Dongheng Lin
Jan 30th, 2024

CS 549: Computer Vision
Assignment 1

Question 1

1. Shift Invariance

To show that the softmax function is invariant to a constant shift, we would like to verify that:

$$\text{softmax}(z) = \text{softmax}(z - C1), \text{ where } C1 = \{C, C, C \dots C\}^T \tag{1}$$

By definition we have:

$$\begin{aligned}
\text{softmax}(\mathbf{z} - C\mathbf{1})_i &= \frac{\exp(z_i - C)}{\sum_{j=1}^{k} \exp(z_j - C)} \\
&= \frac{\exp(z_i)\exp(-C)}{\sum_{j=1}^{k} \exp(z_j)\exp(-C)} \\
&= \frac{\exp(z_i)}{\sum_{j=1}^{k} \exp(z_j)} \\
&= \text{softmax}(\mathbf{z})
\end{aligned}$$

Hence we proved the statement.

2. Derivative

When $i = j$:

$$\frac{\partial y_i}{\partial z_i} = \frac{\exp(z_i)\sum_{j=1}^{k} \exp(z_j) - \exp(z_i)^2}{\left(\sum_{j=1}^{k} \exp(z_j)\right)^2}$$

we know that $y_i = \frac{\exp(z_i)}{\sum_{j=1}^{k} \exp(z_j)}$, plug that in gives:

$$\frac{\partial y_i}{\partial z_i} = y_i(1 - y_i)$$

Similarly, when $i \neq j$:

$$\begin{aligned}
\frac{\partial y_i}{\partial z_j} &= \frac{-\exp(z_i)\exp(z_j)}{\left(\sum_{j=1}^{k} \exp(z_j)\right)^2} \\
&= -y_i y_j
\end{aligned}$$

3. Chain Rule

Given $\mathbf{z} = W^\top \mathbf{x} - \mathbf{u}$, we can express each element of $\mathbf{z}$ as $z_i = \mathbf{w}_i^\top \mathbf{x} - u_i$, where $\mathbf{w}_i$ is the i-th column of W.

Use the result from above, we have $\frac{\partial y_i}{\partial z_j}$, which we can use in the chain rule:

$$\frac{\partial y_i}{\partial \mathbf{x}} = \sum_{j=1}^{k} \frac{\partial y_i}{\partial z_j} \frac{\partial z_j}{\partial \mathbf{x}}$$
$$\frac{\partial z_j}{\partial \mathbf{x}} = \mathbf{w}_j \left(\text{ since } z_j = \mathbf{w}_j^\top \mathbf{x} - u_j \right)$$

Therefore,

$$\frac{\partial y_i}{\partial \mathbf{x}} = \sum_{j=1}^{k} \frac{\partial y_i}{\partial z_j} \mathbf{w}_j$$

Similarly,

$$\frac{\partial y_i}{\partial \mathbf{w}_j} = \sum_{l=1}^{k} \frac{\partial y_i}{\partial z_l} \frac{\partial z_l}{\partial w_j}$$

Since $\frac{\partial z_l}{\partial \mathbf{w}_j}$ is $\mathbf{x}$ when $l = j$ and 0 otherwise,

$$\frac{\partial y_i}{\partial w_j} = \frac{\partial y_i}{\partial z_j} \mathbf{x}$$

Finally,

$$\frac{\partial y_i}{\partial \mathbf{u}} = \sum_{j=1}^{k} \frac{\partial y_i}{\partial z_j} \frac{\partial z_j}{\partial \mathbf{u}}$$

Since $\frac{\partial z_j}{\partial \mathrm{u}}$ is -1 at the j-th component and 0 elsewhere,

$$\frac{\partial y_i}{\partial \mathbf{u}} = -\frac{\partial y_i}{\partial z_i}$$

Question 2

1. Matrix Multiplication

Given $V = \begin{bmatrix} \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix}$. Multiplying V by $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$ gives:

$$\begin{bmatrix} \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{bmatrix}$$

And compute $V^T \begin{bmatrix} 0 \\ 1 \end{bmatrix}$ Multiplying V^T by $\begin{bmatrix} 0 \\ 1 \end{bmatrix}$ gives:

$$\begin{bmatrix} \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix} = \begin{bmatrix} -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{bmatrix}$$

It is clear that Vx rotates $x \in \mathcal{R}^2$ by 45 degrees counterclockwise.

