

Mathematical Optimization Models and Applications II

Yinyu Ye

Department of Management Science and Engineering

Stanford University

Stanford, CA 94305, U.S.A.

<http://www.stanford.edu/~yyye>

Chapters 1, 2.1-2, 6.1-2, 7.2, 11.3, 11.6

Unconstrained Optimization: Logistic Regression I

Similar to SVM, given the two-class **discrimination training data** points $\mathbf{a}_i \in \mathbb{R}^n$, according to the logistic model, the probability that it's in a class C , say in **Red**, is represented by a linear/affine function with slope-vector \mathbf{x} and intersect scalar x_0 :

$$\frac{e^{\mathbf{a}_i^T \mathbf{x} + x_0}}{1 + e^{\mathbf{a}_i^T \mathbf{x} + x_0}}.$$

Thus, for some training data points, we like to determine intercept x_0 and slope vector $\mathbf{x} \in \mathbb{R}^n$ such that

$$\frac{e^{\mathbf{a}_i^T \mathbf{x} + x_0}}{1 + e^{\mathbf{a}_i^T \mathbf{x} + x_0}} = \begin{cases} 1, & \text{if } \mathbf{a}_i \in C \\ 0, & \text{otherwise} \end{cases}.$$

Then the probability to give a “right classification answer” for all training data points is

$$\left(\prod_{\mathbf{a}_i \in C} \frac{e^{\mathbf{a}_i^T \mathbf{x} + x_0}}{1 + e^{\mathbf{a}_i^T \mathbf{x} + x_0}} \right) \left(\prod_{\mathbf{a}_i \notin C} \frac{1}{1 + e^{\mathbf{a}_i^T \mathbf{x} + x_0}} \right)$$

Logistic Regression II

Therefore, we like to maximize the probability when deciding intercept x_0 and slope vector $\mathbf{x} \in R^n$

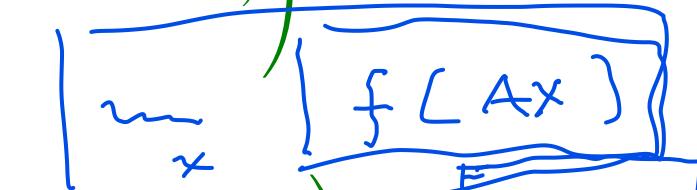
$$\left(\prod_{\mathbf{a}_i \in C} \frac{e^{\mathbf{a}_i^T \mathbf{x} + x_0}}{1 + e^{\mathbf{a}_i^T \mathbf{x} + x_0}} \right) \left(\prod_{\mathbf{a}_i \notin C} \frac{1}{1 + e^{\mathbf{a}_i^T \mathbf{x} + x_0}} \right) = \left(\prod_{\mathbf{a}_i \in C} \frac{1}{1 + e^{-\mathbf{a}_i^T \mathbf{x} - x_0}} \right) \left(\prod_{\mathbf{a}_i \notin C} \frac{1}{1 + e^{\mathbf{a}_i^T \mathbf{x} + x_0}} \right),$$

which is equivalently to maximize

$$-\left(\sum_{\mathbf{a}_i \in C} \ln(1 + e^{-\mathbf{a}_i^T \mathbf{x} - x_0}) \right) - \left(\sum_{\mathbf{a}_i \notin C} \ln(1 + e^{\mathbf{a}_i^T \mathbf{x} + x_0}) \right).$$

Or

$$\min_{x_0, \mathbf{x}} \left(\sum_{\mathbf{a}_i \in C} \ln(1 + e^{-\mathbf{a}_i^T \mathbf{x} - x_0}) \right) + \left(\sum_{\mathbf{a}_i \notin C} \ln(1 + e^{\mathbf{a}_i^T \mathbf{x} + x_0}) \right).$$



This is an **unconstrained optimization** problem, where the objective is a convex function of decision variables: intercept x_0 and slope vector $\mathbf{x} \in R^n$.

Sparse Linear Regression Problems

We want to find a sparsest solution to fit exact data measurements, that is, to minimize the number of non-zero entries in \mathbf{x} such that $A\mathbf{x} = \mathbf{b}$:

$$\begin{aligned} & \text{minimize} && \|\mathbf{x}\|_0 = |\{j : x_j \neq 0\}| \\ & \text{subject to} && A\mathbf{x} = \mathbf{b}. \end{aligned}$$

Sometimes this objective can be accomplished by LASSO:

$$\begin{aligned} & \text{minimize} && \|\mathbf{x}\|_1 = \sum_{j=1}^n |x_j| \\ & \text{subject to} && A\mathbf{x} = \mathbf{b}. \end{aligned}$$

It can be equivalently represented by (?)

$$\begin{aligned} & \text{minimize} && \sum_{j=1}^n y_j \\ & \text{subject to} && A\mathbf{x} = \mathbf{b}, -\mathbf{y} \leq \mathbf{x} \leq \mathbf{y}; \end{aligned}$$

or

$$\begin{aligned} & \text{minimize} && \sum_{j=1}^n (x'_j + x''_j) \\ & \text{subject to} && A(\mathbf{x}' - \mathbf{x}'') = \mathbf{b}, \mathbf{x}' \geq \mathbf{0}, \mathbf{x}'' \geq \mathbf{0}. \end{aligned}$$

Both are **linear programs**!

Sparsest Data Fitting continued

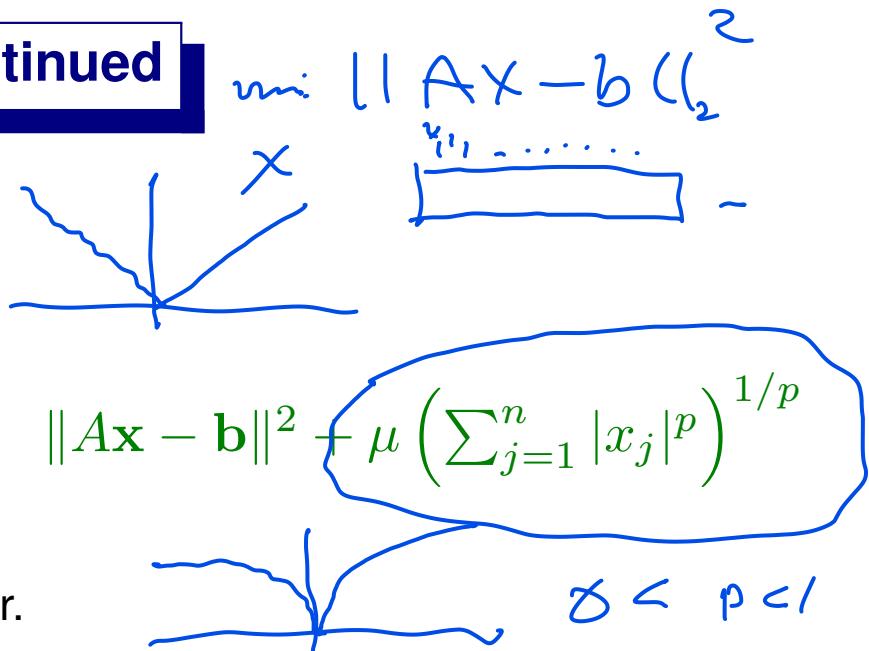
A better approximation of the objective can be accomplished by

$$\begin{aligned} \text{minimize} \quad & \| \mathbf{x} \|_p := \left(\sum_{j=1}^n |x_j|^p \right)^{1/p} \\ \text{subject to} \quad & A\mathbf{x} = \mathbf{b}; \end{aligned}$$

or

$$\text{minimize} \quad \|A\mathbf{x} - \mathbf{b}\|^2 + \mu \left(\sum_{j=1}^n |x_j|^p \right)^{1/p}$$

for some $0 < p < 1$, where $\mu > 0$ is a regularization parameter.



