

Lecture 13: February 13, 2025

Lecturer: Madhur Tulsiani

1 Matrix Scaling

We will consider an application of I-projections to a problem known as matrix scaling. Say we are given two nonnegative matrices $M, N \in \mathbb{R}_+^{n \times n}$ such that for all i, j , $M_{ij} = 0 \Leftrightarrow N_{ij} = 0$. The goal is to multiply (scale) each row i of M by a number r_i and each column j by c_j , such that the resulting matrix M' has the same row and column sums as the target matrix N . Another way of stating this is that we want to find diagonal matrices D_1 and D_2 such that for $M' = D_1 M D_2$, we have

$$\sum_j M'_{ij} = \sum_j N_{ij} \quad \forall i \in [n] \quad \text{and} \quad \sum_i M'_{ij} = \sum_i N_{ij} \quad \forall j \in [n].$$

We will show a special case when the goal is to scale M so that the resulting matrix M' is doubly stochastic i.e.,

$$\sum_j M'_{ij} = \sum_i M'_{ij} = 1 \quad \forall i, j \in [n].$$

First, note that by a *global* scaling of $1 / \sum_{i,j} M_{ij}$, we can assume that $\sum_{i,j} M_{ij} = 1$, and the goal is instead to scale it to have row and column sums equal to $1/n$ i.e.,

$$\sum_j M'_{ij} = \sum_i M'_{ij} = \frac{1}{n} \quad \forall i, j \in [n].$$

We can now think of this as a problem of going from one distribution to another. Assume that $M_{ij} > 0$ for all i, j , and think of the target matrix N with $N_{ij} = 1/n^2$ for all i, j . Since the entries of M are positive and sum to 1, we can think of it as a probability distribution Q with $\text{Supp}(Q) = [n] \times [n]$ (where $q(i, j) = M_{ij}$). We consider the linear family of distributions on $[n] \times [n]$ (written as matrices) with the required row and column sums.

$$\mathcal{L} := \left\{ P \mid \sum_j p(i, j) = \sum_i p(i, j) = \frac{1}{n} \quad \forall i, j \in [n] \right\}$$

Note that the above is a linear family as defined in the previous lecture, since we can consider functions f_1, \dots, f_n and g_1, \dots, g_n defined as

$$f_{i_0}(i, j) = \begin{cases} 1 & \text{if } i = i_0 \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad g_{j_0}(i, j) = \begin{cases} 1 & \text{if } j = j_0 \\ 0 & \text{otherwise} \end{cases}.$$

Then, the above family can be written in terms of the expectations of the functions f_i and g_j for all $i, j \in [n]$. Moreover, we know from the previous lecture that the I-projection P^* of Q onto \mathcal{L} is of the form

$$\begin{aligned} p^*(i, j) &= c_0 \cdot q(i, j) \cdot \exp \left(\sum_{i_0} \lambda_{i_0} \cdot f_{i_0}(i, j) + \sum_{j_0} \mu_{j_0} \cdot g_{j_0}(i, j) \right) \\ &= c_0 \cdot q(i, j) \cdot \exp(\lambda_i + \mu_j) \\ &= (\sqrt{c_0} \cdot \exp(\lambda_i)) \cdot M_{ij} \cdot (\sqrt{c_0} \cdot \exp(\mu_j)) . \end{aligned}$$

Thus, we can define the diagonal matrices D_1 and D_2 as

$$(D_1)_{ii} = \sqrt{c_0} \cdot \exp(\lambda_i) \quad \text{and} \quad (D_2)_{jj} = \sqrt{c_0} \cdot \exp(\mu_j) .$$

We then have that the distribution p^* given by the resulting matrix $M' = D_1 M D_2$, belongs to the linear family \mathcal{L} . Thus, the row and column sums of M' are $1/n$. Combining this with another global scaling (replace $\sqrt{c_0}$ by $\sqrt{c_0 \cdot n}$) we can also get all the row and column sums to be 1 (i.e., make the matrix doubly stochastic).


