

DECISION MAKING BY INTELLIGENT AGENTS IN UNFAMILIAR SETTINGS: THE ROLE OF TEAM COMMUNICATION

M. Laura Frigotto and Alessandro Rossi
ROCK (Research on Organizations, Coordination & Knowledge)
Department of Computer and Management Science
University of Trento
Via Inama 5, Trento I-38100, Italy
E-mail: laura.frigotto@economia.unitn.it; arossi@cs.unitn.it

KEYWORDS

Decision making, constraint satisfaction, neural networks, explanatory coherence.

ABSTRACT

We model individual decision making in unfamiliar settings using constraint satisfaction networks. We investigate to what extent team communication might overcome the limits of partially informed and heterogeneous agents, allowing them to improve their choices in dyadic decision making settings (in terms of compatibility with the choices that they would undertake if fully informed). We show that communication has a non-monotonic effect: while initial increases in communication strength result in better performance, when one exceeds an optimal degree, performance declines up to a point in which independent agents perform better than agents communicating in teams. We show that this is largely due to the fact that too much communication confuses agents, blocking the process of sorting alternatives out. Similar considerations on non-monotonicity still holds if we substitute independent choice under communication with group choice based on voting. However, the latter rule, compared to the former one, shows lower performance for all communication strengths.

INTRODUCTION

In familiar settings, decision making is usually based on probabilities drawn from experienced frequencies and, in the connectionist tradition, choice has been typically modeled as a pattern matching activity in which criteria for selecting one alternative among others refer to the closeness of an initial stimulus to patterns stored in memory.

Conversely, unfamiliar settings challenge intelligent agents since no previously encountered event is considerably similar to the one at point and no previously memorized pattern provides an effective solution. Focusing on this kind of issues, we build on Thagard (1989, 1992a, 1992b, 2000) and we propose a decision making model in which agents, in absence of conditional probabilities derived from direct observation, interpret novel settings ascribing causal links between observed events and explanatory hypotheses. In our model, the criteria for selecting one

among many alternatives results from assessing the internal consistency of arguments (a.k.a. “explanatory coherence”), rather than from closeness to patterns formerly and deliberately stored in memory (Keil 2006).

In this paper we propose an extension for modeling decision makers in a team and we investigate the role of group communication on the decision making process. While other contributions have explored this topic using constraint satisfaction modeling, (Hopfield 1982, Marchiori and Warglien 2005), in this paper we suggest the use of such models to represent the process of inferential reasoning through causality rather than pattern matching, thus reproducing a different and broader class of cognitive phenomena.

THE MODEL

We model decision making in unfamiliar settings building on a peculiar class of connectionist networks, namely constraint satisfaction networks, originally proposed by Rumelhart et al. (1986). We refer to Thagard’s variation of this model, the ECHO model (1992b), in which connections are not initialized to store in memory a series of patterns but, on the contrary, result from the interplay of various causality assumptions between specific units of the network, along a line that dates back to Pierce’s abduction theory (Hartshorne et al. 1931-1958).

Agents are denoted by their schemata, which represents how they make sense of the world on the basis of their available data and causal assumptions: in the decision making process, agents consider dyadic theories or interpretations of a setting – e.g.: a go/no go decision to enter in a new market – each composed by a series of hypotheses causally explaining bits of evidence. It is assumed that these competing theories have been elaborated on the basis of the agent’s logical reasoning, knowledge and access to evidence (the process by which the agent collects evidence and elicits hypothetical statements goes beyond the interest of this paper).

Through an iterative algorithm of relaxation, the model displays a general assessment of the concurrent theories and selects the most coherent and convincing interpretation among the two.