2. Matrix Transpose

To verify that $V^{-1} = V^T$, For a matrix V, if $V^{-1} = V^T$, by definition we have:

$$VV^T = \begin{bmatrix} \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix} \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ -\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

Since $VV^T = I, V^{-1} = V^T$. As we have shown above, It is clear that $V^T x$ rotates $x \in \mathcal{R}^2$ by 45 degrees clockwise.

3. Diagnal Matrix

The transformed vectors, after applying $\Sigma V^\top$, are as follows:

$$\Sigma V^\top \begin{bmatrix} \frac{1}{\sqrt{2}} \\ 0 \end{bmatrix} = \begin{bmatrix} 1.5 \\ -2.5 \end{bmatrix}$$

$$\Sigma V^\top \begin{bmatrix} 0 \\ \frac{1}{\sqrt{2}} \end{bmatrix} = \begin{bmatrix} 1.5 \\ 2.5 \end{bmatrix}$$

$$\Sigma V^\top \begin{bmatrix} -\frac{1}{\sqrt{2}} \\ 0 \end{bmatrix} = \begin{bmatrix} -1.5 \\ 2.5 \end{bmatrix}$$

$$\Sigma V^\top \begin{bmatrix} 0 \\ -\frac{1}{\sqrt{2}} \end{bmatrix} = \begin{bmatrix} -1.5 \\ -2.5 \end{bmatrix}$$

These vectors represent the corners of a transformed square. The original square, formed by the vectors, undergoes a 45-degree clockwise rotation (due to V^T) and then a scaling transformation (due to Σ). The scaling stretches the square by a factor of 3 along the x-axis and a factor of 5 along the y-axis, resulting in a rectangle.

4. Geometric Interpretation

According to the problem definition, for a general squared matrix $B \in \mathcal{R}^{n \times n}$, a similar geometric interpretation would involve finding its singular value decomposition (SVD) where $B = U\Sigma V^T$, and then interpreting the transformations in terms of rotations (by U and V^T) and scaling (by Σ).

5. Homogeneous Equation

Given:

$$A = U\Sigma V^T = \begin{bmatrix} -\frac{1}{\sqrt{5}} & \frac{2}{\sqrt{5}} \\ -\frac{2}{\sqrt{5}} & -\frac{1}{\sqrt{5}} \end{bmatrix} \begin{bmatrix} 2 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix}^\top$$

From Σ, we see that the second singular value is zero. Therefore, the null space of A is spanned by the second column of V.

The matrix V is:

$$V = \begin{bmatrix} \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix}$$

The second column of V is:

$$v_2 = \begin{bmatrix} -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{bmatrix}$$

This vector v_2 is already of unit length (i.e., $\|v_2\| = 1$), as it satisfies [1]:

$$\left\| \begin{bmatrix} -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{bmatrix} \right\| = \sqrt{\left(-\frac{1}{\sqrt{2}}\right)^2 + \left(\frac{1}{\sqrt{2}}\right)^2} = \sqrt{\frac{1}{2} + \frac{1}{2}} = \sqrt{1} = 1$$

Therefore, the solution to the homogeneous equation $Ax = 0$ with $\|x\| = 1$ is:

$$x = \begin{bmatrix} -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{bmatrix}$$

