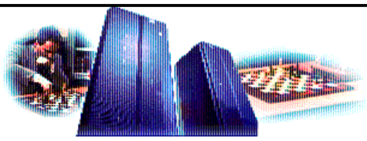

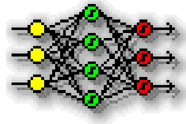

Sistemi Intelligenti
Affective computing

Alberto Borghese

Università degli Studi di Milano
Laboratorio di Sistemi Intelligenti Applicati (AIS-Lab)
Dipartimento di Informatica
Alberto.borghese@unimi.it

Affective computing

It is a branch of Artificial Intelligence that aims at realizing agents that can recognize and express emotional states.

At the crossing of computer science, psychology and cognitive science.

Picard (1995) Affective Computing, TR 321. “Emotions have a stigma in science; they are believed to be inherently non-scientific”

Marvin Minsky, related emotions to the broader issues of machine intelligence stating in The Emotion Machine that emotion is "not especially different from the processes that we call 'thinking.'"

It involves **multi-modal** interfaces and interfacing modalities that incorporate emotional states.

Emotions processing

- Implement modules for **human emotion recognition**, based on physiological parameters or on non-verbal communication.
- Design systems for **simulating internal emotional states**, which can communicate emotions readable by the human user.
- **Models of emotional dynamics**, to explain how human emotional intelligence works and **to reproduce this faculty in the machines**.

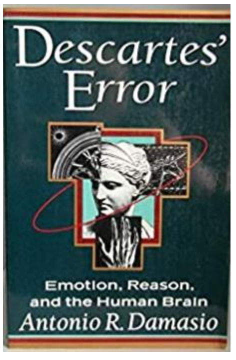

Emotions and Reasoning are tightly interconnected

Why?

Essere gentile, bellissima, amichevole, avere iniziativa, essere spiritosa, distinguere il bene dal male, innamorarsi, essere golosa di fragole alla panna, fare innamorare qualcuno, imparare dall'esperienza, avere proprietà di linguaggio, riflettere su sé stessa, avere una varietà di comportamenti, essere veramente creativa" (A. Turing, 1940).

Emotional Artificial intelligence

- To get truly intelligent machines: emotions are an important part of our intellective faculties.
- To improve human-machine interaction, making it a bit closer to human-human interaction.
- Application domains: enterteinment (videogames, social robots), health care,…

Emotions representation

Continuos (Russel’s circumplex model, 1980)

Discrete (Ekman’s basic emotions: *Joyful, Sad, Surprised, Angry, Fearful, Disgusted*, 1974)

Valence: the degree of attractiveness an individual, activity, or thing possesses as a behavioral goal.

Arousal: is a state of heightened physiological activity. This includes having strong emotions like anger and fear. For example the fight, flight or freeze response is a state of emotional arousal.

High level of arousal can be associated to **stress**. Perceived inability to cope with the expected tasks.

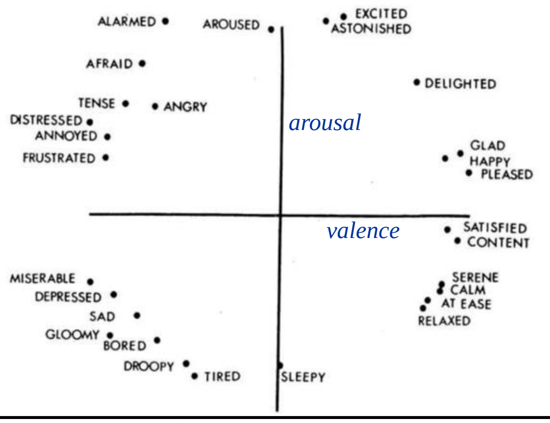

Emotions Input/output

Covert communication features.