It is worth to stress the main difference of this model with respect to “standard” pattern matching models based on constraints satisfaction. According to the latter

models, a series of patterns X_i , for $i=1, \dots, n$, are stored into the network through a careful selection of the weights and the aim of the relaxation process is to assess the ability of the model to recognize the correct pattern X_i when the signal coming from the environment is noisy (the initial activation is $X_i + E$). Conversely, in the ECHO model, each weight is assigned on the basis of the existence of an explanatory relationship that links two different units of the network according to a rule that will be explained in the next Subsection, and, similarly to other constraint satisfaction models, such as Axelrod (1997), the aim of the relaxation process is to observe what solution the model displays.

Structure of the model

Agents.

An agent is modeled as a constraint satisfaction network of n units representing either scraps of available evidence (a.k.a. *evidence units*) or hypotheses giving causal explanations for one or several evidences (a.k.a. *units of hypothesis, explanatory units* or *explainers*); moreover, *special evidence units* (units that are only connected with evidence units) are introduced in the model in order to clamp evidence units to positive activation values (see below for the details).

Explanatory units are grouped into two competing sets (in the following: theories “A” and “B”), representing alternative interpretations of the problem setting. Thus, the network can be imagined as a three-layer graph, in which the top layer represents explanatory units belonging to theory A, the middle layer represents evidence units, while the bottom layer collects the explanatory units belonging to theory B.

Units' Activation.

Units' initial activation represents the agent's original beliefs regarding the units, that is, the agent's preliminary confidence on the environmental evidence and on the various theories' features.

Units' activations, that may take values in the $[-1, 1]$ interval, are updated overtime according to the relaxation rule (see below); the fixed point that is reached represents the final belief of the agent regarding the units of the model.

This steady state may highlight that the agent favors one theory over the other if all explanatory units from one theory - say A - are positive while vice versa occurs for the units of the other theory - say B. If such a configuration does not occur, the model does not give a clear indication in terms of choice of one theory over the other, suggesting a case in which the agent does not judge the collected evidence and/or the supporting hypotheses conclusive.

Connections.

Connections or weights in each agent's network w_{ij} are set in order to reflect the competitive or cooperative relationship that exists between two units of the network (see Thagard 1992b for the full rationale). For sake of

simplicity, we restrict our analysis to simple direct causality, thus longer causal chains are not considered.

Positive connections between units of evidence and explanatory units represent direct causal relations (e.g., event x is causally explained by hypothesis y) and their intensity is coded through positive weights, such that higher weights correspond to higher causal relationships.

In order to introduce positive feedback between co-hypotheses, positive connections are also introduced between explanatory units (from the same theory) that jointly explain the same set of evidence units.

Finally, competitive relationships are modeled by introducing negative connections between explanatory units from opposing theories, that jointly explain the same set of units of evidence.

Formal Procedure for Connections' Initialization.

Let $\mathbf{s} = (s_1, \dots, s_i, \dots, s_n) = (a_1, \dots, a_k, b_1, \dots, b_l, e_1, \dots, e_m) \in \mathbb{R}^n$ be the vector of the activations of all the units of the network, where k and l represent the number of explanatory units belonging, respectively, to theory A and B, m is the number of evidence units (note that $n = k + l + m$). Let also $\mathbf{f} = (f_1, \dots, f_i, \dots, f_n) \in \mathbb{R}^n$ (with $f_i = 0$ for $i = 1, \dots, k + l$) be the vector of the activations of the special evidence units.

Let W be a $n \times n$ null matrix. Define α as the excitatory default value for assigning positive connections among units and β as the inhibition default value for assigning negative connections among units. Then, the weights w_{ij} (for $i = 1, \dots, n, j = 1, \dots, n$) are assigned according to the following steps:

Step 1. *Positive connection between an explanatory unit and a unit of evidence*: for each unit of type e that is causally explained by one or more explanatory units of type a :

- i. let i corresponds to the position of the unit e in \mathbf{s} ;
- ii. compute the number r of explanatory units of type a that explain s_i ;
- iii. for each one of the r explanatory units of type a that explain s_i :
 - a. let j correspond to the position of the unit a in \mathbf{s} ;
 - b. set $w_{ij} = w_{ji} = \alpha / r$;
- iv. repeat step 1. for theory B.