Or simply

$$\begin{aligned} \text{minimize} \quad & \| \mathbf{x} \|_p^p := \left(\sum_{j=1}^n |x_j|^p \right)^p \\ \text{subject to} \quad & A\mathbf{x} = \mathbf{b}; \end{aligned} \quad \text{or} \quad \text{minimize} \quad \|A\mathbf{x} - \mathbf{b}\|^2 + \beta \left(\sum_{j=1}^n |x_j|^p \right);$$

where the former is a **linearly constrained (nonconvex) optimization problem** and the latter is an **unconstrained (nonconvex) optimization problem**

Akret

Quadratic Programming (QP): Portfolio Management

For expected return vector \mathbf{r} and co-variance matrix \mathbf{V} of an investment portfolio, one management model is:

minimize $\mathbf{x}^T \mathbf{V} \mathbf{x} - \mathbf{r}^T \mathbf{x}$
 subject to $\mathbf{r}^T \mathbf{x} \geq \mu$,
 $\mathbf{e}^T \mathbf{x} = 1$, $\mathbf{x} \geq \mathbf{0}$,

where \mathbf{e} is the vector of all ones.

This is a (convex) quadratic program.

or simply

minimize $\mathbf{x}^T \mathbf{V} \mathbf{x} - \rho \mathbf{r}^T \mathbf{x}$
 subject to $\mathbf{r}^T \mathbf{x} \geq \mu$,
 $\mathbf{e}^T \mathbf{x} = 1$,
 $x_1 + \dots + x_n = 1$

More CLP Examples: Robust Portfolio Management

In applications, \mathbf{r} and V may be estimated under various scenarios, say \mathbf{r}_i and V_i for $i = 1, \dots, m$. Then, we like

$$\begin{array}{ll}
 \text{minimize} & \max_i \mathbf{x}^T V_i \mathbf{x} \\
 \text{subject to} & \min_i \mathbf{r}_i^T \mathbf{x} \geq \mu, \\
 & \mathbf{e}^T \mathbf{x} = 1, \mathbf{x} \geq \mathbf{0}.
 \end{array}
 \Rightarrow
 \begin{array}{ll}
 \text{minimize} & \alpha \\
 \text{subject to} & \mathbf{r}_i^T \mathbf{x} \geq \mu, \forall i \\
 & \sqrt{\mathbf{x}^T V_i \mathbf{x}} \leq \alpha, \forall i \\
 & \mathbf{e}^T \mathbf{x} = 1, \mathbf{x} \geq \mathbf{0}.
 \end{array}$$

This is a quadratically constrained quadratic program (QCQP). If factorize $V_i = R_i^T R_i$ and let $\mathbf{y}_i = R_i \mathbf{x}$, we can rewrite the problem as

$$\begin{array}{ll}
 \text{minimize} & \alpha \\
 \text{subject to} & \mathbf{r}_i^T \mathbf{x} \geq \mu, \mathbf{y}_i - R_i \mathbf{x} = \mathbf{0}, \forall i \\
 & \|\mathbf{y}_i\| \leq \alpha, \forall i, \mathbf{e}^T \mathbf{x} = 1, \mathbf{x} \geq \mathbf{0},
 \end{array}$$

MDSEK

which is an SOCP with additional benefits.

Portfolio Selection Problem

If no more than k stocks can be selected into your portfolio as a policy constraint?

$$\begin{aligned} & \text{minimize} && \mathbf{x}^T V \mathbf{x} \\ & \text{subject to} && \mathbf{r}^T \mathbf{x} \geq \mu, \\ & && \mathbf{e}^T \mathbf{x} = 1, \\ & && \mathbf{0} \leq \mathbf{x} \leq \mathbf{y}, \mathbf{e}^T \mathbf{y} \leq k, \mathbf{y} \in \{0, 1\}^n \end{aligned}$$

$n \gg k$

This is a mixed-integer quadratic program (MIP).

If the integer variables are restricted 0 or 1, it is also names as the binary optimization problem.

Graph Realization and Sensor Network Localization

Given a graph $G = (V, E)$ and sets of non-negative **weights**, say $\{d_{ij} : (i, j) \in E\}$, the goal is to compute a **realization** of G in the **Euclidean space** \mathbf{R}^d for a **given low dimension** d , where the distance information is preserved.

More precisely: given anchors $\mathbf{a}_k \in \mathbf{R}^d$, $d_{ij} \in N_x$, and $\hat{d}_{kj} \in N_a$, find $\mathbf{x}_i \in \mathbf{R}^d$ such that

$$\mathbf{x} \in \mathbf{R}^2 \quad \left\{ \begin{array}{l} \|\mathbf{x}_i - \mathbf{x}_j\|^2 = d_{ij}^2, \forall (i, j) \in N_x, i < j, \\ \|\mathbf{a}_k - \mathbf{x}_j\|^2 = \hat{d}_{kj}^2, \forall (k, j) \in N_a. \end{array} \right.$$

This is a set of **Quadratic Equations**, which can be represented as an optimization problem:

$$\min_{\mathbf{x}_i \forall i} \sum_{(i, j) \in N_x} (\|\mathbf{x}_i - \mathbf{x}_j\|^2 - d_{ij}^2)^2 + \sum_{(k, j) \in N_a} (\|\mathbf{a}_k - \mathbf{x}_j\|^2 - \hat{d}_{kj}^2)^2.$$

Does the system have a **localization** or **realization** of all \mathbf{x}_j 's? Is the localization **unique**? Is there a **certification** for the solution to make it **reliable** or **trustworthy**? Is the system **partially** localizable with a **certification**?

It can be relaxed to **SOCP** (change “=” to “ \leq ”) or **SDP**.

