Bayesact II - Theory and Model

Jesse Hoey
School of Computer Science
University of Waterloo

Feb 4th, 2014

Readings:

- Jesse Hoey, Areej Alhothali and Tobias Schröder. *Bayesian Affect Control Theory*. In Proc. Humaine Association Conference on Affective Computing and Intelligent Interaction (ACII), 2013.
- longer version on the webpage `http://www.bayesact.ca`

Two examples

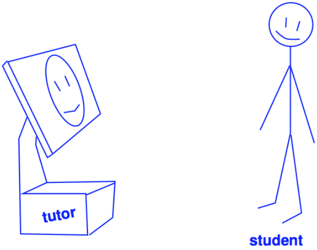

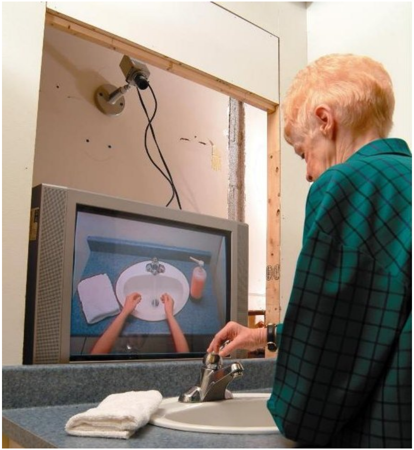

intelligent tutoring

assisted handwashing

Affective “EPA” Space

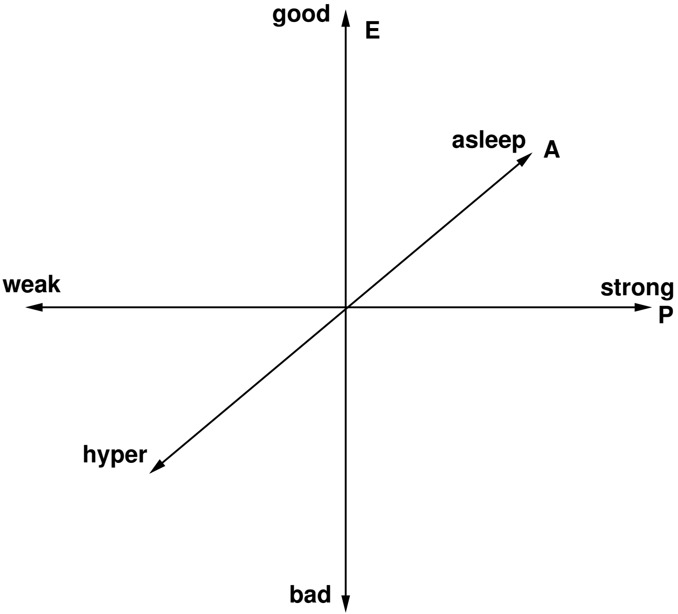

Affective “EPA” Space

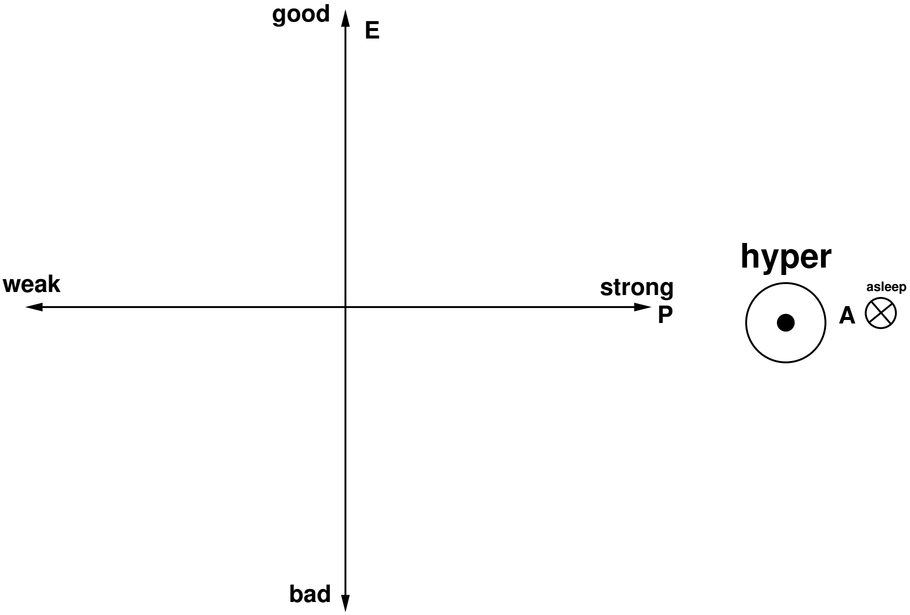

Affective "EPA" Space

