

theory and technology of reputation

BUSINESS
INSTITUTIONS
MARKET REGULATION
GRID SYSTEMS
ONLINE COLLABORATIVE ENVIRONMENTS

SOCIAL KNOWLEDGE for
eGOVERNANCE

erep

::: http://megatron.iiia.csic.es/eRep :::

ISBN 978-88-85059-26-9

Consiglio Nazionale delle Ricerche
Istituto di Scienze e Teconologie della Cognizione
via S.Martino della Battaglia, 44
00185 Rome, Italy
This work was supported by the European Community under the FP6 programme (eRep project, contract number CIT5-028575).

CONTENTS

Reputation is defined as social knowledge that allows for the accomplishment of various social decisions. From the dawn of humankind, reputation has helped regulate society, but it has become even more crucial in this modern age of connectivity, characterized by a dramatically enlarging range of interactions and the continual generation of new types of aggregation. Reputation thus gets applied, under several names, to regulate new societal ties, just as it used to regulate the old ones. But despite this critical role, reputation generation, transmission, and use remain unclear.

This booklet presents the outcomes of a scientific research project pertaining to reputation, carried out by a cross-disciplinary research team known as eRep. The project approached reputation as a complex phenomenon related to the formation and circulation of social evaluations and attempted to consider its role and impact on the maintenance of social order.

The theoretical framework for this project grounded reputation within a social and cognitive perspective. Thus, the analysis focused on how reputational dynamics might be exploited to achieve desired outcomes. We applied this approach to three concrete cases. First, in an electronic auction context, we studied the salience of competing signals, both reputational and objective, through laboratory experiments. Second, for Internet services, we confirmed the validity of a reputational system for selecting dependable service providers in a simulated grid. Third, in a conceptual experiment describing a market in which good sellers are rare and volatile, we explored the role of false reputation. This booklet reports briefly on the findings, as well as the methodology and technology that produced them.

Finally, we suggest some practical implications and suggestions aimed at specific interest groups.

Reading path // Content aimed at

- business
- readers interested or involved in local policymaking
- research groups
- readers interested or involved in management of online communities

An ancient artifact for modern challenges

If we were to list the most influential and recurrent social constructions over time, reputation would undoubtedly appear. In human societies, the exchange of social evaluation dictates partner selection, social control, and coalition formation—to name just a few of the main functions of reputation. From the very moment a community targets and evaluates an agent, the agent's life changes. Reputation represents the immaterial equivalent of a scarlet letter sewn onto one's clothes but is even more powerful than any physical designation, because the individual displaying a particular reputation may not even perceive the evaluation, nor can he or she necessarily control or manipulate it.

Agents?

We use the term "agent" in this booklet to refer to an entity that is able to perform actions autonomously in a given context. Agents include users (i.e., persons, but also companies and institutions) and artificial entities capable of (limited) autonomous action, called software agents.

stakeholders' decisions often depend on a firm's reputation. In everyday life, reputation works as a compass to help people avoid dangerous partnerships and find reliable collaborators, in both small and large social groups. Reputation remains a social artifact based on ancestral activities, such as word of mouth, chatty talk, and grooming. These apparently frivolous occupations, which have kept us busy for an important share of our lifetimes, actually enable us to make better decisions and have important effects on establishing what constitutes acceptable or unacceptable behavior in a society. From an evolutionary point of view, gossip and reputation complement social norms: Reputation evolves along with implicit norms to encourage socially desirable conducts, such as benevolence or altruism, and discourage socially unacceptable ones, like cheating.

Gossip and the origins of languages

Amongst the most fascinating theories on the origins of language is that proposed by Robin Dunbar. He explains the beginning and uses of language as grooming and gossip, highlighted by the abilities and limits of language as part of human life.

Building and maintaining a good reputation is paramount and, in some contexts, essential to survival, as in competitive markets, in which

Online reputation

As soon as the Internet became a context for social interaction, a digitized version of reputation appeared. Before the Internet era, reputational information was locally nested; its reach was limited by geography and social boundaries.

With the advent of the Internet, with its online social networks, online product review sites, and powerful online search engines, reputational information about a target began to spread at a faster pace and to a wider range.

The more the Internet diffuses online social problems that once were limited to the brick-and-mortar world, the more new implementations of that ancient artifact emerge. At the same time, the Internet continues to shape that old artifact, adapting it to online settings to perform a regulatory role.

Principles derived from the use of reputation can regulate and improve efficiency in a variety of contexts that involve interactions among individuals or among individuals, institutions, and organizations. This booklet presents a particular theory of reputation; in the next section, we describe research conducted under the auspices of this project, which addresses issues of fraud in Internet auctions, Internet services, and false reputations. We draw on these results to present some practical implications and suggestions aimed at specific interest groups.

Reputation as a means to avoid adverse selection

Information asymmetry—a situation in which one party in a transaction has more or better information than the other—can create a lemons market, as the economist Robert Akerlof (1970) showed in a classic paper. To alleviate the lemons problem, information about the seller needs to circulate. Such information might be transmitted, in the form of a direct image, from an agent that previously has interacted with the target to other important agents.

theory

Image and reputation: two levels of information

From a social-cognitive perspective, understanding reputation requires distinguishing it from another social artifact, which we call image. Both pertain to the evaluation of a given object (the target), which may be an individual or a group, as developed by another social agent (the evaluator).

Image is the output of the process of evaluation of another agent, fed by trusted communications, direct experience, or both. In their social lives, people continuously assess their colleagues, friends, and partners on their personal features, competence, behaviors, and so on. These evaluations reflect the social images of those agents. In other words, image is an assessment of the positive or negative qualities of a target according to a norm or a competence. Thus, image is attributable; the identity of the evaluator is always clearly expressed, such as in the sentence, "I believe John is a good guy" (i.e., the evaluator, I, is clearly identified).

Reputation interrelates with but also differs from image. Image is the set of evaluative beliefs about a given target; reputation pertains to the process and effect of transmitting social evaluations. Image is assumed to be true by the agent that holds it; reputation is the voice the agent considers spreading. Thus, reputation focuses on the transmission of social evaluations, not the truth of their content, as perceived by agents. For example, if a certain company has a good public reputation, I might still have a bad image of it, because of my negative experiences with the company. Note that the agent who holds a negative image still might diffuse the good reputation, at least in some social settings (e.g., to be considered an insider, because he or she works for the company).