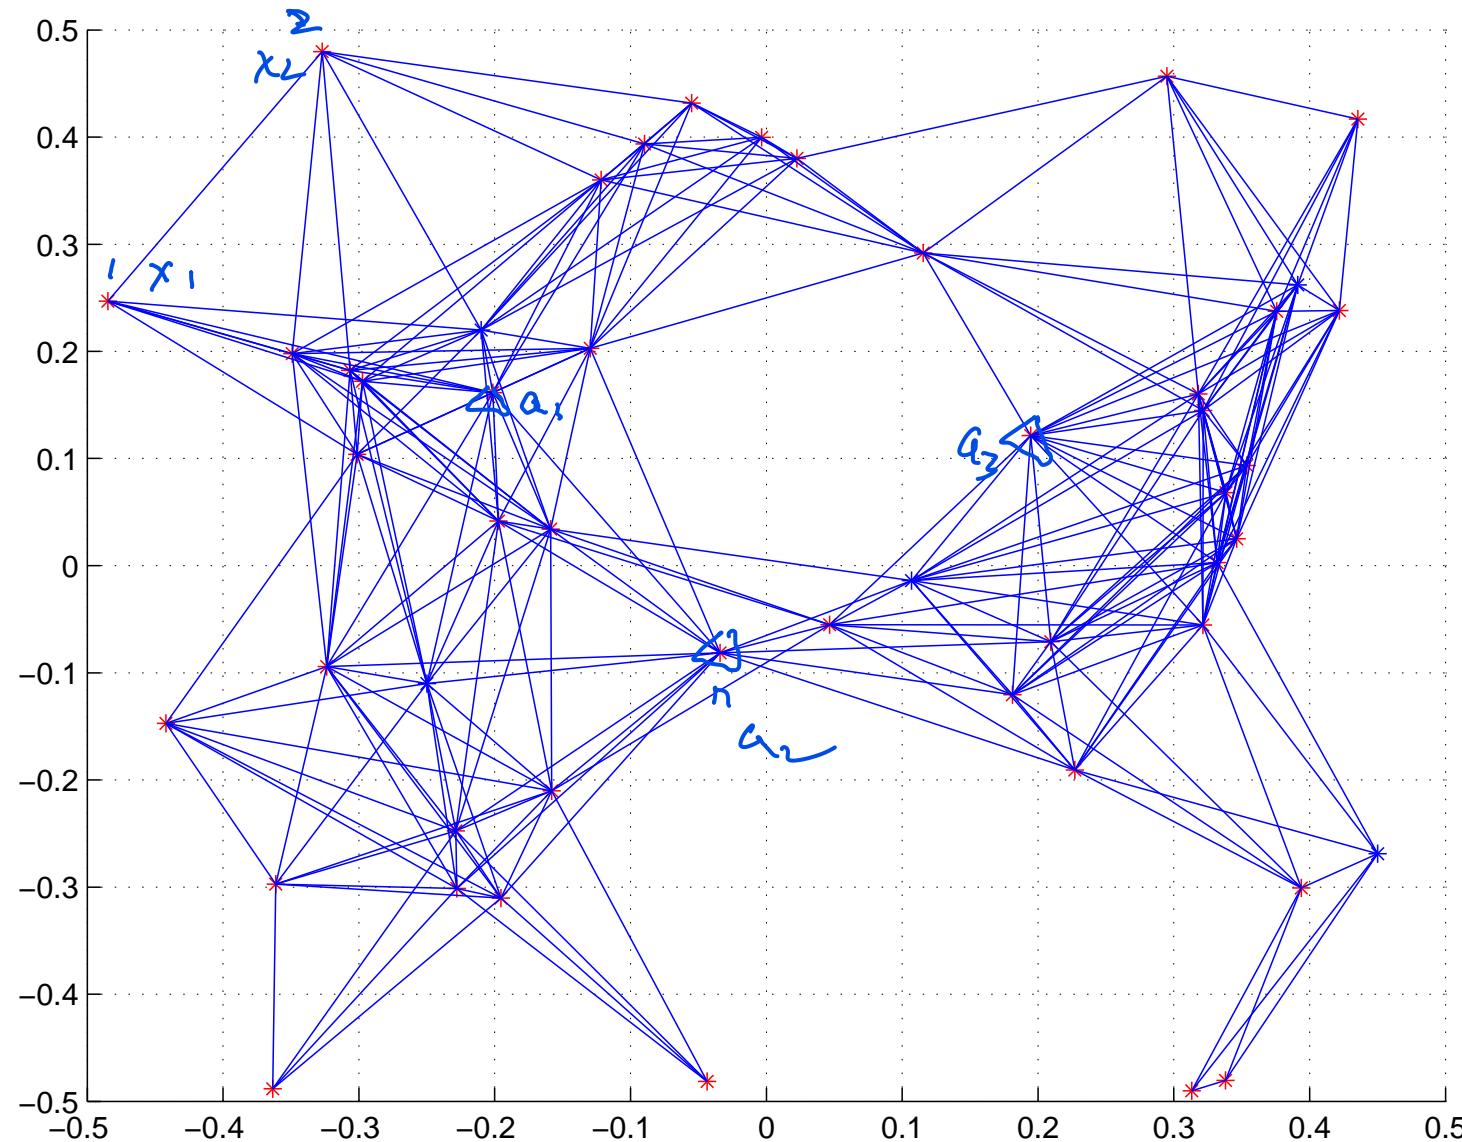


Figure 1: 50-node 2-D Sensor Localization.

Matrix Representation of SNL and SDP Relaxation

Let $X = [\mathbf{x}_1 \ \mathbf{x}_2 \ \dots \ \mathbf{x}_n]$ be the $d \times n$ matrix that needs to be determined and \mathbf{e}_j be the vector of all zero except 1 at the j th position. Then

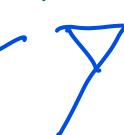
$$\underbrace{\mathbf{x}_i - \mathbf{x}_j}_{\text{wavy line}} = X(\mathbf{e}_i - \mathbf{e}_j) \quad \text{and} \quad \mathbf{a}_k - \mathbf{x}_j = [I \ X](\mathbf{a}_k; -\mathbf{e}_j)$$

so that

$$\Rightarrow \|\mathbf{x}_i - \mathbf{x}_j\|^2 = (\mathbf{e}_i - \mathbf{e}_j)^T X^T X (\mathbf{e}_i - \mathbf{e}_j)$$

$$\|\mathbf{a}_k - \mathbf{x}_j\|^2 = (\mathbf{a}_k; -\mathbf{e}_j)^T [I \ X]^T [I \ X] (\mathbf{a}_k; -\mathbf{e}_j) =$$

$$(\mathbf{a}_k; -\mathbf{e}_j)^T \begin{pmatrix} I & X \\ X^T & \underbrace{X^T X}_{\text{circled}} \end{pmatrix} (\mathbf{a}_k; -\mathbf{e}_j).$$



Or, equivalently, \rightsquigarrow *EVOS*

$$(\mathbf{e}_i - \mathbf{e}_j)^T Y (\mathbf{e}_i - \mathbf{e}_j) = d_{ij}^2, \forall i, j \in N_x, i < j,$$

$$(\mathbf{a}_k; -\mathbf{e}_j)^T \begin{pmatrix} I & X \\ X^T & Y \end{pmatrix} (\mathbf{a}_k; -\mathbf{e}_j) = \hat{d}_{kj}^2, \forall k, j \in N_a,$$

$$Y \preceq X^T X.$$

Relax $Y = X^T X$ to $Y \succeq X^T X$, which is equivalent to matrix inequality:

$$-Y - X^T X \succeq 0. \quad \begin{pmatrix} I & X \\ X^T & Y \end{pmatrix} \succeq 0. \quad \text{rank}(Y) = 2$$

This matrix has **rank** at least d ; if it's d , then $Y = X^T X$, and the converse is also true.

The problem is now an SDP problem: when the SDP relaxation is **exact**?

Algorithm: **Convex relaxation** first and **steepest-descent-search** second strategy?

Stochastic Optimization and Learning

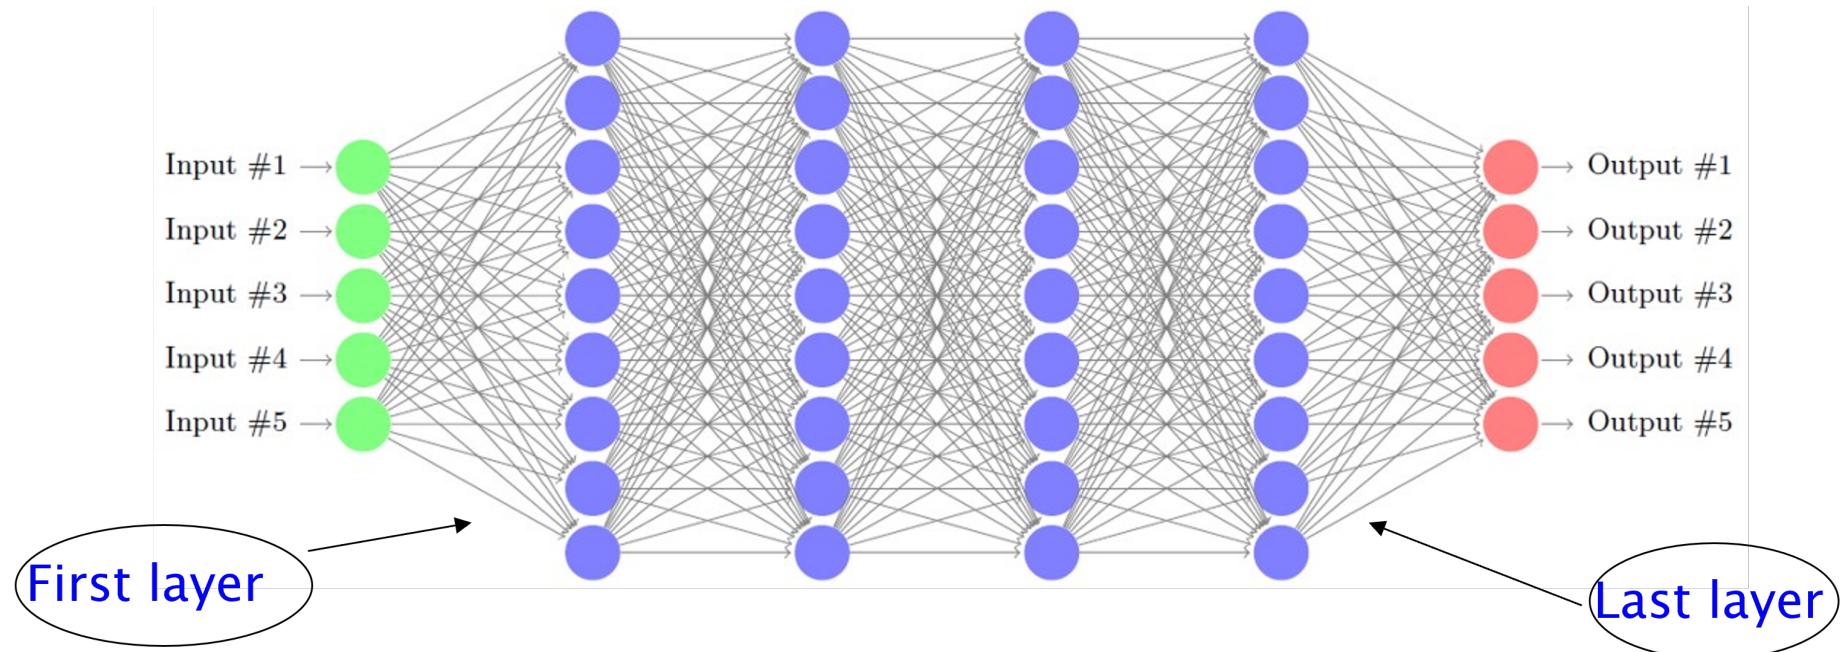
In real world, we most often do

$$f(x) = \underbrace{E_{\xi} \text{Loss}(x, \xi)}_{\text{mimimize}_{x \in X} E_{F_\xi}[h(x, \xi)]} \quad (1)$$

where ξ represents random variables with the joint distribution F_ξ .