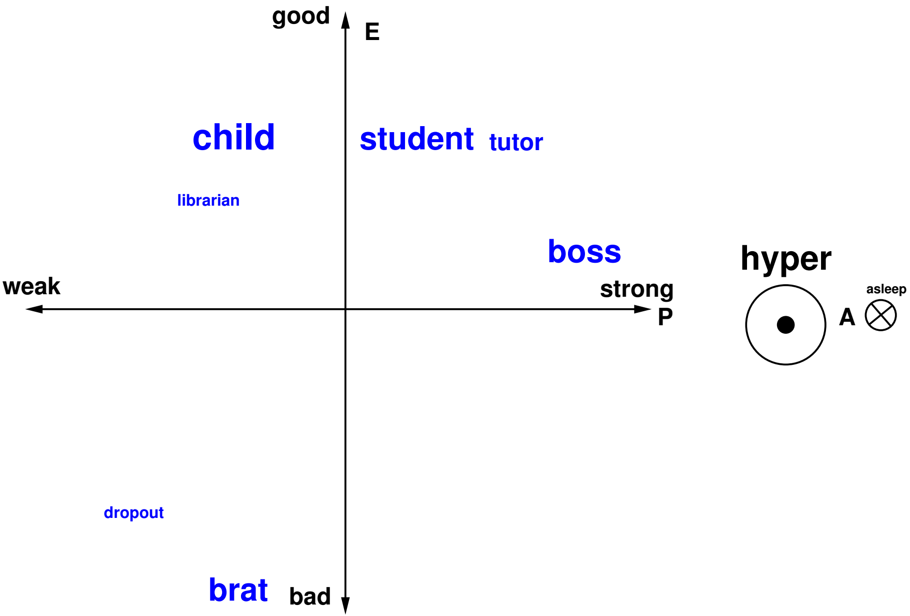

Affective “EPA” Space

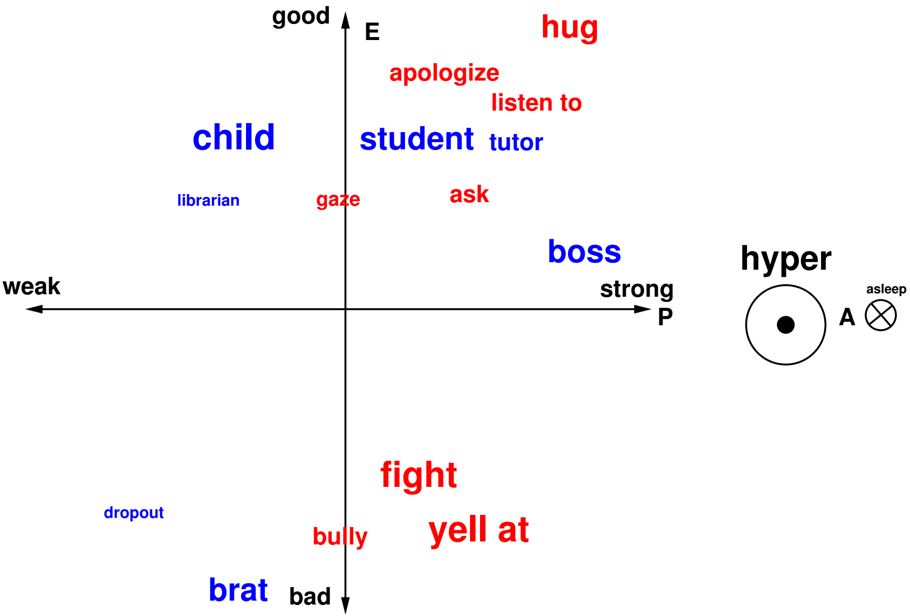

Affect Control Theory

- Actor-Behaviour-Object
- fundamental sentiments: $\mathbf{F} \in [-4.3, 4.3]^9$
- transient impressions: $\mathbf{T} \in [-4.3, 4.3]^9$
- deflection $D = \sum_i w_i(\mathrm{f}_i - \tau_i)^2$
- prediction $\mathbf{T}_{t+1} = \boldsymbol{M}\mathscr{G}(\mathbf{F}_t, \mathbf{T}_t)$ **learned empirically**
- **Affect Control Principle**: actors work to experience transient impressions that are consistent with their fundamental sentiments

ACT Examples

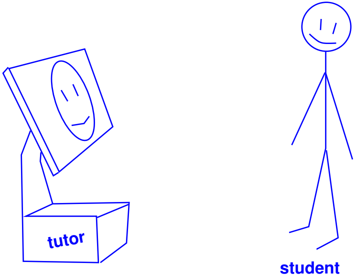

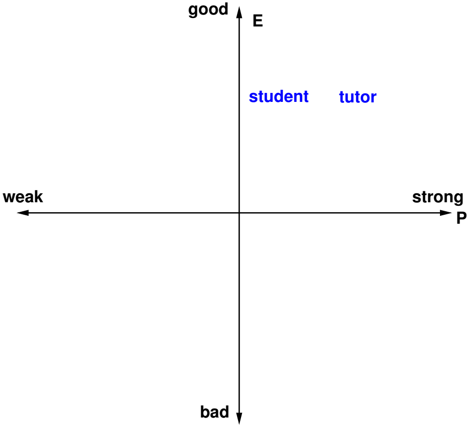

Deflection:

ACT Examples

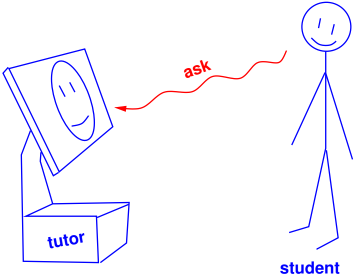

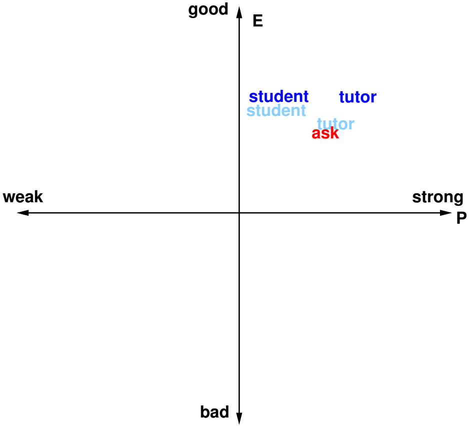

Deflection: 1.0

ACT Examples

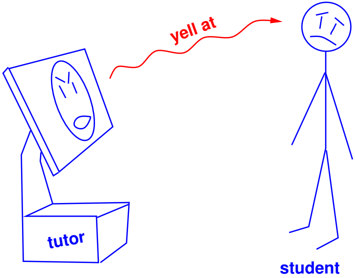

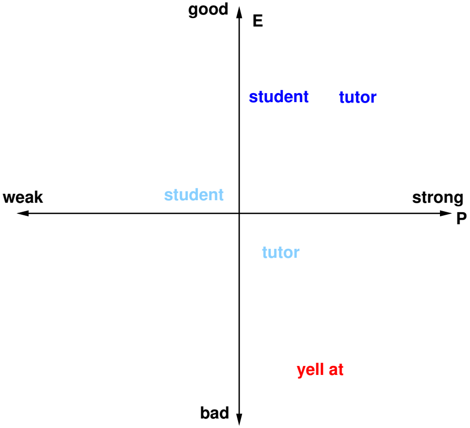

Deflection: 8.0

ACT Examples

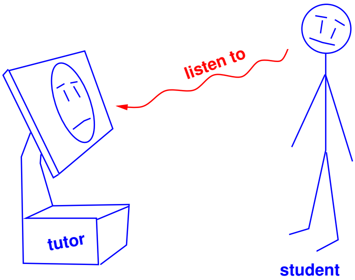

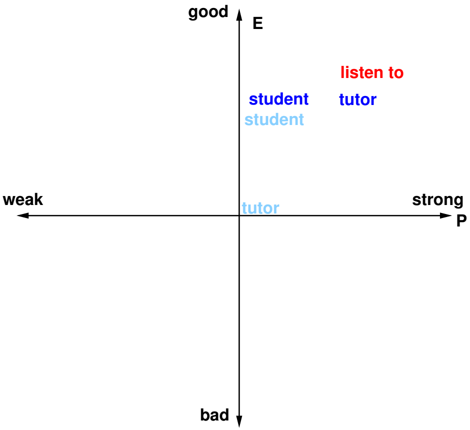

Deflection: 6.0

ACT Examples

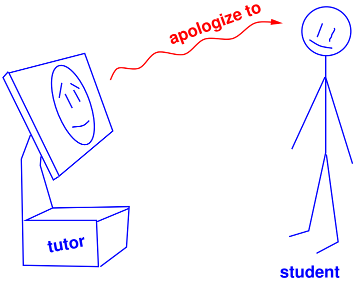

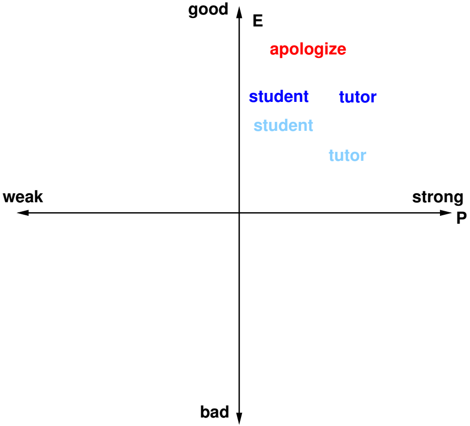

Deflection: 4.0

Interact

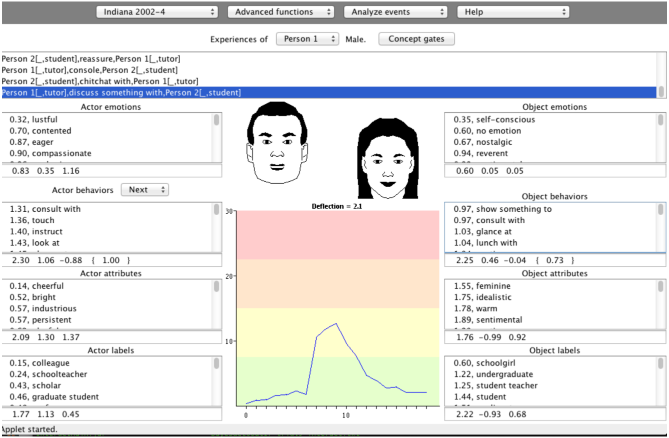

http://www.indiana.edu/~socpsy/ACT/interact.htm

POMDPs for HCI and AI

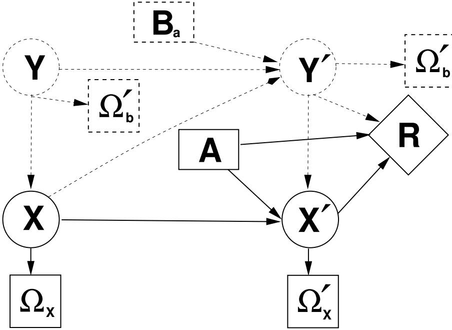

Used for:

- Tutoring systems
- Assistive Technology
- Dialogue systems

- Ω_x, A, X: "propositional"
- Ω_f, B_a, Y: "affective"