Step 2. *Positive connections between explanatory units that belong to the same theory*: for each couple of units of type a that are co-hypotheses (they jointly explain one or more units of evidence):

- i. let i correspond to the position of one unit of type a in \mathbf{s} and j correspond to the position of the other unit in \mathbf{s} ;
- ii. set $w_{ij} = w_{ji} = \sum_{e=k+l+1}^n w_{ie} \cdot \mathbf{I}$,
where $\mathbf{I} = \begin{cases} 1 & \text{if } w_{ie} = w_{je} \\ 0 & \text{otherwise} \end{cases}$;
- iii. repeat step 2. for theory B.

Step 3. *Negative connections between explanatory units belonging to different theories*: for each couple of units, one belonging to theory A and the other one belonging to theory B, that competitively explain one or more units of evidence:

- i. let i correspond to the position of the unit of type a in \mathbf{s} and j correspond to the position of the unit of type b in \mathbf{s} ;
- ii. set p as the number of units of evidence that are jointly explained by s_i and s_j ;
- iii. set q as the overall number of co-hypotheses (of type a and b) that jointly explain the units of evidence at step ii;
- iv. set $w_{ij} = w_{ji} = \beta p / (q/2)$.

Note that W is symmetric, $w_{ii} = 0$ for $i = 1, \dots, n$ and $w_{ij} = 0$ for $i = 1 + k + l, \dots, n$ and $j = 1 + k + l, \dots, n$.

Relaxation rule.

Units' activation values are updated through an algorithm that increases the degree of coherence of the network in the sense that it performs a gradient-descent path towards levels of activation of the units that better satisfy constraints (see McClelland and Rumelhart 1989, Hertz et al. 1991 for a treatment of similar rules). At each iteration, units' activation levels are synchronously updated according to the following rule (Thagard 1992b, 2000):

$$s_j(t+1) = (1-d)s_j(t) + \begin{cases} \text{net}_j(\max - s_j(t)) & \text{if } \text{net}_j > 0, \\ \text{net}_j(s_j(t) - \min) & \text{otherwise} \end{cases} \quad (1)$$

where $s_j(t)$ is the activation of unit j at time t , d is a decay parameter that, at each iteration, weakens the activation value of every unit. Min and max represent, the lower and upper boundaries of the units' activation and are generally set, respectively, at -1 and 1. Finally,

$$\text{net}_j = \sum_i w_{ij} s_i(t) + f_j \quad (2)$$

is the net input to unit j , computed as the sum of the activation of all the units weighted by the connections w_{ij} linking each of these units with unit j . Note also that, in the case of evidence units, the net input also includes the value of the corresponding special evidence unit. It is worth to mention that the formal treatment of this model is still incomplete. In particular, there is no proof of convergence of the system to a stationary state, nor of coherence maximization, since through relaxation the system might settle on a local maximum. However, there is a considerable body of literature (Thagard 1989, 1991, 1992a, 199b, Nowak and Thagard 1992a, 1992b, Eliasmith and Thagard 1997) that has shown convergence towards fixed points in finite time. In the simulations reported in the next Section, we will employ a choice of parameters consistent with previous

literature and we will study the issues of convergence and of multiplicity of local maxima (fixed points).

Group interaction.

A group is a set of p agents modeled as a "network of agents' networks" (Hutchins 1995; Marchiori and Warglien 2005). Communication between two agents is modeled by linking each unit s_j of one agent with the corresponding unit s_j of the other agent. In this respect, communication is here intended as a parsimonious activity of beliefs exchange, in which only the activation of the units, and not the schemata is shared. In this group model, the vector of units is represented by the union of the p agents' vector of units (\mathbf{s}), while the weight matrix contains both the individual weight matrices (that are arranged as $n \times n$ blocks along the main diagonal) and the communication matrices (that are arranged as $n \times n$ blocks outside the main diagonal). The strength of communication is modeled through the communication intensity parameter $\delta \geq 0$; note that for $\delta = 0$ no communication occurs and the model reduces to a mere union of independent agents. Also, we assume the simplest form of communication: each agent communicates with everyone else with the same strength (each communication matrix has $\alpha\delta$ over the main diagonal and 0 elsewhere). This model is still a constraint satisfaction network and we apply the same relaxation rule for modeling individual decision making under team communication.