- Pros: In many cases, the expected value is a good measure of performance
- Cons: One has to know the exact distribution of ξ to perform the stochastic optimization so that we most frequently use sample distribution. Then, deviant from the assumed distribution may result in sub-optimal solutions. Even know the distribution, the solution/decision is generically risky.

Deep-Learning on Neural-Network I



The **input vector** is denoted by \mathbf{x} and the **output vector** of layer l is denoted by \mathbf{y}^l . The **edge-weights** of layer l are denoted by $w_{i,j}^l$ where the relation of input-output is

$$y_j^l = \max\{0, w_{0,j}^l + \sum_i w_{i,j}^l y_j^{l-1}\}, \forall j, l = 1, \dots, L.$$

where the formula is called **ReLU operator/function** and $\mathbf{y}^0 = \mathbf{x}$.

Deep-Learning on Neural-Network II

The Deep-Learning is to use massive **sample images/inputs** \mathbf{x} to optimize/train (or learn edge-weights $w_{i,j}^l$) such that a (classification) sample-average error function is minimized. In other words, for this example, the outputs of images/inputs of Panda and Gibbon are **distinguishable/separable**, or they belong to different regions in the output space.

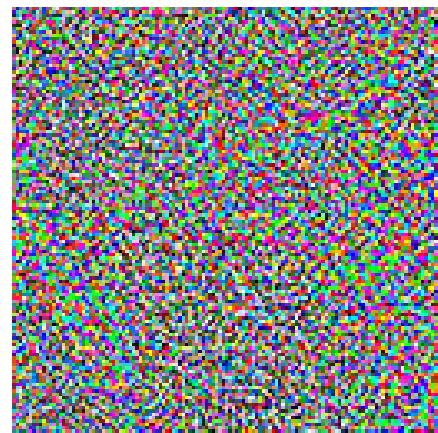
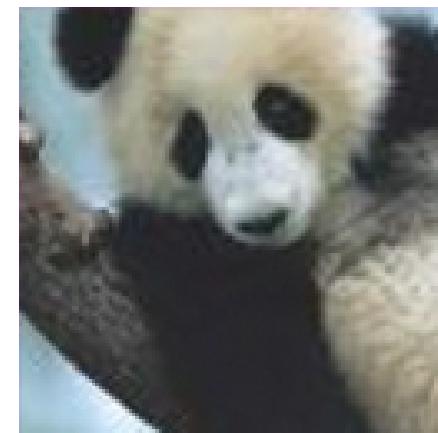
When all weights are determined, then the last-layer output vector of the neural-network, denoted by $\mathbf{y}^L(\mathbf{x})$, is a **vector function/mapping** of an input vector \mathbf{x} .

The **neural network verification**, for this example, is to find the smallest distortion of a given typical Panda image such that its output is in the output-region of normal Gibbon images, that is,

$$\text{minimize}_{\mathbf{x}} \quad \|\mathbf{x} - \hat{\mathbf{x}}\|^2$$

$$\text{subject to} \quad \mathbf{y}^L(\mathbf{x}) \in \text{a (convex) region outside of } \mathbf{y}^L(\hat{\mathbf{x}}).$$

Learning with Noises/Distortions

 $+\epsilon$  $=$ 

“panda”

57.7% confidence

“gibbon”

99.3% confidence



Goodfellow et al. [2014]

Distributionally Robust Optimization and Learning

On the other hand: Why does error occur? Believing that the sample distribution is the true distribution...

In practice, although the exact distribution of the random variables may not be known, people usually know certain observed samples or training data and other statistical information. Thus, we can consider an enlarged distribution set \mathcal{D} that confidently containing the sample distribution, and do

$$\underset{\mathbf{x} \in X}{\text{minimize}} \quad \left\{ \max_{F_\xi \in \mathcal{D}} \mathbf{E}_{F_\xi} [h(\mathbf{x}, \xi)] \right\} \quad (2)$$

In DRO, we consider a set of distributions \mathcal{D} and choose one to minimize the expected value for the worst distribution in \mathcal{D} . When choosing \mathcal{D} , we need to consider the following:

- Tractability
- Practical (Statistical) Meanings
- Performance (the potential loss comparing to the benchmark cases)

This is a nonlinear Saddle-Point Min-Max optimization/zero-sum-game problem

Reinforcement Learning: Markov Decision/Game Process



- RL/MDPs provide a mathematical framework for modeling sequential decision-making in situations where outcomes are partly random and partly under the control of a decision maker.
- Markov game processes (MGPs) provide a mathematical framework for modeling sequential decision-making of two-person turn-based zero-sum game.
- MDGPs are useful for studying a wide range of optimization/game problems solved via dynamic programming, where it was known at least as early as the 1950s (cf. Shapley 1953, Bellman 1957).
- Modern applications include dynamic planning under uncertainty, reinforcement learning, social networking, and almost all other stochastic dynamic/sequential decision/game problems in Mathematical, Physical, Management and Social Sciences.