Input / Output

- Facial expressions (FACS – Ad ogni emozione evocata è associato un insieme di micro-movimenti che la caratterizzano)
- Gestures (azioni del corpo)
- Proximity
- Speech prosody (EML Markup language)

```
<emotion xmlns="http://www.w3.org/2009/10/emotionml" category-set="http://.../xml#everyday-categories">
<category name="afraid" value="0.4"/>
<reference role="expressedBy“ uri="#sent1"/>
</emotion>
```

- Speech rhythm (pauses, speed…)

Measurements through physiological data (e.g. skin conductance - Empathica platform)

Graphical models

A **graphical model** o **modello probabilistico su grafo (PGM)** è un modello probabilistico che evidenzia le dipendenze tra le variabili randomiche (può evolvere eventualmente in un albero). Viene utilizzato nell'inferenza statistica.

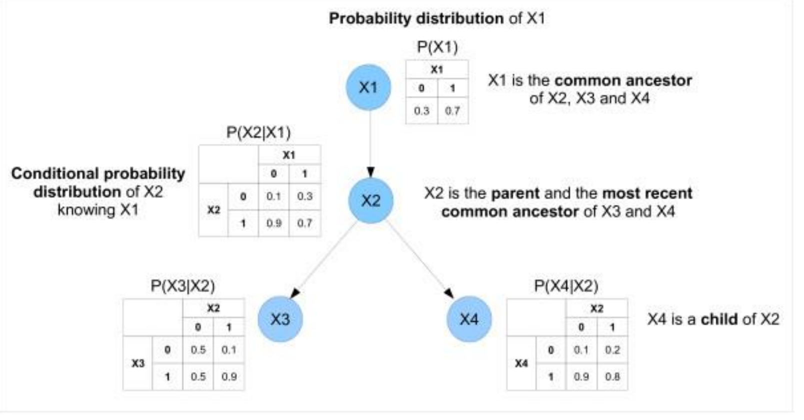

Il teorema di Bayes si può rappresentare come un modello grafico a 2 passi.

AIBOs graphical model of the interation

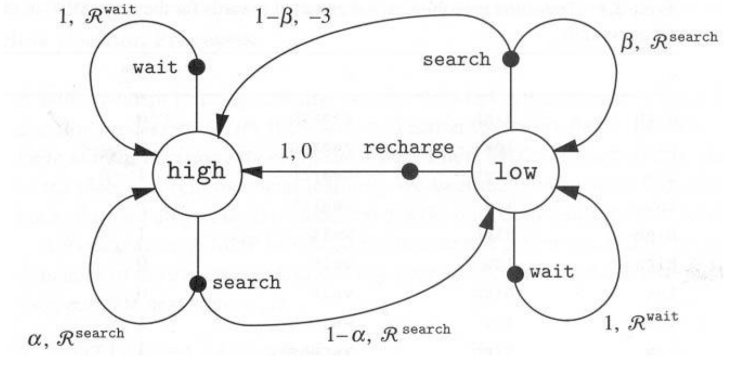

Emotions processing

- It can be well represented through graphical models
- Agent’s emotional states can be decided through **Probabilistic Finite State Machines** working on the **probabilistic graph**.

Interacting with an artificial partner: modeling the role of emotional aspects

I. Cattinelli, M. Goldwurm and N.A. Borghese (2008) Biological Cybernetics, pp.254-259.

Probabilistic Finite State Machine

- States: $\{s_1, s_2, \ldots s_N\}$ the set of **emotional states** of the agent
- Inputs: $\{u_1, u_2, \ldots u_M\}$ emotions provided to the agent **by the partner**
- Outputs: $\{y_1, y_2 \ldots y_P\}$ emotions **displayed by the agent** (equal to its emotional state)
- Next state: probability of the transition to a next state at time t. It can be seen as a 3D matrix where the triple: dim1 (state), dim2 (input), dim3 (next state), $P_t : S \times U \times S \rightarrow [0\ 1]$, provides the probability of the transition to next state from actual state with actual action.
- Output function: y= g(s). Hypothesis y = s.
- s_0 is the initial state

At each step:

1) Input
2) Processing
3) Output

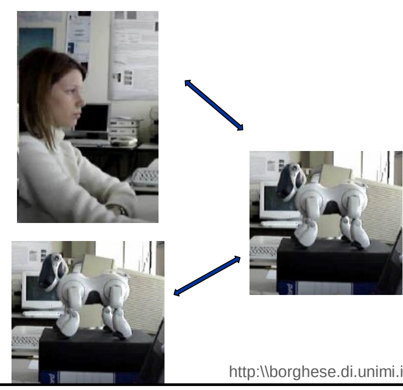

A basic model of an emotional agent

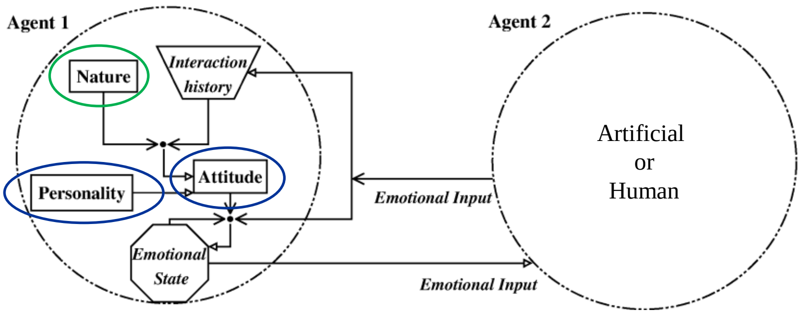

Personality is the initial transition matrix. P_0.

Attitude is the actual transition matrix P_t.

Nature is the criterion that drives the update of the current **Attitude** of the agent.

Agent 2 can be also a Human partner

Emotion Input

Emotion input are grouped into K categories (e.g. “nice” inputs = positive valence and positive arousal)

For each category, c_k, we introduce:
- An **eligibility trace** $e_t(c_k)$
- A set of **Target States** $TS(c_k)$ defined by the agent’s nature.

Example: if the nature is imitative,
If c_k corresponds to a joyful input,
TS can be “Happy”

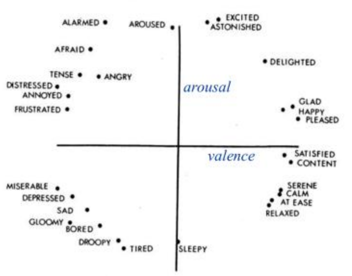

Attitude towards a target state is incremented when the eligibility of the category to which the input belongs is over a pre-defined threshold:

$P_{t+1}(s,u,ts) = P_t(s,u,ts) + \Delta \quad \forall s \in S,\ u \in U,\ ts \in TS(c_k(u))$
$P_{t+1}(s,u,ts) = P_t(s,u,ts) - k\Delta \quad$ otherwise

Eligibility trace on categories

$e_t(\underline{c}_k) = \alpha\, e_{t-1}(c_k) + h(c_k, u_j)$ if u_j is clustered in c_k

$e_t(\underline{c}_k) = \alpha\, e_{t-1}(c_k)$ $e_t(\underline{c}_k)$ decays otherwise

h(.) represents the affinity between the input and the category

Human-robot interaction

Robot agent has 4 emotional states: *Neutral*, *Joyful*, *Sad*, *Angry*.

User produces 7 emotional outputs that are input to the robot's agent: *Joyful, Sad, Surprised, Angry*, *Fearful*, *Disgusted* (defined by P. Ekman) and *Neutral*. These are recognised by real-time image processing.

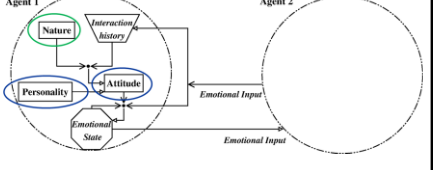

Variability in robot's response to Human partner emotion.

Variability depends on **Personality** and **Nature**, besides the input from the user partner.

