{\partial \mathbf{z}^{L+1,i}}{\partial (\mathbf{z}^{m,i})^T} = \frac{\partial \mathbf{z}^{L+1,i}}{\partial (\mathbf{s}^{m,i})^T} W^m$$



Jacobian Evaluation: Fully-connected Layer III

- For the layer $L + 1$, if using a linear activation function with

$$\mathbf{z}^{L+1,i} = \mathbf{s}^{L,i},$$

then we have

$$\frac{\partial \mathbf{z}^{L+1,i}}{\partial (\mathbf{s}^{L,i})^T} = \mathcal{I}_{n_{L+1}}.$$



Gradient versus Jacobian I

- Operations for gradient

$$\begin{aligned} & \frac{\partial \xi_i}{\partial \text{vec}(S^{m,i})^T} \\ &= \left(\frac{\partial \xi_i}{\partial \text{vec}(Z^{m+1,i})^T} \odot \text{vec}(I[Z^{m+1,i}])^T \right) P_{\text{pool}}^{m,i} \\ & \frac{\partial \xi_i}{\partial W^m} = \frac{\partial \xi_i}{\partial S^{m,i}} \phi(\text{pad}(Z^{m,i}))^T \\ & \frac{\partial \xi_i}{\partial \text{vec}(Z^{m,i})^T} = \text{vec} \left((W^m)^T \frac{\partial \xi_i}{\partial S^{m,i}} \right)^T P_{\phi}^m P_{\text{pad}}^m \end{aligned}$$



Gradient versus Jacobian II

- For Jacobian we have

$$\frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(S^{m,i})^T} = \left(\frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(Z^{m+1,i})^T} \odot (\mathbf{1}_{n_{L+1}} \text{vec}(I[Z^{m+1,i}])^T) \right) P_{\text{pool}}^{m,i}$$

$$\frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(W^m)^T} = \begin{bmatrix} \text{vec}\left(\frac{\partial \mathbf{z}_1^{L+1,i}}{\partial S^{m,i}} \phi(\text{pad}(Z^{m,i}))^T\right)^T \\ \vdots \\ \text{vec}\left(\frac{\partial \mathbf{z}_{n_{L+1}}^{L+1,i}}{\partial S^{m,i}} \phi(\text{pad}(Z^{m,i}))^T\right)^T \end{bmatrix}$$



Gradient versus Jacobian III

$$\begin{aligned}
 & \frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(\mathbf{Z}^{m,i})^T} \\
 = & \begin{bmatrix} \text{vec} \left((\mathbf{W}^m)^T \frac{\partial \mathbf{z}_1^{L+1,i}}{\partial \mathcal{S}^{m,i}} \right)^T \mathbf{P}_{\phi}^m \mathbf{P}_{\text{pad}}^m \\ \vdots \\ \text{vec} \left((\mathbf{W}^m)^T \frac{\partial \mathbf{z}_{n_{L+1}}^{L+1,i}}{\partial \mathcal{S}^{m,i}} \right)^T \mathbf{P}_{\phi}^m \mathbf{P}_{\text{pad}}^m \end{bmatrix}. \quad (2)
 \end{aligned}$$



Implementation I

- For gradient we did

$$\Delta \leftarrow \text{mat}(\text{vec}(\Delta)^T P_{\text{pool}}^{m,i})$$

$$\frac{\partial \xi_i}{\partial W^m} = \Delta \cdot \phi(\text{pad}(Z^{m,i}))^T$$

$$\Delta \leftarrow \text{vec} \left((W^m)^T \Delta \right)^T P_{\phi}^m P_{\text{pad}}^m$$

$$\Delta \leftarrow \Delta \odot I[Z^{m,i}]$$

- Now for Jacobian we have similar settings but there are some differences



Implementation II

- We do not really store the Jacobian:

$$\frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(W^m)^T} = \begin{bmatrix} \text{vec}\left(\frac{\partial z_1^{L+1,i}}{\partial S^{m,i}} \phi(\text{pad}(Z^{m,i}))^T\right)^T \\ \vdots \\ \text{vec}\left(\frac{\partial z_{n_{L+1}}^{L+1,i}}{\partial S^{m,i}} \phi(\text{pad}(Z^{m,i}))^T\right)^T \end{bmatrix}$$

- Recall Jacobian is used for matrix-vector products

$$G^S \mathbf{v} = \frac{1}{C} \mathbf{v} + \frac{1}{|S|} \sum_{i \in S} ((J^i)^T (B^i(J^i \mathbf{v}))) \quad (3)$$



Implementation III

- The form

$$\frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(W^m)^T} = \begin{bmatrix} \text{vec}\left(\frac{\partial z_1^{L+1,i}}{\partial S^{m,i}} \phi(\text{pad}(Z^{m,i}))^T\right)^T \\ \vdots \\ \text{vec}\left(\frac{\partial z_{n_{L+1}}^{L+1,i}}{\partial S^{m,i}} \phi(\text{pad}(Z^{m,i}))^T\right)^T \end{bmatrix} \quad (4)$$

is like **the product of two things**



Implementation IV

- If we have

$$\frac{\partial z_1^{L+1,i}}{\partial S^{m,i}}, \dots, \frac{\partial z_{n_{L+1}}^{L+1,i}}{\partial S^{m,i}}, \text{ and } \phi(\text{pad}(Z^{m,i}))$$

probably we can do the matrix-vector product without multiplying these two things out

- We will talk about this again later



Implementation V

- We already know how to obtain

$$\phi(\text{pad}(Z^{m,i}))$$

so the remaining issue is on obtaining

$$\frac{\partial z_1^{L+1,i}}{\partial S^{m,i}}, \dots, \frac{\partial z_{n_{L+1}}^{L+1,i}}{\partial S^{m,i}}$$