Exercise 1.1. Where did we use the fact that $M_{ij} > 0$ for all $i, j \in [n]$?

Exercise 1.2. Use this the above techniques to solve the matrix scaling problem for an arbitrary target matrix N (assuming $M_{ij} = 0 \Leftrightarrow N_{ij} = 0$).

Matrix scaling and its generalization, known as operator scaling have found a variety of applications in combinatorial optimization, complexity theory and analysis. Please take a look at the recent tutorial by Wigderson [Wig17] for an introduction to many of these connections.

2 Error-Correcting Codes

Over the next few lectures, we will switch to the “coding theory” part of the course and see how to construct (and work with) error correcting codes. We first define the model(s) we will be working with, and consider the tasks that we will aim for.



As illustrated above, the goal is to communicate a message, which we will take to be one of M possible values in $[M] = \{1, \dots, M\}$. The problem is that the transmission goes through a “channel” which introduces some form of noise. The goal is then to “encode” the message as a length- n string in a finite alphabet (say) \mathcal{X} , in such a way that even after the noise, the received noisy transmission (say in \mathcal{Y}^n) can be decoded to the intended message (maybe with high probability). A code is thus specified by two (not necessarily efficiently computable) maps $Enc: [M] \rightarrow \mathcal{X}^n$ and $Dec: \mathcal{Y}^n \rightarrow [M]$. The specific requirements from these maps, depend on the error model we consider.

Shannon model. In this setting, we will think of the message being a random variable W chosen uniformly randomly in $[M]$, and the errors introduced by the channel as being the result of a probabilistic process. The goal will be design (deterministic) maps Enc and Dec such that

$$\mathbb{P} [\text{Dec}(\text{Noisy-Transmission}(\text{Enc}(W))) = W] \rightarrow 1,$$

where the probability is over the choice of W , and the errors introduced by the noisy transmission through the channel. In particular, we will limit our discussion to *discrete memoryless channels* where the “discrete” part refers to \mathcal{X} and \mathcal{Y} being finite, and the “memoryless” property refers to the fact that the noise acts independently on each of the n symbols transmitted, and does not depend on what happens to the previous symbols. The channel is thus specified by a collection of probability distributions $P(Y|X)$ where X and Y refer to the input and output of the channel for a *single* transmission.

Hamming model. Here we think of the message as an arbitrary element of $[M]$, and the errors introduced by the channel as adversarial. Of course, we need some assumption, as otherwise all transmissions can be mapped to a single output string with no hope of recovery. We will thus assume a bound on the *number of positions* where $\bar{x} \in \mathcal{X}^n$ is corrupted. In particular, we take $\mathcal{Y} = \mathcal{X}$, but assume $\bar{x} \in \mathcal{X}^n$ is corrupted in at most t positions, to give the output $\bar{y} \in \mathcal{X}^n$. We will want codes which satisfy

$$\forall w \in [M] \quad \text{Dec}(\text{Noisy-Transmission}(\text{Enc}(w))) = w,$$

if the noisy transmission corrupts the string $\bar{x} = \text{Enc}(W)$ in at most t (arbitrary) positions. We will first discuss the Shannon model of errors for a couple of lectures, before considering codes for the Hamming model.

2.1 Discrete memoryless channels

As discussed above, a discrete memoryless channel is specified by a collection of probability distributions, or a transition matrix, specifying the probabilities $\mathbb{P}[Y = y|X = x]$, where $x \in \mathcal{X}$, $y \in \mathcal{Y}$, and X and Y are random variables corresponding respectively to the input and the output. The limits on how efficiently we can transmit data through a channel, will be specified by a quantity called the channel capacity, defined as

$$C := \max_{P(X)} I(X;Y),$$

where we maximize over all the distributions $P(X)$ for the input X , and the joint distribution of (X, Y) is then given by the transition matrix corresponding to the channel. Consider the following examples.