THE SIMULATION MODEL

Let us first describe the baseline model, that is, a model in which an agent has an accurate knowledge of all the relevant environmental evidence (e_i), explanatory units (a_i and b_i), and causal connections (w_{ij}). We refer to this baseline model as the *full information and knowledge* (FIK) treatment. For tractability purposes, we restrict, in running all subsequent simulations, to the case in which $k = l = m = 4$. Table 1 shows the FIK treatment.

Table 1: List of Explainers in the FIK Treatment

| unit of evidence | explainers in theory A | explainers in theory B |
|------------------|------------------------|------------------------|
| e_1 | a_1, a_2, a_3 | b_1 |
| e_2 | a_3 | b_2, b_3 |
| e_3 | a_2, a_3, a_4 | b_3, b_4 |
| e_4 | a_3, a_4 | b_4 |
| tot. nr. | 9 | 6 |

Figure 1 represents the schemata and shows the cooperative (solid lines) and the competitive (dashed lines) connections (cooperative links among explainers and actual weights w_{ij} , that are omitted for clarity).

In the next Section we will use this treatment as a benchmark for comparing the performance of agents having an incomplete description of the schemata represented in Figure 1; this treatment will be referred as *partial information and knowledge* (PIK).

These bounded agents are modeled using a table of explainers that represents a subset of the corresponding table in the FIK treatment.

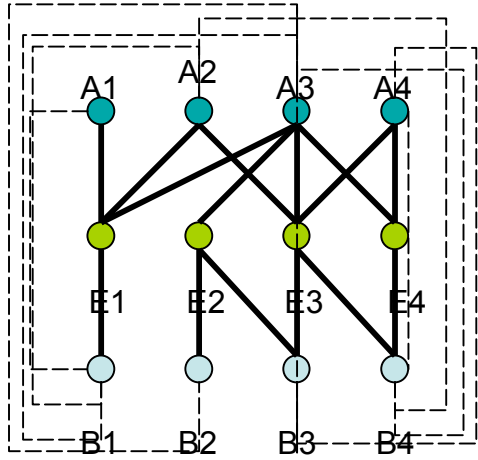


Figure 1: The Schemata in the FIK Treatment

Table 2 summarizes $p=5$ different partially informed agents. Note that agents have full knowledge over theory B, but only partial knowledge over theory A. The reason for this design will be given in the next Section.

Table 2: List of Explainers in the PIK Treatment

| ev. u. | explainers in theory A | | |
|----------|------------------------|------------|------------|
| | agent 1 | agent 2 | agent 3 |
| e_1 | a_1, a_2, a_3 | a_1, a_2 | – |
| e_2 | – | – | – |
| e_3 | a_2, a_4 | a_2, a_3 | a_2, a_3 |
| e_4 | – | a_4 | – |
| tot. nr. | 5 | 5 | 2 |
| ev. u. | agent 4 | agent 5 | |
| | e_1 | a_1, a_2 | a_1, a_2 |
| | e_2 | a_3 | – |
| | e_3 | a_2, a_4 | a_2, a_4 |
| | e_4 | – | a_3, a_4 |
| tot. nr. | 5 | 6 | |

For each setting that we are going to analyze, we derive our results via Monte Carlo simulations with randomly generated initial activation values, from a uniform distribution in the interval $[-1, 1]$. Random initial values are employed as a way to study to what extend final outcomes are affected by initial beliefs of the agent: by randomly initializing the activation of the units we obtain an estimate of the numbers of fixed points and of the size of their basins of attraction.