The difference between image and reputation appears even clearer if we consider that agents who spread reputation do not need to commit to the truth of the information. On the contrary, transmitting an image implies the commitment of the transmitting agent to the evaluation content.

With regard to reputation, agents are more likely to transmit uncertain information, and a given positive or negative reputation may circulate across a population of agents even if the majority of those agents do not believe its content.

In terms of the social sets involved, an image refers to three sets of agents:

-➔ Those who share the evaluation, or evaluators

-➔ an evaluation Target (T)

-➔ a set of Beneficiaries (B), or agents sharing the goal, or the norm on which bases the evaluation takes place.

In addition, reputation involves third parties, or gossipers. A third party is an agent in the position to transmit reputation without being responsible for that evaluation and who thus enlarges the social network.

Benevolence or prudence

When circulating information, gossipers may follow different strategies, depending on the direction of their benevolence. According to the agents' autonomy and self-interest, these strategies might be the following:

1. In the case of benevolence toward the beneficiaries, gossiping tends to be negative and critical. The aggregate reputation then follows some prudence rule, such as "pass on negative evaluations even if uncertain; pass on positive evaluations only if certain." This rule promotes the circulation of a reputation that exacerbates the negative characteristics of the target. Voices circulating among pupils (in both B and E roles) about their teachers (T role) often exhibit such characteristics. The evaluators E tend to be benevolent toward themselves (B) but have no intersection with the teachers (T) and thus sense no benevolence toward them.

The distinction we make between image and reputation is not simply that found in the literature between **direct experience** and **communicated information**. An agent can strongly believe that the target exhibits a certain characteristic (e.g., reliable contract fulfillment) even if it lacks any direct experience with the target and instead depends on reported experiences. Alternatively, an agent might decide to send a piece of information that comes from his or her experience but mask it as anonymous or claim to have just heard the information, to hide the actual amount of involvement with the target.

Why transmit reputation?

Agents transmit reputation, acting as gossipers, to demonstrate to other agents that they have access to information and are willing to share it. They transmit this information for reasons such as altruism, status improvements (i.e., to be considered a good guy), and building a social system that functions according to their preferences. Active gossipers generally constitute "in-groups," because they are perceived to share the criteria for image formation, to be interested in the spread of reputation, and therefore to adopt beneficiaries' goal(s) and cooperate with evaluators. Consequently, it is a good idea for agents to spread information about reputation as soon as they receive it. However, several factors may affect the convenience of contributing to this transmission, including certainty about and acceptance of the evaluation, the reputation of the source of the information, a sense of responsibility and accountability for the effects of distributing the evaluation to others, and benevolence toward the beneficiary versus toward the target or none at all.

theory

BENEVOLENCE OR PRUDENCE

Overlapping	M ∩ E	E ∩ T	E ∩ B	B ∩ T
HIGH	underprovision	Underprovision Overestimation	Provision Underestimation	Underprovision overestimation
LOW	provision	Provision Underestimation	Underprovision	Provision Underestimation

2. When benevolence of transmitters is more target-oriented, a courtesy rule, such as "pass on positive evaluations even if uncertain; pass on negative evaluations only if certain," becomes likely. In a courtesy equilibrium, nobody expresses critiques because of their fear of negative reciprocation (i.e., retaliation). The prevalence of positive evaluations on eBay might be attributed to this effect. Recent (March 2008) changes in the eBay feedback policy (sellers' evaluations of buyers now are restricted to positive values) could represent a tentative attempt to minimize overlap between the sets of evaluators and targets and thereby dampen the courtesy rule that resulted in a striking 98% of positive evaluations.

3. If gossipers do not have strong benevolence toward either group (B or T), theory predicts scarce reputation transmission. In such cases, information may be induced by some institutional reinforcement. For example, teachers' grades of students likely fall in this category. In electronic contexts, feedback could be required to complete a transaction or prompted by sending e-mail reminders to buyers.

The systematic application of a courtesy or prudence rule in reputation spreading may induce aggregate circulation of selected forms of evaluations, whether positive or negative.
Because this selective transmission depends on the specific motivations of the gossiper (and evaluator) agents, we want to understand the social and cognitive conditions that determine the application of those rules.

Consider two significant cases:
1. general overlapping of evaluators (E), target (T), beneficiaries (B).
2. overlapping of E and B, while T has a separate status.

On case 1, we expect positive evaluations to prevail, and therefore surpass by far the number of critic evaluations.
On case 2, inversely, we expect the emergence of a sort of "social alarm", useful to warn the community of beneficiaries-evaluators about a possible danger coming from an external target.

How Reputation mechanisms can reduce Internet fraud

The role of reputation in Internet commerce, especially for Internet auctions, has received extensive attention, both from researchers and in the industry. The rise of the Internet has changed the theater in which consumers make consumptive decisions; they increasingly purchase goods using auction sites and obtain information about products from review sites and Weblogs. Rather than engaging in face-to-face contacts with a salesperson or expert, an increasing number of consumers rely on information provided by anonymous others on the Internet. Whereas online purchasing has made life easier for many consumers, many others have suffered misleading recommendations from sellers and so-called experts.

Unfortunately, fraud has become a well-known risk in online purchasing, especially because sellers can easily hide their real identity. Moreover, on Weblogs and review sites, information sometimes gets manipulated fraudulently.

> A Dutch company posted fake questions on an advice site, which it then answered using a fake expert who recommended that particular company. Extensive research into the IP addresses of posters can ultimately reveal their identity, and enable legal action, but such laborious investigations can commence only after a fraud has been committed and detected; many consumers have no idea how to start such a process.

Thus, it might be more effective and practical to create a system that indicates the trustworthiness of a person on the Internet, before the transaction begins. If it were possible to mobilize and combine the experiences of many consumers, the proportion of fraudulent information and misinformed people on the Internet might decline. This challenge is substantive, especially with the rise of Web 2.0 applications, which likely will increase the frequency of interactions among anonymous persons.

Many companies, such as auction sites, realized early in their evolution that consumer trust would be essential to their business and therefore developed reputation systems that allow consumers to rate a salesperson or transaction easily. The increasing number of online purchases suggests such reputation systems will become ever more important indicators of trust on the Internet, not only for auction sites but also on sites that offer advice from expert consumers or various Web 2.0 applications in different sectors that require users to accept or make contacts. Even in the context of negotiating about computing resources (e.g., auctions for disk space or CPU time),

findings

HOW REPUTATION MECHANISMS CAN REDUCE INTERNET FRAUD

reputation systems may help lower the proportion of fraudulent activities, increasing the robustness and efficiency of Internet-based applications.

