CAP6671 Intelligent Systems

Lecture 8:

Agent Reputation and Trust Testbed

Instructor: Dr. Gita Sukthankar
Email: gitars@eecs.ucf.edu

Schedule: T & Th 9:00-10:15am

Location: HEC 302

Office Hours (in HEC 232):
T & Th 10:30am-12

Trust Decisions in Reputation Exchange Networks

- **Agents perform transactions to obtain needed resources**
  - **Transactions have risk because partners may be untrustworthy**
  - **Agents must learn whom to trust and how trustworthy to be**
- **When agents can exchange reputations**
  - **Agents must also learn when to request reputations and what reputations to tell**
  - **Agents' trust decisions affect each other**
    - Difficult to learn each decision independently

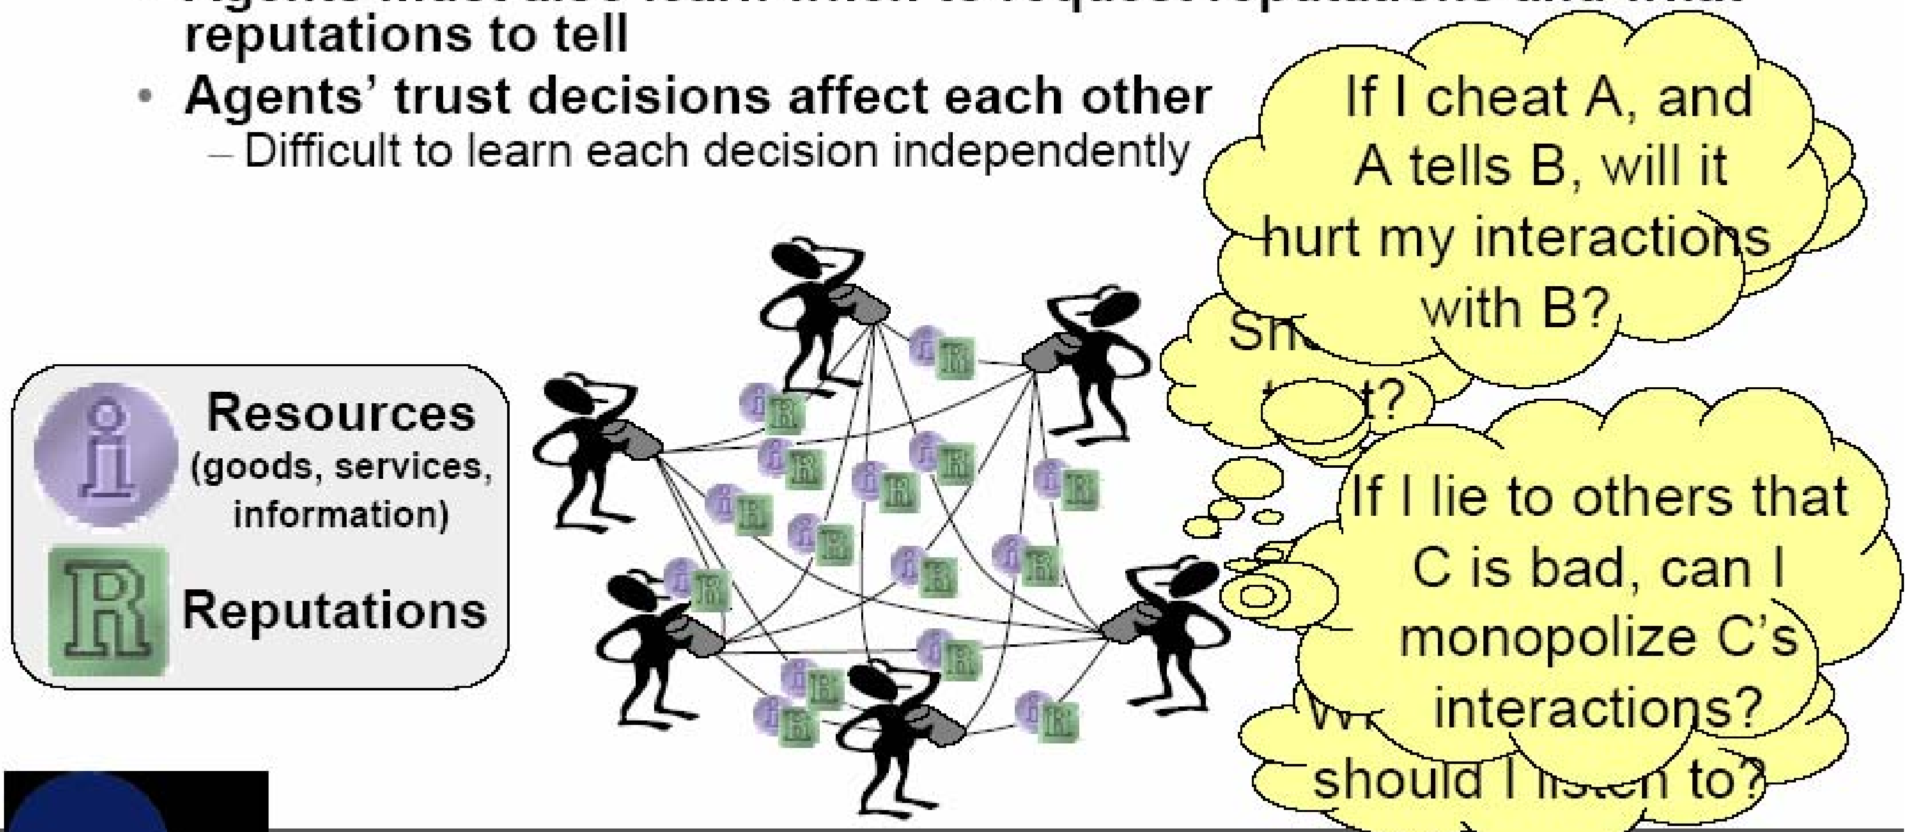

Reputation & Trust

- Models realistic economies in which agents depend on each other but are still self-interested
- Tradeoffs between selfish and altruistic behavior
- Reputation helps the agent make money by increasing the opinion requests
- But expending too many resources increasing one's reputation might not be worth it....
- Modeling trust allows the agent not to be exploited by other agents
- Agent must be adaptive to avoid exploitation

Enumerating Decisions in a Trust Strategy

Num agents = a
Num transaction types = e
Num choices/decision = n

Truster

Trustee

Transaction

Agent Role	Fundamental	Reputation
Trustee	How trustworthy should I be? $\text{combinations} = n^{ae}$	Should I tell an accurate reputation? $\text{combinations} = n^{a^2 e}$
Truster	Should I trust? $\text{combinations} = n^{ae}$	Should I believe this reputation? $\text{combinations} = n^{a^2 e}$

If these decisions affect each other, there are $n^{(2ae(1+a))}$ possible strategies!

How to learn the best strategy with so many choices?

Agent Reputation and Trust (ART) Testbed

- An open-source tool for
  - Experimentation: Easily-repeatable experiments in a common environment
  - Competitions: Most promising trust technologies are identified