A final remark regarding the parameterization of the model: we have run all the instances of the model according to a choice of parameters ($d=0.05$, $\alpha=0.04$, $\beta=-0.06$, $s_j=0.01$ for $j=1, \dots, n$, $f_j=0.1$ for $j=k+l+1, \dots, n$) that has been commonly employed in the previous literature on ECHO, and it has shown a remarkable capability of fitting data from various empirical domains. While some sensitivity analysis on these parameters have shown us that qualitative results do not

change over a large parametric space, a deeper robustness analysis has not been performed, and, in the following, we report sensitivity analysis results only on communication intensity.

RESULTS

Full Information and Knowledge Treatment

Figure 2 shows the results of a simulation of 50 agents with random initial activation values under the FIK treatment. Axes measure the normalized Euclidean distance of the vector of activations from the two “ideal” vectors of activations represented by $\mathbf{s}_A=(1, 1, 1, 1, -1, -1, -1, -1, 1, 1, 1, 1)$ and $\mathbf{s}_B=(-1, -1, -1, -1, 1, 1, 1, 1, 1, 1, 1, 1)$. Each line represents the path that one agent follows during the relaxation process while the dots represent the fixed-points.




Figure 2: Paths of Convergence of 50 Randomly Initialized Agents, FIK Treatment

Statistics from 10000 random trials are collected in Table 4.

Table 4: Number of Fixed Points and Size (=Frequency) of the Basins of Attractions in the FIK Treatment

| Fixed Points | Th. A | Th. B | Inconclusive outcome |
|--------------|-------|-------|----------------------|
| 2 | 67% | 33% | 0 |

The relaxation rule converges in less than 500 iterations into one of the following two fixed points:

$\mathbf{s}_1=(0.4296, 0.6149, 0.7470, 0.6017, -0.3441, -0.3278, -0.4221, -0.3054, 0.6877, 0.6967, 0.6907, 0.6965)$;

$\mathbf{s}_2=(-0.2989, -0.5105, -0.4857, -0.4113, 0.5689, 0.4767, 0.6331, 0.6499, 0.6784, 0.6727, 0.6813, 0.6837)$;

corresponding, respectively, to selecting theory A and theory B. Overall, the simulation tells us that in the baseline problem the final outcome is affected by initial activations, and that theory A has a larger basin of attraction than theory B (with an observed frequency of 67% and 33% each). Note also that in this setting, the

model never relaxes toward an “inconclusive outcome”. We take these results as a reference point for evaluating performance of agents in all the PIK treatments.

Partially Informed Agents

Partially Informed Agents without Communication.

Partially informed agents have been *ad hoc* designed in order to induce choices that are at some degree non-consistent with the FIK case. Figure 3 reports the outcome of a simulation of 50 agents of type 1 with random initial activation values (results for agents 2-5 are very similar and hence are omitted).




Figure 3: Paths of Convergence of 50 Randomly Initialized Agents, FIK Treatment, No Communication

In this case, all trials converge to the fixed-point: $s_3=(-0.2447, -0.4658, -0.2447, -0.3893, 0.5284, 0.3413, 0.5917, 0.6144, 0.6843, 0.7035, 0.6816, 0.7136)$; that corresponds to theory B. Similar results hold for agents 2-5. A simple metric for measuring the performance of partially informed agents is represented by the frequency in which the choice under PIK is equal to the one under FIK (given the same initial activation). Only 33% of times agents in PIK select the same theory selected in FIK. In the next Subsection, we will use this as a benchmark to evaluate if communication in teams can increase this level of performance.

The Role of Communication.

