

# The Agent Reputation and Trust (ART) Testbed

Karen K. Fullam<sup>1</sup>, Tomas Klos<sup>2</sup>, Guillaume Muller<sup>3</sup>, Jordi Sabater-Mir<sup>4</sup>,  
K. Suzanne Barber<sup>1</sup>, and Laurent Vercouter<sup>3</sup>

<sup>1</sup> The University of Texas at Austin, USA  
{kfullam, barber}@lips.utexas.edu

<sup>2</sup> Center for Mathematics and Computer Science (CWI), Amsterdam, The Netherlands  
tomas.klos@cwi.nl

<sup>3</sup> École Nationale Supérieure des Mines, Saint-Étienne, France  
{muller, Laurent.Vercouter}@emse.fr

<sup>4</sup> Artificial Intelligence Research Institute (IIIA-CSIC), Barcelona, Spain  
jsabater@iiaa.csic.es

**Abstract.** The Agent Reputation and Trust (ART) Testbed initiative has been launched with the goal of establishing a testbed for agent reputation- and trust-related technologies. The ART Testbed serves in two roles: (1) as a competition forum in which researchers can compare their technologies against objective metrics, and (2) as a suite of tools with flexible parameters, allowing researchers to perform customizable, easily-repeatable experiments. In the Testbed’s artwork appraisal domain, agents, who value paintings for clients, may purchase opinions and reputation information from other agents to produce accurate appraisals. The ART Testbed features useful data collection tools for storing, downloading, and replaying game data for experimental analysis.

## 1 Introduction

The Agent Reputation and Trust (ART) Testbed [1] serves two roles: (1) as a competition forum for comparing technologies against objective metrics, independent from participating researchers, and (2) as an environment for performing customizable, easily-repeatable experiments. In the testbed’s art appraisal domain, agents value paintings for clients and gather opinions from other agents to produce accurate appraisals. As a versatile, universal experimentation site, the ART Testbed scopes relevant trust research problems and unites researchers toward solutions via unified experimentation methods. Through objective, well-defined metrics, the testbed provides researchers with tools for comparing and validating their approaches. The testbed also serves as an objective means of presenting technology features—both advantages and disadvantages—to the community. In addition, the ART Testbed places trust research in the public spotlight, improving confidence in the technology and highlighting relevant applications.

## 2 Testbed Domain Problem

The ART Testbed provides functionality for researchers of trust and reputation in multi-agent systems. As such, it operates in two modes: competition and experimentation. In competition mode, each participating researcher controls a single agent, which works in competition against every other agent in the system. At the 2006 iTrust conference, we will demonstrate the ART Testbed with a variety of agents (e.g., participants in the First ART Testbed Competition, at AAMAS 2006). To utilize the testbed's experimentation mode, the Testbed is downloadable for researcher use independent of the competition [5]: results may be compared among researchers for benchmarking purposes, since the testbed provides a well-established environment for easily-repeatable experimentation.

The testbed operates in an art appraisal domain (see [2] for a detailed justification), where researchers' agents function as painting appraisers with varying levels of expertise in different artistic eras. Clients request appraisals for paintings from different eras; if an appraiser does not have the expertise to complete the appraisal, it may purchase opinions from other appraisers. Other appraisers estimate the accuracy of opinions they send by the cost they choose to invest in generating an opinion, and opinion providers may lie about the estimated accuracy of their opinions. Appraisers produce appraisals using their own opinion and opinions received from other appraisers, receiving more clients, and thus more profit, for producing more accurate appraisals. They may also purchase reputation information from each other about third-party agents. Appraisers attempt to accurately value their assigned paintings; their decisions about which opinion providers to trust directly impact the accuracy of their final appraisals. In competition mode, the winning agent is selected as the appraiser with the highest bank account balance.