Figure 1: Reordered Alma Statue

Question 3

3.1 Reorder

My solution employs a greedy algorithm to reassemble an image from its shredded pieces. This is achieved by computing the pairwise distances between the shreds based on pixel similarity and iteratively selecting the best fitting pairs to reconstruct the original image. Specifically, for each pair of shreds, the distance function calculates the 'distance' between them, representing how well they fit together. Specifically, it computes the L2 norm between the last column of one shred and the first column of another. Mathematically, this can be represented as:

$$\text{dist}(I_j, I_i) = \sum_{k=1}^{K} \left(L_j(k) - F_i(k)\right)^2$$

Where $\text{dist}(I_j, I_i)$ is the distance measure between the i-th and j-th image strips, I_j, I_i denotes two different image strips. $L_j(k)$ is The pixel values in the last column of strip I_j at row k. $F_i(k)$ The pixel values in the first column of strip I_i at the same row index k. It is manifest that smaller distance indicates larger similarity.

3.2 Align & Reorder

The part 3.2 adapts most data structures and greedy-matching algorithms in 3.1, where the core adjustments lies in the similarity calculation.. Instead of a straightforward comparison, this function now accounts for vertical sliding, considering up to 20% of the strip's height. It calculates the similarity over various alignments by sliding one strip relative to another, both upwards and downwards, and selects the maximum similarity from these alignments that will be used for reordering task. At the same time, it will also record the best slide offset between slides for final composition. There is also a normalization step incorporated to improve quality of final alignment.