Figure 4 shows the results (for agent 1) from a simulation with 50 teams of 5 agents each, with random initial activation values and $\delta = 0.5$ (results for agents 2-5 are not shown because they are very similar). Statistics from 2000 random trials (with groups of 5 agents) show that the relaxation process converges in less than 1500 iterations into one of the following three fixed points (again, we only report agent 1 fixed-points, results for agent 2-5 are qualitatively indistinguishable): $s_4=(0.6626, 0.7741, 0.6745, 0.7053, -0.5896, -0.3934, -0.6300, -0.5356, 0.7683, 0.735, 0.7704, 0.7338)$; $s_5=(0.6483, 0.5915, 0.5704, 0.2623, -0.5443, 0.6014, 0.6112, 0.6233, 0.7658, 0.7895, 0.8048, 0.7902)$; $s_6=(-0.6018, -0.7199, -0.6166, -0.6542, 0.7449, 0.6211, 0.7641, 0.7682, 0.7683, 0.7922, 0.7670, 0.7950)$.

According to team members’ initial activations, agents can choose theory A (s_4), theory B (s_6) or be unable to choose (s_5) and basins of attraction have sizes, respectively, of about 19.36%, 71.80% and 8.84%.




Figure 4: Paths of Convergence of 50 Randomly Initialized Agents, PIK Treatment with Communication

Despite the emergence of this inconclusive outcome, which signals that communication might sometimes confuse agents instead of helping decision making, performance under this communication setting is around 45%, a value considerably higher than the performance with no communication (33%).

Finally, sensitivity over communication strength are reported in Tables 4 and 5.

Table 4: Number of Fixed Points, Size of the Basins of Attractions in the PIK Treatment, for various Communication Strength (Data Pooled over Agents 1-5)

| δ | Fixed Points | Th. A | Th. B | Inconclusive outcomes |
|----------|--------------|--------|--------|-----------------------|
| 0 | 1 | 0 | 100% | 0 |
| 0.25 | 2 | 2.73% | 97.27% | 0 |
| 0.5 | 3 | 19.36% | 71.80% | 8.84% |
| 0.75 | 12 | 14.19% | 55.06% | 30.75% |
| 1 | 41 | 6.88% | 28.56% | 64.56% |

Table 4 shows that increases in communication strength have a direct effect on the number and frequency of inconclusive outcomes.

Table 5: Performance Measures (Frequency of Compatibility of One Agent’s Choice under PIK–FIK Treatments. Data Pooled over Agents 1-5)

| δ | Ind. Choice | Maj. Rule |
|----------|-------------|-----------|
| 0 | 33% | 19.19% |
| 0.25 | 33.98% | 21.73% |
| 0.5 | 45.00% | 38.94% |
| 0.75 | 35.90% | 32.38% |
| 1 | 19.39% | 18.38% |

Table 5 collects performance of individual decision making at various communication strength, highlighting a clear non-monotonic relationship between the two variables: while, at the beginning, increases in the communication strength result in improvements, when one exceeds an optimal threshold, performance declines until independent agents perform better than agents communicating in teams. Table 5 offers an alternative measure of performance, based on a group voting procedure (majority rule). In this case performance is always lower than the ones reported with the previous metric, suggesting that peer pressure or conformity to the majority might be less effective for implementing effective decision making in teams of interacting agents.

CONCLUSIONS

In this paper, we investigated the impact of imperfect information on decision making performance by individual agents in unfamiliar settings and we assessed to what extent communication might better outcomes. We showed a non-monotonic effect of communication. For low levels of communication, group interaction results in better individual performance. However, once the optimal degree is passed, performance declines up to a point in which individuals are better alone than in teams. We claim that this phenomenon is due to the confounding effects of too much communication, which makes agents unable to sort out alternatives. The introduction of a voting procedure based on the majority rule, highlights similar non-monotonic results, but with lower performance for all communication levels.