Some pictures of the interaction

Agent-agent emotional interaction

Output of one agent is the input to the other one. **We have two PFSA.**

$A^1 = <S, U, Y, P^1(S,U)\ S^1_0>$ $g(Y) = I$
$A^2 = <S, U, Y, P^2(S,U), S^2_0>$ $g(Y) = I$

The set of the emotional states, S, is shared by the two agents.

Output (= emotional state) of one agent is the input to the other one.
$u^1 = s^2$
$u^2 = s^1$

Initial states are generally different
$S^1_0 \neq S^2_0$

Personality of the two agents is generally different
$P^1_0 \neq P^2_0$

Study of this situation

Attitude evolution

The two agents need to adapt to each other to have a successful emotional interaction. How?

They have to **adapt their Attitude run-time**.

Learning attitude

One of the two agents act as the environment: partner agent (e.g. Ghosts in the Pacman game), and the other as the learning agent (e.g. the Pacman):

Partner agent

- The partner agent emotional state can be observed by the learning agent through the partner's agent output (equal to the state).
- The partner agent emotional state, can either belong or not belong to TS.
- The learning agent can modify the emotional state of the partner agent through its output.

Learning agent

- Observes the partner emotional state.
- Receives a positive reward each time the partner enters in a target state.
- Has to maximize the long-time reward (maximize the time in which the partner is in a target emotional state).

Attitude Q-learning

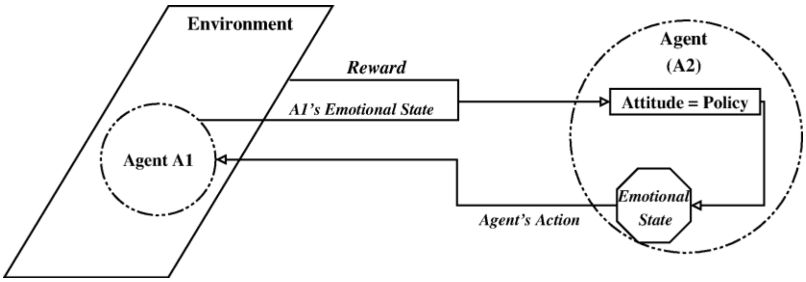

For the **learning agent**, initial $Q_0(s,a) = P^2_0(s,a)$ is equal to agent's personality: the more the value of the action the higher is the probability of choosing that action (pursuit).

Q-learning cycle

1) The learning agent observes state s and choses a according to the actual attitude $P^2_t(s,a,s')$, that is $Q_t(s,a)$.
2) The learning agent observes the next state, s', and the reward r_{t+1}. $r_{t+1} = +1$ if s' belongs to TS; $r_{t+1} = 0$ otherwise.
3) Update: $Q_t(s,a) = P_{t+1}(s,u,s') = P_t(s,u,s') + \Delta \quad \forall s \in S, u \in U, s' = ts \in TS(c_k(u))$
$Q_t(s,a) = P_{t+1}(s,u,s') = P_t(s,u,s') - k\Delta$

Interaction graph - I

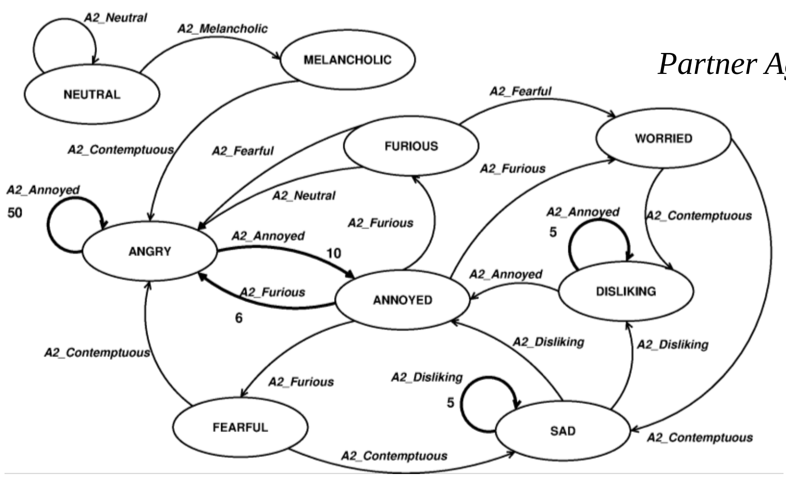

Partner Agent, a_1

- A^1 and A^2 have the personality of "friendly" agents.
- Goal of A^2 is to make A^1 **angry**: TS = {***Annoyed, Angry, Furious***}
- It succeeds for 78% of the time.