BayesAct

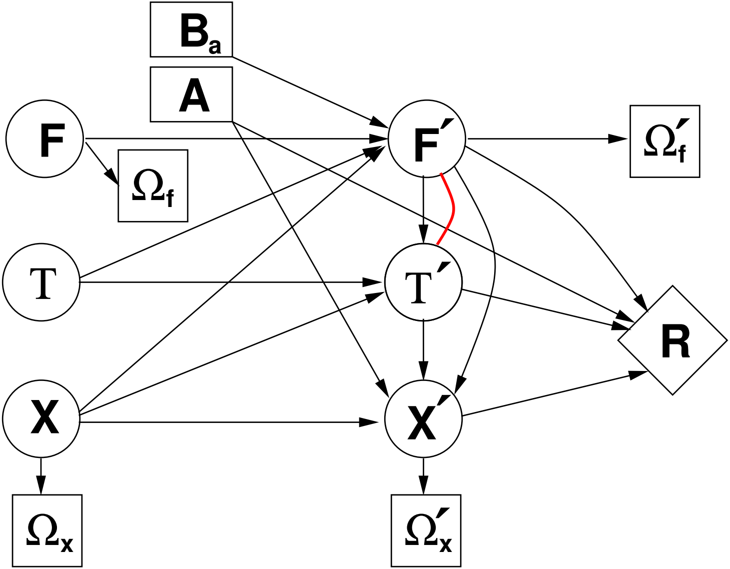

BayesAct Basics

- Fundamental sentiments $\mathbf{F} = \{\mathrm{F}_{ij}\}$ where $\mathrm{F}_{ij}, i \in \{a, b, c\}, j \in \{e, p, a\}$
- Transient impressions $\mathbf{T} = \{\mathrm{T}_{ij}\}$
- Transient dynamics $Pr(\boldsymbol{\tau}'|\boldsymbol{\tau}, \mathbf{f}', \mathbf{x}) = \delta(\boldsymbol{\tau}' - \boldsymbol{M}\mathscr{G}(\mathbf{f}', \boldsymbol{\tau}, \mathbf{x}))$
- Affect Control Potential $\varphi(\mathbf{f}', \boldsymbol{\tau}') \propto e^{-(\mathbf{f}'-\boldsymbol{\tau}')^T \Sigma^{-1} (\mathbf{f}'-\boldsymbol{\tau}')}$
- Fundamental dynamics $Pr(\mathbf{f}'|\mathbf{f}, \boldsymbol{\tau}, \mathbf{x}, \mathbf{b_a}, \varphi)$
- Reward function $R(\mathbf{f}, \boldsymbol{\tau}, \mathbf{x}) = R_x(\mathbf{x}) + R_s(\mathbf{f}, \boldsymbol{\tau})$
- Application dynamics $Pr(\mathbf{x}'|\mathbf{x}, \mathbf{f}', \boldsymbol{\tau}', a)$
- Observation functions $Pr(\omega_{\boldsymbol{f}}|\mathbf{f}), Pr(\omega_{\boldsymbol{x}}|\mathbf{x})$

Bayesact Hybrid DDN - detail

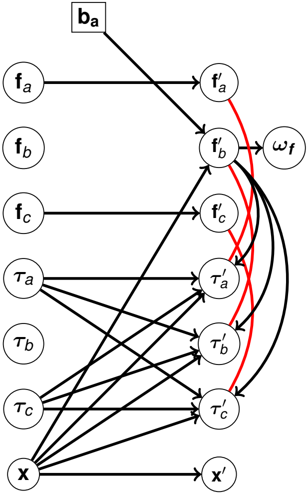

(*a* and some links to $\mathbf{x}'$ left out for clarity)

- $\mathbf{f}'_b$ is set by $\mathbf{b_a}$ when $x_w = agent$ (turn)
- $\mathbf{f}'_b$ is unconstrained (but has evidence from ω_f) when $x_w = client$
- τ_b represents transients about the *last behaviour*
- behaviours are ephemeral, identity persists
- affect control potential φ

Transient Dynamics

$$\boldsymbol{\tau}' = \boldsymbol{M}(\mathbf{x})\mathscr{G}(\mathbf{f}', \boldsymbol{\tau}, \mathbf{x})$$

$$\mathbf{t} \leftarrow \mathscr{G}(\mathbf{f}', \boldsymbol{\tau}, \mathrm{x}_w = \texttt{agent})$$

X_w is the "turn":

- can be `agent` or `client`
- must be part of **X**
- is **partially observable**

$$\mathscr{G}(\mathbf{f}', \boldsymbol{\tau}, \mathrm{x}_w = \texttt{agent}) = [1 \quad \tau_{ae} \quad \tau_{ap} \quad \tau_{aa} \quad \mathrm{f}'_{be} \quad \mathrm{f}'_{bp} \quad \mathrm{f}'_{ba} \quad \tau_{ce} \quad \tau_{cp} \quad \tau_{ca} \quad \tau_{ae}\mathrm{f}'_{be} \quad \tau_{ae}\mathrm{f}'_{bp} \quad \tau_{ae}\tau_{ce} \quad \tau_{ae}\tau_{cp} \quad \tau_{ap}\mathrm{f}'_{be} \quad \tau_{ap}\mathrm{f}'_{bp} \quad \tau_{ap}\tau_{cp} \quad \tau_{ap}\tau_{ca} \quad \tau_{aa}\mathrm{f}'_{ba} \quad \tau_{ce}\mathrm{f}'_{be} \quad \tau_{cp}\mathrm{f}'_{be} \quad \tau_{ce}\mathrm{f}'_{bp} \quad \tau_{cp}\mathrm{f}'_{bp} \quad \tau_{ca}\mathrm{f}'_{bp} \quad \tau_{cp}\mathrm{f}'_{ba} \quad \tau_{ae}\tau_{ce}\mathrm{f}'_{be} \quad \tau_{ae}\tau_{cp}\mathrm{f}'_{bp} \quad \tau_{ap}\tau_{cp}\mathrm{f}'_{bp} \quad \tau_{ap}\tau_{ca}\mathrm{f}'_{bp}]^T$$