- Further we need to take all data (or data in the selected subset) into account
- In the end what we have is the following procedure



Implementation VI

- In the beginning we have

$$\Delta \in R^{d^{m+1} a^{m+1} b^{m+1} \times n_{L+1} \times l} \quad (5)$$

This corresponds to

$$\frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(Z^{m+1,i})^T} \odot (\mathbf{1}_{n_{L+1}} \text{vec}(I[Z^{m+1,i}])^T), \forall i = 1, \dots, l$$



Implementation VII

- We then calculate

$$\Delta \leftarrow \text{mat} \left(\begin{array}{c} \left[(P_{\text{pool}}^{m,1})^T \text{vec}(\Delta_{:,:,1}) \right] \\ \vdots \\ \left[(P_{\text{pool}}^{m,l})^T \text{vec}(\Delta_{:,:,l}) \right] \end{array} \right)_{d^{m+1} \times a_{\text{conv}}^m b_{\text{conv}}^m n_{L+1} l}$$

- Recall that the pooling matrices are **different across instances**



Implementation VIII

- The above operation corresponds to

$$\frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(S^{m,i})^T} = \left(\frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(Z^{m+1,i})^T} \odot (\mathbb{1}_{n_{L+1}} \text{vec}(I[Z^{m+1,i}])^T) \right) P_{\text{pool}}^{m,i}$$

- Now we get

$$\begin{bmatrix} \frac{\partial z_1^{L+1,1}}{\partial S^{m,1}} & \cdots & \frac{\partial z_{n_{L+1}}^{L+1,1}}{\partial S^{m,1}} & \cdots & \frac{\partial z_{n_{L+1}}^{L+1,l}}{\partial S^{m,l}} \end{bmatrix} \in R^{d^{m+1} \times a_{\text{conv}}^m b_{\text{conv}}^m n_{L+1} l}$$



Implementation IX

- For gradient, the next step is to calculate

$$\frac{\partial \xi_i}{\partial W^m} = \dots$$

but here for Jacobian we have mentioned that we do not explicitly get

$$\frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(W^m)^T}$$

- Therefore, the next operation is

$$V \leftarrow \text{vec}((W^m)^T \Delta) \in \mathbb{R}^{hhd^m a_{\text{conv}}^m b_{\text{conv}}^m n_{L+1} l \times 1}$$



Implementation X

- This is same as

$$\text{vec} \left((W^m)^T \left[\frac{\partial z_1^{L+1,1}}{\partial S^{m,1}} \quad \cdots \quad \frac{\partial z_{n_{L+1}}^{L+1,1}}{\partial S^{m,1}} \quad \cdots \quad \frac{\partial z_{n_{L+1}}^{L+1,l}}{\partial S^{m,l}} \right] \right).$$

- Now V is a big vector like

$$\begin{bmatrix} v_1^1 \\ \vdots \\ v_{n_{L+1}}^1 \\ \vdots \\ v_{n_{L+1}}^l \end{bmatrix}$$



Implementation XI

Note that “ \mathbf{v} ” here is not the vector in matrix-vector products. We happen to use the same symbol

- From (2), we then calculate

$$\begin{aligned}
 & (\mathbf{v}_1^1)^T P_\phi^m P_{\text{pad}}^m \\
 & \quad \vdots \\
 & (\mathbf{v}_{n_{L+1}}^1)^T P_\phi^m P_{\text{pad}}^m \\
 & \quad \vdots \\
 & (\mathbf{v}_{n_{L+1}}^l)^T P_\phi^m P_{\text{pad}}^m
 \end{aligned}$$



Implementation XII

- For each resulting vector, we convert it to

$$\text{mat} \left(\mathbf{v}^T \mathbf{P}_{\phi}^m \mathbf{P}_{\text{pad}}^m \right)_{d^m \times a^m b^m}$$

This corresponds to

$$\frac{\partial z_1^{L+1,i}}{\partial Z^{m,i}}, \dots, \frac{\partial z_{n_{L+1}}^{L+1,i}}{\partial Z^{m,i}}, i = 1, \dots, l$$



Implementation XIII

- Finally,

$$\Delta \leftarrow \Delta \odot \left[\underbrace{I[Z^{m,1}] \cdots I[Z^{m,1}]}_{n_{L+1}} \cdots \underbrace{I[Z^{m,l}] \cdots I[Z^{m,l}]}_{n_{L+1}} \right] \quad (6)$$

This is equivalent to

$$\frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(Z^{m,i})^T} \odot (\mathbf{1}_{n_{L+1}} \text{vec}(I[Z^{m,i}])^T), \forall i = 1, \dots, l$$



Implementation XIV

- Note that in the beginning of the calculation, we assume that in (5)

$$\frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(Z^{m+1,i})^T} \odot (\mathbf{1}_{n_{L+1}} \text{vec}(I[Z^{m+1,i}])^T), \forall i = 1, \dots, l$$

is available. The calculation here is to provide information for the previous layer



MATLAB Implementation I

```
dzdS{m} = vTP(model, net, m, num_data,  
              dzdS{m}, 'pool_Jacobian');  
  
dzdS{m} = reshape(dzdS{m},  
                  model.ch_input(m+1), []);  
  
V = model.weight{m}' * dzdS{m};  
dzdS{m-1} = vTP(model, net, m, num_data,  
                 V, 'phi_Jacobian');  
  
% vTP_pad
```



MATLAB Implementation II

```
dzdS{m-1} = reshape(dzdS{m-1},  
    model.ch_input(m), model.ht_pad(m),  
    model.wd_pad(m), []);  
p = model.wd_pad_added(m);  
dzdS{m-1} = dzdS{m-1}(:, p+1:p+model.ht_input(m),  
    p+1:p+model.wd_input(m), :);  
  
dzdS{m-1} =  
    reshape(dzdS{m-1}, [], nL, num_data)  
    .* reshape(net.Z{m} > 0, [], 1, num_data);
```



MATLAB Implementation III

- In the last line for doing (6), we do not need to repeat each $I[Z^{m,i}]$ n_{L+1} times. For `.*`, MATLAB does the expansion automatically



Discussion I

- For doing several CG steps, we should store

$$\frac{\partial z_1^{L+1,i}}{\partial S^{m,i}}, \dots, \frac{\partial z_{n_{L+1}}^{L+1,i}}{\partial S^{m,i}}, i = 1, \dots, l$$

in (4).