Questions about reputation in purchase decisions

To what extent do such reputation systems affect the purchase process of consumers? Do consumers really appreciate and use these systems? Do consumers rely more on reputation scores when they purchase an expensive product? How important is trust when product quality tends to be heterogeneous (e.g., variations in wear and tear)? Is a good description of the product just as important as, for example, a high-quality picture? In the bidding process, might people forget about reputation issues when they enter a bidding competition?

To unravel and respond to these important questions, we conducted a series of experiments in which people bid on different products in an experimentally controlled Web auction setting.

Experimental results: Effects of reputation and image

The results from the laboratory experiments show that the image and reputation of the seller are the dominant choice criteria in the selection of sellers.

The expression "Web 2.0" describes changing trends in the use of the World Wide Web technology, aimed at enhancing creativity, communications, information sharing, collaboration. It refers to the increased interconnectivity and interactivity of Web-delivered content. Web 2.0 concepts include the development and evolution of Web culture communities, such as social networking sites (Facebook), video-sharing sites (YouTube), wikis, and blogs. Most Web 2.0 platforms rely on some reputational mechanisms.

That is, the effect of image is more pronounced than the effect of reputation, but reputation nonetheless plays an important role in shaping bidders' decisions. Obtaining a good reputation and high image score is therefore essential for success. However, trustworthiness results not only from feedback provided by other buyers but also from other cues, such as photos and product descriptions, traceability and location of seller, and offered payment methods (e.g., PayPal). Sellers also need to realize that their reputation goes beyond the auction site; consumers can trace sellers and find additional reputational information on bulletin boards or Web sites that offer personal information

The **laboratory experiments** involved human subjects who interacted on a Web auction platform that had been specifically designed for this research. The platform uses the features of popular auction sites, such as the display of several reputation scores, based on our distinction between image and reputation:

-» **eBay's feedback score**: Based on the direct experiences of unknown others, or image-based information.

-» **Friend's experiences score**: Based on the direct experiences of friends, another form of image-based information.

-» **Hearsay/Rumour score**: Evaluations of the target collected by an artificial Web site that searches the Web (based on e-mail addresses of the sellers) without telling the original source; this score represents reputation-based information:

	Less expensive	More expensive
Homogeneous	Lonely Planet guide	New LCD television
Heterogeneous	Second-hand design chair	Second-hand LCD television

about their identities (e.g., Rapleaf).

Reputation seemingly should play a more important role for products that vary in quality relative to those products that are more homogeneous, yet we find that this is not the case. Even for standardized products, such as a Lonely Planet guide or DVD, reputational feedback scores strongly determine seller selection. However, reputation plays a somewhat more pronounced role when the price of the product is higher. Bidders extend their search activities and are more likely to incorporate feedback scores in these situations, which implies that sellers of valuable products should take even more care of their reputation.

Finally, consumers pay attention to differences among sellers. When there is little or no variation in a category, that criterion becomes less significant, because it does not provide any means to differentiate. Consumers search more and rely more on photos when the sellers vary greatly in terms of the quality of the photos they provide (see figure on next page for details). When all sellers exhibit high- (low-) quality photos, consumers pay less attention to the photos. However, when some have high-quality photos and some have low-quality photos, consumers consider it worthwhile to extend their search activities and closely investigate the photos. If all sellers have good reputation scores, they will search for other criteria as a means to differentiate (e.g., photo quality)—which does not mean that reputation becomes less important.

findings

HOW REPUTATION MECHANISMS CAN REDUCE INTERNET FRAUD

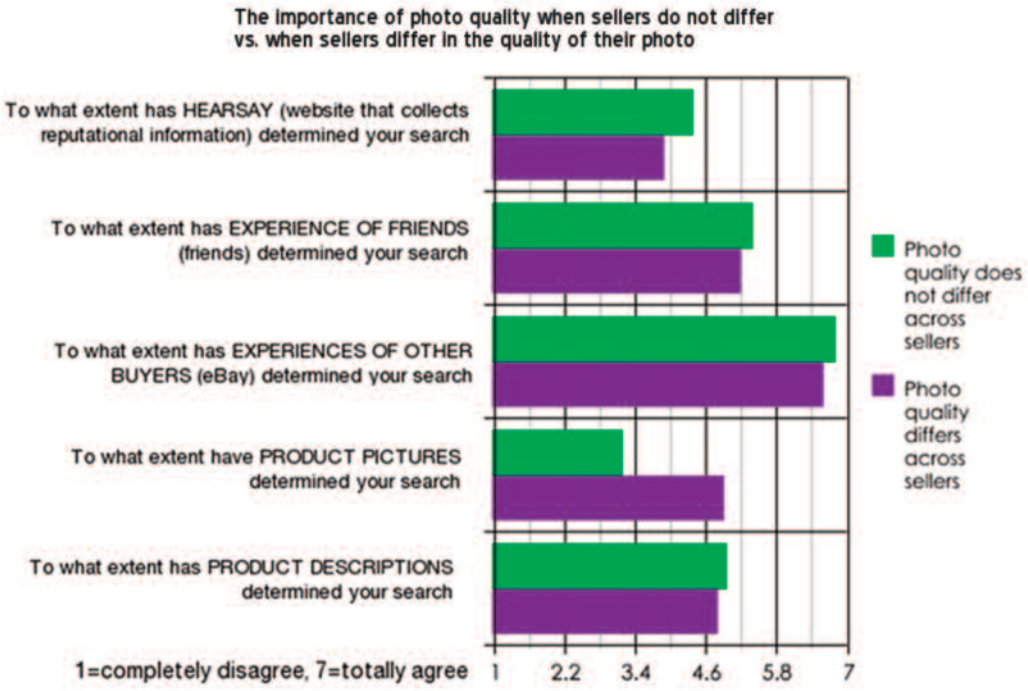

Experiment results: Impact of photo quality and product descriptions

Sellers thus can promote their merchandise through high-quality pictures and clear product descriptions. The importance of photo quality is greatest for second-hand products that are likely to differ in quality (e.g., chairs); without a good photo, it is impossible for consumers to determine product quality. For products such as televisions and Lonely Planet tour guides, photo quality is less important, because these items are highly standardized, and a negative reputation score cannot be mitigated by a high-quality picture. The results also demonstrate that using a catalog picture instead of an actual picture is not harmful when the seller is trustworthy. In the case of an untrustworthy seller

though, a catalog picture is harmful; it would be better for this seller to provide high-quality photos.

