Week 9

Anomoly Detection & Recommender Systems

Anomoly Detection

- Motivation:
  - Remove outliers based on probabilistic methods

Anomoly Detection (con't)

- Procedure:
  - Use some "normalized" function, p(x,y,...), to fit known, good data
  - E.g. gaussian, multivariate gaussian
  - Assign some threshold, ε, s.t. if $p(x_0,y_0,\ldots)<\varepsilon$, for a given $x_0,y_0,\ldots$, label outlier
  - Adjust ε for best results

Recommender Systems

- Motivation:
  - Given user preferences on some small set, predict preferences on larger set

Recommender Systems (con't)

- Procedure:
  - Collect data on some subset of user preferences & classifiers:

Movie	Alice (1)	Bob (2)	Carol (3)	Dave (4)	x_1 (romance)	x_2 (action)
Love at last	5	5	0	0	0.9	0
Romance forever	5	?	?	0	1.0	0.01
Cute puppies of love	?	4	0	?	0.99	0
Nonstop car chases	0	0	5	4	0.1	1.0
Swords vs. karate	0	0	5	?	0	0.9

[Andrew Ng]

Recommender Systems (con’t)

- Not how it’s done anymore...
- Method nowadays is physics-inspired

Restricted Boltzman Machines

- Consider an ANN with the following layout:

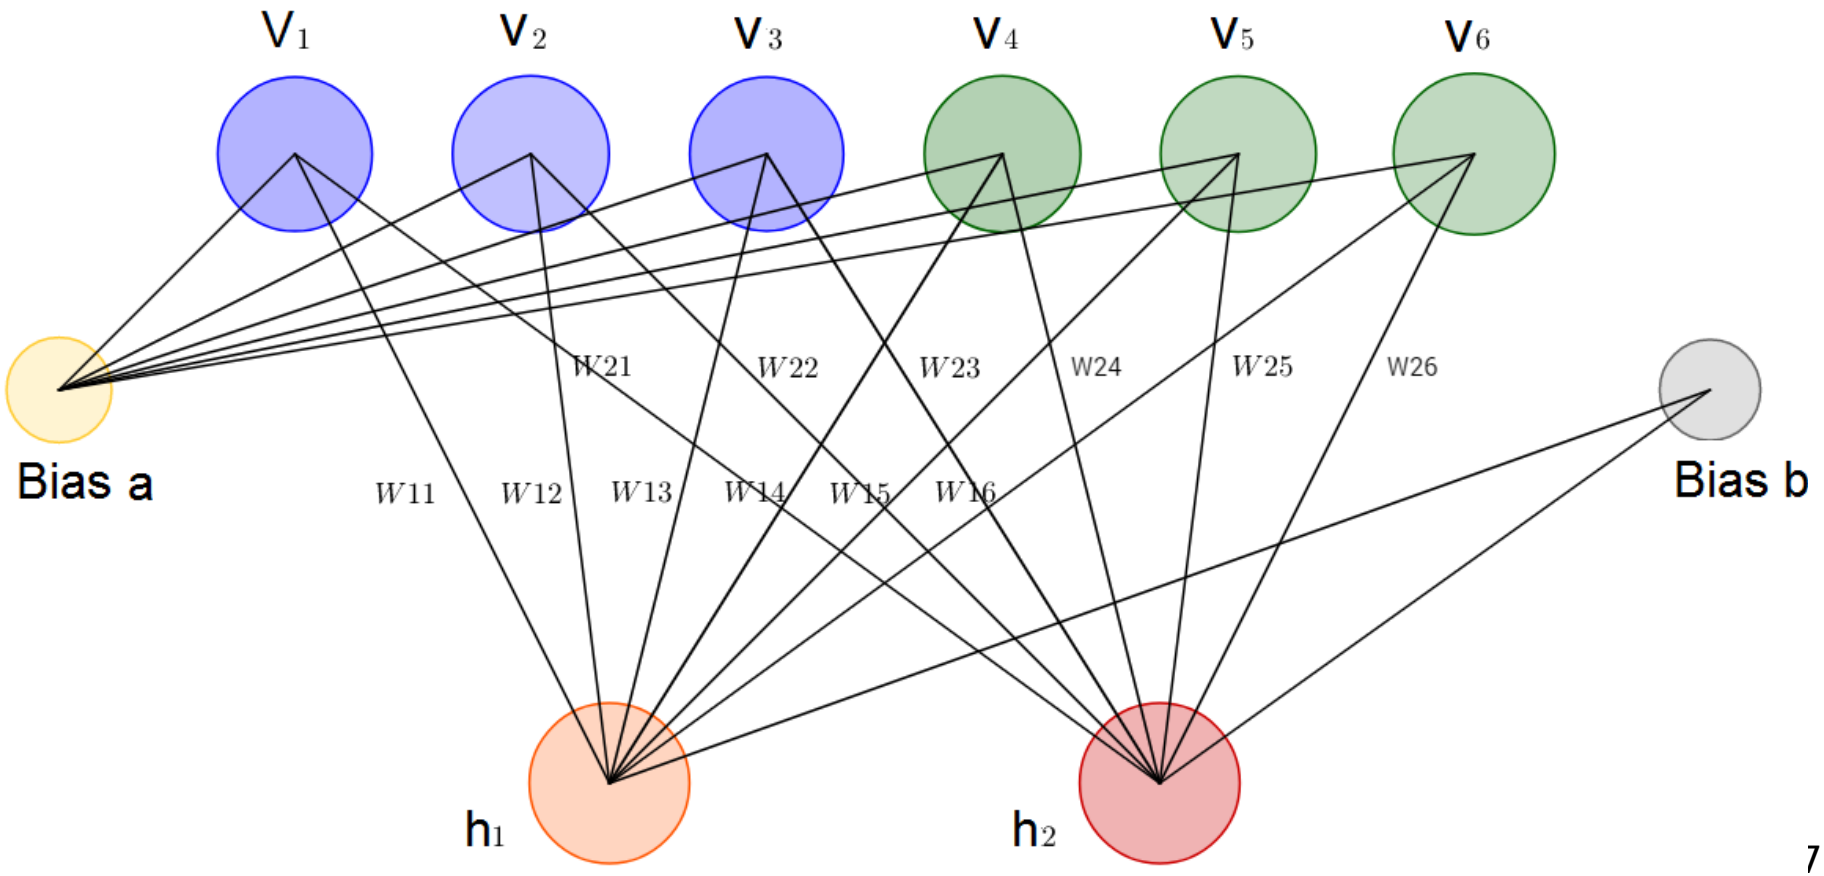

https://towardsdatascience.com/deep-learning-meets-physics-restricted-boltzmann-machines-part-i-6df5c4918c15

Restricted Boltzman Machines (Con't)

- Treat system as having some total, finite, Energy:

$$E(\vec{v}, \vec{h}) = -\sum_i a_i v_i - \sum_i b_i h_i - \sum_{ij} v_i h_j w_{ij}$$

- W_{ij} are values which minimize the energy
- Relies on binary 0,1:
  - Netflix, Youtube, etc all changed to binary...

Restricted Boltzman Machines (Con't)

- Treat likes/dislikes probabilistically:

$$p(v_i = 1|\vec{h}) = \frac{1}{1 + e^{-(a_i + \sum h_i w_{ij})}}$$

Restricted Boltzman Machines (Con't)