$$\mathscr{G}(\mathbf{f}', \boldsymbol{\tau}, \mathrm{x}_w = \texttt{client}) = [1 \quad \tau_{ce} \quad \tau_{cp} \quad \tau_{ca} \quad \mathrm{f}'_{be} \quad \mathrm{f}'_{bp} \quad \mathrm{f}'_{ba} \quad \tau_{ae} \quad \tau_{ap} \quad \tau_{aa} \quad \tau_{ce}\mathrm{f}'_{be} \quad \tau_{ce}\mathrm{f}'_{bp} \quad \tau_{ce}\tau_{ae} \quad \tau_{ce}\tau_{ap} \quad \tau_{cp}\mathrm{f}'_{be} \quad \tau_{cp}\mathrm{f}'_{bp} \quad \tau_{cp}\tau_{ap} \quad \tau_{cp}\tau_{aa} \quad \tau_{ca}\mathrm{f}'_{ba} \quad \tau_{ae}\mathrm{f}'_{be} \quad \tau_{ap}\mathrm{f}'_{be} \quad \tau_{ae}\mathrm{f}'_{bp} \quad \tau_{ap}\mathrm{f}'_{bp} \quad \tau_{aa}\mathrm{f}'_{bp} \quad \tau_{ap}\mathrm{f}'_{ba} \quad \tau_{ce}\tau_{ae}\mathrm{f}'_{be} \quad \tau_{ce}\tau_{ap}\mathrm{f}'_{bp} \quad \tau_{cp}\tau_{ap}\mathrm{f}'_{bp} \quad \tau_{cp}\tau_{aa}\mathrm{f}'_{bp}]^T$$

Matrix "M" dynamics

tdynamics-male.txt gives M^T:

Z000000000	-0.25	-0.09	0.07	-0.15	0.03	-0.02	-0.09	-0.38	-0.03
Z100000000	0.44	-0.02	0.05	0.11	0.03	-0.01	0.01	0.00	-0.01
Z010000000	0.00	0.59	-0.05	0.03	0.15	-0.07	0.00	-0.06	0.00
Z001000000	0.01	0.07	0.65	0.00	0.02	0.29	-0.02	0.02	0.01
Z000100000	0.41	-0.09	-0.08	0.54	-0.14	-0.08	0.11	0.20	0.05
Z000010000	-0.04	0.47	0.11	-0.05	0.72	0.14	-0.01	-0.14	0.00
Z000001000	-0.10	-0.05	0.28	-0.12	-0.04	0.62	-0.02	0.05	0.04
Z000000100	0.02	0.02	0.00	0.05	0.02	-0.01	0.62	-0.07	0.00
Z000000010	-0.02	-0.04	0.01	-0.04	0.00	-0.03	-0.01	0.67	-0.04
Z000000001	-0.01	0.01	0.00	0.00	0.02	0.04	0.04	0.08	0.67
Z100100000	0.05	0.01	0.00	0.02	0.01	-0.01	0.04	0.02	0.00
Z100010000	-0.03	0.02	0.00	0.02	-0.01	0.01	-0.01	0.02	0.01
Z100000100	0.00	0.00	-0.01	0.01	-0.01	-0.01	0.00	0.01	-0.01
Z100000010	0.01	0.00	0.00	-0.01	0.00	0.02	0.00	0.01	0.03
Z010100000	0.01	0.04	-0.02	0.00	0.00	0.01	-0.01	-0.01	-0.02
Z010010000	0.00	-0.07	0.02	-0.01	-0.01	0.01	0.00	0.03	0.00
Z010000010	0.02	0.01	0.00	0.03	0.01	0.00	0.00	-0.01	-0.03
Z010000001	-0.01	0.02	0.01	-0.02	0.05	0.01	0.01	-0.01	0.01
Z001001000	0.00	-0.02	-0.06	0.00	0.00	-0.02	0.03	0.00	0.01
Z000100100	0.13	0.02	0.00	0.12	0.03	-0.01	0.05	0.03	0.01
Z000100010	-0.06	-0.03	-0.01	-0.04	-0.01	-0.02	0.00	0.02	0.00
Z000010100	-0.06	0.01	0.00	-0.05	0.00	0.01	-0.03	-0.01	0.01
Z000010010	0.07	0.04	0.02	0.06	0.01	0.03	-0.01	-0.04	0.01
Z000010001	0.01	0.02	0.01	0.02	0.02	0.00	-0.02	-0.02	-0.01
Z000001010	0.03	-0.01	-0.02	0.04	0.00	-0.01	-0.01	0.00	-0.02
Z100100100	0.03	0.01	0.00	0.03	0.01	0.00	0.01	0.01	0.00
Z100010010	0.02	0.01	0.01	0.03	0.00	-0.02	-0.01	-0.01	-0.01
Z010010010	-0.02	0.00	0.01	-0.03	-0.01	0.01	-0.01	0.00	0.01
Z010010001	0.02	-0.01	-0.04	0.00	-0.02	-0.01	0.00	0.01	-0.02

Transient Impressions

Grouping terms together:

$$\boldsymbol{\tau}' = \boldsymbol{M}(\mathbf{x})\mathscr{G}(\mathbf{f}', \boldsymbol{\tau}, \mathbf{x}) = \mathscr{H}(\boldsymbol{\tau}, \mathbf{x})\mathbf{f}'_b + \mathscr{C}(\boldsymbol{\tau}, \mathbf{x})$$