Interaction graph - II

Partner Agent, a_1

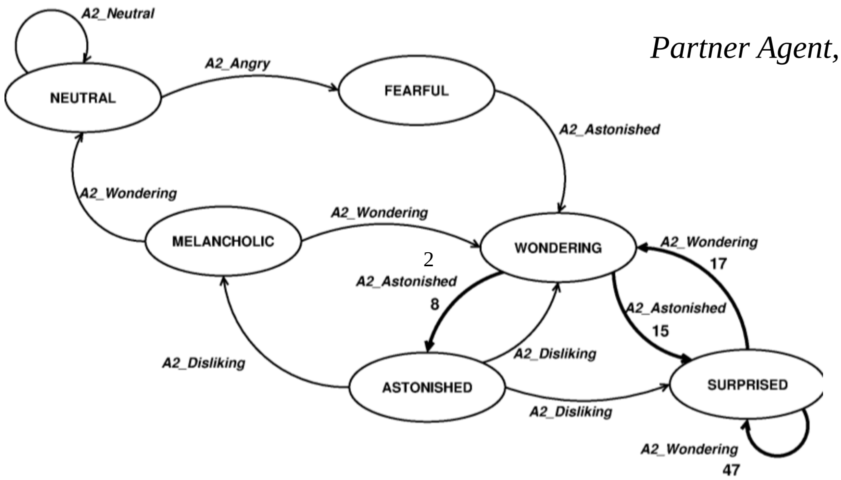

- A^1 and A^2 have the personality of "friendly" agents.
- Goal of A^2 is to make A^1 **surprised**: TS = {***Wondering, Surprised, Astonished***}
- It succeeds for 95% of the time.

Interaction graph - III

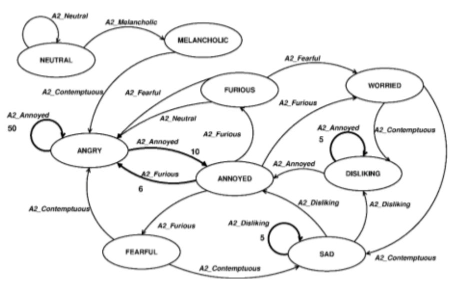

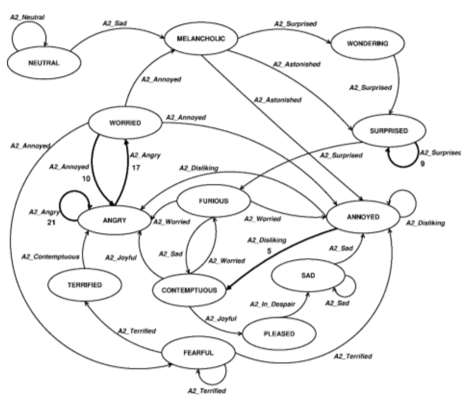

Partner Agent, a_1

- A^1 and A^2 have the personality of "friendly" agents. However, A2 has a **different personality** with respect to the previous situation.
- Goal of A^2 is to make A^1 **angry**: TS = {***Annoyed, Angry, Furious***}
- It succeeds for 51% of the time now.

Quantitative analysis

How good is a model?

Which states will be the most frequent ones?

How long would it take to go from state i to state j?

….

We will use Markov Random Chains theory

Markov chain

- A set of states, $S = \{X_i\}$, for which the probability of a state, X_i, next to X_j, depends only on the transition probability from X_i to X_j : $p(X_i, X_j)$.