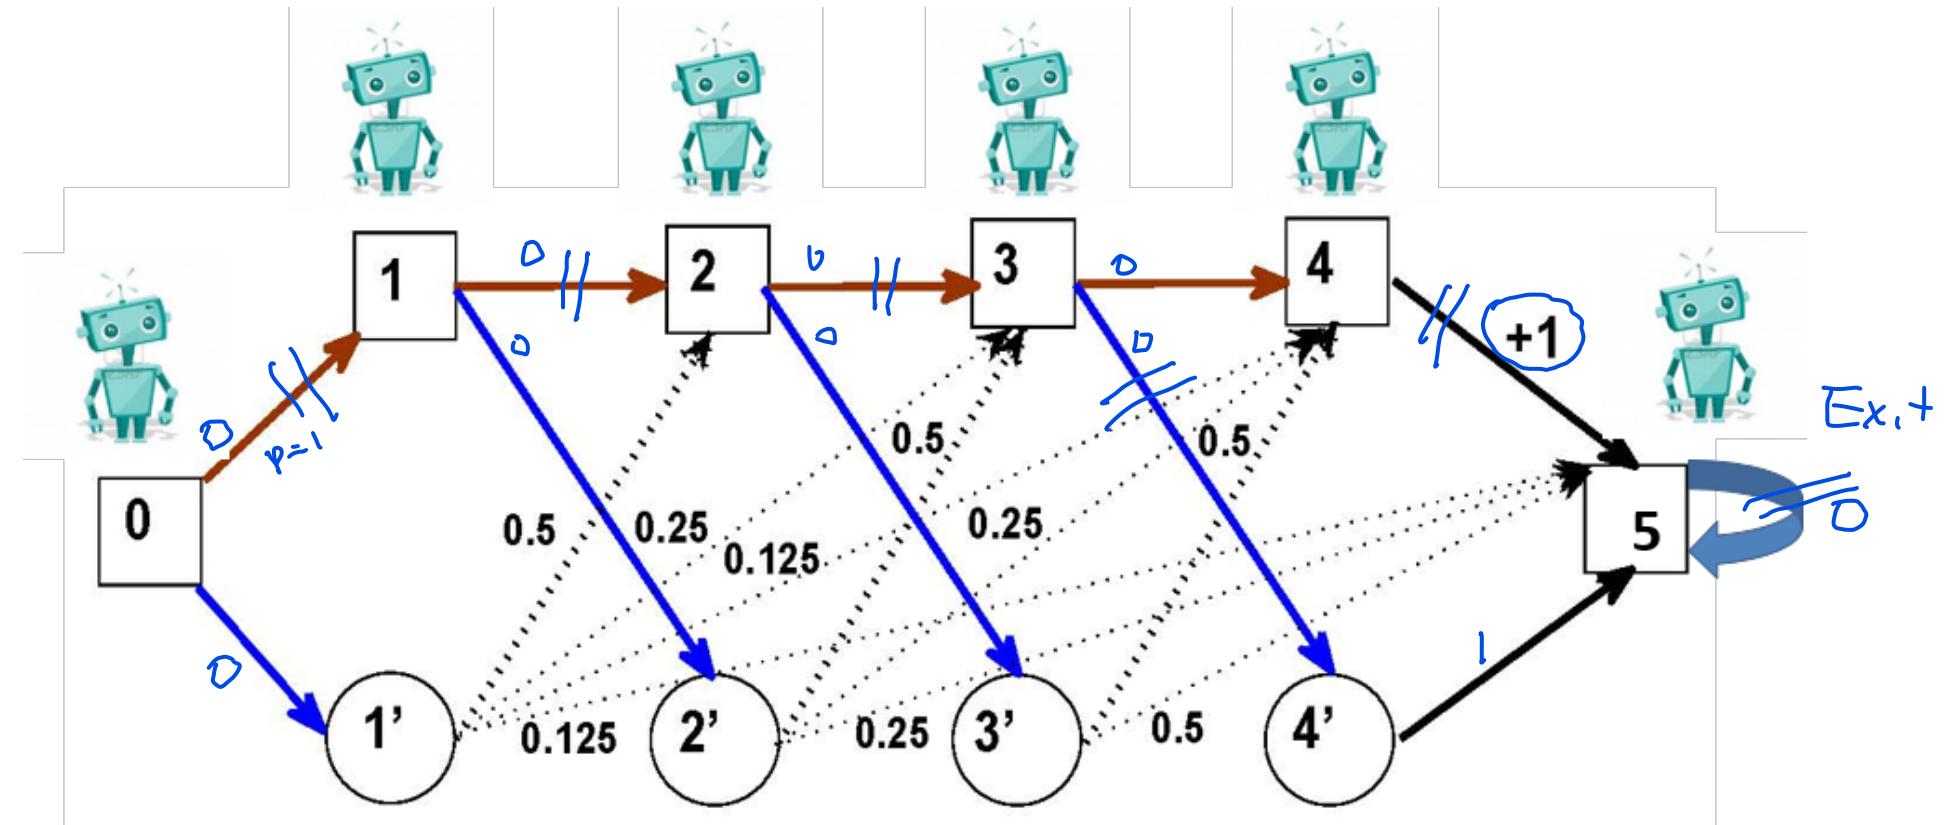
MDP Stationary Policy and Cost-to-Go Value

- An MDP problem is defined by a given number of states, indexed by i , where each state has a number of actions, A_i , to take. Each action, say $j \in A_i$, is associated with an (immediate) cost c_j of taking, and a probability distribution p_j to transfer to all possible states at the next time period. *total $s \dots$*
total $c \dots$
- A **stationary** policy for the decision maker is a function $\pi = \{\pi_1, \pi_2, \dots, \pi_m\}$ that specifies an action in each state, $\pi_i \in A_i$, that the decision maker will take at any time period; which also lead to a **cost-to-go** value for each state.
- The MDP is to find a stationary policy to minimize/maximize the expected discounted sum over the infinite horizon with a discount factor $0 \leq \gamma < 1$:

$$\sum_{t=0}^{\infty} \gamma^t E[c^{\pi_{i^t}}(i^t, i^{t+1})].$$

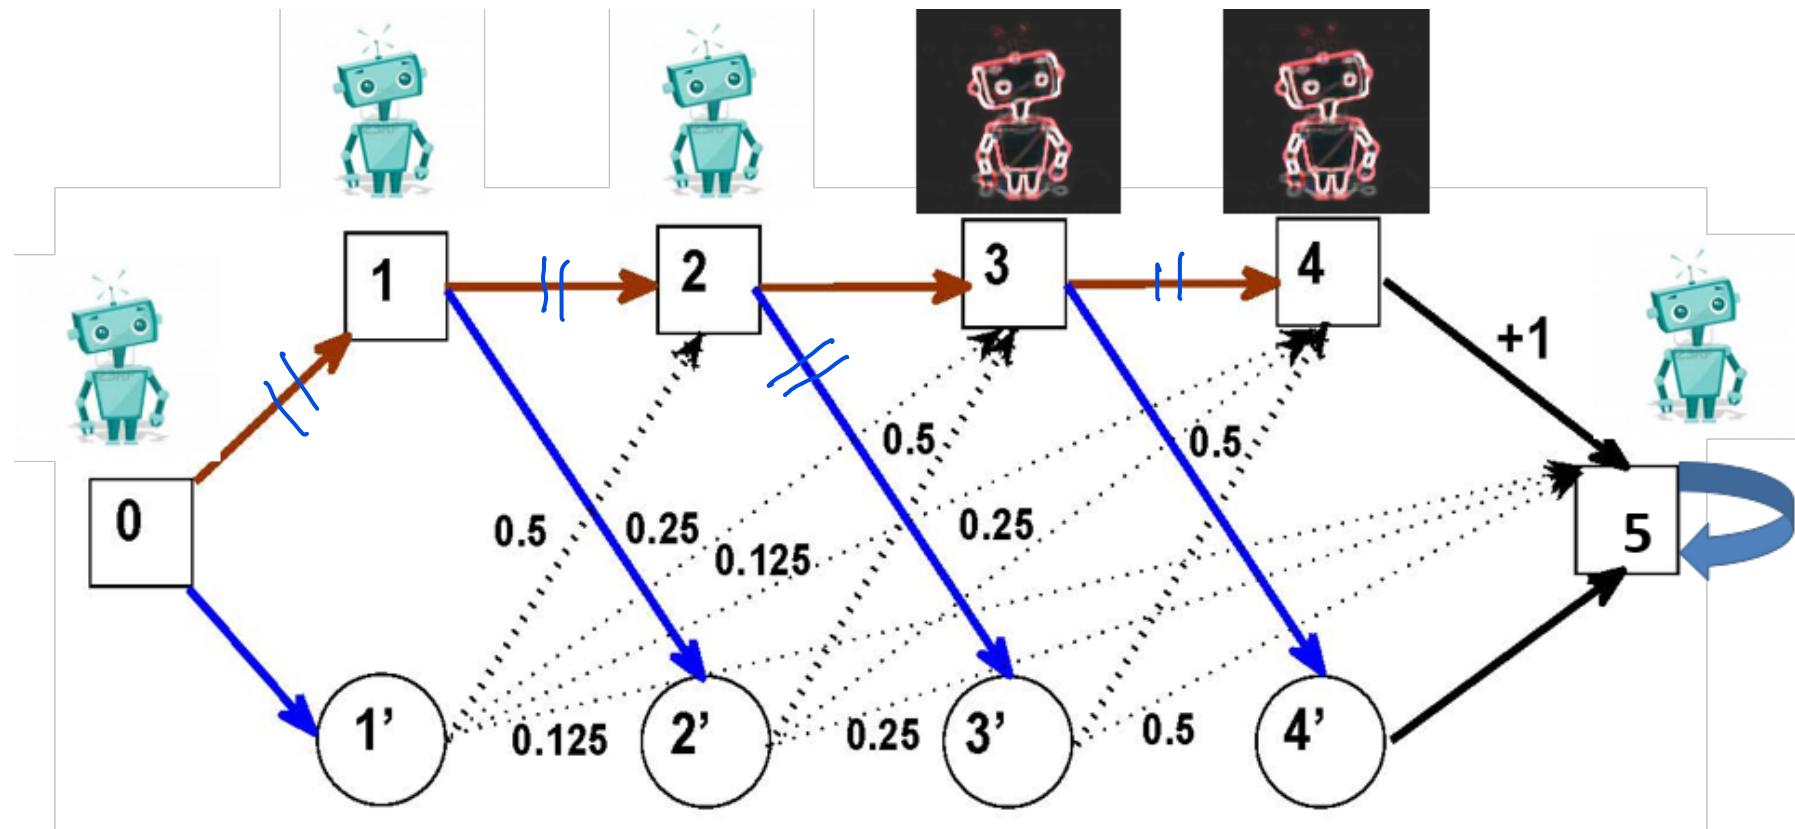
- If the states are partitioned into two sets, one is to minimize and the other is to maximize the discounted sum, then the process becomes a two-person turn-based zero-sum **stochastic game**.

