1 MATLAB Example

Download the ICA MATLAB code from Moodle. Open and run the `ica_image_mix.m` file. The code loads a pair of images, superimposes them, and then unmixes the images. Answer the following questions, and then modify the code to test that your reasoning is correct.

1. Take two greyscale images. Pick one of them and generate an inverted version of it (white becomes black and vice versa). Will this affect the result of ICA when you try to unmix the inverted image with the second image? What happens if you mix the inverted and the original greyscale image?

   Solution

   The MATLAB code is presented in listing 1.

   Listing 1: The MATLAB code for part one.

```
% 1-1-1
f2 = 255 - f2;

% 1-1-2
f1 = 255 - f2;
```

   Inverting one of the images will not significantly affect the outcome of ICA. This is because ICA is sign-invariant; meaning that it is indifferent to whether the signals are positive or negative. ICA's primary goal is to separate statistically independent signals, not to focus on their specific values or signs. Since ICA does not distinguish between a signal and its inverted version (i.e., if the source signal is s_i, ICA treats s_i and $-s_i$ as equivalent), inverting a greyscale image (changing white to black and vice versa) does not impact the separation process. Therefore, when an inverted greyscale image is mixed with a different image, ICA will still be able to successfully separate them. The inversion does not alter the statistical independence between the two images; therefore, ICA will function as expected.

   On the other hand, if you mix an inverted greyscale image with its original version and apply ICA, it will affect the result. Inverting a greyscale image is essentially a linear transformation where all pixel values are multiplied by -1, producing the negative of the original image. Since the inverted and original images are closely related, they are highly correlated. This violates a key assumption of ICA, which requires that the independent components be statistically independent. Therefore, ICA will struggle to properly unmix these two signals due to their inherent dependence.

   ■

2. Is ICA insensitive to the choice of correlation created in the mixing matrix? To test this, modify the seed of the random mixing matrix.

   Solution

   The MATLAB code is shown in listing 2.

   Listing 2: The MATLAB code for part two.

```
% 1-2
rand('seed', 3);
```

   In ICA, the choice of the mixing matrix plays a crucial role in the correlation between observed signals. When using a randomly generated mixing matrix, the resulting observed signals are typically uncorrelated, allowing ICA to effectively separate the independent source signals.

However, in a special case where the mixing matrix introduces strong correlation, this correlation can hinder ICA's ability to accurately separate the signals, making it difficult for the algorithm to recover the original independent components.

■

3. Will ICA be affected if you modify one of the images by adding a vector of random white noise entries?

Solution

The MATLAB code is provided in listing 3.

Listing 3: The MATLAB code for part three.

```
% 1-3
f1 = imnoise(f1);
```

Modifying one of the images by adding a vector of random white noise will affect ICA's performance by increasing the complexity of the observed signals and potentially creating correlations that obscure the underlying independent components. This can lead to difficulties in accurately separating the original signals, resulting in misidentification or failure in recovery. Even if the noise image is statistically independent of the original image, the added variability from the noise complicates the separation process. While ICA may still perform reasonably well with low levels of noise, higher noise levels can significantly diminish the accuracy of the separation.

■

2 Whitening

Recall that in ICA, whitening is done by projecting a zero-mean distribution x through the matrix $V = D^{-\frac{1}{2}}E^\top$, where E is the matrix consisting of the eigenvectors of the covariance matrix of x, and D is a diagonal matrix composed of the eigenvalues of the corresponding eigenvectors in E.

1. Explain how such a projection whitens the data, i.e., how it ensures that the data, once projected, is uncorrelated and has unit variance.

Solution

Whitening is a preprocessing step in ICA that seeks to eliminate the correlations in the data. To this end, consider the projection $y = Vx$, where $V = D^{-\frac{1}{2}}E^\top$ as defined in the question. Denoting the expected value operation by $E\{\cdot\}$, the covariance matrix of the zero-mean projection can be computed as

$$E\{yy^\top\} = E\{Vxx^\top V^\top\} = VE\{xx^\top\}V^\top = D^{-\frac{1}{2}}E^\top E\{xx^\top\}ED^{-\frac{1}{2}}.$$

Using the eigendecomposition $E\{xx^\top\} = EDE^\top$, it can be further deduced that

$$E\{yy^\top\} = D^{-\frac{1}{2}}\underbrace{E^\top E}_{=I}D\underbrace{E^\top E}_{=I}D^{-\frac{1}{2}} = D^{-1}D = I.$$

This shows that the covariance matrix of y is equal to the identity matrix. More specifically, the projected data has unit variance (diagonal elements of I) and is uncorrelated because the non-diagonal elements of I are all zero.

■

2. While this projection is done in one time step, show that $z = Vx$ is a stationary point of the iterative learning rule $\Delta V = \gamma(I - zz^\top)V$.