  $p(X_i, X_j) = P(s_{t+1} = X_j \mid s_t = X_i)$

- $\mu^0(X_i)$ is the a-priori probability of being in a state $X_i = P(s^0 = X_i)$
- $P(s^n = X_j \mid s^0 = X_i) = (P^n)_{ij}$ – probability of going from state i to state j in n steps.
- $P(s^n = X_j) = (P^n)_j$ – probability of going to state j in n steps from any state.
- $\mu^n(X_j) = P(s^n = X_j)$ is the probability of being in state X_j in n steps.
- $(P^n)_{ij} = \sum \mu^0(X_i)\ (P^n)_{ij}$

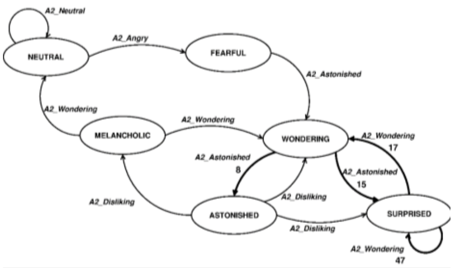

Primitive Markov chains

Primitive Markov chains if there is a probability greater than zero to enter in all states in some steps: $(P^n)_i > 0$ for some value of n, that is $(P^n)_{ij} > 0$ for all $i,j \in S$. To this aim the graph should be:

- **Strongly connected** (each vertex should be reachable from any other vertex).
- **Aperiodic.** Given the set of cycles, the MCD of their lengths is not greater than 1.

We focus on the **essential components of the graph**, that is components that support primitive Markov chains, in which once entered, the agent never leaves.

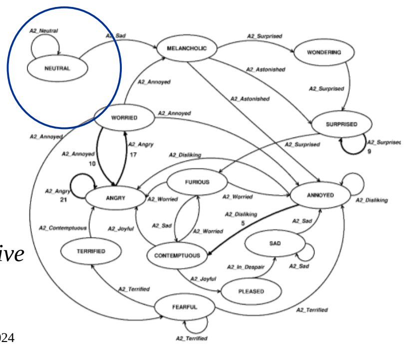

Not primitive

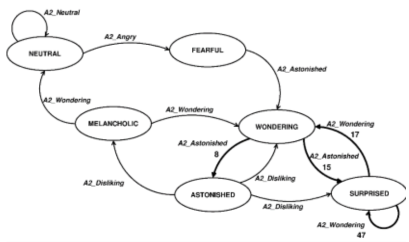

Why Markov Chains?

Markov chains have **no inputs**

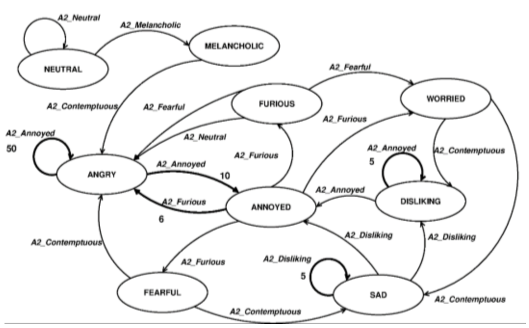

We can put inside the same matrix M the two transitions matrixes of the learning agent and its partner

M(i,j) provides the probability to go from state i= (a,b), a for agent A_1 and b for agent A_2, to state j = (a',b'), a' for agent A_1 and b' for agent A_2.

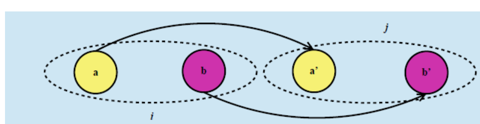

$M(i,j) = P_1(a,b,a') \times P_2(b,a', b')$

We have no external input. We have a closed system of which we can study the bhavior.

Proprietà delle catene di Markov primitive - I

1) Esiste ed è unica una **distribuzione statistica stazionaria** (che non varia nel tempo), $\pi(X_j)$, tale che:

$$\pi' P = \pi'$$

dove π' è l'autovettore di P che corrisponde all'autovalore = 1.