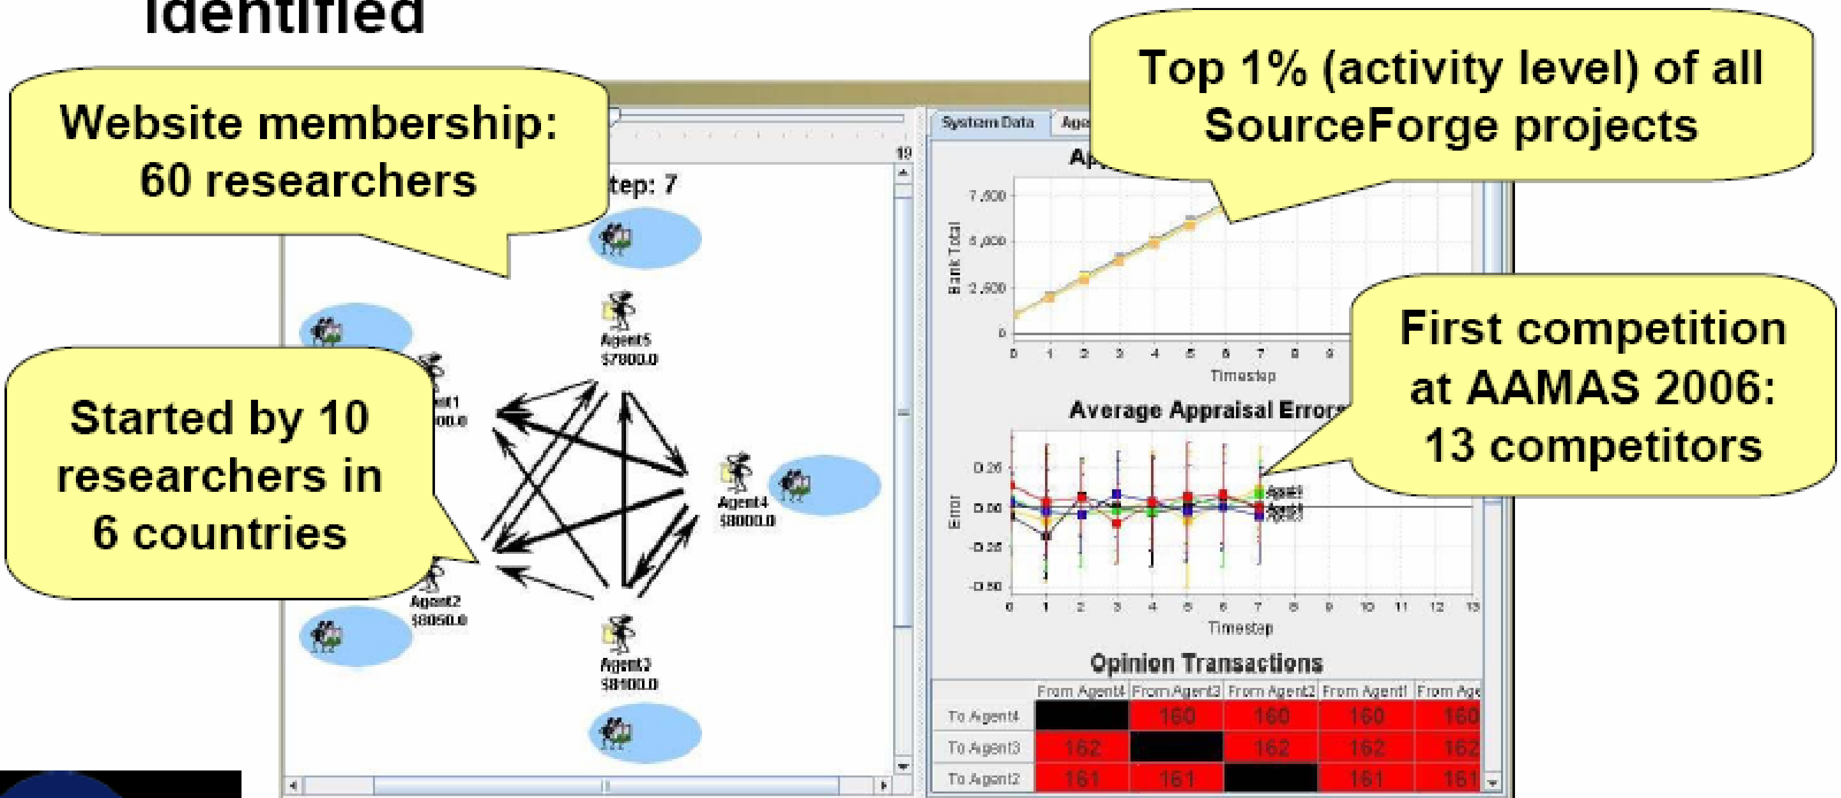

ART vs. TAC?