Managing disputes and the impact of retaliation

Sellers must remain aware of potential negative evaluations from buyers. Products delivered with scratches or not delivered at all almost automatically lead to negative evaluations that lower the seller's reputation score. However, most consumers do not provide negative feedback when the product simply is delivered late, though they also do not provide positive feedback. Sellers clearly should inform customers as soon as possible about potential delays to avoid serious norm violations and negative evaluations.

The research project also investigated the impact of a system in which the evaluators (buyers) get rated by the sellers (increasing the possibility of retaliation) compared with a system in which evaluators are not rated (no retaliation). When retaliation is possible, we posit that consumers will be more inclined to rate sellers positively. The laboratory experiments confirm this assumption: Consumers provide more positive ratings and less negative ratings when they are evaluated by the seller. However, in severe cases (i.e., no product delivery), people negatively evaluate the other party, irrespective of potential retaliation. Thus, the retaliation effect exists, especially in cases of moderate problems (delivery too late, scratches).

According to our analyses, the sellers' trustworthiness depends most on the eBay feedback score (purchase evaluations from unknown bidders), followed by friends' feedback score (purchase evaluations by friends), and then hearsay (evaluations from the Web site that collects seller information). Consumers grant special status to the first-hand experiences of other bidders. Most respondents (69%) systematically choose a reputable seller, that is, one that earns at least 90% positive scores on both image-based measures (eBay and friends').

findings

The effects of reputation on the Internet of Services

Even after four decades of rapid advances, computing continues to experience revolutionary changes at all levels, including hardware, middleware, and network infrastructure, but perhaps more important at the level of intelligent applications. Emerging technologies such as the semantic Web and Web services have been transforming the Internet from a network of information to a network of knowledge and services. The number of services offered on the Internet is expected to rise dramatically. These intelligent services can communicate and negotiate with one another over the Internet infrastructure to build the so-called Internet of Services (IoS).

The Internet of Services (IoS) requires an efficient allocation mechanism to match the demand and supply of resources. It needs a market. Markets can collect existing resources and service supplies, as well as corresponding demands, and thereby achieve an even utilization by leveling heterogeneous user behavior. Similar to other utilities, services traded in markets in huge numbers tend to be simple in nature. They are distinguishable by their service quality characteristics but equivalent otherwise. In turn, given their equal characteristics, competition occurs by signaling lower prices.

However, with this vision of the IoS, several questions emerge, especially regarding the risks involved in IoS market transactions. The bilateral economic exchanges envisioned for IoS markets will involve risks resulting from strategic and parametric uncertainties. A key mechanism for mitigating such risk, reducing uncertainties, and increasing trust in IoS markets is reputation.

Reputation reduces uncertainty in the IoS by conveying cooperation

Reputation works as a signaling device that distinguishes between trustworthy and untrustworthy transaction partners. Furthermore, it changes the long-term utility functions of the market participants (by introducing potential profit losses for those identified as cheaters) and thereby encourages transaction partners to cooperate. Our results in this area show that reputation is effective in fully automated environments like the IoS. Our simulations of a typical Internet structure, with complex services built upon simpler services running on network nodes support this hypothesis: More jobs will be completed correctly if a trust and reputation mechanism is introduced in the market.

Information spreading; faster reputation effects

A primary influence on reputation is the circulation rate of information in the system. We assume that the farther information spreads in a system, the faster the system expresses reputation information and the more information a single market participant can obtain about a potential transaction partner. For example, if a market only uses direct experience, every buyer must trade at least once with a seller before he or she can assess its trustworthiness and react accordingly. In contrast, if information circulates throughout the whole system, buyers obtain information about a seller faster, and the effects of reputation occur earlier. The goal is finding a suitable trade-off between spending time gathering more reputation information and using that time to produce a service. In the simulation results that we show as an example, contract fullfilment rates are much better with reputation for both kinds of agents (complex service agents and basic service agents).

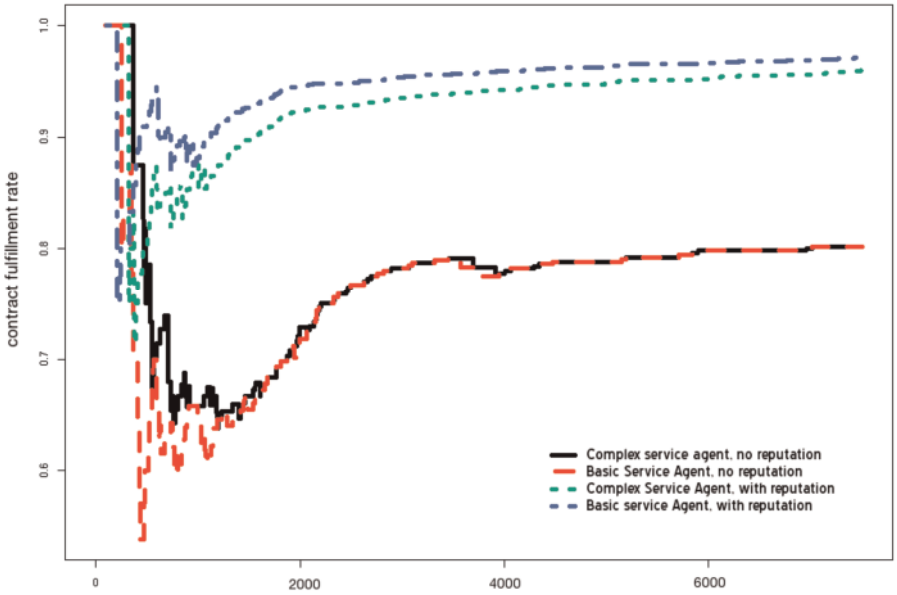

The simulated environment models a two-layer market in which complex service agents buy the needed basic services from other agents, which in turn trade with agents selling low level "resources". Communicatation takes place among complex service agents and basic service agents about the performance of their sellers.
The results show that a sensibly higher fulfillment rate is achieved in presence of reputational information.

findings

When false reputation spreads

Social information helps reduce uncertainty in common decision-making situations in everyday life. Consider the great amount of uncertainty associated with decisions in, for example, the commercial or electoral environments. The presence of image or reputation information helps reduce such uncertainty and facilitates decisions. What happens to a society characterized by a substantial proportion of false social information? This question is not trivial; such a scenario can apply to a wide range of circumstances.