Where $\mathscr{H}$ is 9×3 matrix, and $\mathscr{C}$ is 9×1 vector
We postulate this is deterministic:

$$Pr(\boldsymbol{\tau}'|\boldsymbol{\tau}, \mathbf{f}', \mathbf{x}) = \delta(\boldsymbol{\tau}' - \mathscr{H}(\boldsymbol{\tau}, \mathbf{x})\mathbf{f}'_b - \mathscr{C}(\boldsymbol{\tau}, \mathbf{x}))$$

where:

$$\delta(\mathbf{x}) = \begin{cases} 1 & \text{if} \quad \mathbf{x} = 0 \\ 0 & \text{otherwise} \end{cases}$$

Bayesact Hybrid DDN - detail

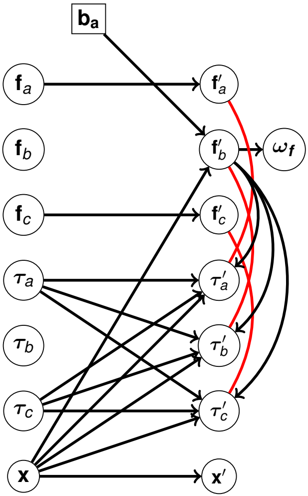

(*a* and some links to $\mathbf{x}'$ left out for clarity)

- affect control potential φ prevents us from sampling from $Pr(\mathbf{f}'|\mathbf{f}, \mathbf{b_a}, \mathbf{x})$ (since we don't know how to predict $\mathbf{f}'_b$ on client turns)
- Instead, we need to sample from $Pr(\mathbf{f}'|\mathbf{f}, \mathbf{b_a}, \mathbf{x}, \varphi)$

Hybrid BN "trick"

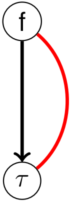

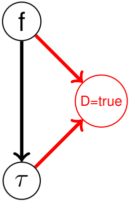

$P(\tau|f)$

$\varphi(f,\tau) \propto e^{-(f-\tau)^2}$

$P(D = true|f,\tau) \propto e^{-(f-\tau)^2}$

$P(\tau|f,\varphi) = \sum_{D=true} P(\tau, D|f)$

$= P(D = true|\tau, f)P(\tau|f)$

$= \varphi(f,\tau)P(\tau|f)$

Sentiment Dynamics

To compute $Pr(\mathbf{f}'|\mathbf{f}, \mathbf{b_a}, \mathbf{x}, \varphi)$, sum over all values of $\boldsymbol{\theta_f} = Pr(\mathbf{f}'|\mathbf{f}, \boldsymbol{\tau}, \mathbf{b_a}, \mathbf{x})$ and $\boldsymbol{\tau}'$:

$$Pr(\mathbf{f}'|\mathbf{f}, \boldsymbol{\tau}, \mathbf{x}, \mathbf{b_a}, \varphi) \propto \int_{\boldsymbol{\theta_f}, \boldsymbol{\tau}'} Pr(\boldsymbol{\theta_f}, \mathbf{f}', \mathbf{f}, \boldsymbol{\tau}', \boldsymbol{\tau}, \mathbf{x}, \mathbf{b_a}, \varphi)$$

$$= \int_{\boldsymbol{\theta_f}, \boldsymbol{\tau}'} Pr(D = true|\mathbf{f}', \boldsymbol{\tau}')Pr(\boldsymbol{\tau}'|\mathbf{f}', \boldsymbol{\tau}, \mathbf{x})Pr(\mathbf{f}'|\mathbf{f}, \boldsymbol{\tau}, \mathbf{x}, \mathbf{b_a}, \boldsymbol{\theta_f})Pr(\boldsymbol{\theta_f}|\mathbf{x})Pr(\mathbf{f}, \boldsymbol{\tau}, \mathbf{b_a}, \mathbf{x})$$

$Pr(\boldsymbol{\tau}'|...)$ is deterministic and selects one value for $\boldsymbol{\tau}'$

$$\propto \varphi(\mathbf{f}', \boldsymbol{\tau}, \mathbf{x}) \int_{\boldsymbol{\theta_f}} \boldsymbol{\theta_f}(\mathbf{f}'; \mathbf{f}, \boldsymbol{\tau}, \mathbf{x}, \mathbf{b_a})Pr(\boldsymbol{\theta_f}|\mathbf{x})$$

$$= \varphi(\mathbf{f}', \boldsymbol{\tau}, \mathbf{x}) \left[\mathbb{E}_{Pr(\boldsymbol{\theta_f}|\mathbf{x})}(\boldsymbol{\theta_f})\right]$$

and affective "inertia", and setting of $\mathbf{f}'$ by $\mathbf{b_a}$ (on agent turn) encoded in the prior:

$$\mathbb{E}_{Pr(\boldsymbol{\theta_f}|\mathbf{x})}(\boldsymbol{\theta_f}) \propto e^{-(\mathbf{f}' - \langle \mathbf{f}, \mathbf{b_a} \rangle)^T \boldsymbol{\Sigma}_f^{-1}(\mathbf{x})(\mathbf{f}' - \langle \mathbf{f}, \mathbf{b_a} \rangle)}$$

where

$$\boldsymbol{\Sigma}_f(\mathbf{x}) = \begin{bmatrix} \boldsymbol{I}_3\beta_a^2 & 0 & 0 \\ 0 & \boldsymbol{I}_3\beta_b^2(\mathbf{x}) & 0 \\ 0 & 0 & \boldsymbol{I}_3\beta_c^2 \end{bmatrix}$$ and $\langle \mathbf{f}, \mathbf{b_a} \rangle = \{\mathbf{f}_a, \mathbf{b_a}, \mathbf{f}_c\}$