ART Testbed Game Rules

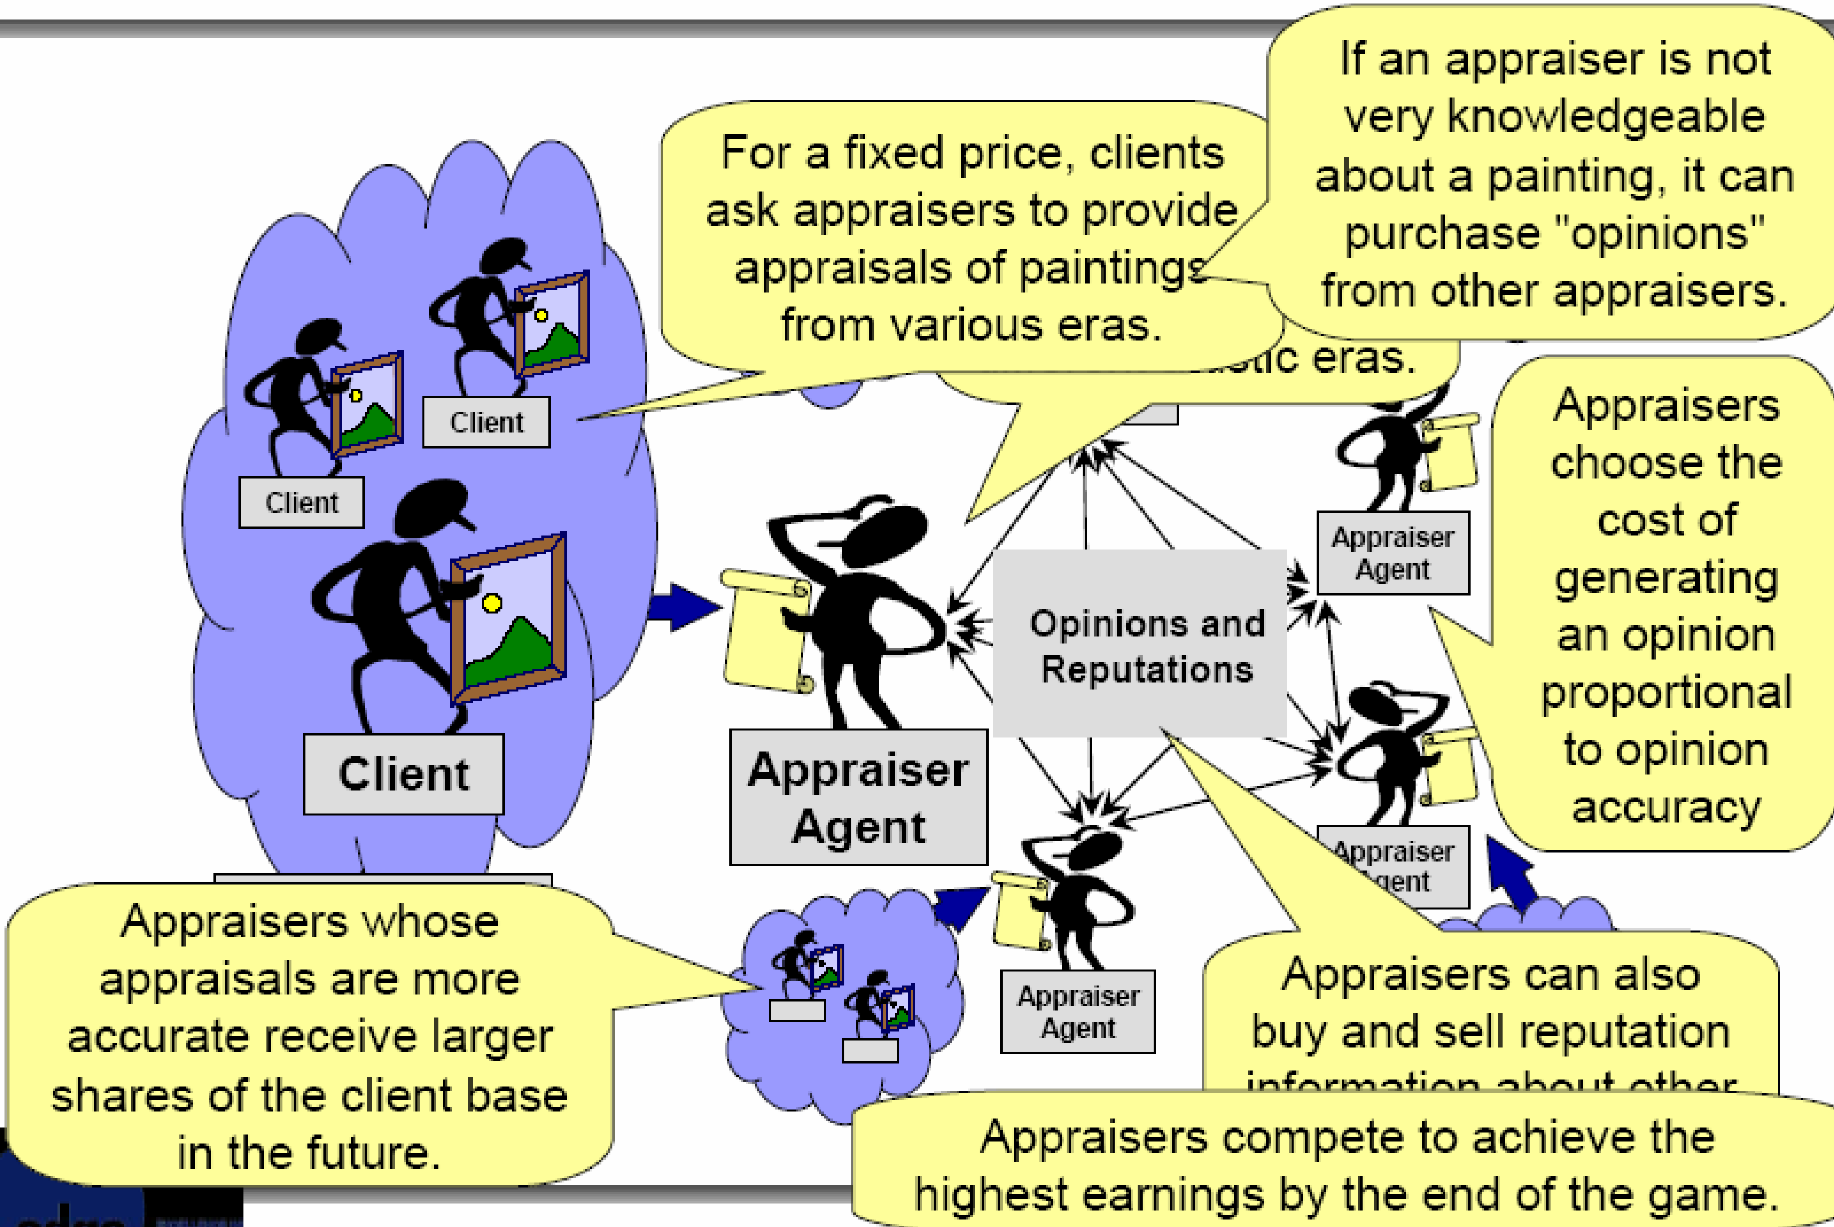

ART Competition Rules

- Game is an unknown number of iterations
- Agents (acting as art appraisers) make money by:
  - Performing appraisals for clients ($100)
  - Generating opinions for other agents ($10)
  - Answering reputation requests for other agents ($1)
- Agents can spend money:
  - On their appraisals (either for their client or for opinions)
  - Buying opinions and reputation requests
- Agents get new clients based on performance from previous timesteps.