We therefore performed a series of simulation-based explorations of the effect of false information on the formation and revision of social evaluation, employing a model of reputation processing (RepAge) that we developed. The experiment creates an abstract market populated by buyers and sellers who trade goods of different quality. Buyers transmit information about sellers to fellow buyers, but limited stock makes good sellers a scarce resource, which creates a motivation for agents not to distribute accurate information. We studied the reaction of the market as the number of cheaters increases (i.e., lying buyers who provide answers that are opposite to what they believe). The study also includes considerations of image only versus the communication of image and reputation together. Reputation is more efficient for providing high-quality contracts when the information is reliable. As the number of cheaters increases though, information becomes corrupted, and the reputation mechanism fails, as quality plunges. With just image, cheating is ineffective, and quality does not depend on the number of cheaters. In this case, performance becomes worse than that achieved by using reputation when the number of cheaters is moderate, but it is better otherwise.

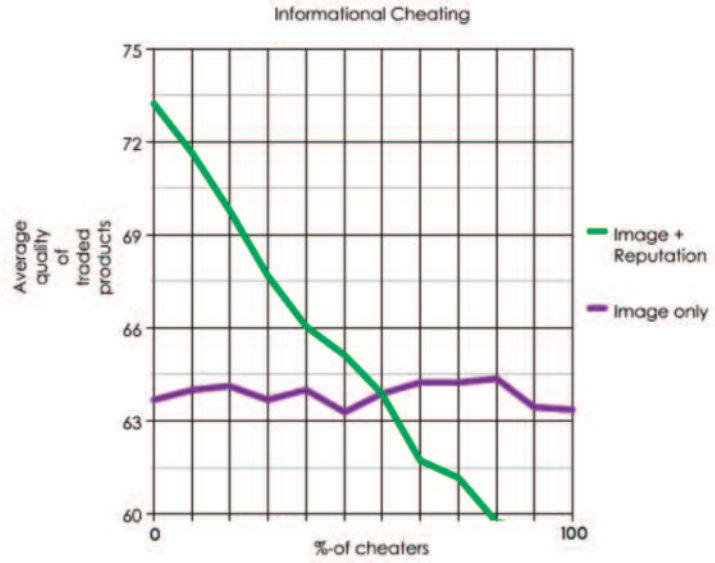

How the research was carried out: the process

findings

This booklet results from a cross-disciplinary collaborative project among several research institutions, namely, the eRep project, which benefitted from synergies among partners who bring solid and diversified disciplinary backgrounds to the table.

The scientific background underlying this research includes cognitive science (CNR-ISTC), organization science and economy (RuG), and computer science and ICT (CSIC-IIIA and UBT). All partners have demonstrated competitive research capacity in highly innovative fields, including agent and multi-agent theory and technology and agent-based social simulation.

During the project, we investigated several research questions connected to the usage and effects of reputation in different environments. These environments range from e-commerce, where humans interact, to hybrid systems and fully automated systems, such as IoS in which software agents interact. When users interact in a social environment with the mediation of carefully designed software tools, the result is a complex system that a single discipline cannot hope to understand or improve. Moreover, communication problems across the different disciplines often make this task even more demanding. To achieve purposeful communication then, we relied on a set of recommendations from established information systems research frameworks.

The cross-disciplinary and cross-methodological foundation of the project enabled us to investigate the research question in complementary ways. For example, our investigation of the use of reputation in e-commerce environments relies on behavioral science. Through laboratory experiments and social simulation studies, we pursued answers to the research questions pertaining to the design of policies, as we describe in this booklet. These findings in turn helped us propose reasonable policies for the IoS environment, which are based on solid knowledge that consists of findings from reputation theory. Moreover, some hypotheses deduced from reputation theory received further support and specification from the project work, which offered a refinement of our theoretical work.

Our research also was characterized by continuous feedback between the development of models or software and their evaluation. This feedback loop enabled us to integrate findings from behavioral research and social simulations into the design of our technology. Knowledge from both fields—empirical

HOW THE RESEARCH WAS CARRIED OUT: THE PROCESS

findings and simulation—helped us refine the software prototype and enhanced our general knwledge base. As a result, we also can offer insights about and support for managing cross-disciplinary and cross-methodological teams to attain beneficial synergy effects.

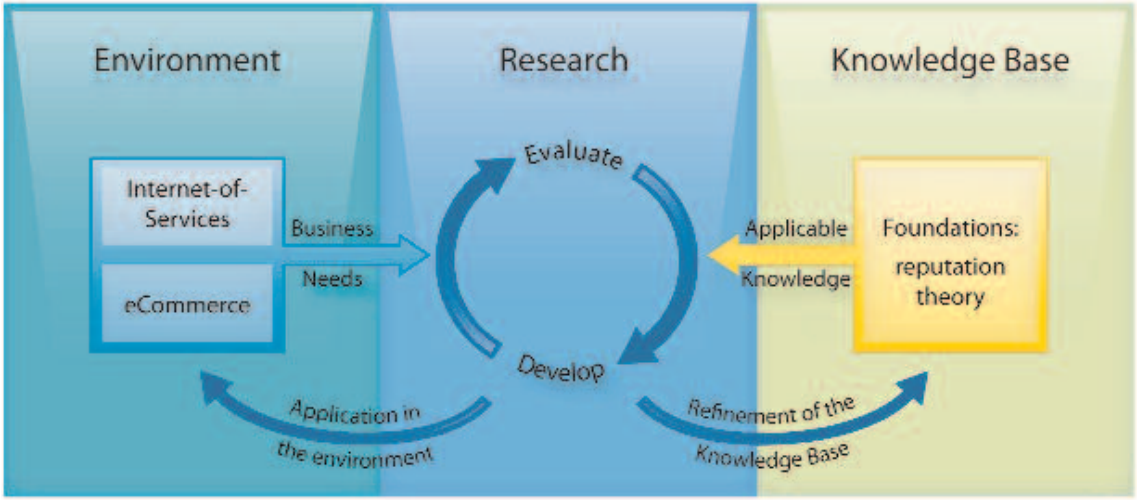

Information Systems Framework (adapted from Hevner, 2004)

How the research was carried out: the technology

The eRep technology advances three features related to computational reputation models:

-➔ The model; that is, the development of new technology to build better models.
-➔ The integration of the model into the agent architecture.
-➔ The use and integration of reputation in virtual environments.