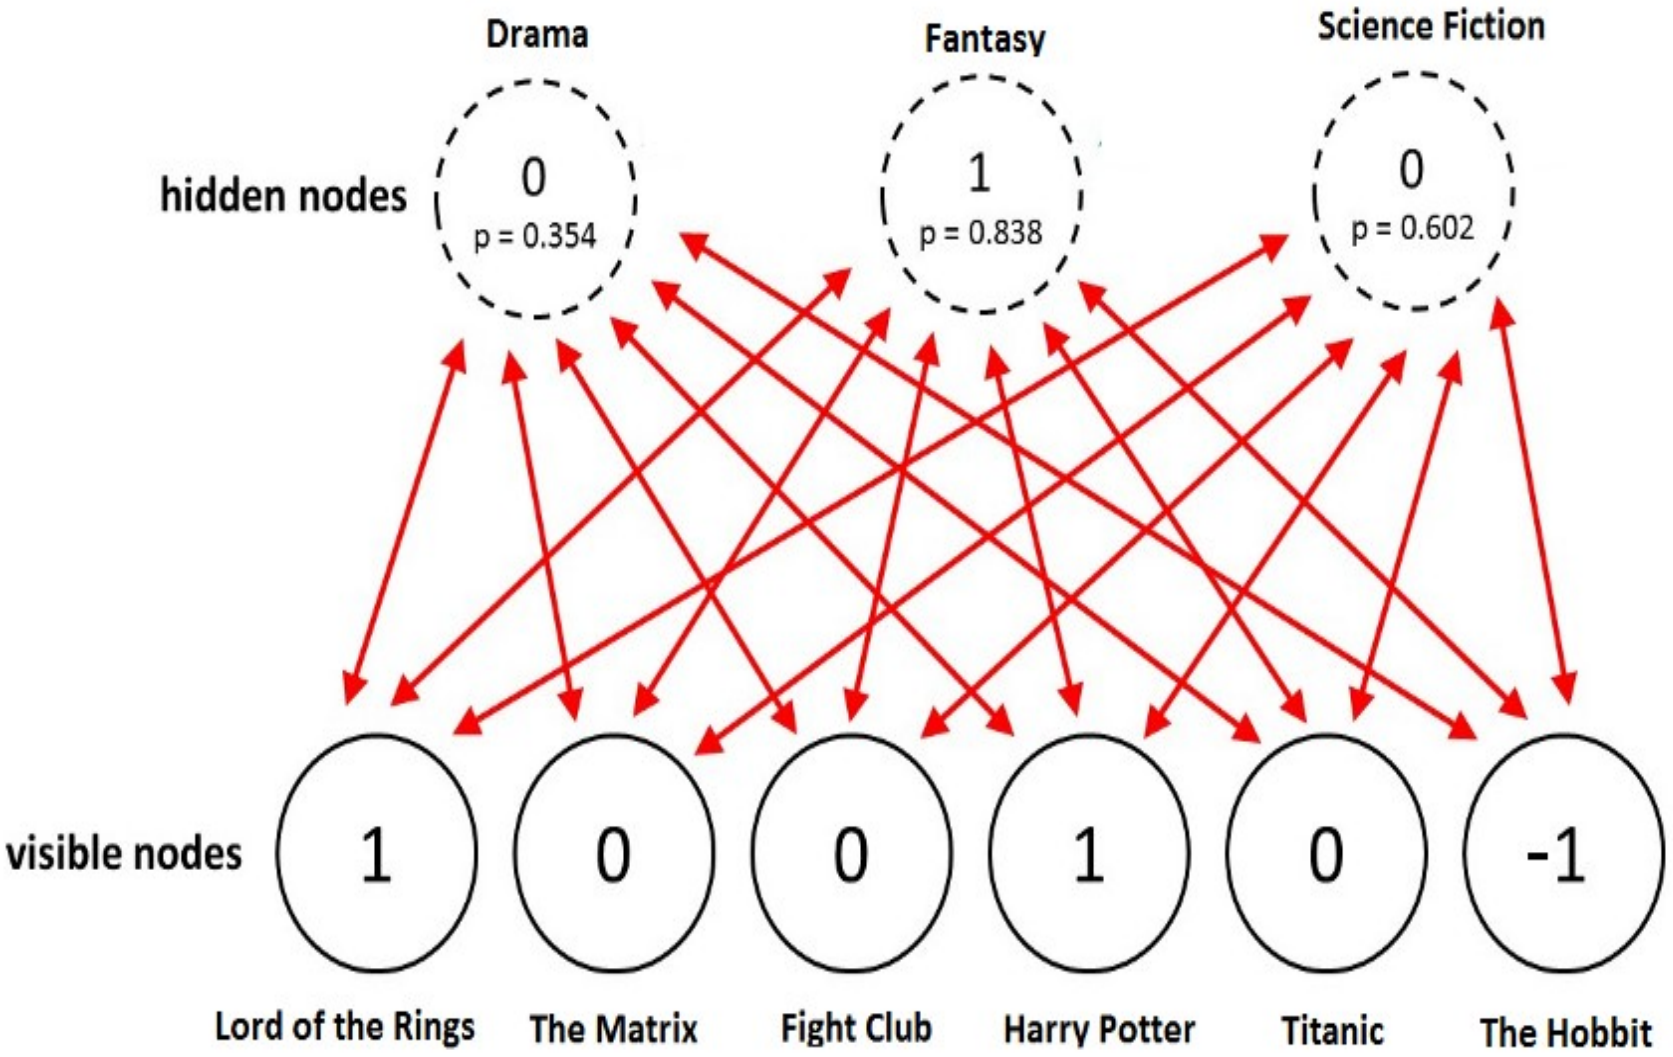

https://towardsdatascience.com/deep-learning-meets-physics-restricted-boltzmann-machines-part-i-6df5c4918c15

Restricted Boltzman Machines (Con't)