- The reason is that it's used for all CG steps (Jacobian matrix remains the same)
- Recalculating them at each CG step is too expensive



Discussion II

- The memory cost is

$$I \times n_{L+1} \times \left(\sum_{m=1}^{L^c} d^{m+1} a_{\text{conv}}^m b_{\text{conv}}^m + \sum_{m=L^c+1}^L n_{m+1} \right) \quad (7)$$

- It is proportional to
 - Number of classes
 - Number of data for the subsampled Hessian
- This memory cost is high
- Thus later we will consider a different approach to reduce the memory consumption



Outline

- 1 Backward setting
 - Jacobian evaluation
 - Gauss-Newton Matrix-vector products
- 2 Forward + backward settings
 - R operator
 - Gauss-Newton matrix-vector product
- 3 Complexity analysis



Gauss-Newton Matrix-Vector Products I

- We check

$$Gv$$

though the situation of using G^S (i.e., a subset of data) is the same

- The Gauss-Newton matrix is

$$G = \frac{1}{C} \mathcal{I} + \frac{1}{I} \sum_{i=1}^I \begin{bmatrix} (J^{1,i})^T \\ \vdots \\ (J^{L,i})^T \end{bmatrix} B^i [J^{1,i} \quad \dots \quad J^{L,i}]$$



Gauss-Newton Matrix-Vector Products II

- The Gauss-Newton matrix vector product is

$$\begin{aligned}
 G\mathbf{v} &= \frac{1}{C}\mathbf{v} + \frac{1}{I} \sum_{i=1}^I \begin{bmatrix} (J^{1,i})^T \\ \vdots \\ (J^{L,i})^T \end{bmatrix} B^i [J^{1,i} \quad \dots \quad J^{L,i}] \begin{bmatrix} \mathbf{v}^1 \\ \vdots \\ \mathbf{v}^L \end{bmatrix} \\
 &= \frac{1}{C}\mathbf{v} + \frac{1}{I} \sum_{i=1}^I \begin{bmatrix} (J^{1,i})^T \\ \vdots \\ (J^{L,i})^T \end{bmatrix} \left(B^i \sum_{m=1}^L J^{m,i} \mathbf{v}^m \right),
 \end{aligned}$$

(8)



Gauss-Newton Matrix-Vector Products III

where

$$\mathbf{v} = \begin{bmatrix} \mathbf{v}^1 \\ \vdots \\ \mathbf{v}^L \end{bmatrix}$$

- Each \mathbf{v}^m , $m = 1, \dots, L$ has the same length as the number of variables (including bias) at the m th layer.



Jacobian-vector Product I

- For the convolutional layers,

$$\begin{aligned}
 & J^{m,i} \mathbf{v}^m \\
 = & \begin{bmatrix} \text{vec} \left(\frac{\partial z_1^{L+1,i}}{\partial S^{m,i}} \left[\phi(\text{pad}(Z^{m,i}))^T \mathbb{1}_{a_{\text{conv}}^m b_{\text{conv}}^m} \right] \right)^T \mathbf{v}^m \\ \vdots \\ \text{vec} \left(\frac{\partial z_{n_{L+1}}^{L+1,i}}{\partial S^{m,i}} \left[\phi(\text{pad}(Z^{m,i}))^T \mathbb{1}_{a_{\text{conv}}^m b_{\text{conv}}^m} \right] \right)^T \mathbf{v}^m \end{bmatrix} \\
 & \in R^{n_{L+1} \times 1}
 \end{aligned}$$

- By this formulation, we need



Jacobian-vector Product II

- a for loop to generate n_{L+1} vectors

$$\text{vec} \left(\frac{\partial z_1^{L+1,i}}{\partial \mathcal{S}^{m,i}} \left[\phi(\text{pad}(Z^{m,i}))^T \mathbf{1}_{a_{\text{conv}}^m b_{\text{conv}}^m} \right] \right)^T$$

$$\vdots$$

$$\text{vec} \left(\frac{\partial z_{n_{L+1}}^{L+1,i}}{\partial \mathcal{S}^{m,i}} \left[\phi(\text{pad}(Z^{m,i}))^T \mathbf{1}_{a_{\text{conv}}^m b_{\text{conv}}^m} \right] \right)^T$$

- the product between the above matrix and a vector \mathbf{v}^m



Jacobian-vector Product III

- Is there a way to avoid a for loop?
- For a language like MATLAB/Octave, we hope to avoid for loops
- Also we hope the code can be simpler and shorter
- We use the following property

$$\text{vec}(AB)^T \text{vec}(C) = \text{vec}(A)^T \text{vec}(CB^T)$$



Jacobian-vector Product IV

- The first element is

$$\begin{aligned}
 & \text{vec} \left(\underbrace{\frac{\partial z_1^{L+1,i}}{\partial S^{m,i}}}_A \underbrace{[\phi(\text{pad}(Z^{m,i}))^T \mathbb{1}_{a_{\text{conv}}^m b_{\text{conv}}^m}]}_B \right)^T \underbrace{\mathbf{v}^m}_{\text{vec}(C)} \\
 &= \frac{\partial z_1^{L+1,i}}{\partial \text{vec}(S^{m,i})^T} \times \\
 & \text{vec} \left(\text{mat}(\mathbf{v}^m)_{d^{m+1} \times (h^m h^m d^{m+1})} \begin{bmatrix} \phi(\text{pad}(Z^{m,i})) \\ \mathbb{1}_{a_{\text{conv}}^m b_{\text{conv}}^m}^T \end{bmatrix} \right).
 \end{aligned}$$