Cognitive computational reputation models

The technology development effort focuses on the design and implementation of cognitive models of reputation. There are several computational models of reputation, but very few have the support of solid theories. A cognitive model of reputation (and the complexity associated with it) is not always necessary, but it is essential in complex mixed societies that include not only virtual entities but also humans. The virtual entities in such societies must deal with human reputation, and thus, cognitive models become relevant.

Reputation models and the agent architecture

The integration of cognitive reputation models with the rest of the components of the agent architecture somehow, in current research, has been neglected. The reputation model appears as a black box, with very few connections to other elements of the architecture (e.g., agent memory, planners). In contrast, the eRep project is oriented toward a tight integration between the reputation model and the components of the architecture. This integration suggests new uses of the reputation model, which is no longer a simple reactive element but instead becomes a proactive element of the architecture.

Adding reputation to virtual societies

Institutions can regulate a complex society and guarantee, up to a certain point, that the society remains robust against improper or unethical behaviors, because they define interaction protocols and norms. The concept of an electronic institution derives from such human institutions. In open multi-agent systems, autonomous entities interact to achieve individual goals, and their behavior cannot be guaranteed. Therefore, similar to what happens in human societies, we require mechanisms to avoid the collapse of the system, despite any local behaviors.

Reputation mechanisms can provide control where institutions cannot, and thus, the integration of reputation models into electronic institutions is one of the key contributions of the eRep project. We

HOW THE RESEARCH WAS CARRIED ON: THE TECHNOLOGY

have designed several software components to support the use of different reputation mechanisms (both centralized and distributed) in the context of an electronic institution. In particular, we have designed and implemented an ontology of reputation that enables agents to exchange information about reputation, even when they are using different models.

Technology that includes both electronic institutions and reputation models then suggests a prototype application in the IoS context. In the prototype, different autonomous nodes, placed in an Internet-like network structure, host agents that are negotiating with one another. Agents can use reputation models to increase their own utility and avoid being cheated when buying resources. These resources (e.g., CPU power, disk space) then may be recombined into high-level services and delivered to external users.

The market itself consists of four types of players that act on two interrelated markets. The first is a resource market, centered on trading computational and data resources, such as processors or memory, between resource agents (sellers of the resources) and basic service agents (buyers of the resources). The second market, a basic service market, involves trades of basic application services between basic service agents (sellers) and complex service agents (buyers of basic services). Compared with consumer markets, these envirnoments exhibit some differences; for example, strong time restrictions usually apply to the purchase of a service or resource in an IoS.

Institutions are a mechanism to regulate a complex society and guarantee to a certain point that this society is robust against wrong behaviours by defining interaction protocols and norms. The concept of electronic institution is inspired from these human institutions. In open multi-agent systems you have autonomous entities that interact to achieve individual goals. The behaviour of these entities cannot be guaranteed. Therefore, and similar to what happens in human societies, you need mechanisms to avoid the collapse of the system in spite of the local behaviours.

Opportunities and Challenges in a Connected World

With the ongoing evolution of the Internet and the continued development of online social networks, many new business and opportunities are emerging and growing. Some have occupied niches of limited amplitude; others have grown massively and with unprecedented speed.

For example, with Facebook and MySpace, people can connect and share information easily. It took Facebook just two years to reach a market audience of 50 million, and by the beginning of 2009, it could count more than 150 million active users. Flickr, a photo management and sharing platform, hosts approximately 3 billion photographs, ranked by an (undisclosed) algorithm according to a parameter the site calls "interestingness." The extent to which this system effectively creates a shared "taste" and propagates it through patterns of imitation of successful photographers remains to be investigated.

The Reddit and Digg social news platforms, widely used in the EU, similarly have been growing their user bases exponentially. Several million users collectively filter and discuss current news items, enough that the sites started presenting themselves as grassroots challengers to professional mainstream media news desks. Yet criticisms challenged the allegedly "democratic" nature of these media when it became clear that only a small proportion of "power users" actually was responsible for determining the front page stories on Digg. Such problems of democracy reflect the probably inherent biases that affect these platforms.

Finally, on eBay, anybody can create an account to buy and sell items, and its overall success results from the success of each of the transactions taking place. Without a reputation mechanism, it would be impossible to predict whether transaction partners will act honestly or are solely attempting to make money by selling items that they never intend to deliver.

In terms of IT governance, these new systems pose an essential problem: Because of their openly distributed nature, anybody can participate in them, and their outcomes depend on the individual decisions and actions of participants. Yet these individuals' own specific goals are opaque and difficult to predict, as well as subject to change in response to various external factors.

To ease interactions and enhance trust and cooperation in online social environments, institutions or distributed mechanisms are necessary to constrain individual behavior with the

OPPORTUNITIES AND CHALLENGES IN A CONNECTED WORLD

help of formal and informal rules. Such institutions and mechanisms also can create indirect incentives for participants to change or improve their behavior without compromising their autonomy.

If they can use reputation, participants in a market can rely on more information than just their own experiences; then, they can use information provided by others to decide whether to trust an unknown transaction partner. In turn, transaction partners are encouraged to comply with their transaction agreements to gain a good reputation and encourage more transaction offers in the future, a form of positive reward. Each transaction partner also must try to avoid a negative reputation, which could even result in a market exclusion as a punishment for fraudulent activities.

Reputation is a distributed mechanism, but online reputation can be regulated by institutions that provide the tools for managing it. The choices that institutions make regarding how reputation gets created, collected, and presented cannot be considered neutral. A theory of reputation, as outlined here, is the key to successful reputation design. Only by consulting a theory of reputation can designers find hints of the answers to questions such as:

-➜ What kind of dynamics does a specific reputation process trigger? How do cultural evolution, opinion dynamics, and idea circulation take place in such an environment? Do these systems move toward conformism, or do they favor innovation?

-➜ When circulating a voice, do gossiping agents follow a prudence rule or a courtesy rule?

-➜ What is the responsibility of the transmitting agent? What are the effects of broadcasting evaluations instead of spreading them through direct links?

When it features a specific focus on online communities and social networks and their distributed nature, the theory of reputation accounts for cognitive problems, such as the problem that occurs when individuals lie about the trustworthiness of others. On the basis of the knowledge we have obtained through this project, we offer some theory-driven recommendations for decision makers who want to regulate simulated societies as Internet environments, collaboration platforms, or telecommunication networks.