ART Testbed

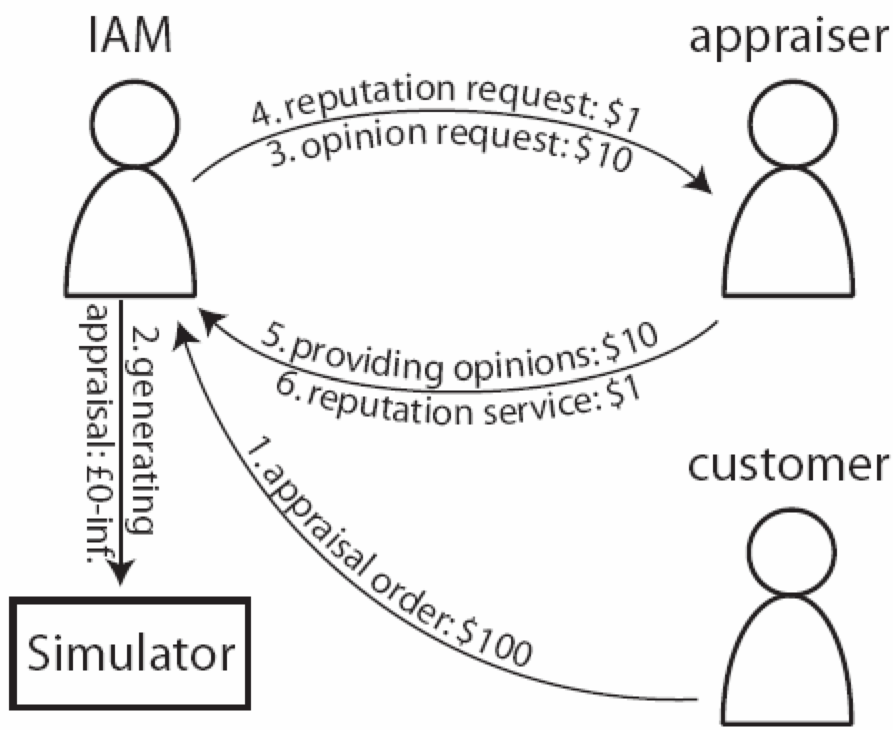

IAM
appraiser
4.reputation request: $1
3.opinion request: $10
5.providing opinions: $10
6.reputation service: $1
customer
2.generating
appraisal: £0-inf.
1.appraisal order: $100
Simulator

Appraisals

- Agents are initialized with a private expertise vector
- Expertise vector determines the variance of error when appraising a painting while spending a certain amount of money
- Simulator generates appraisals according to the formula:

$$var(e_i) = v^2 \left(s_i + \frac{\alpha}{C_i} \right)^2$$

v=painting value
s=expertise value
C=money expended

Combining Appraisals

- All appraisals are calculated by the testbed according to this formula:

$$\overline{e} = \sum_{i=0}^{q} w_i \cdot e_i$$

E=estimate
w=weight

- Agents need to decide:
  - Which opinions (if any) to solicit
  - How to weight opinions
  - No exploration in function space is permitted

IAM Client-Service Decisions

1. How much it should spend on its own appraisal given its expertise vector
2. Whether it needs to ask for external opinions for an order and which agents opinions it should choose
3. Whether it should ask for reputation values and how these values can be incorporated into the choice of agent opinions
4. Setting the weights on the opinions it receives

IAM Agent Services

- How much money to spend on generating opinions for other agents
- How to generate reputation values for other agents

What strategy would you use?