Jacobian-vector Product V

- If all elements are considered together

$$\begin{aligned}
 & J^{m,i} \mathbf{v}^m \\
 &= \frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(S^{m,i})^T} \times \\
 & \text{vec} \left(\text{mat}(\mathbf{v}^m)_{d^{m+1} \times (h^m h^m d^{m+1})} \left[\begin{array}{c} \phi(\text{pad}(Z^{m,i})) \\ \mathbb{1}_{a_{\text{conv}}^m}^T b_{\text{conv}}^m \end{array} \right] \right). \tag{9}
 \end{aligned}$$

This involves

- One matrix-matrix product
- One matrix-vector product



Transposed Jacobian-vector Products I

- After deriving (9), from (8), we sum results of all layers

$$\sum_{m=1}^L J^{m,i} \mathbf{v}^m$$

- Next we calculate

$$\mathbf{q}^i = B^i \left(\sum_{m=1}^L J^{m,i} \mathbf{v}^m \right). \quad (10)$$

- This is usually easy



Transposed Jacobian-vector Products II

- We mentioned earlier that if the squared loss is used

$$B^i = \begin{bmatrix} 2 & & \\ & \vdots & \\ & & 2 \end{bmatrix}$$

is a diagonal matrix



Transposed Jacobian-vector Products III

- Finally, we calculate

$$\begin{aligned}
 & (J^{m,i})^T \mathbf{q}^i \\
 &= \left[\text{vec} \left(\frac{\partial \mathbf{z}_1^{L+1,i}}{\partial \mathbf{S}^{m,i}} [\phi(\text{pad}(Z^{m,i}))^T \mathbb{1}_{a_{\text{conv}}^m b_{\text{conv}}^m}] \right) \cdots \right. \\
 & \quad \left. \text{vec} \left(\frac{\partial \mathbf{z}_{n_{L+1}}^{L+1,i}}{\partial \mathbf{S}^{m,i}} [\phi(\text{pad}(Z^{m,i}))^T \mathbb{1}_{a_{\text{conv}}^m b_{\text{conv}}^m}] \right) \right] \mathbf{q}^i
 \end{aligned}$$



Transposed Jacobian-vector Products IV

$$\begin{aligned}
 &= \sum_{j=1}^{n_{L+1}} q_j^i \text{vec} \left(\frac{\partial z_j^{L+1,i}}{\partial \mathcal{S}^{m,i}} [\phi(\text{pad}(Z^{m,i}))^T \mathbb{1}_{a_{\text{conv}}^m b_{\text{conv}}^m}] \right) \\
 &= \text{vec} \left(\sum_{j=1}^{n_{L+1}} q_j^i \left(\frac{\partial z_j^{L+1,i}}{\partial \mathcal{S}^{m,i}} [\phi(\text{pad}(Z^{m,i}))^T \mathbb{1}_{a_{\text{conv}}^m b_{\text{conv}}^m}] \right) \right) \\
 &= \text{vec} \left(\left(\sum_{j=1}^{n_{L+1}} q_j^i \frac{\partial z_j^{L+1,i}}{\partial \mathcal{S}^{m,i}} \right) [\phi(\text{pad}(Z^{m,i}))^T \mathbb{1}_{a_{\text{conv}}^m b_{\text{conv}}^m}] \right)
 \end{aligned}$$



Transposed Jacobian-vector Products V

$$\begin{aligned}
 &= \text{vec} \left(\text{mat} \left(\left(\frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(\mathbf{S}^{m,i})^T} \right)^T \mathbf{q}^i \right) \right) \times \\
 &\quad \left[\phi(\text{pad}(\mathbf{Z}^{m,i}))^T \mathbb{1}_{a_{\text{conv}}^m b_{\text{conv}}^m} \right]. \tag{11}
 \end{aligned}$$

This needs a matrix-vector product and then a matrix-matrix product



Fully-connected Layers I

- Similar to the results of the convolutional layers, for the fully-connected layers we have

$$J^{m,i} \mathbf{v}^m = \frac{\partial \mathbf{z}^{L+1,i}}{\partial (\mathbf{s}^{m,i})^T} \text{mat}(\mathbf{v}^m)_{n_{m+1} \times (n_m+1)} \begin{bmatrix} \mathbf{z}^{m,i} \\ \mathbf{1}_1 \end{bmatrix}.$$

$$(J^{m,i})^T \mathbf{q}^i = \text{vec} \left(\left(\frac{\partial \mathbf{z}^{L+1,i}}{\partial (\mathbf{s}^{m,i})^T} \right)^T \mathbf{q}^i [(\mathbf{z}^{m,i})^T \mathbf{1}_1] \right).$$



Implementation I

- As before, we must handle all instances together
- We discuss only

$$\begin{bmatrix} \sum_{m=1}^L J^{m,1} \mathbf{v}^m \\ \vdots \\ \sum_{m=1}^L J^{m,l} \mathbf{v}^m \end{bmatrix} \in \mathbb{R}^{n_{L+1} \times 1}$$