2) π consente di determinare la frequenza con cui uno stato X_j viene visitato. Per ogni coppia i,j $\in$ S, vale:

$$\lim_{n\to\infty} (P)^n{}_{ij} = \lim_{n\to\infty} Prob(X \quad = j) = \pi_j$$

La distribuzione limite degli X_n è indipendente dallo stato iniziale della catena ed è coincidente con la distribuzione stazionaria.

3) $\pi(X_j)$ quindi ci dice quanto tempo dovremo aspettare, τ_j. per entrare una seconda volta nello stato X_j (T = 1/f).

$$E_j(\tau_j) = 1 / \pi_j \quad \text{for each } j \in S$$

Haggstrom 2002

Proprietà delle catene di Markov primitive - II

$E_i(\tau_j)$ = ? Quanto tempo dobbiamo aspettare per passare dallo stato i allo stato j?

a) Costruiamo la matrice G(z) = I – Pz, con z vettore di k variabili.
b) Costruiamo per ogni elemento i,j la matrice aggiunta $G_{ij}^+(z)$ ottenuta come determinante della matrice derivate eliminando la i-esima riga e la j-esima colonna di G(z).
c) Calcoliamo $r_{ij}(z) = (-1)^{(i+j)} \det(G_{ij}^+(z))$
d) Poniamo z = 1 e otteniamo:

$$E_i(\tau_j) = \frac{r'_{ij} r_{jj} - r_{ij} r'_{jj}}{r_{jj}^2}$$

where r' is computed on the transpose matrix.

Errore di approssimazione

$\mu^n(X_j)$ – probability of entering in the state X_j after n step (relative frequency)

$\pi(X_j)$ – asintotic probability of entering in the state X_j (stationary case)

$\forall \varepsilon > 0 \quad |\mu^n(X_j) - \pi(X_j)| < \varepsilon$ *total variation error*

if $n \geq t\left(1 + \frac{log_2 k - log_2 \varepsilon - 1}{-log_2 m(P^t)}\right)$

- t è l'intero più piccolo per cui $P^t > 0$
- k è la cardinalità di S
- $m(P^t)$ è il massimo tra la differenza, misurata mediante TV, tra due righe di P^t

Quantitative behavioural analysis - I

- A^1 and A^2 have the personality of "friendly" agents.
- Goal of A^2 is to make A^1 **angry**: TS = {***Annoyed, Angry, Furious***}
- It succeeds for 78% of the time.
- 15 states in total between A_1 and A_2

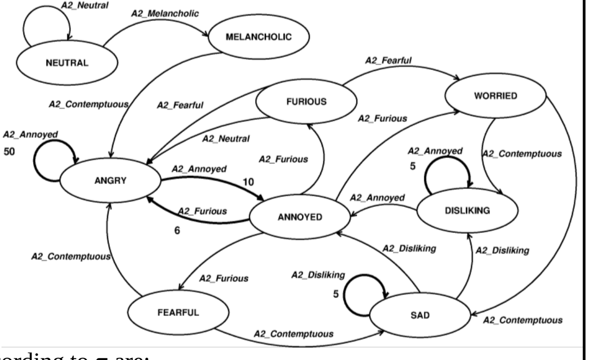

The most probable pairs of states according to π are:
- (Angry Annoyed), with p = 0.5148
- (Annoyed, Furious) with p = 0.1548
- (Sad, Disliking), with p = 0.0973

Quantitative behavioural analysis- II

- A^1 and A^2 have the personality of "friendly" agents.
- Goal of A^2 is to make A^1 **surprised**: TS = {***Wondering, Surprised, Astonished***}
- It succeeds for 95% of the time.
- 10 states in total between A_1 and A_2