Figure 2: Reordered Person

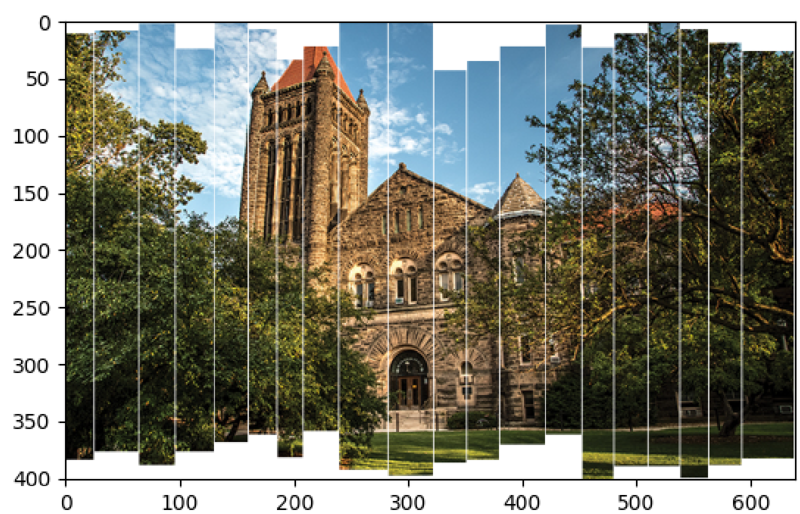

Figure 3: Aligned Altgeld

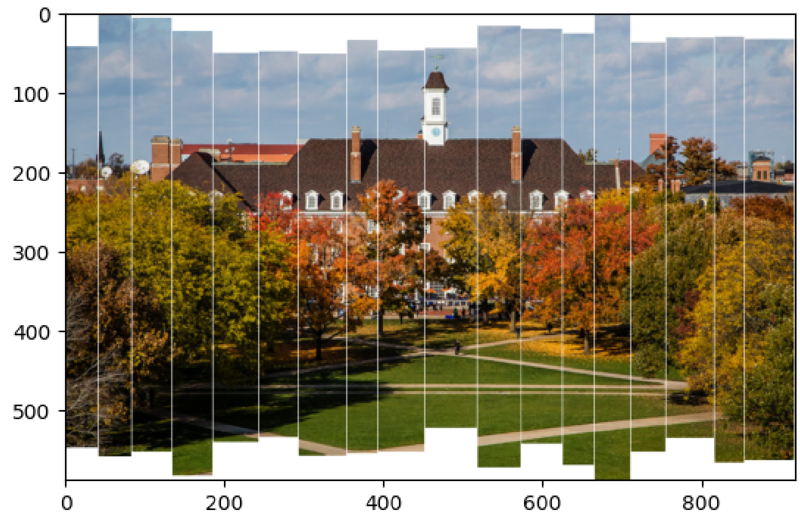

Figure 4: Aligned Campus

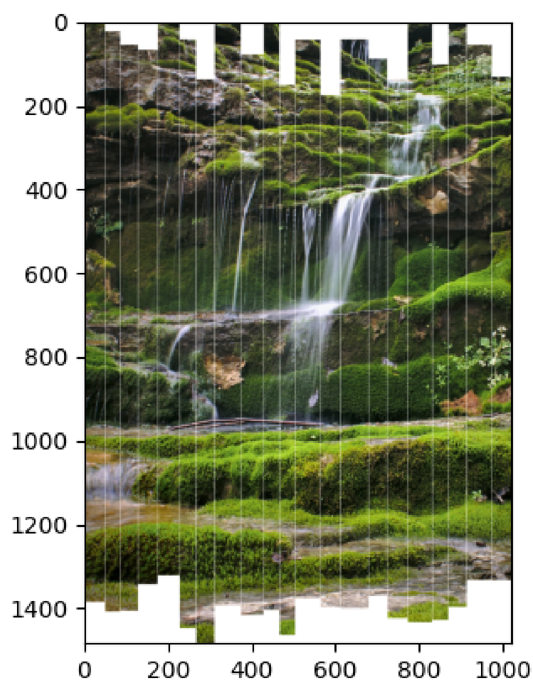

Figure 5: Aligned Nature

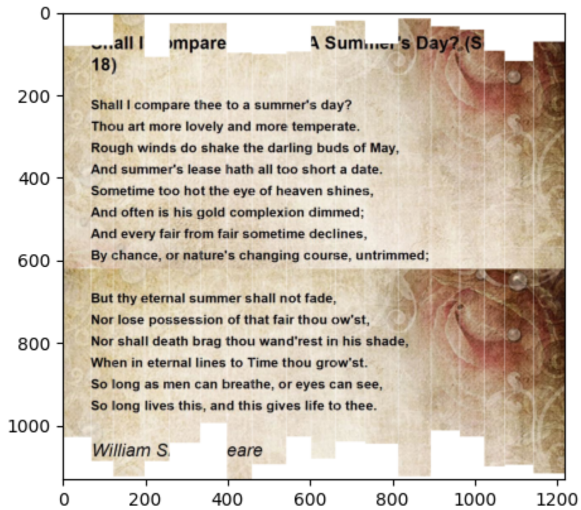

Figure 6: Aligned Sonnet

3.3 Noisy Data

The key modification involved redefining the similarity function. Moving away from the sum of squared distances used previously, we implemented zero mean normalized cross-correlation (ZNCC). This measure, which is essentially the dot product between normalized pixel vectors, is more robust against variations in lighting and coloration. To further counteract noise, we modified our approach to consider more than just a single column of pixels. By averaging the last three columns of each strip, we aimed to mitigate the impact of irregular boundaries and enhance the robustness of our similarity calculations.

As it is done in 3.2, the similarity function was thus adapted to normalize the pixel values of the strips, accounting for zero mean and unit norm, before computing the ZNCC. This modification was necessary to achieve a more accurate assessment of similarity under noisy conditions. The distance function continued to play a crucial role, now leveraging the revised similarity calculations to ascertain the best alignment and order of the strips, even with the added noise and irregularities.

Implementing these changes, we managed to stitch the image strips from the scan dataset together. While the result was not perfect as it is shown in 7. While most shards are rearragend correctly, the left most shred in the Figure is not correctly matched with the body of the statue. Since the current distance metric works perfect on all the other pieces and another test shred set. We suspect that the failure is due to the inherent limitation of greedy-method based matching that failed to find the global optimal order of the shreds. Our observations revealed that certain regions of the image were more challenging to stitch accurately, likely due to more pronounced noise or discoloration, highlighting areas for future improvement in our approach.