ACKNOWLEDGMENTS

The authors wish to thank Roberto Gabriele, James March, Ricardo Alberto Marques Pereira, Jay McClelland, Paul Thagard, Massimo Warglien, Enrico Zaninotto, three anonymous referees and the participants to a Cifrem seminar for helpful comments on a previous draft. Financial support from MIUR (projects FIRB03 and PRIN05) is gratefully acknowledged. The usual disclaimer applies.

REFERENCES

- Axelrod, R. 1997. *The Complexity of Cooperation. Agent-based Models for Competition and Collaboration*. Princeton, NJ, Princeton University Press.
- Eliasmith, C. and P. Thagard. 1997. Waves, particles, and explanatory coherence. *British Journal for the Philosophy of Science*, 48, 1-19.
- Keil, F.C. 2006. Explanation and Understanding. *Annual Review of Psychology*, 57, 227-254.
- Hartshorne C., Weiss P. and Burks A. Eds. 1931-1958. *Collected Papers of Charles Sanders Peirce*. 8 vols. Cambridge, MA, Harvard University Press.
- Hertz J.; A. Krogh; and R.G. Palmer. 1991. *Introduction to the theory of neural computation*. Santa Fe Institute studies in the sciences of complexity, Lecture Notes: Vol. 1, Redwood City, CA, Addison-Wesley Publishing Company.

- Hopfield, J.J. 1982. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences of the USA*, 79, 8, 2554-2558.
- Hutchins, E. 1995. *Cognition in the wild*. Cambridge, MA, The MIT Press.
- Marchiori, D. and M. Warglien. 2005. Constructing shared interpretations in a team of intelligent agents: the effects of communication intensity and structure. In *Agent-Based Simulation: From Modeling Methodologies to Real-World Applications. Post-Proceedings of the 3rd International Workshop on Agent-Based Approaches in Economic and Social Complex Systems 2004*, T. Terano (Ed.). Berlin, Germany, Springer Verlag.
- McClelland, J.L. and D.E. Rumelhart. 1989. Explorations in Parallel Distributed Processing. Cambridge, MA, The MIT Press.
- Nowak, G. and P. Thagard. 1992a. Copernicus, Ptolemy, and explanatory coherence. In *Cognitive models of science*, R. Giere (Ed.). Vol. 15, Minneapolis, MN, University of Minnesota Press, 274-309.
- Nowak, G. and P. Thagard. 1992b. Newton, Descartes, and explanatory coherence. In *Philosophy of science, cognitive psychology and educational theory and practice*, R. Duschl and H.R. Hamilton (Eds.). Albany, NY, SUNY Press, 69-115.
- Rumelhart, D.E.; P. Smolensky; J.L. McClelland; and G.E. Hinton. 1986. Schemata and sequential thought processes in PDP models. In *Parallel distributed processing: Explorations in the microstructure of cognition*, D.E. Rumelhart, J. L. McClelland, and the PDP Research Group (Eds.). Vol. 2. Cambridge, MA, MIT Press.
- Thagard, P. 1989. Explanatory coherence. *Behavioral and Brain Sciences*, 12, 435-467.
- Thagard, P. 1992a. Adversarial problem solving: Modelling an opponent using explanatory coherence. *Cognitive Science*, 16, 123-149.
- Thagard, P. 1992b. *Conceptual revolutions*. Princeton, NJ, Princeton University Press.
- Thagard, P. 2000. *Coherence in thought and action*. Cambridge, MA, The MIT Press.

AUTHOR BIOGRAPHIES



LAURA FRIGOTTO is Research Fellow at ROCK (Research on Organizations, Coordination and Knowledge, University of Trento) and Ph. D. Candidate in Economics and Management at Ca' Foscari University of Venice. At present, she is Visiting Scholar at Stanford University. She is interested in organizational decision making.



ALESSANDRO ROSSI holds a Ph. D. in Organization and Management (University of Udine) and is Assistant Professor of Business Economics and Management at the University of Trento. He is interested in behavioral and computational decision making. He is member of ROCK (<http://rock.cs.unitn.it>).