A theory for understanding and driving reputation dynamics in the society

Reputation dynamics drive or adopt social machineries, often in implicit and not easily recognizable manners. Questioning and understanding the role of reputation leads to exploitations of the characteristics of an artifact to orientate the appropriate machineries. The theoretical framework that forms the base of the eRep project has proven robust in predicting the outcomes of a number of real-world settings, interpreting certain social phenomena, and providing a powerful tool for the design of theory-driven reputation architectures in various areas. We briefly discuss the cases of public governance and discrimination as examples.

Public Governance

The style of public governance in Europe is moving from a pure top-down approach to a decentralized, more inclusive method, seeking to overcome the historical opposition between centralization and total deregulation.

Governance refers to a set of mechanisms for regulating complex social systems, which must be characterized by:

-➜ decentralization

-➜ dynamism

-➜ bidirectionality (both top-down, from institutions to citizens, and viceversa)

-➜ a mix of spontaneous and deliberate behavior

We argue that reputation provides one of the most effective, spontaneous, and efficient versions of such mechanisms. Reputation is a social artifact evolved precisely to achieve social order in the absence of a central authority. Reputation dynamics can be exploited to work together with designed institutions and thus achieve the level of decentralization that a modern approach to governance demands.

Accountability of institutions, decentralized regulations, and grassroots involvement represent the pillars of this new style of governance; they simultaneously are the long-term desiderata, yet to be achieved. Assessments of the performance of public offices by heterogeneous actors (single citizens, organizations, other institutions), such that these evaluations effectively drive institutional behavior, require ad hoc designed tools and practices, along with a theory-driven approach. Client satisfaction with a public service, often assessed by asking the user for direct feedback, can be very misleading, because general perceptions of the quality of a service often come from reported evaluations.

UNDERSTANDING AND DRIVING REPUTATION DYNAMICS IN THE SOCIETY

A special case of reputation in the scientific community: Peer review

Peer review, the standard procedure that journals and granting agencies use to ensure the scientific quality of the papers they publish, represents a reciprocal and symmetric type of evaluation, according to the social cognitive theory of reputation. Thus, it offers narrow access and transparency to the target. Because peer review is characterized by target accessibility and bidirectionality, the theory expects it to lead to a courtesy equilibrium (target accessibility leads to reciprocation/retaliation, and bidirectionality leads to leniency), which in turn promotes less rigorous evaluations. Other factors might change this situation, such as when the reviewer wants to publish in the same outlet.

THE REPUTATIONAL STRUCTURE OF PREJUDICE

Rioting spread in southern Italy in July 2008: Small Gypsy communities were being accused of "child stealing" and attacked by angry mobs. The accusations proved to be false, and the situation thus revealed the power of gossip/reputation. The evaluation targets (i.e., excluded communities of Gypsies) were totally separated from the evaluators, with little understanding or knowledge crossing the community boundaries, so the prudence rule applied: The worst possible evaluation is accepted quickly and spreads like flame.

The roots of discrimination: the spread of unfair evaluations

A general theory of reputation can help predict the phenomena of marginalization, exclusion, and prejudice, which often refer to a particular reputational structure of the societies involved.

The phenomena of prejudice and exclusion are linked to the dynamics of social evaluation, so a general framework that accounts for both the cognitive (micro) and social (macro) aspects of reputation would be a helpful tool for designing policies and tools aimed at including segregated/marginalized communities. If communities gather first-hand (image-based) information about other groups, for example, they might minimize the corruptible nature of transmitted evaluations. Local government bodies should build the chances of interaction among otherwise segregated communities to encourage first-hand knowledge gathering. Even this step might not be sufficient, because a positive image can coexist with a negative reputation; a clear case where any policy not grounded in theory easily could fail.

Reputation for business and institutions

Implication for business

People need to trust others before they will trade; this long-standing assumption is strongly supported by our results. The more information that is available about a seller, the more trust consumers can develop in that seller, and the more likely they are to start trading. Because bidders base their decisions on sellers' reputation feedback scores, sellers have a strong motivation to maintain their trustworthiness and engage in socially accepted behavior. This drive enhances the system's effectiveness and fairness.

Bidders also can estimate the trustworthiness of sellers by evaluating their product descriptions, the presence and quality of pictures, minimum opening prices, the presence of previous bids, the availability of payment methods such as escrow services, reputation feedback mechanisms, and the presence of other labels (e.g., Power Seller). Our study's results show that consumers pay most attention to the feedback provided by the seller's previous buyers, which means these buyers need to be motivated to evaluate sellers for the system to work.

These previous buyers generally provide feedback according to three criteria: the expected benefits, the ease of evaluation, and the potential for repercussions.

Even a slight modification in the design of a reputation system can create substantial changes in the market structure. For example, the implementation of a bidirectional reputation system instead of a unidirectional one leads to very different outcomes in terms of the total number of users in the market. In the lifespan of an online marketplace then, reputation scores should be frequently adapted according to changing economical and social dynamics. A brand-new marketplace benefits from a leniency effect, a positive evaluation bias that occurs frequently in bidirectional rating systems when consumers are kind to others by reciprocating their positive evaluations, which increases the user base. However, when there is a critical mass of users, a prudence effect should be elicited to encourage less positively biased evaluations, which improves the robustness of the market. That is, consumers will offer positive evaluations only when they are certain of the performance of a seller and negative ones when they are uncertain.

Auction sites might increase trust in sellers by generating "system trust," or trust in the auction site, which in turn engenders trust in sellers. When auction sites do their utmost to check the behaviors of sellers (e.g., confirming locations and telephone

numbers), educate bidders and sellers, and resolve disputes, consumer trust in the auction site and in sellers should increase. Because feedback provision about the seller's performance (social data) helps consumers select their seller of interest, links to social networks (e.g., Facebook) might be helpful as well, turning auction sites into a small world of social interactions online.

The results of our study also should apply to online review Web sites, whose enormous growth has prompted the spread of product evaluations through "word-of-mouse." These sites also require effective reputation feedback systems that help consumers evaluate the quality of the product and of the evaluations provided by others.