- Following earlier derivation



Implementation II

$$\begin{aligned}
 \begin{bmatrix} J^{m,1} \mathbf{v}^m \\ \vdots \\ J^{m,l} \mathbf{v}^m \end{bmatrix} &= \begin{bmatrix} \frac{\partial \mathbf{z}^{L+1,1}}{\partial \text{vec}(S^{m,1})^T} \text{vec} \left(\text{mat}(\mathbf{v}^m) \begin{bmatrix} \phi(\text{pad}(Z^{m,1})) \\ \mathbb{1}_{a_{\text{conv}}^m}^T b_{\text{conv}}^m \end{bmatrix} \right) \\ \vdots \\ \frac{\partial \mathbf{z}^{L+1,l}}{\partial \text{vec}(S^{m,l})^T} \text{vec} \left(\text{mat}(\mathbf{v}^m) \begin{bmatrix} \phi(\text{pad}(Z^{m,l})) \\ \mathbb{1}_{a_{\text{conv}}^m}^T b_{\text{conv}}^m \end{bmatrix} \right) \end{bmatrix} \\
 &= \begin{bmatrix} \frac{\partial \mathbf{z}^{L+1,1}}{\partial \text{vec}(S^{m,1})^T} \mathbf{p}^{m,1} \\ \vdots \\ \frac{\partial \mathbf{z}^{L+1,l}}{\partial \text{vec}(S^{m,l})^T} \mathbf{p}^{m,l} \end{bmatrix},
 \end{aligned}$$



Implementation III

- where

$$\text{mat}(\mathbf{v}^m) \in R^{d^{m+1} \times (h^m h^m d^m + 1)}$$

and

$$\mathbf{p}^{m,i} = \text{vec} \left(\text{mat}(\mathbf{v}^m) \begin{bmatrix} \phi(\text{pad}(Z^{m,i})) \\ \mathbf{1}^T \\ a_{\text{conv}}^m \\ b_{\text{conv}}^m \end{bmatrix} \right). \quad (12)$$



Implementation IV

- To get $\mathbf{p}^{m,i}$, a matrix-matrix product is needed. For all $i = 1, \dots, l$ the calculation can be done by a matrix-matrix product

$$\text{mat}(\mathbf{v}^m) \begin{bmatrix} \phi(\text{pad}(Z^{m,1})) & \cdots & \phi(\text{pad}(Z^{m,l})) \\ \mathbf{1}_{a_{\text{conv}}^m b_{\text{conv}}^m}^T & \cdots & \mathbf{1}_{a_{\text{conv}}^m b_{\text{conv}}^m}^T \end{bmatrix}$$

$$\in R^{d^{m+1} \times a_{\text{conv}}^m b_{\text{conv}}^m l}$$



Implementation V

- To get

$$\begin{bmatrix} \frac{\partial \mathbf{z}^{L+1,1}}{\partial \text{vec}(S^{m,1})^T} \mathbf{p}^{m,1} \\ \vdots \\ \frac{\partial \mathbf{z}^{L+1,l}}{\partial \text{vec}(S^{m,l})^T} \mathbf{p}^{m,l} \end{bmatrix},$$

we need l matrix-vector products

- There is no good way to transform it to matrix-matrix operations



Implementation VI

- To avoid a for loop over all data, here we implement the matrix-vector product

$$J^{m,i} \mathbf{v}^m = \frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(S^{m,i})^T} \mathbf{p}^{m,i} \quad (13)$$

by a summation of all rows of the following matrix

$$\begin{bmatrix} \frac{\partial z_1^{L+1,i}}{\partial \text{vec}(S^{m,i})} & \cdots & \frac{\partial z_{n_{L+1}}^{L+1,i}}{\partial \text{vec}(S^{m,i})} \end{bmatrix}_{d^{m+1} a_{\text{conv}}^m b_{\text{conv}}^m \times n_{L+1}} \odot \begin{bmatrix} \mathbf{p}^{m,i} & \cdots & \mathbf{p}^{m,i} \end{bmatrix}_{d^{m+1} a_{\text{conv}}^m b_{\text{conv}}^m \times n_{L+1}} \cdot$$



Implementation VII

- For example, summing up all elements of the first column is the inner product between the first row of

$$\frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(\mathbf{S}^{m,i})^T} \text{ and } \mathbf{p}^{m,i}$$

- Then all the l matrix-vector products

$$\mathbf{J}^{m,i} \mathbf{v}^m = \frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(\mathbf{S}^{m,i})^T} \mathbf{p}^{m,i}, \quad i = 1, \dots, l.$$



Implementation VIII

can be done in one line by

$$\begin{bmatrix} \dots & \frac{\partial z_1^{L+1,i}}{\partial \text{vec}(S^{m,i})} & \dots & \frac{\partial z_{n_{L+1}}^{L+1,i}}{\partial \text{vec}(S^{m,i})} & \dots \end{bmatrix} \odot \begin{bmatrix} \dots & \mathbf{p}^{m,i} & \dots & \mathbf{p}^{m,i} & \dots \end{bmatrix}$$

- The code (convolutional layers) is like



Implementation IX

```
for m = LC : -1 : 1
    var_range = var_ptr(m) : var_ptr(m+1) - 1;
    ab = model.ht_conv(m)*model.wd_conv(m);
    d = model.ch_input(m+1);

    p = reshape(v(var_range), d, []) *
        [net.phiZ{m}; ones(1, ab*num_data)];
    p = sum(reshape(net.dzdS{m}, d*ab, nL,
        []) .*
        reshape(p, d*ab, 1, []), 1);
    Jv = Jv + p(:);
```



Implementation X

end

- Note that

`sum(:,1);`

sums up all rows

- For $\mathbf{p}^{m,i}$ we do not duplicate it n_{L+1} times. Instead, for `.*`, MATLAB does the expansion automatically



Outline

- 1 Backward setting
 - Jacobian evaluation
 - Gauss-Newton Matrix-vector products
- 2 Forward + backward settings
 - R operator
 - Gauss-Newton matrix-vector product
- 3 Complexity analysis