An MDGP Toy Example: Maze Robot Runners (Simplified)

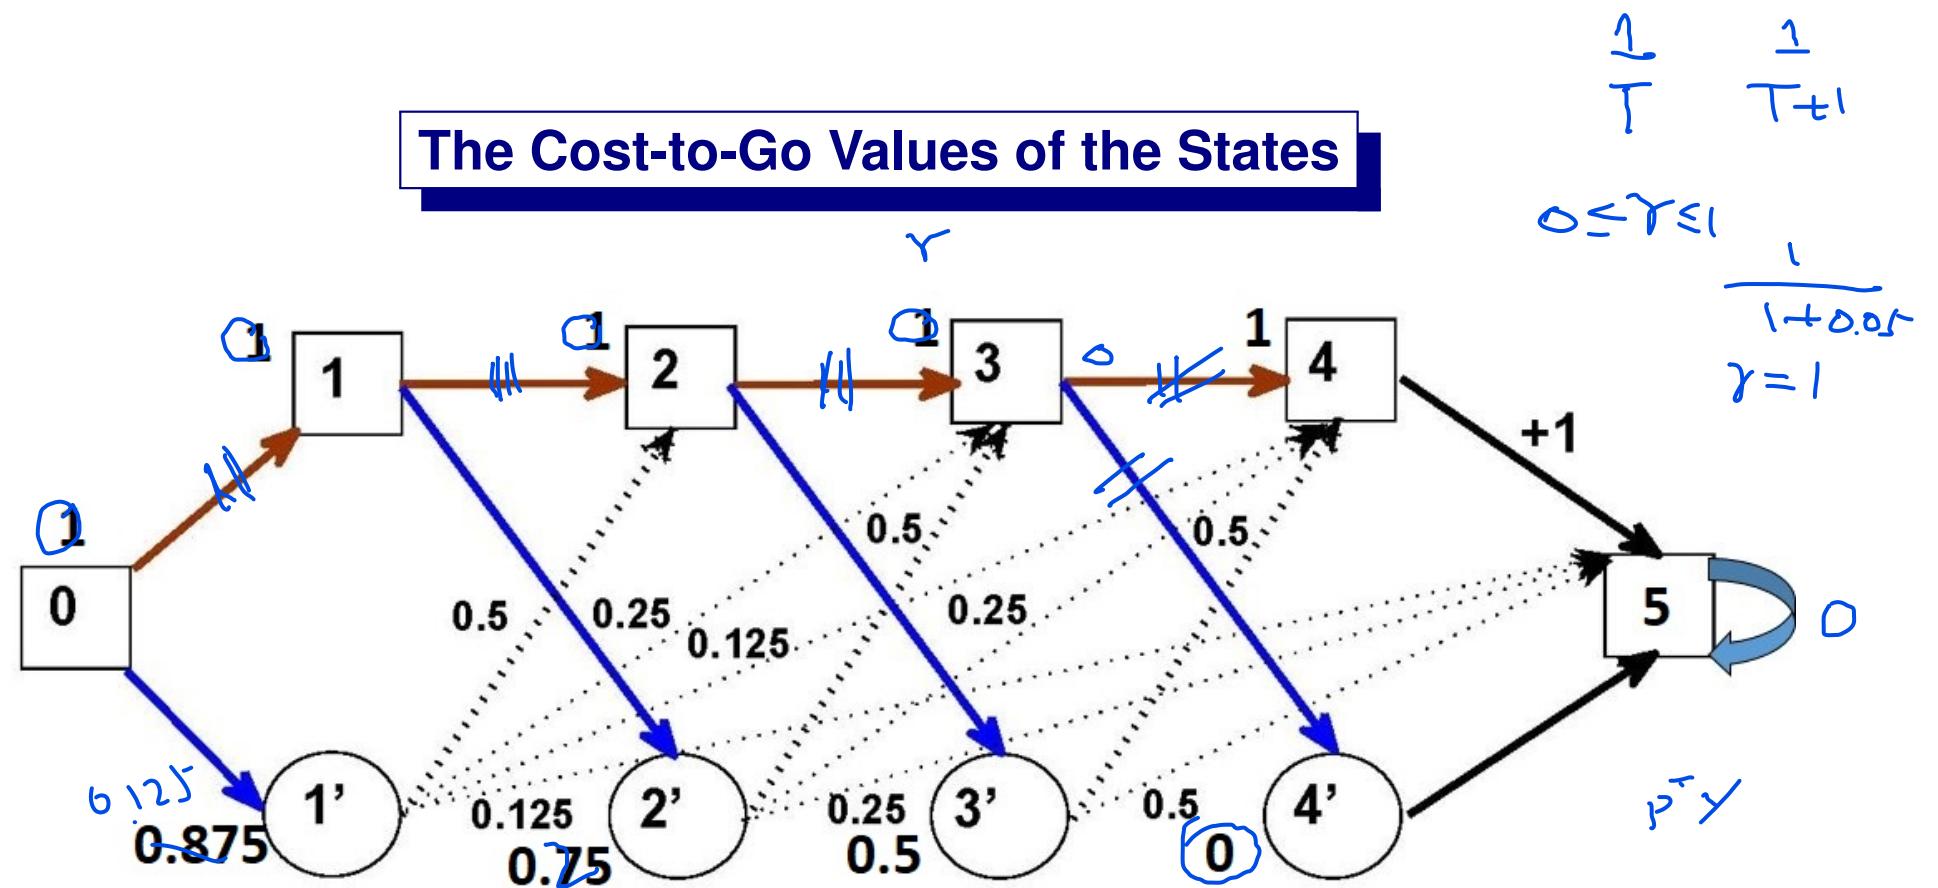


Actions are in red, blue and black; and all actions have zero cost except the state 4 to the exit/termination state 5. Which actions to take from every state to minimize the total cost (called optimal policy)?

Toy Example: Game Setting



States $\{0, 1, 2, 5\}$ minimize, while States $\{3, 4\}$ maximize.



Cost-to-go values on each state when actions in red are taken: the current policy is not optimal since there are better actions to choose to minimize the cost.

The Optimal Cost-to-Go Value Vector

Let $\mathbf{y} \in \mathbf{R}^m$ represent the **cost-to-go** values of the m states, i th entry for i th state, of a given policy.

The MDP problem entails choosing an optimal policy where the corresponding cost-to-go value vector \mathbf{y}^* satisfying:

$\max \mathbf{y}_i$

\Leftarrow

$$y_i^* = \min_{j \in \mathcal{A}_i} \{c_j + \gamma \mathbf{p}_j^T \mathbf{y}^*, \forall j \in \mathcal{A}_i\}, \forall i,$$

$$\mathbf{y} = \mathcal{T}(\mathbf{y})$$

with optimal policy

$$\pi_i^* = \arg \min \{c_j + \gamma \mathbf{p}_j^T \mathbf{y}^*, \forall j \in \mathcal{A}_i\}, \forall i.$$

In the Game setting, the conditions becomes:

$$\left\{ \begin{array}{l} y_i^* = \min_{j \in \mathcal{A}_i} \{c_j + \gamma \mathbf{p}_j^T \mathbf{y}^*, \forall j \in \mathcal{A}_i\}, \forall i \in I^-, \\ y_i^* = \max_{j \in \mathcal{A}_i} \{c_j + \gamma \mathbf{p}_j^T \mathbf{y}^*, \forall j \in \mathcal{A}_i\}, \forall i \in I^+ \end{array} \right.$$

and

$$\left\{ \begin{array}{l} y_i^* = \min_{j \in \mathcal{A}_i} \{c_j + \gamma \mathbf{p}_j^T \mathbf{y}^*, \forall j \in \mathcal{A}_i\}, \forall i \in I^-, \\ y_i^* = \max_{j \in \mathcal{A}_i} \{c_j + \gamma \mathbf{p}_j^T \mathbf{y}^*, \forall j \in \mathcal{A}_i\}, \forall i \in I^+ \end{array} \right.$$

$\{1, \dots, m\}$

They both are **fix-point** or **saddle-point** optimization problems. The MDP problem can be cast as a linear program; see next page.