Sentiment Dynamics

Probabilistic generalisation of the affect control principle:

$$\psi(\mathbf{f}',\boldsymbol{\tau},\mathbf{x}) = (\mathbf{f}' - \boldsymbol{M}(\mathbf{x})\mathscr{G}(\mathbf{f}',\boldsymbol{\tau},\mathbf{x}))^T\boldsymbol{\Sigma}^{-1}(\mathbf{f}' - \boldsymbol{M}(\mathbf{x})\mathscr{G}(\mathbf{f}',\boldsymbol{\tau},\mathbf{x}))$$

Affective "inertia":

$$\xi(\mathbf{f}',\mathbf{f},\mathbf{b_a},\mathbf{x}) \equiv (\mathbf{f}' - \langle\mathbf{f},\mathbf{b_a}\rangle)^T\boldsymbol{\Sigma}_f^{-1}(\mathbf{x})(\mathbf{f}' - \langle\mathbf{f},\mathbf{b_a}\rangle)$$

Fundamental Dynamics:

$$Pr(\mathbf{f}'|\mathbf{f},\boldsymbol{\tau},\mathbf{x},\mathbf{b_a},\varphi) \propto e^{-\psi(\mathbf{f}',\boldsymbol{\tau},\mathbf{x})-\xi(\mathbf{f}',\mathbf{f},\mathbf{b_a},\mathbf{x})}$$

Rearranging and completing the squares, we get:

$$Pr(\mathbf{f}'|\mathbf{f},\boldsymbol{\tau},\mathbf{x},\mathbf{b_a},\varphi) \propto e^{-(\mathbf{f}'-\boldsymbol{\mu_n})^T\boldsymbol{\Sigma}_n^{-1}(\mathbf{f}'-\boldsymbol{\mu_n})}$$

where

$$\boldsymbol{\mu_n} = \boldsymbol{\Sigma}_n\mathscr{K}^T(\boldsymbol{\tau},\mathbf{x})\boldsymbol{\Sigma}^{-1}\mathscr{C}(\boldsymbol{\tau},\mathbf{x}) + \boldsymbol{\Sigma}_n\boldsymbol{\Sigma}_f^{-1}(\mathbf{x})\langle\mathbf{f},\mathbf{b_a}\rangle$$

$$\boldsymbol{\Sigma}_n = (\mathscr{K}^T(\boldsymbol{\tau},\mathbf{x})\boldsymbol{\Sigma}^{-1}\mathscr{K}(\boldsymbol{\tau},\mathbf{x}) + \boldsymbol{\Sigma}_f^{-1}(\mathbf{x}))^{-1}.$$

Emotions

Emotion: a shared communication mechanism to help resolve deflection

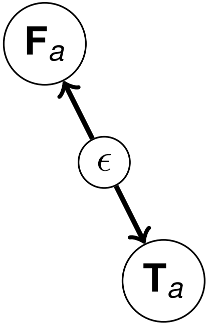

- Emotion $\epsilon = \{\epsilon_e, \epsilon_p, \epsilon_a\}$
- $\epsilon = f(\mathbf{f}_a, \boldsymbol{\tau}_a)$
- $\epsilon \propto (\boldsymbol{\tau}_a - R\mathbf{f}_a - d)$

- When fundamental identity is confirmed by transient, *"characteristic emotion"* occurs
- characteristic $\epsilon \propto (\mathbf{f}_a(1 - R) - d)$

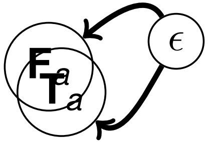

Belief Monitoring

$$b(\boldsymbol{s}_t) \equiv Pr(\boldsymbol{s}_t|\omega_0,\ldots,\omega_t,\mathbf{b_a}_0,\ldots,\mathbf{b_a}_t)$$

which can be written as

$$\begin{aligned} b(\boldsymbol{s}_t) &= \int_{\boldsymbol{s}_{t-1}} Pr(\boldsymbol{s}_t,\boldsymbol{s}_{t-1}|\omega_0,\ldots,\omega_t,\mathbf{b_a}_0,\ldots,\mathbf{b_a}_t) \\ &\propto Pr(\omega_t|\boldsymbol{s}_t)\int_{\boldsymbol{s}_{t-1}} Pr(\boldsymbol{s}_t|\boldsymbol{s}_{t-1},\mathbf{b_a}_t)b(\boldsymbol{s}_{t-1}) \\ &= Pr(\omega_t|\boldsymbol{s}_t)\mathbb{E}_{b(\boldsymbol{s}_{t-1})}\left[Pr(\boldsymbol{s}_t|\boldsymbol{s}_{t-1},\mathbf{b_a}_t)\right] \end{aligned} \quad (1)$$

where $Pr(\boldsymbol{s}_t|\boldsymbol{s}_{t-1},\mathbf{b}_t)$ factored:

$$Pr(\boldsymbol{s}_t|...) = Pr(\mathbf{x}'|\mathbf{x},\mathbf{f}',\tau',a)Pr(\tau'|\tau,\mathbf{f}',\mathbf{x})Pr(\mathbf{f}'|\mathbf{f},\tau,\mathbf{x},\mathbf{b_a}) \quad (2)$$

Particle Filter Updates

Represent $b(\boldsymbol{s}_t)$ with a set of samples

$$b(\boldsymbol{s}) \propto \sum_{i=1}^{N} w_i \delta(\boldsymbol{s} - \boldsymbol{s}_i)$$

iterate:

1. $\{a, \mathbf{b_a}\} \sim \pi^{\dagger}(\mathbf{f}'_b)$ (computed with $b(\boldsymbol{s}_t)$)).
2. Take action $\{a, \mathbf{b_a}\}$ and receive observation ω.
3. Sample unweighted samples $\boldsymbol{s}_i$, from $b(\boldsymbol{s})$
4. Draw a new samples, $\boldsymbol{s}'_i$ from $Pr(\cdot|\boldsymbol{s}_i, \mathbf{b_a})$:
5. Compute new weights $w_i = Pr(\omega|\boldsymbol{s}'_i)$
6. (Resample)
7. New state is $b(\boldsymbol{s}')$, start again

Policy

POMCP (Silver and Veness, NIPS 2010)

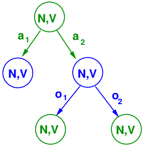

Policy

POMCP (Silver and Veness, NIPS 2010)

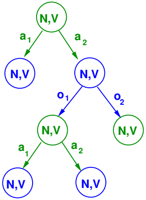

Policy

POMCP (Silver and Veness, NIPS 2010)

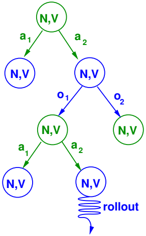

Policy

POMCP (Silver and Veness, NIPS 2010)

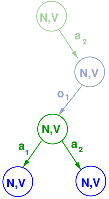

Policy

POMCP (Silver and Veness, NIPS 2010)

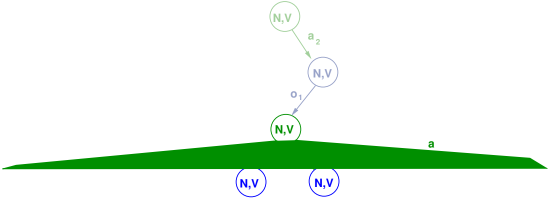

Policy

POMCP (Silver and Veness, NIPS 2010)

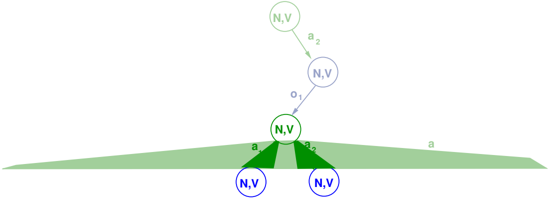

$$\pi^\dagger(\mathbf{f}'_b) = \int_{\mathbf{f}'_a,\mathbf{f}'_c} \int_{\boldsymbol{s}} Pr(\mathbf{f}'|\mathbf{f},\boldsymbol{\tau},\mathbf{x},\varphi) b(\boldsymbol{s})$$
$$= \int_{\mathbf{f}'_a,\mathbf{f}'_c} \int_{\boldsymbol{s}} e^{-(\mathbf{f}'-\boldsymbol{\mu}^\dagger_{\boldsymbol{n}})^T(\boldsymbol{\Sigma}^\dagger_n)^{-1}(\mathbf{f}'-\boldsymbol{\mu}^\dagger_{\boldsymbol{n}})} b(\boldsymbol{s})$$

where

$$\boldsymbol{\mu}^\dagger_{\boldsymbol{n}} = \boldsymbol{\Sigma}^\dagger_n \mathscr{K}^T(\boldsymbol{\tau},\mathbf{x})\boldsymbol{\Sigma}^{-1}\mathscr{C}(\boldsymbol{\tau},\mathbf{x}) + \boldsymbol{\Sigma}^\dagger_n(\boldsymbol{\Sigma}^\dagger_f(\mathbf{x}))^{-1}\mathbf{f}$$
$$\boldsymbol{\Sigma}^\dagger_n = (\mathscr{K}^T(\boldsymbol{\tau},\mathbf{x})\boldsymbol{\Sigma}^{-1}\mathscr{K}(\boldsymbol{\tau},\mathbf{x}) + (\boldsymbol{\Sigma}^\dagger_f(\mathbf{x}))^{-1})^{-1}$$

$\boldsymbol{\Sigma}^\dagger_f$ is the same as $\boldsymbol{\Sigma}_f$ but with unconstrained behaviours

Implementation

Python class `Agent`, subclassed by defining:

- `initXvar`: initialise **X**
- `sampleXvar`: draw a sample from **X**
- `sampleXObservation`: sample an observation from **X**
- `evalSampleXvar`: evalute a sample from **X** given an observation
- input/output mappings

For using with POMCP:

- `reward`
- γ discount factor
- observation and action comparison functions

Key Parameters

param.	default	meaning
α	1.0	variance of a diagonal uniform Σ, the deflection potential covariance. (smaller means the affect control principle is stronger)
β_a	0.01	identity inertia for agent (larger means agent shifts identities more)
β_c	0.01	identity inertia for client (larger means agent thinks client will be shifting identities more)
β_a^0	0.01	initial identity variance for agent (larger means agent is more uncertain of its own identity)
β_c^0	0.01	initial identity variance for client (larger means agent is more uncertain of client's identity)
γ	1.0	model environment noise variance
γ_d	0.9	discount factor (if needed)
N	300	number of samples (use smallest number possible)

Testing

1. Analytical reduction of *BayesAct* to ACT
2. Run *BayesAct* (fixed ids) alongside *Interact*
3. Simulate with unknown identities
4. Simluate adaptation to changing identities
5. proof of concept system (tutoring, handwashing)

Next

- Simulations and Trials of Bayesact
- Signal processing
- Facial Expressions