Outline

- 1 Backward setting
 - Jacobian evaluation
 - Gauss-Newton Matrix-vector products
- 2 Forward + backward settings
 - R operator
 - Gauss-Newton matrix-vector product
- 3 Complexity analysis



Reverse versus Forward Autodiff I

- We mentioned before that two types of autodiff are forward and reverse modes
- For the Jacobian evaluation, at layer m ,

$$J^{m,i} = \left[\frac{\partial \mathbf{z}^{L+1,i}}{\partial \text{vec}(W^m)^T} \quad \frac{\partial \mathbf{z}^{L+1,i}}{\partial (\mathbf{b}^m)^T} \right],$$

naturally we follow the gradient calculation to use the reverse mode

- But this may not be a good decision



Reverse versus Forward Autodiff II

- In particular, we must store $J^{m,i}, \forall i$, or more precisely,

$$\frac{\partial z_1^{L+1,i}}{\partial S^{m,i}}, \dots, \frac{\partial z_{n_{L+1}}^{L+1,i}}{\partial S^{m,i}}, i = 1, \dots, l,$$

where the memory cost is

$$l \times n_{L+1} \times \left(\sum_{m=1}^{L^c} d^{m+1} a_{\text{conv}}^m b_{\text{conv}}^m + \sum_{m=L^c+1}^L n_{m+1} \right)$$

This memory cost is higher than other stored information



Reverse versus Forward Autodiff III

- For example, the $Z^{m,i}, \forall i$ stored from the forward process takes

$$l \times \left(\sum_{m=1}^{L^c} d^m a^m b^m + \sum_{m=L^c+1}^{L+1} n_m \right),$$

which is independent to the number of classes.

- We will show a solution to address this memory difficulty
- First, for the Jacobian-vector product, we will use the forward mode of automatic differentiation



Reverse versus Forward Autodiff IV

- Recall earlier we said that by the forward mode, the Jacobian-vector product can be done in just one pass



R Operator I

- Consider $g(\boldsymbol{\theta}) \in R^{k \times 1}$. Following Pearlmutter (1994), we define

$$\mathcal{R}_{\mathbf{v}}\{g(\boldsymbol{\theta})\} \equiv \frac{\partial g(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}^T} \mathbf{v} = \begin{bmatrix} \nabla g_1(\boldsymbol{\theta})^T \mathbf{v} \\ \vdots \\ \nabla g_k(\boldsymbol{\theta})^T \mathbf{v} \end{bmatrix}. \quad (14)$$

- Note that

$$\begin{bmatrix} \nabla g_1(\boldsymbol{\theta})^T \\ \vdots \\ \nabla g_k(\boldsymbol{\theta})^T \end{bmatrix}$$

is the **Jacobian** of $g(\boldsymbol{\theta})$



R Operator II

- This definition can be extended to a matrix $M(\boldsymbol{\theta}) \in R^{k \times t}$ by

$$\begin{aligned} \mathcal{R}_{\mathbf{v}}\{M(\boldsymbol{\theta})\} &\equiv \text{mat}(\mathcal{R}_{\mathbf{v}}\{\text{vec}(M(\boldsymbol{\theta}))\})_{k \times t} \\ &= \text{mat}\left(\frac{\partial \text{vec}(M(\boldsymbol{\theta}))}{\partial \boldsymbol{\theta}^T} \mathbf{v}\right)_{k \times t} = \begin{bmatrix} \nabla M_{11}^T \mathbf{v} & \cdots & \nabla M_{1t}^T \mathbf{v} \\ \vdots & \ddots & \vdots \\ \nabla M_{k1}^T \mathbf{v} & \cdots & \nabla M_{kt}^T \mathbf{v} \end{bmatrix} \end{aligned}$$

- Clearly,

$$\mathcal{R}_{\mathbf{v}}\{M(\boldsymbol{\theta})\} = (\mathcal{R}_{\mathbf{v}}\{M(\boldsymbol{\theta})^T\})^T. \quad (15)$$

R Operator III

- If $h(\cdot)$ is a scalar function, we let

$$h(M(\boldsymbol{\theta})) = \begin{bmatrix} h(M_{11}) & \cdots & h(M_{1t}) \\ \vdots & \ddots & \vdots \\ h(M_{k1}) & \cdots & h(M_{kt}) \end{bmatrix}$$

and

$$h'(M(\boldsymbol{\theta})) = \begin{bmatrix} h'(M_{11}) & \cdots & h'(M_{1t}) \\ \vdots & \ddots & \vdots \\ h'(M_{k1}) & \cdots & h'(M_{kt}) \end{bmatrix} .$$



R Operator IV

- Because

$$\nabla(h(M_{ij}(\boldsymbol{\theta})))^T \mathbf{v} = h'(M_{ij}) \nabla(M_{ij})^T \mathbf{v},$$

we have

$$\begin{aligned} \mathcal{R}_{\mathbf{v}}\{h(M(\boldsymbol{\theta}))\} &= \begin{bmatrix} \nabla h(M_{11})^T \mathbf{v} & \cdots & \nabla h(M_{1t})^T \mathbf{v} \\ \vdots & \ddots & \vdots \\ \nabla h(M_{k1})^T \mathbf{v} & \cdots & \nabla h(M_{kt})^T \mathbf{v} \end{bmatrix} \\ &= h'(M(\boldsymbol{\theta})) \odot \mathcal{R}_{\mathbf{v}}\{M(\boldsymbol{\theta})\}, \end{aligned} \tag{16}$$



R Operator V

where \odot stands for the Hadamard product (i.e., component-wise product).