Solution

Denoting the expected value operation by $E\{\cdot\}$, it holds that

$$\begin{aligned}
E\{\Delta V\} &= E\{\gamma(I - zz^\top)V\} \\
&= \gamma(I - E\{zz^\top\})V \\
&= \gamma(I - E\{Vxx^\top V^\top\})V \\
&= \gamma(I - VE\{xx^\top\}V^\top)V \\
&= \gamma(I - D^{-\frac{1}{2}}E^\top E\{xx^\top\}ED^{-\frac{1}{2}})V.
\end{aligned}$$

Using the eigendecomposition $E\{xx^\top\} = EDE^\top$ in conjunction with the above result will yield

$$E\{\Delta V\} = \gamma(I - D^{-\frac{1}{2}}\underbrace{E^\top E}_{=I}D\underbrace{E^\top E}_{=I}D^{-\frac{1}{2}})V = \gamma(I - D^{-1}D)V = \gamma(I - I)V = 0.$$

■

3 Negentropy

ICA uses a fundamental property of Gaussian distributions to estimate the independent components. This property is that the entropy of a Gaussian distribution is larger than any other distribution with the same mean and variance.

1. For a random variable x and the two associated probability density functions $g(x)$ (corresponding to Gaussian distribution) and $f(x)$ with the same mean and variance, show that the above property is true.

   Hint: Use the fact that the relative entropy $D(f||g)$ of two probability density functions $f(x)$ and $g(x)$ is non-negative, i.e., $D(f||g) = \int_{-\infty}^{\infty} f(x) \ln\left(\frac{f(x)}{g(x)}\right) dx \geq 0$.

   Solution

   Consider $g(x)$ as a Gaussian probability density function (PDF) with mean μ and variance σ^2. Therefore, it holds that

$$g(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}. \tag{1}$$

   Let $f(x)$ be an arbitrary PDF with the same mean and variance. The differential entropy of a random variable x with $f(x)$ PDF is defined as

$$h(f) = -\int_{-\infty}^{\infty} f(x) \ln(f(x)) dx. \tag{2}$$

   Using eqs. (1) and (2), the entropy of the Gaussian distribution can be derived as

$$\begin{aligned}
h(g) &= -\int_{-\infty}^{\infty} g(x) \ln(g(x)) dx \\
&= -\int_{-\infty}^{\infty} g(x) \ln\left(\frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}\right) dx \\
&= -\int_{-\infty}^{\infty} g(x) \left(\ln\left(\frac{1}{\sqrt{2\pi\sigma^2}}\right) - \frac{(x-\mu)^2}{2\sigma^2}\right) dx \\
&= -\ln\left(\frac{1}{\sqrt{2\pi\sigma^2}}\right) \underbrace{\int_{-\infty}^{\infty} g(x) dx}_{=1} + \frac{1}{2\sigma^2} \underbrace{\int_{-\infty}^{\infty} g(x)(x-\mu)^2 dx}_{=\sigma^2} \\
&= \frac{1}{2}\ln(2\pi\sigma^2) + \frac{1}{2} \\
&= \frac{1}{2}\ln(2\pi e\sigma^2).
\end{aligned} \tag{3}$$

   The definition of the relative entropy given in the hint can be further expanded into

$$\begin{aligned}
& 0 \leq D(f||g) = \int_{-\infty}^{\infty} f(x) \ln\left(\frac{f(x)}{g(x)}\right) dx \\
\Leftrightarrow\ & 0 \leq D(f||g) = \int_{-\infty}^{\infty} f(x) \ln(f(x)) dx - \int_{-\infty}^{\infty} f(x) \ln(g(x)) dx \\
\Leftrightarrow\ & 0 \leq D(f||g) = -h(f) - \int_{-\infty}^{\infty} f(x) \ln(g(x)) dx.
\end{aligned} \tag{4}$$

Moreover, note that

$$
\begin{aligned}
\int_{-\infty}^{\infty} f(x)\ln(g(x))dx &= \int_{-\infty}^{\infty} f(x)\ln\left(\frac{1}{\sqrt{2\pi\sigma^2}}e^{-\frac{(x-\mu)^2}{2\sigma^2}}\right)dx \\
&= \int_{-\infty}^{\infty} f(x)\left(\ln\left(\frac{1}{\sqrt{2\pi\sigma^2}}\right)-\frac{(x-\mu)^2}{2\sigma^2}\right)dx \\
&= \ln\left(\frac{1}{\sqrt{2\pi\sigma^2}}\right)\underbrace{\int_{-\infty}^{\infty} f(x)dx}_{=1} - \frac{1}{2\sigma^2}\underbrace{\int_{-\infty}^{\infty} f(x)(x-\mu)^2dx}_{=\sigma^2 \text{ (same variance as } g(x))} \\
&= -\frac{1}{2}\ln(2\pi\sigma^2)-\frac{1}{2} \\
&= -\frac{1}{2}\ln(2\pi e\sigma^2) \\
&= -h(g).
\end{aligned}
\tag{5}
$$

Putting eqs. (4) and (5) together, it follows that

$$0 \leq D(f||g) = h(g) - h(f); \tag{6}$$

meaning that Gaussian distribution has the maximum entropy among all the other distributions with the same mean and variance.

■

2. Show that the negentropy is thus always non-negative and discuss what this means for ICA.

Solution

The negentropy is given by $J(x) = h(g) - h(f)$, and is always non-negative as shown in eq. (6). In ICA, the goal is to maximize the non-Gaussianity among the independent components. In other words, the aim is to find the maximum negentropy because, the larger the negentropy, the further away the distribution is from that of a Gaussian.

■