IAM Agent

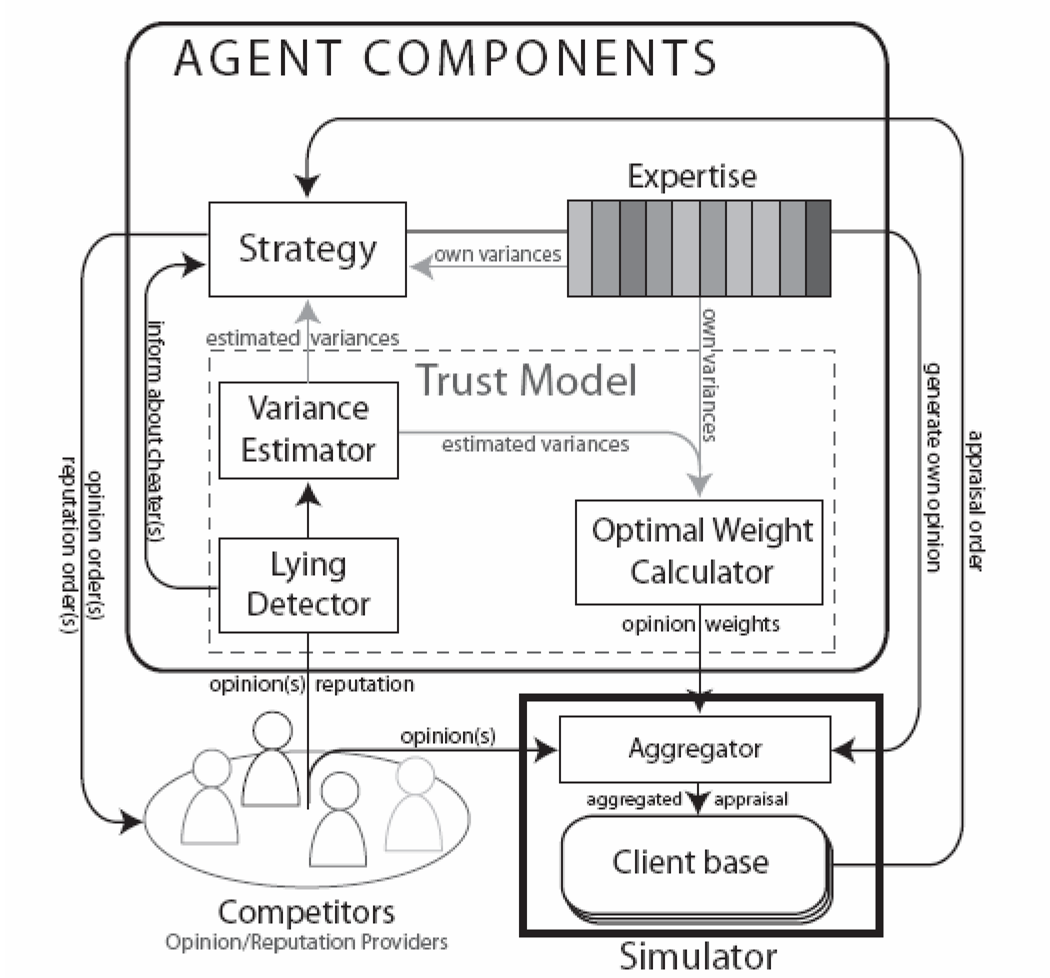

AGENT COMPONENTS
Expertise
Strategy
own variances
estimated variances
inform about cheater(s)
Trust Model
own variances
Variance Estimator
estimated variances
Optimal Weight Calculator
Lying Detector
opinion weights
opinion order(s)
reputation order(s)
generate own opinion
appraisal order
opinion(s)
reputation
opinion(s)
Aggregator
aggregated appraisal
Client base
Competitors
Opinion/Reputation Providers
Simulator

Methods

- IAM uses several statistical and optimization technqiues to improve its performance
- Calculate optimal weights based on variances of the opinions
- Estimate variances based on a Bayesian analysis to determine the most likely values of (s,C) used by other agents
- Use chi-square statistical test to identify liars
- Minimizing amount of money spent in estimates

Calculating Optimal Weights

- Must aggregate opinions from different agents
- Weighting function must be linear
- Minimize mean square error
- Best Linear Unbiased Estimates uses information about variance to determine weight

$$w_i = \frac{1/var(e_i)}{\sum_{i=0}^{q} 1/var(e_i)}$$

- IAM doesn't use reputation value in weight estimates.

Estimating Truthful Variances

- Calculate expected variance by marginalizing over values of (s,C) used by the agents

$$E[var(e_i/v)] = \sum_{s_i \in \mathcal{S}} \sum_{C_i \in \mathcal{C}} P(s_i, C_i) \left(s_i + \frac{\alpha}{C_i} \right)^2$$

- Maintain a conditional probability table and update it according to agent experience

Identifying Liars

- Use chi-squared test to test the hypothesis that the agents last k opinions are truthful
- Calculate the maximum mse based on max and min bounds on (s,C)

$$Q_k = \frac{1}{2.25} \sum_{i=1}^{k} (e_i / v_i - 1)^2$$

- Calculate probability that last k opinions were generated truthfully

Generating Appraisals

- Determine how much money it is worth spending on each appraisal
- Even for different values of s \$4 is a good break point based on the simulator parameters