The Equivalent LP Formulation for MDP

This model can be reformulated as an LP:

$$\begin{aligned}
 & \text{maximize}_{\mathbf{y}} && \sum_{i=1}^m y_i \\
 & \text{subject to} && \left. \begin{array}{l} y_1 - \gamma \mathbf{p}_j^T \mathbf{y} \leq c_j, \quad j \in \mathcal{A}_1 \\ \vdots \\ y_i - \gamma \mathbf{p}_j^T \mathbf{y} \leq c_j, \quad j \in \mathcal{A}_i \\ \vdots \\ y_m - \gamma \mathbf{p}_j^T \mathbf{y} \leq c_j, \quad j \in \mathcal{A}_m. \end{array} \right\}
 \end{aligned}$$

Theorem 1 When \mathbf{y} is maximized, there must be at least one inequality constraint in \mathcal{A}_i that becomes equal for every state i , that is, maximal \mathbf{y} is a fixed point solution.

The Maze Runner Example

The Fixed-Point formulation:

$$\begin{aligned}y_0 &= \min\{0 + \gamma y_1, 0 + \gamma(0.5y_2 + 0.25y_3 + 0.125y_4 + 0.125y_5)\} \\y_1 &= \min\{0 + \gamma y_2, 0 + \gamma(0.5y_3 + 0.25y_4 + 0.25y_5)\} \\y_2 &= \min\{0 + \gamma y_3, 0 + \gamma(0.5y_4 + 0.5y_5)\} \\y_3 &= \min\{0 + \gamma y_4, 0 + \gamma y_5\} \\y_4 &= 1 + \gamma y_5 \\y_5 &= 0 \text{ (or } y_5 = 0 + \gamma y_5)\end{aligned}$$

The LP formulation:

$$\text{maximize}_y \quad y_0 + y_1 + y_2 + y_3 + y_4 + y_5$$

subject to change each equality above into inequality

The Interpretations of the LP Formulation

The LP variables $\mathbf{y} \in \mathbf{R}^m$ represent the expected present **cost-to-go** values of the m states, respectively, for a given policy.

The LP problem entails choosing variables in \mathbf{y} , one for each state i , that maximize $\mathbf{e}^T \mathbf{y}$ so that it is the **fixed point**

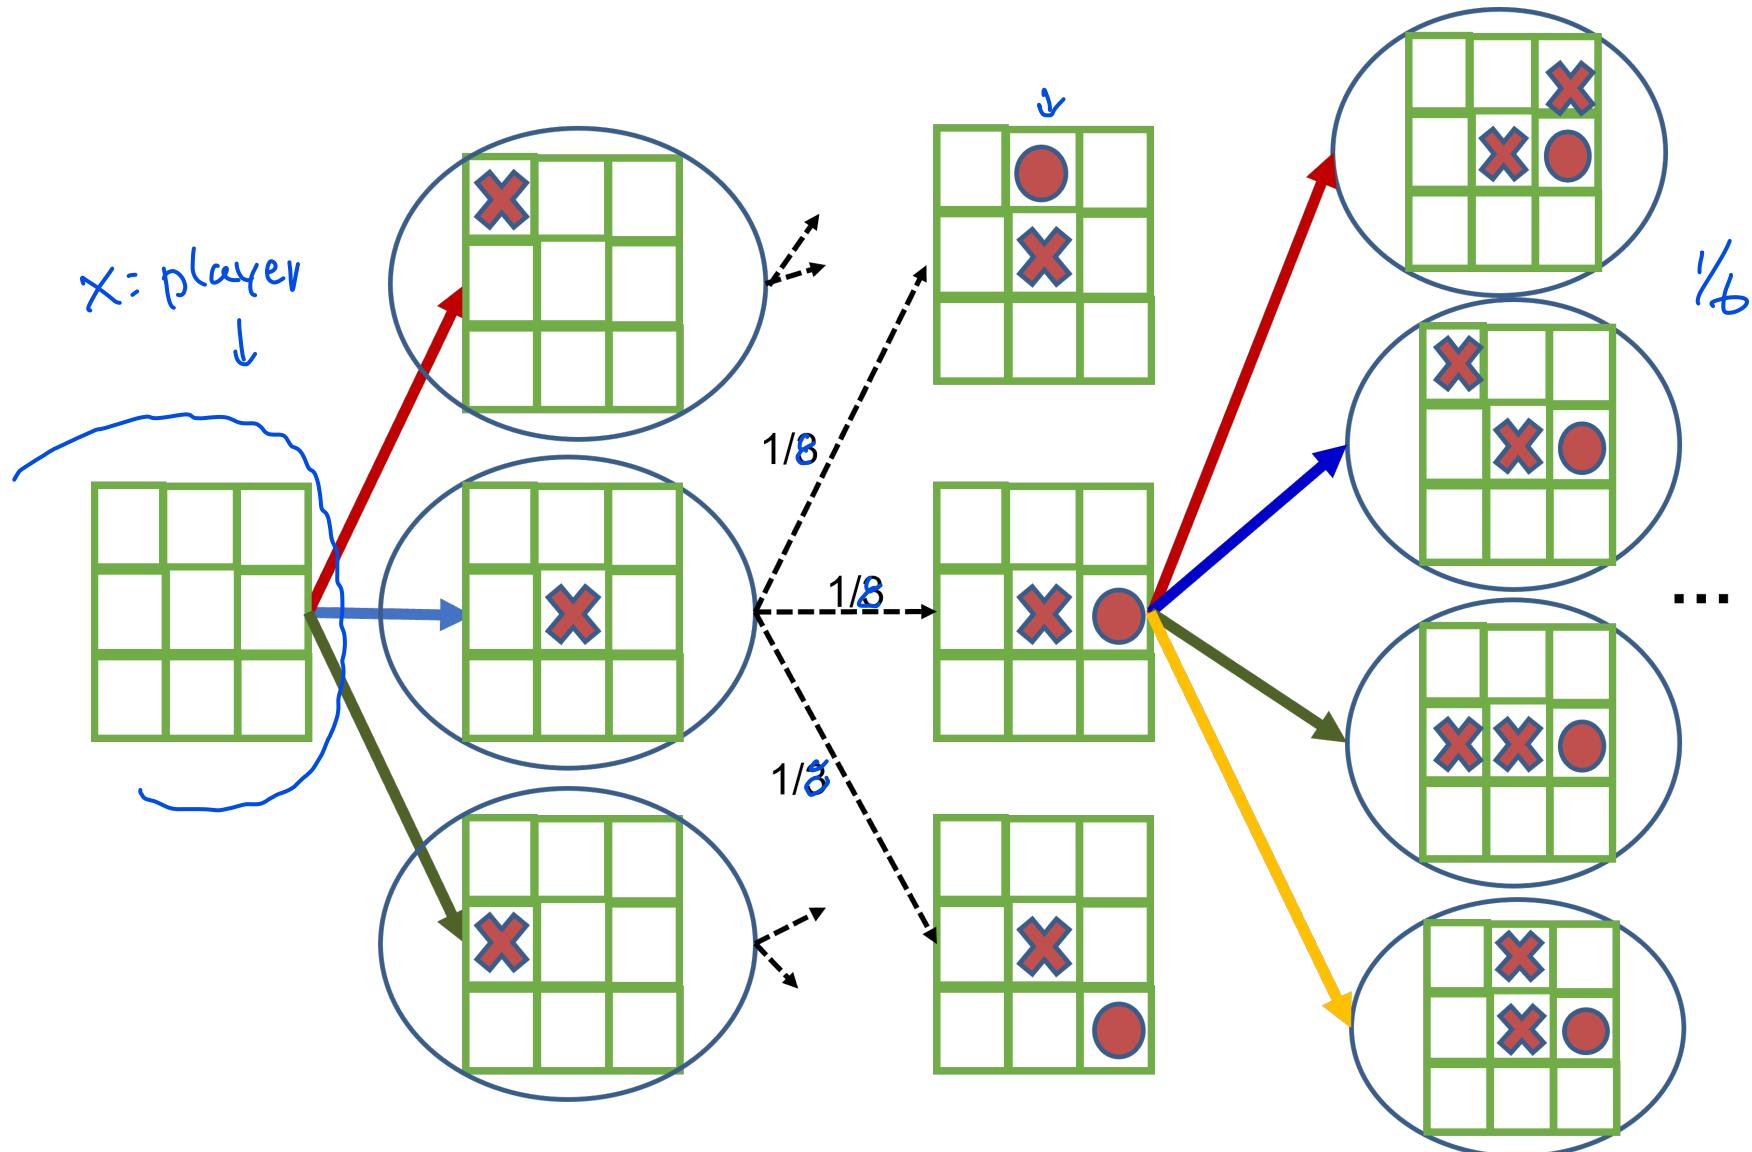
$$y_i^* = \min_{j \in \mathcal{A}_i} \{ \mathbf{c}_{j_i} + \gamma \mathbf{p}_{j_i}^T \mathbf{y} \}, \quad \forall i,$$