References

[1] yes (https://math.stackexchange.com/users/263664/yes). Homogeneous system of equations and svd. Mathematics Stack Exchange. URL:https://math.stackexchange.com/q/4373721 (version: 2022-02-04).

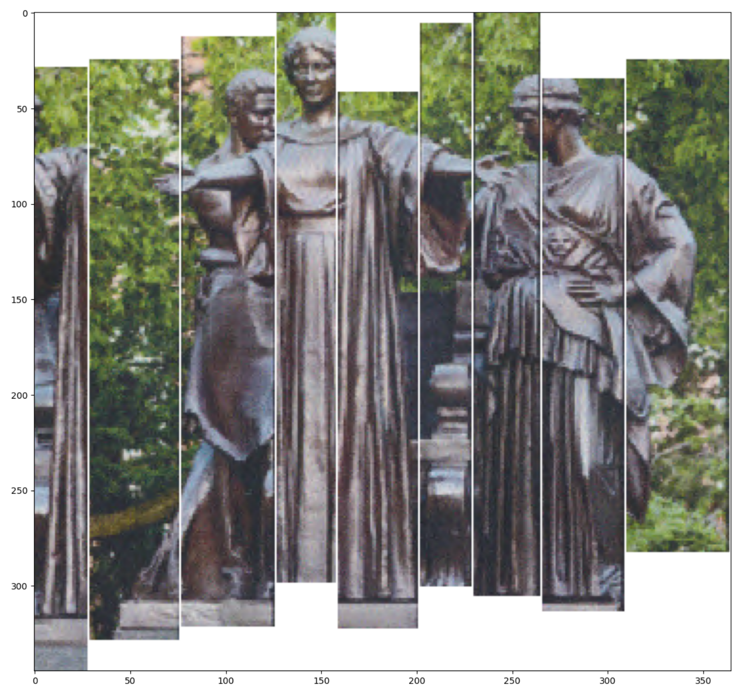

Figure 7: Imperfect Matching Result

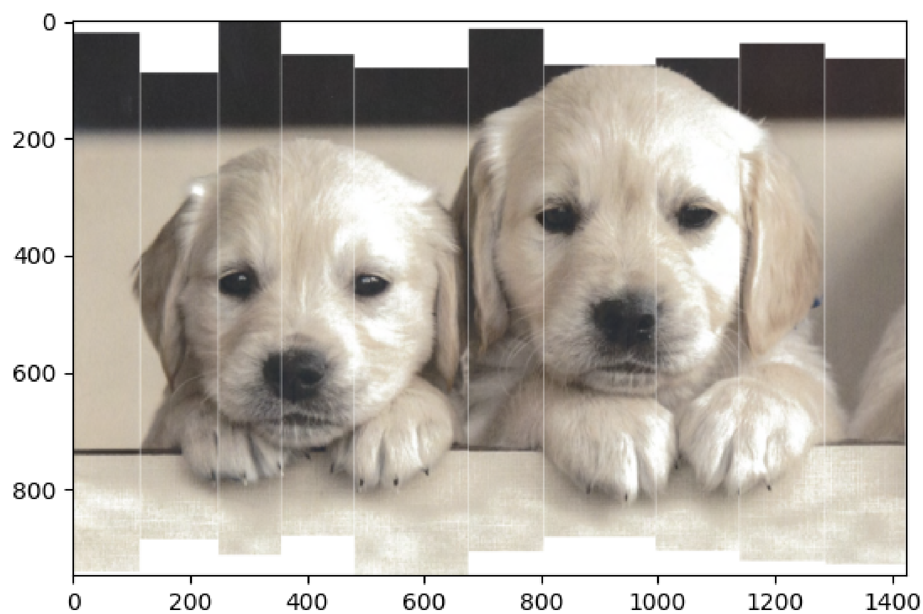

Figure 8: Successed in Matching Noisy Dogs