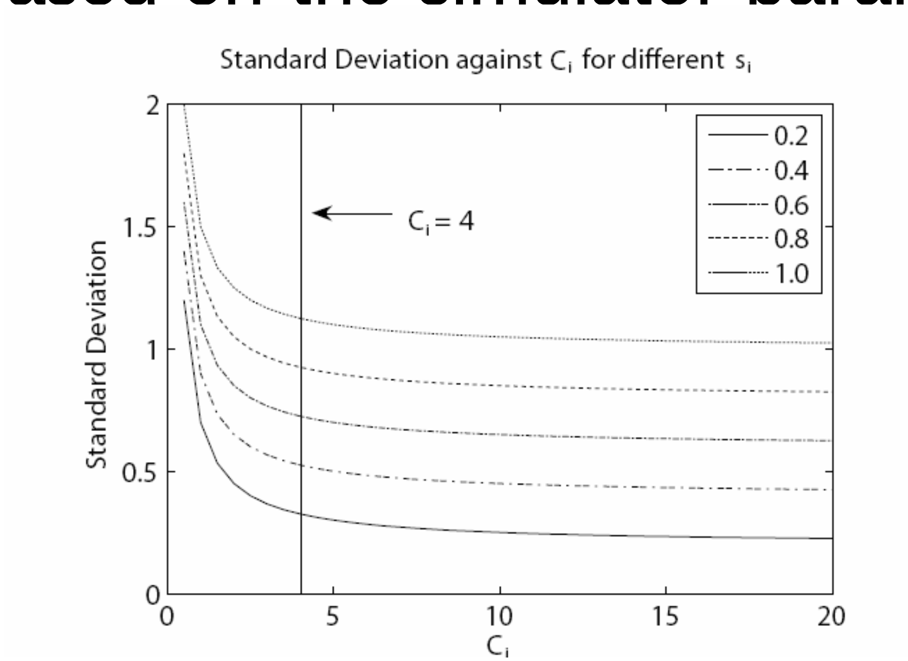

Combining Opinions

- Always use your own opinion regardless of your expertise value
- Sort other agents into estimated variances for the art era of the painting
- Eliminate agents if they have $p > 0.6$ of cheating on chi-square test
- Calculate reduction in variance for adding each agent from the list
- If combined variance of final appraisal is reduced less than 15% stop the selection process

Earning thru Honesty

- Generate good appraisals ($4) for non-cheating agents
- Retaliate against cheating agents by spending a fraction of that money (0.01)
- Generate fair reputation values
- For cheating agents provide random reputation values

Competition Results

Agent	Affiliation	Revenue	Cost	Profit
IAM	University of Southampton	149812	18299	131583
Neil	Nanyang Technological University	116764	13741	103023
Frost	Bogazici University	120753	18176	102577
Sabatini	Universidad Carlos III de Madrid	127137	25726	101411
Joey	University of Nebraska-Lincoln	111985	19506	92479
	mean	*125290*	*19076*	*106215*

Cost Percentages

Agent	Opinion Costs	Opinion Generation Costs	Reputation Costs
IAM	59.46	40.54	0.00
Neil	6.80	92.03	1.16
Frost	31.65	68.35	0.00
Sabatini	38.85	61.15	0.00
Joey	0.00	100.00	0.00

Revenue Percentages

Agent	Client Payments	Opinion Payments	Reputation Payments
IAM	96.09	3.89	0.03
Neil	98.63	1.37	0.00
Frost	98.45	1.52	0.03
Sabatini	88.25	11.72	0.03
Joey	96.96	3.00	0.04

Discussion

- Agents can benefit from 3rd party info if it is easier to establish the information as reliable (e.g. combining appraisals)
- Generally more economical to purchase opinions from 3rd parties than to invest heavily in own opinion
- For reputation it was easier to learn reputation models for other agents than to purchase them because of:
  - Small number of agents
  - Problems with reputation semantics

References

- K. Fullam slides from

  https://webspace.utexas.edu/fullamkk/