- If $M(\boldsymbol{\theta})$ and $T(\boldsymbol{\theta})$ have the same size,

$$\mathcal{R}_v\{M(\boldsymbol{\theta}) + T(\boldsymbol{\theta})\} = \mathcal{R}_v\{M(\boldsymbol{\theta})\} + \mathcal{R}_v\{T(\boldsymbol{\theta})\}. \quad (17)$$

- Lastly, we have

$$\mathcal{R}_v\{U(\boldsymbol{\theta})M(\boldsymbol{\theta})\} = \mathcal{R}_v\{U(\boldsymbol{\theta})\}M(\boldsymbol{\theta}) + U(\boldsymbol{\theta})\mathcal{R}_v\{M(\boldsymbol{\theta})\} \quad (18)$$



R Operator VI

Proof: Note that

$$(\mathcal{R}\{U(\boldsymbol{\theta})M(\boldsymbol{\theta})\})_{ij} = \nabla ((U(\boldsymbol{\theta})M(\boldsymbol{\theta}))_{ij})^T \mathbf{v}.$$

With

$$(U(\boldsymbol{\theta})M(\boldsymbol{\theta}))_{ij} = \sum_{p=1}^m U_{ip}M_{pj},$$

we have both $U_{ip} \in R^1$ and $M_{pj} \in R^1$. Then,

$$\nabla (U_{ip}M_{pj})^T \mathbf{v} = ((\nabla U_{ip})^T \mathbf{v}) M_{pj} + U_{ip} ((\nabla M_{pj})^T \mathbf{v}).$$



R Operator VII

The summation

$$\sum_{p=1}^m ((\nabla U_{ip})^T \mathbf{v}) M_{pj}$$

leads to the (i, j) component of

$$\mathcal{R}_{\mathbf{v}}\{U(\boldsymbol{\theta})\}M(\boldsymbol{\theta})$$

Thus we have (18)

- For simplicity, subsequently we use $\mathcal{R}\{g(\boldsymbol{\theta})\}$ to denote $\mathcal{R}_{\mathbf{v}}\{g(\boldsymbol{\theta})\}$



R Operator for $J^i \mathbf{v}$ I

- We have

$$J^i \mathbf{v} = \frac{\partial \mathbf{z}^{L+1,i}}{\partial \boldsymbol{\theta}^T} \mathbf{v} = \mathcal{R}\{\mathbf{z}^{L+1,i}\}.$$

- We consider the following forward operations by assuming that

$$\mathcal{R}\{\mathbf{z}^{m,i}\}$$

is available from the previous layer



R Operator for $J^i v$ II

- From (18), we have

$$\begin{aligned}
 & \mathcal{R}\{\phi(\text{pad}(Z^{m,i}))\} \\
 &= \mathcal{R}\left\{\text{mat}\left(P_{\phi}^{m,i} P_{\text{pad}}^{m,i} \text{vec}(Z^{m,i})\right)\right\} \\
 &= \text{mat}\left(\mathcal{R}\left\{P_{\phi}^{m,i} P_{\text{pad}}^{m,i} \text{vec}(Z^{m,i})\right\}\right) \\
 &= \text{mat}\left(P_{\phi}^{m,i} P_{\text{pad}}^{m,i} \mathcal{R}\{\text{vec}(Z^{m,i})\}\right) \quad h^m h^m d^m \times a_{\text{Conv}}^m b_{\text{Conv}}^m
 \end{aligned}$$



R Operator for $J^i \mathbf{v}$ III

- From (17) and (18), we have

$$\begin{aligned}
 & \mathcal{R}\{S^{m,i}\} \\
 &= \mathcal{R}\{W^m \phi(\text{pad}(Z^{m,i})) + \mathbf{b}^m \mathbb{1}_{a_{\text{conv}}^m b_{\text{conv}}^m}^T\} \\
 &= \mathcal{R}\{W^m \phi(\text{pad}(Z^{m,i}))\} + \mathcal{R}\{\mathbf{b}^m \mathbb{1}_{a_{\text{conv}}^m b_{\text{conv}}^m}^T\} \\
 &= \mathcal{R}\{W^m\} \phi(\text{pad}(Z^{m,i})) + W^m \mathcal{R}\{\phi(\text{pad}(Z^{m,i}))\} + \\
 & \quad \mathcal{R}\{\mathbf{b}^m\} \mathbb{1}_{a_{\text{conv}}^m b_{\text{conv}}^m}^T \\
 &= V_W^m \phi(\text{pad}(Z^{m,i})) + W^m \mathcal{R}\{\phi(\text{pad}(Z^{m,i}))\} + \\
 & \quad \mathbf{v}_b^m \mathbb{1}_{a_{\text{conv}}^m b_{\text{conv}}^m}^T,
 \end{aligned}$$



R Operator for $J^i \mathbf{v}$ IV

where we have

$$\begin{aligned}\mathcal{R}\{W^m\} &= V_W^m, \\ \mathcal{R}\{\mathbf{b}^m\} &= \mathbf{v}_b^m.\end{aligned}$$

- Note that

$$\mathbf{v} = \begin{bmatrix} \mathbf{v}^1 \\ \vdots \\ \mathbf{v}^L \end{bmatrix},$$

and each \mathbf{v}^m , $m = 1, \dots, L$ has the same length as the number of variables (including bias) at the m th layer.