Example 2.1 (Noiseless channel). We take $\mathcal{X} = \mathcal{Y} = \{0, 1\}$, and have $\mathbb{P}[Y = b_1|X = b_2] = 1$ if $b_1 = b_2$ and 0 otherwise i.e., symbols are transmitted without any corruption. It is easy to see that $Y = X$ and

$$C = \max_{P(X)} I(X;Y) = \max_{P(X)} H(X) = 1.$$

Example 2.2 (Noiseless random channel). Let $\mathcal{X} = \{0, 1\}$, $\mathcal{Y} = \{a, b, c, d\}$ with $\mathbb{P}[Y = a|X = 0] = \mathbb{P}[Y = b|X = 0] = 1/2$, $\mathbb{P}[Y = c|X = 0] = \mathbb{P}[Y = d|X = 0] = 1/2$, and all other probabilities being 0. Thus, the value of Y still uniquely specifies X , and we have

$$C = \max_{P(X)} I(X;Y) = \max_{P(X)} \{H(X) - H(X|Y)\} = \max_{P(X)} H(X) = 1.$$

Example 2.3 (Binary symmetric channel: BSC(p)). We take $\mathcal{X} = \mathcal{Y} = \{0, 1\}$ and

$$\mathbb{P}[Y = b_1|X = b_2] = \begin{cases} 1 - p & \text{if } b_1 = b_2 \\ p & \text{if } b_1 \neq b_2 \end{cases}.$$

From the symmetry of the above transition matrix, we can say that $Y = X + Z \pmod 2$ where $Z = 1$ with probability p and 0 with probability $1 - p$. Using this, we get

$$I(X;Y) = H(Y) - H(Y|X) = H(X + Z) - H((X + Z)|X) = H(X + Z) - H_2(p).$$

Since $H(X + Z) \leq 1$ and is maximized when $P(X)$ is uniform over $\{0, 1\}$, we get that for the binary symmetric channel, $C = 1 - H_2(p)$.

Exercise 2.4 (Binary erasure channel). Take $\mathcal{X} = \{0, 1\}$, $\mathcal{Y} = \{0, 1, \perp\}$ and define the transition probabilities as

$$\mathbb{P}[Y = b_1 | X = b_2] = \begin{cases} 1 - p & \text{if } b_1 = b_2 \\ p & \text{if } b_1 = \perp \end{cases},$$

i.e., each symbol in $\{0, 1\}$ is transmitted without error with probability $1 - p$, and is converted to \perp (erasure) with probability p (but 0 is never converted to 1 , and vice-versa). Prove that $C = 1 - p$.

2.2 Channel coding theorem

We next discuss Shannon's channel coding theorem which says that the maximum achievable efficiency for any code is given by the channel capacity. To measure the efficiency of a code, we need a few definitions.

Recall that a code is specified by maps $\text{Enc} : [M] \rightarrow \mathcal{X}^n$ and $\text{Dec} : \mathcal{Y}^n \rightarrow [M]$. We will also take the map Enc to be injective so that no two messages have the same encoding. We can think of this process as the following Markov chain

$$W \rightarrow (\text{Enc}(W) = \bar{\mathbf{X}}) \rightarrow \bar{\mathbf{Y}} \rightarrow (\hat{W} = \text{Dec}(\bar{\mathbf{Y}})),$$

where $\bar{\mathbf{X}} \in \mathcal{X}^n$ and $\bar{\mathbf{Y}} \in \mathcal{Y}^n$ denote the input and output sequences for the channel. We take the probability of error for a code to be $p_e = \mathbb{P}[W \neq \hat{W}]$, where the probability is over the choice of the message, and the noise in the channel.

Definition 2.5. We define the rate of a code as above to be

$$R := \frac{\log M}{n \cdot \log |\mathcal{X}|} \quad (\text{bits per transmission}).$$

We say that a rate R is achievable for a channel, if there exists a sequence of codes for $n \geq n_0$ with rates at least R and error probabilities $p_e^{(n)}$ such that $p_e^{(n)} \rightarrow 0$ as $n \rightarrow \infty$. We define the maximum achievable rate for a channel as $R^* = \sup \{R \mid R \text{ is achievable}\}$.