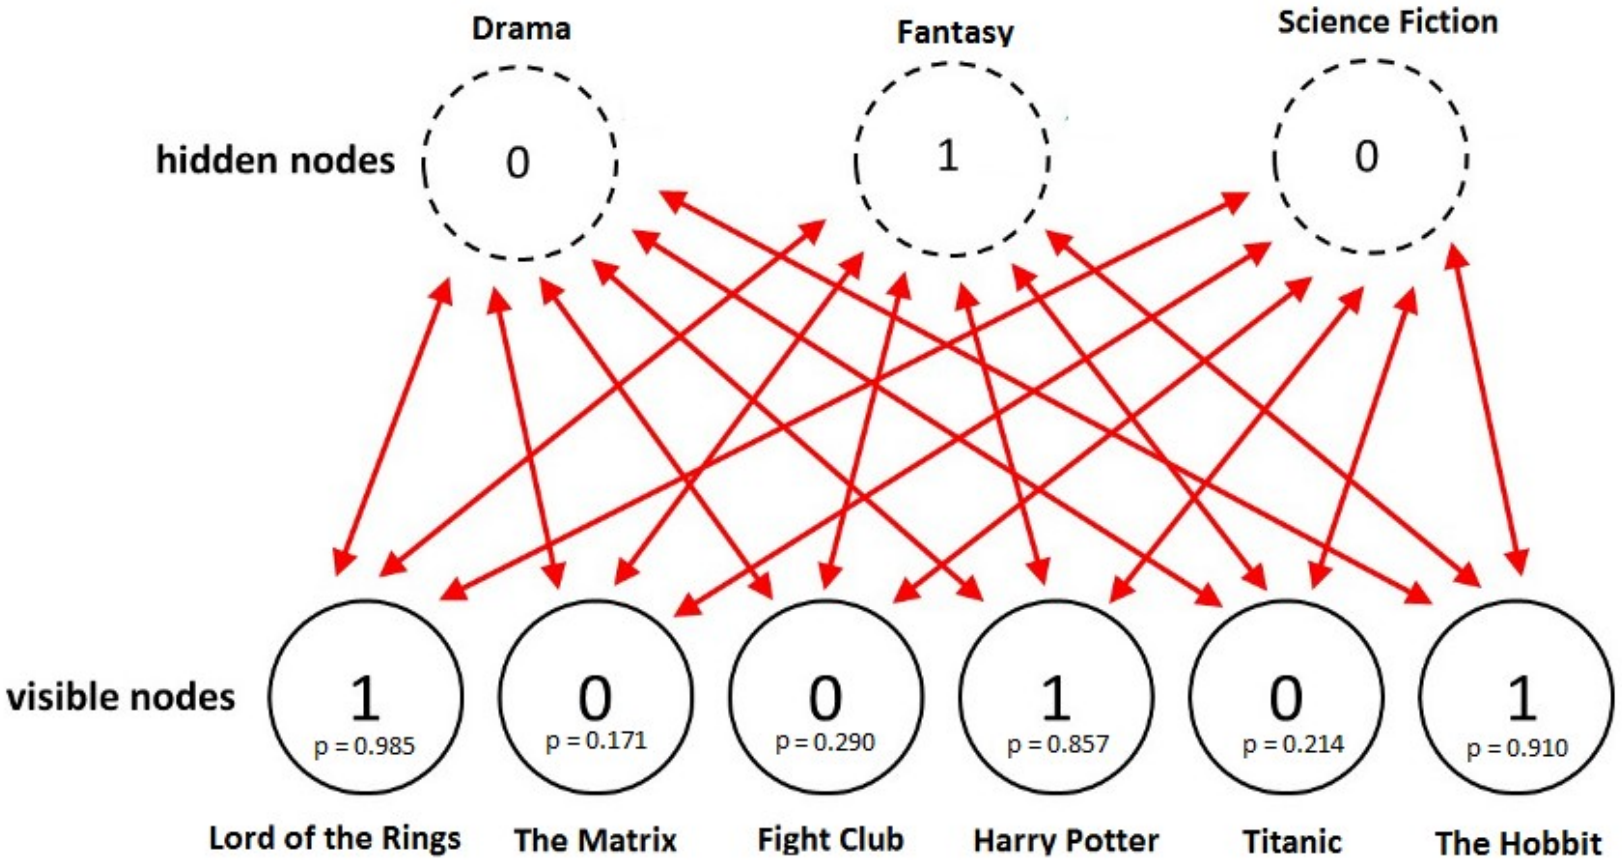

https://towardsdatascience.com/deep-learning-meets-physics-restricted-boltzmann-machines-part-i-6df5c4918c15