R Operator for $J^i \mathbf{v} \mathbf{V}$

- We further split \mathbf{v}^m to V_W^m (a matrix form) and \mathbf{v}_b^m
- From (16), we have

$$\mathcal{R}\{\sigma(S^{m,i})\} = \sigma'(S^{m,i}) \odot \mathcal{R}\{S^{m,i}\}. \quad (19)$$

- From (18), we have

$$\begin{aligned} & \mathcal{R}\{Z^{m+1,i}\} \\ &= \mathcal{R}\{\text{mat}(P_{\text{pool}}^{m,i} \text{vec}(\sigma(S^{m,i})))\} \\ &= \text{mat}(\mathcal{R}\{P_{\text{pool}}^{m,i} \text{vec}(\sigma(S^{m,i}))\}) \\ &= \text{mat}\left(P_{\text{pool}}^{m,i} \mathcal{R}\{\text{vec}(\sigma(S^{m,i}))\}\right) \end{aligned}$$



R Operator for $J^i \mathbf{v}$ VI

- We can continue this process until we get

$$J^i \mathbf{v} = \mathcal{R}\{\mathbf{z}^{L+1,i}\}.$$

- Clearly, we do not need to store

$$\frac{\partial \mathbf{z}_1^{L+1,i}}{\partial \mathcal{S}^{m,i}}, \dots, \frac{\partial \mathbf{z}_{n_{L+1}}^{L+1,i}}{\partial \mathcal{S}^{m,i}}$$

as before, so the memory issue is solved

- But how about

$$(J^i)^T(\cdot)?$$

We will explain later that they are not needed



Outline

- 1 Backward setting
 - Jacobian evaluation
 - Gauss-Newton Matrix-vector products
- 2 Forward + backward settings
 - R operator
 - Gauss-Newton matrix-vector product
- 3 Complexity analysis



Gauss-Newton Matrix-vector Product I

- From the above discussion, we have known how to calculate

$$J^i \mathbf{v}$$

- Calculate

$$B^i(J^i \mathbf{v})$$

is known to be easy



Gauss-Newton Matrix-vector Product II

- Now for

$$(J^i)^T (B^i J^i \mathbf{v}),$$


if we define

$$\mathbf{u} = B^i J^i \mathbf{v},$$

then

$$(J^i)^T \mathbf{u} = \left(\frac{\partial \mathbf{z}^{L+1,i}}{\partial \boldsymbol{\theta}^T} \right)^T \mathbf{u}.$$

- But earlier the gradient calculation is

$$(J^i)^T \nabla_{\mathbf{z}^{L+1,i}} \xi(\mathbf{z}^{L+1,i}; \mathbf{y}^i, \mathbf{z}^{1,i}) = \left(\frac{\partial \mathbf{z}^{L+1,i}}{\partial \boldsymbol{\theta}^T} \right)^T \frac{\partial \xi_i}{\partial \mathbf{z}^{L+1,i}}$$


Gauss-Newton Matrix-vector Product III

- Thus the same backward procedure can be used
- All we need is to replace

$$\frac{\partial \xi_i}{\partial \mathbf{z}^{L+1,i}}$$

with

u

- Therefore, we do not need to explicitly derive J^i at all.



Gauss-Newton Matrix-vector Product IV

- Thus for $(J^i)^T \mathbf{u}$, there is no need to store

$$\frac{\partial z_1^{L+1,i}}{\partial S^{m,i}}, \dots, \frac{\partial z_{n_{L+1}}^{L+1,i}}{\partial S^{m,i}}$$



Outline

- 1 Backward setting
 - Jacobian evaluation
 - Gauss-Newton Matrix-vector products
- 2 Forward + backward settings
 - R operator
 - Gauss-Newton matrix-vector product
- 3 Complexity analysis



Complexity Analysis I

- We have known from past slides that matrix-matrix products are the bottleneck (though in our cases some slow MATLAB functions are also bottlenecks in practice)
- For simplicity, in our analysis we just count the number of matrix-matrix products without worrying about their sizes



Complexity Analysis II

- For approaches solely by backward settings, if

$$\frac{\partial z_1^{L+1,i}}{\partial S^{m,i}}, \dots, \frac{\partial z_{n_{L+1}}^{L+1,i}}{\partial S^{m,i}}$$

are stored, then the complexity of a Newton iteration is proportional to

$$(n_{L+1} + 1) + \#CG \times 2,$$

where $\#CG$ is the number of CG steps in that iteration



Complexity Analysis III

- If not, then

$$\#CG \times ((n_{L+1} + 1) + 2)$$

Note that here we assume that $Z^{m,i}$ are not stored either, so at each CG step, a forward process is needed

- Therefore, “1” of “ $n_{L+1} + 1$ ” comes from one product in the forward process. In the backward process we need n_{L+1} products

$$\text{vec} \left((W^m)^T \begin{bmatrix} \frac{\partial z_1^{L+1,1}}{\partial S^{m,1}} & \cdots & \frac{\partial z_{n_{L+1}}^{L+1,1}}{\partial S^{m,1}} & \cdots & \frac{\partial z_{n_{L+1}}^{L+1,l}}{\partial S^{m,l}} \end{bmatrix} \right)$$



Complexity Analysis IV

- The situation is slightly different from the Gradient calculation, which needs “3” products (one in forward and two in backward).

The reason is that now we do not need

$$\frac{\partial \xi_i}{\partial W^m} = \Delta \cdot \phi(\text{pad}(Z^{m,i}))^T$$

- For “# CG \times 2”, the “2” is from (9) and (11)



Complexity Analysis V

- If using R operators, then

$$\#CG \times (3 + 2)$$

products are needed, where “3” are from the forward process

$$W^m \phi(\text{pad}(Z^{m,i})),$$

and

$$V_W^m \phi(\text{pad}(Z^{m,i})), W^m \mathcal{R}\{\phi(\text{pad}(Z^{m,i}))\},$$

and “2” are from the backward process



Complexity Analysis VI

- Clearly, under the same memory consumption, the one using R operators is much more efficient



Discussion I

- At this moment in the Python code we are not using the forward mode for Jv
- It was not available before
- However, since version 2.10 released in January 2020, this functionality is provided:
https://www.tensorflow.org/api_docs/python/tf/autodiff/ForwardAccumulator
- It will be interesting to do the implementation and make a comparison