Implications for EU

As more consumers sell to consumers, the number of fraudulent sellers has increased. The European Union might facilitate seamless and secure transactions by educating consumers (e.g., through econsumer.gov or infomercials) and building platforms on which individual consumers can report complaints and determine the trustworthiness of sellers. Regulations and consumer rights should be uniform to facilitate international trades. Because of the greater importance of reputational systems in Internet markets, we argue that a discussion is needed to determine how reputational systems might be implemented on a global scale. Such systems could be implemented by large auction sites, though other players (i.e. review sites and blogs) may be hesitant to adopt them. It remains debatable whether legislation is the best option or if an independent (EU-supported) system should be implemented that the Web sites can adopt voluntarily.

This project also offers insights for those who want to implement systems in which reputation provides a means to make better selections. Institutions such as local governments can enrich their current evaluation criteria with reputation scores to select the best companies to participate in public offerings.

Reputation theory and technology for research groups

The eRep project offers research groups a tested technology for deploying, simulating, and evaluating models that involve agents (both human and virtual) and their interactions. From an organizational level, electronic institutions provide the framework for specifying the interaction models and performative structure. On the basis of this framework, the eRep project provides another level of interaction control: reputation mechanisms. At this social level, the proposed technology involves a generic agent architecture that can use distinct reputation mechanisms to communicate with human users through the Web and participate in electronic institutions. In the following subsections, we synthesize the possible contributions of this technology to several research areas.

Cognitive science

Reputation, a cognitive phenomena, plays a crucial role in agents' interactions and may require cognitive agents to exploit its full potential. In this sense, the eRep technology can model cognitive agents and use cognitive reputation models. Because the chosen agent's architecture is based on the classical BDI (belief, desire, intention) agent architecture, flexibility in the design and implementation of cognitive models increase considerably compared with other agent architecture approaches. Furthermore, the simulation platform based on electronic institutions offers tools to design interaction scenarios for testing these cognitive models.

Another important aspect for cognitive science is the possibility of combining computational agents and human users. The eRep technology allows for their combined participation in the same scenario through graphical interfaces on the Web for humans and through a set of API designed by the project for artificial agents.

Multi-agent systems

The field of multi-agent systems (MAS) is broad and covers several areas. The technology associated with the eRep project provides a connection to trust and reputation systems, which itself has produced extensive literature that needs to be tested and compared with other models. The problem of interoperability among agents using different reputation models has prevented such testing, because each model likely uses different representation values and even different ontologies to describe what a reputation value is.

The eRep technology provides a common ontology of reputation and thus avoids these problems. The agent architecture includes all functionalities associated with reputation information in terms of this ontology. Also, the eRep technology entails a me-

REPUTATION THEORY AND TECHNOLOGY FOR RESEARCH GROUPS

chanism for translating information provided by reputation models into the terms of this common ontology. Therefore, researchers can test different reputation models in the same environment using the same reasoning process.
As part of this technology, we provide a repository of implemented reputation models, both centralized and decentralized, with their respective transformation functions.

Social simulation

Social simulation, considered by Robert Axelrod as a "third way of doing science", can be defined as the application of computational methods to problems in the social sciences.

Social simulation can help in the formal definition of social problems, finding new ways to approach them. Depending on the level of detail chosen, social simulation can be applied to abstract issues like norms or cooperation, acting as an "intuition pump", or to specific issues, trying to simulate collective events accurately, both for forecast and for educational purposes.

In the eRep project, we developed a theory-based module for reasoning about reputation and image that can be used in simulative environments. The approach proposed, that innovates on current representations of reputational information, has been successfully employed to different fields: an abstract market with false information, a simulated auction setting, a complex market for the Internet of Services. We are able to collaborate with research projects with our expertise in design and implementation of social simulations involving reputational constructs.

In this booklet we have presented a scientific approach to a fascinating, multi-faceted social artefact: reputation. The approach leaned upon a cognitive theory of social evaluation. It helped formulating research questions, designing experiments and the platforms that supported them; contributed to the interpretation of research results and to the production of hints and clues suggesting how to apply the lessons learned to the specific application contexts, which consist in the areas of electronic commerce and of the Internet of services. We have also tried to provide insights about reputational dynamics in specific social issues, such as discrimination.

Far from pretending to have shown an exhaustive theory, we hope to have painted the road towards a better employ of the artifact of reputation – especially in electronic contexts, where reputation technology is rather ubiquitous, but respondent to engineering constraints more than scientific knowledge.

This booklet tries to propose a reasoned, up-to-date state of the art in the field of reptuation and its application. Already while we are getting this in print, our research groups are busy with further research and experiments. We are expanding the study of reputation mechanisms to C2C web interactions, such as product evaluations on a review site; we are translating our findings in an agent-based model to explore how reputational information may spread through a network of consumers, and how reputational mechanisms may serve as a filter to downsize the impact of inaccurate product recommendations. This line of research fits in projects on word-of-mouth and network dynamics which are currently conducted at the University of Groningen using surveys, laboratory experiments and simulation studies in ensemble.

New simulation experiments are in preparation to investigate in further depth the issue of informational accuracy. The research group at ISTC-CNR will perform these simulation and try to validate them against political scenarios and social networks data.

At the same time, the exploitation of electronic institutions as a tool to perform laboratory experiments where humans participate together with software agents and further research on cognitive models for reputation will be carried on at the IIIA-CSIC.

The Internet of Services will over time become an ubiquitous, important backbone for businesses and consumers. This also creates a growing dependence on the reliability of network links, the availability of services and the trustworthiness of service providers. While reputation mechanisms are one promising way to reduce those risks, the University of Bayreuth group will also investigate into other risk management concepts like SLA negotiations, insurances, or policy enforcement using a mixture of explorative case studies, simulation studies and evaluation of software artifacts.

CONTRIBUTORS

Mario Paolucci (project coordinator)
Torsten Eymann (unit coordinator)
Wander Jager (unit coordinator)
Jordi Sabater-Mir (unit coordinator)

LABSS-ISTC-CNR

http://labss.istc.cnr.it

Theory, social simulation

Mario Paolucci, Rosaria Conte, Samuele Marmo, Stefano Picascia, Walter Quattrociocchi

Universitaet Bayreuth

http://www.bwl7.uni-bayreuth.de/

Internet of services, technology

Torsten Eymann, Tina Balke, Stefan Koenig

Rijksuniversiteit Groningen

http://www.rug.nl/feb/index

Internet fraud, laboratory experiments

Wander Jager, Thijs Broekhuizen, Debra Trampe, Mirjiam Tuk

Artificial Intelligence Research Institute

http://iiia.csic.es

Simulations, technology, platform

Jordi Sabater, Ismel Brito, Isaac Pinyol, Daniel Villatoro