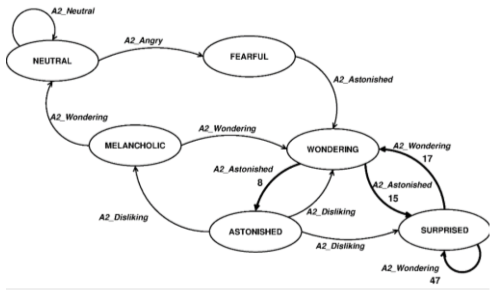

The most probable pairs of states according to π are:

- (Surprised, Wondering), with p = 0.6286
- (Wondering, Astonished) with p = 0.2292
- (Astonished, Disliking), with p = 0.0917

Quantitative behavioural analysis - III

- A^1 and A^2 have the personality of "friendly" agents.
- Goal of A^2 is to make A^1 **angry**: TS = {***Annoyed, Angry, Furious***}
- It succeeds for 78% of the time.
- SS pair: Melancholic, Contemptuos
- 15 states in total between A_1 and A_2

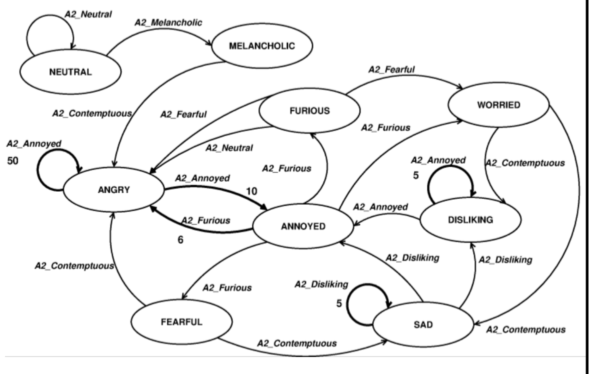

A minimum of 5,91 and a maximum of 213,1 steps, on average, for going from state in SS to states in ES (mean 77,98)

Quantitative behavioural analysis- IV

- A^1 and A^2 have the personality of "friendly" agents.
- Goal of A^2 is to make A^1 **surprised**: TS = {***Wondering, Surprised, Astonished***}
- SS pair: Neutral, Angry.
- It succeeds for 95% of the time.
- 10 states in total between A_1 and A_2

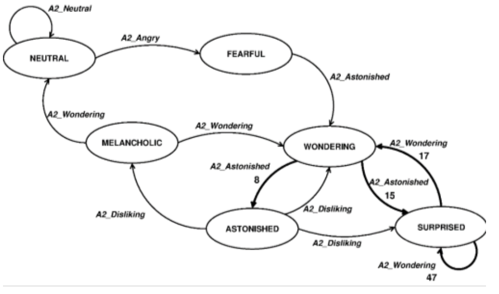

A minimum of 3,86 and a maximum of 12,43 steps, on average, for going from state in SS to states in ES (mean 7,07)

Comments

I valori ottenuti dalla distribuzione stazionaria, sono molto vicini ai valori sperimentali

- La distribuzione stazionaria è un descrittore affidabile del comportamento dei nostri agenti
- L'errore di approssimazione è minore di 0,001 già dopo 38 e 27 passi rispettivamente.

- Q-learning è risultato efficace: i target state visitati dall'agente partner sono effettivamente i più frequentemente visitati.
- Nell'esempio in cui il goal è rendere il partner **surprised** la policy è particolarmente efficace (servono solo 7 passi in media)
- Nell'esempio in cui il goal è rendere il partner **angry** la policy è meno efficace perchè servono in media 78 passi. Tuttavia, occorre considerare anche che i TS angry hanno una probabilità di ingresso molto bassa.

Take home message

Modello di interazione uomo-macchina o agente-agente, basato su natura e personalità.

Ha una natura probabilistica, che rende ricca l'interazione perchè non stereotipata.

Il modello consente di stimare il tempo media per entrare in uno stato emotivo e la frequenza con cui viene visitato.

Agendo su personalità e natura si possono ottenere comportamenti diversi e personalizzati sull'utente.

L'apprendimento con rinforzo può essere utilizzato per ottenre agenti che interagiscano secondo la loro natura pre-definita.