with an optimal policy

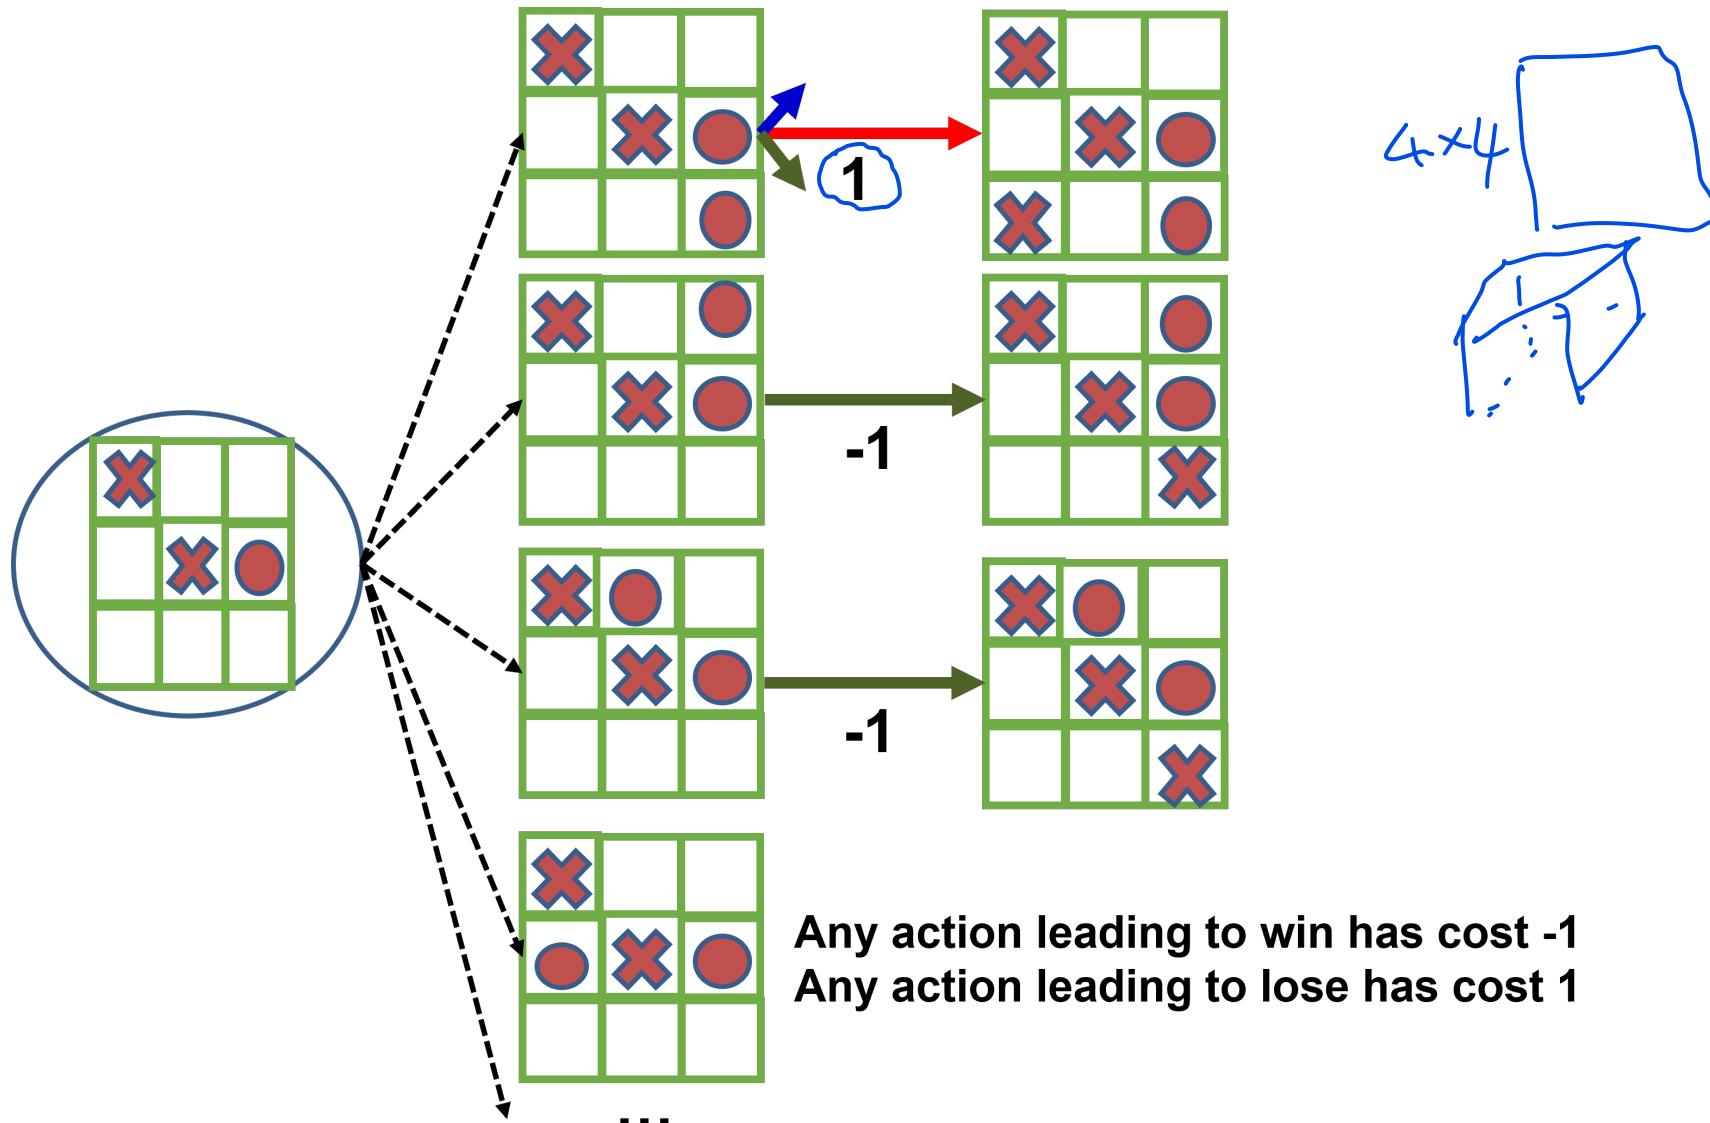
$$\pi_i^* = \arg \min \{ \mathbf{c}_j + \gamma \mathbf{p}_j^T \mathbf{y}, \quad j \in \mathcal{A}_i \}, \quad \forall i.$$

It is well known that there exist a **unique** optimal stationary policy value vector \mathbf{y}^* where, for each state i , y_i^* is the minimum expected present cost that an individual in state i and its progeny can incur.

States/Actions in the Tic-Tac-Toe Game against a Random Player



Action Costs in the Tic-Tac-Toe Game against a Random Player



States/Actions in the Tic-Tac-Toe Game against an Adversary Player?