## 3 Testbed Architecture

As shown in Figure 1, the testbed architecture, implemented in Java, consists of several components (see [3] for a detailed description of the ART Testbed architecture). The Testbed Server manages the initiation of all games by starting a Simulation Engine for each game. The Simulation Engine is responsible for controlling the simulation environment by enforcing chosen parameters. In each timestep, the Simulation Engine assigns clients with paintings to each appraiser. Then appraisers conduct reputation and opinion transactions with each other as described above. Finally, the Simulation Engine assesses each appraiser's accuracy based on the opinions the appraiser purchases and the 'weights' the appraiser places on those opinions. Weights are real values between zero and one that an appraiser assigns, based on its trust model, to another's opinion.

Through the Simulation Engine, the Database collects environment and agent data, such as true painting values, opinions, transaction messages, calculated final appraisals, client share allocations, and bank balances. With access tools for navigating, downloading/uploading, and replaying Database logs, data sets

**Fig. 1.** The ART Testbed architecture.

are made available to researchers after each game session for game re-creation and experimental analysis.

User Interfaces permit researchers to observe games in progress and access information collected in the Database by graphically displaying details. Figure 2 shows the Game Monitor Interface, by which observers can view opinion and reputation transactions between agents on the left and detailed statistics, such as bank balance, about each appraiser agent on the right. The Game Monitor Interface's play-pause buttons permit games to be played and replayed, regardless of whether the game is in progress or completed.

Finally, the abstract Agent class is designed to allow researchers to easily implant customized internal trust representations and trust revision algorithms while permitting standardized communication protocols with entities external to the appraiser agent. Users simply create a class inheriting from the Agent class, implementing a method for each of the agent's necessary strategic decisions. The abstract class Agent handles all required inter-agent communication, as well as communication between agents and the Simulation Engine.

## 4 Conclusions

The ART Testbed 1) provides researchers with a common experimentation environment and 2) allows researchers to compete against each other to determine the most viable technology solutions. Initial experimentation by Fullam and Barber [4] and by participants in the First Annual ART Testbed Competition (held in May, 2006 at AAMAS-06) and the First Spanish ART Competition (held in April, 2006 together with the AgentCities.es school) shows the testbed's

**Fig. 2.** The Game Monitor Interface for viewing game data.

art appraisal problem to be non-trivial. However, in the future, game rules may be changed to keep the competition challenging (for example, not revealing actual painting values to appraisers, changing the format by which reputations are represented or introducing multi-dimensional trust characteristics, such as quality, timeliness, and availability). Further, the Testbed may be expanded to incorporate multiple problem scenarios relevant to a wider range of trust-related research. Possible improvements to the Testbed include logging additional data and developing data processing utilities for forensic analysis of games and wrappers to permit designing agents in languages other than Java. In addition, organizing the development effort to permit contributions from the research community via the Testbed's SourceForge webpage [5] will speed the implementation of these features and the correction of bugs. Development progress can be monitored through the testbed website [1], where updates to testbed development are posted.

## References

1. ART Testbed Team. *Agent Reputation and Trust Testbed Website*. <http://www.art-testbed.net/>, 2006.
2. K. Fullam, T. Klos, G. Muller, J. Sabater-Mir, A. Schlosser, Z. Topol, K. S. Barber, J. S. Rosenschein, L. Vercouter, and M. Voss. A Specification of the Agent Reputation and Trust (ART) Testbed In *Proc. AAMAS*, 2005, pp. 512–518.
3. K. Fullam, T. Klos, G. Muller, J. Sabater-Mir, Z. Topol, K. S. Barber, J. S. Rosenschein, and L. Vercouter. The Agent Reputation and Trust (ART) Testbed Architecture In *Proc. Trust Workshop at AAMAS*, 2005, pp. 50–62.
4. K. Fullam, and K. S. Barber. Learning Trust Strategies in Reputation Exchange Networks In *Proc. AAMAS*, 2006.
5. ART Testbed Team. *ART Testbed SourceForge project page*. <https://sourceforge.net/projects/art-testbed>, 2006.