Remark 2.6. Often achievable rates are defined with respect to the maximum probability of error $\lambda_e = \max_{w \in [M]} \mathbb{P}[\hat{W} \neq W \mid W = w]$, where the probability is taken only over the noise in the

channel. However, it will be slightly more convenient to work with the average error probability p_e for us. Moreover, you can check that R^* does not change according to the two notions. In particular, check that if there exists a sequence of codes with rate at least R such that $p_e^{(n)} \rightarrow 0$, then for every $\varepsilon > 0$, there exists a sequence of codes with rate at least $R - \varepsilon$ such that $\lambda_e^{(n)} \rightarrow 0$ (simply discard messages for which probability of error larger than $2p_e$).

We can now state Shannon's channel coding theorem for discrete memoryless channels.

Theorem 2.7. *For any discrete memoryless channel, we have $R^* = C$.*

While we will prove that $R^* \leq C$ for every channel, we will only prove $R^* \geq C$ for the binary symmetric channel. The proof idea for the case of general channels is similar and (using a random code) but the analysis a bit more cumbersome.

2.3 Channel capacity as an upper bound on achievable rates

We first prove the following.

Proposition 2.8. *Let R be any achievable rate for a given channel with capacity C . Then, $R \leq C$.*

Proof: Since R is achievable, there exists a sequence of codes with encoding lengths (also called block-lengths) $n \geq n_0$ and rate at least R , such that $p_e^{(n)} \rightarrow 0$. Consider any such code with block-length n and rate $\log(M)/n \geq R \Rightarrow M \geq 2^{nR}$. Recall that we think of the transmission process as the Markov chain

$$W \rightarrow (\text{Enc}(W) = \bar{\mathbf{X}}) \rightarrow \bar{\mathbf{Y}} \rightarrow (\hat{W} = \text{Dec}(\bar{\mathbf{Y}})) .$$

Applying Fano's inequality gives that for $p_e^{(n)} = \mathbb{P}[\hat{W} \neq W]$, we have

$$1 + p_e^{(n)} \cdot \log(M) \geq H_2(p_e^{(n)}) + p_e^{(n)} \cdot \log(\text{Supp}(W) - 1) \geq H(W|\hat{W})$$

We can write the RHS of the inequality above as

$$H(W|\hat{W}) = H(W) - I(W; \hat{W}) \geq \log(M) - I(W; \bar{\mathbf{Y}}) = nR - I(\bar{\mathbf{X}}; \bar{\mathbf{Y}}) .$$

Finally, we can analyze the mutual information term as follows (using $Y_{<i}$ as a shorthand for Y_1, \dots, Y_{i-1})

$$\begin{aligned}
 I(\bar{\mathbf{X}}; \bar{\mathbf{Y}}) &= H(\bar{\mathbf{Y}}) - H(\bar{\mathbf{Y}}|\bar{\mathbf{X}}) \\
 &= \sum_{i=1}^n (H(Y_i | Y_{<i}) - H(Y_i | Y_{<i}, \bar{\mathbf{X}})) \\
 &\leq \sum_{i=1}^n (H(Y_i) - H(Y_i | X_i)) \\
 &= \sum_{i=1}^n I(X_i; Y_i) \\
 &\leq n \cdot C
 \end{aligned}$$

Note that we used above that $H(Y_i | Y_{<i}, \bar{\mathbf{X}}) = H(Y_i | X_i)$ by the memoryless property of the channel, and that for all i , $I(X_i; Y_i) \leq C$. Combining the above bounds, we get

$$1 + p_e^{(n)} \cdot \log(M) \geq \log(M) - nC \quad \Rightarrow \quad R \cdot (1 - p_e^{(n)}) \leq C + \frac{1}{n}.$$

Taking the limit $n \rightarrow \infty$ then gives $R \leq C$ as desired. ■

References

[Wig17] Avi Wigderson, *Operator scaling: theory, applications and connections*, 2017, Tutorial given at CCC 2017. 2