How to Trust a Few Among Many

Etuk, A., T. J. Norman, M. Sensoy, and M. Srivatsa (2017), Journal of Autonomous Agents and Multiagent Systems, 31:531-560.

Quiero MAS

Adam Plowcha, Becca Horzewski, Mustafa Ridha

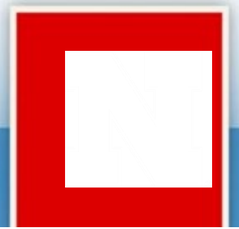

Introduction

- This paper examines a methodology for determining which information sources to sample in order to create the most accurate overall report
- Determining which sources to trust is typically done by:
  - Observation based sampling (OBS)
  - Majority Based Sampling (MBS)
  - Random Based Sampling (RBS)
- In this paper Diversity-Based Sampling which uses the realisation of the Trust in Information through DiversitY (TIDY) framework is introduced

Observation Based Sampling

- Uses assessments of the trustworthiness of individual sources to guide sampling
- Discounts Opinions
- Challenging for the decision-maker to determine the true reliability of sources
- When constrained by budget, OBS selected the most trusted sources according to the budget allowance

Majority Based Sampling

- Inclined towards opinions held by larger groups of sources
- Benefits from the high proportion of honest sources, who are likely to provide reliable reports to filter bogus reports
- Not affected by the dynamic nature of sources in a system
- Filters based only on statistics on the reports, and not source's behavior

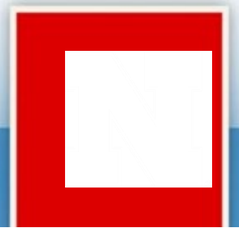

Random Sampling

- Most popular method in conducting surveys
- Each source has an equal probability of being sampled
- Does not perform filtering like MBS
- Does not perform weighting like OBS

Background

- The trustworthiness of information sources is an important factor in making informed and reliable decisions about what is true in the world.
- Typical approaches to estimating a value of some state is to rely on reports from as many sources as possible (exploiting the “wisdom of the crowd effect”)
- However, capturing and distributing evidence can be costly.
- Sources in a complex, dynamic and constrained environments can adopt a model of diversity in order to minimise redundant sampling and mitigate the effect of certain biases.
- The TIDY framework offers a general approach for resource constrained information fusion from variously trusted sources

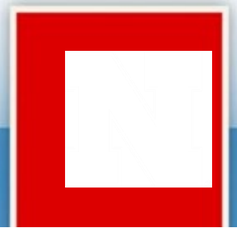

The TIDY Approach - Basic Terminology

- TIDY: “Trust in Information through DiversitY”
- Information Source: an entity that can provide information about the state of the environment
  - Identified by a unique ID number and a vector of feature values
- Feature: observable characteristic of an information source
- Report: a reading of the environmental state by an information source
- Report History: all reports from an information source up to the current time
- Decision Maker: central entity that obtains reports from information sources and combines them to create an aggregate report

The TIDY Approach - Similarity

- Similarity Metric: measure of how similar two information sources are, based on their features
  - Uses M5 model tree algorithm
    - Nodes are features, path taken from node is the difference in values for that feature for the pair of information sources in question
    - Leaves are linear regression models used to estimate similarity score (weighted sum of feature value differences)
- Diversity Structure: stratification of all information sources into exhaustive and disjoint groups
  - Created by hierarchical clustering with the similarity metric as the distance function

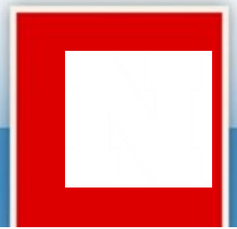

The TIDY Approach - Agreement

- Report Agreement: whether or not reports from two different information sources match/agree
  - Difference between them is less than some threshold
- Source Agreement: degree to which two information sources agree, based on the agreement of their reports
  - Expected value of a Beta distribution based on number of agreeing reports and number of disagreeing reports
- Similarity Metric Revision: update of the similarity metric function based on observed source agreement, and to handle any changes in information sources
  - Re-train the model tree with observed feature distances and corresponding source agreement (how often this occurs is topic of future work)

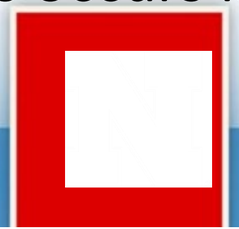

The TIDY Approach - Trustworthiness

- Report Assessment: whether a report is accurate or not, compared to ground truth information obtained at a later time
  - Difference between them is less than some threshold
- Source Trustworthiness: degree of accuracy of a source, based on previous report assessments
  - Expected value of a Beta distribution based on number of accurate reports and number of inaccurate reports
  - Trustworthiness of a group of sources is the average trustworthiness value of the information sources in the group

The TIDY Approach - Sampling

- Sampling Cost: the energy, time, or bandwidth cost of obtaining information from a set of information sources
  - Sum of the sampling costs for each information source contacted
  - Strategy: sample from one information source from each group, prioritizing the groups by trustworthiness, until budget is exhausted (if extra budget, sample more sources from larger groups)

The TIDY Approach - Fusion

- Fusion: function by which the decision maker synthesizes an aggregate report from a set of reports from various information sources
  - First, average the reports from information sources within a group
  - Then, take the trustworthiness-weighted average of the groups' reports

- Information Quality: degree of accuracy of the decision maker's aggregate report compared to the ground truth information
  - Mean absolute error
  - The dependent variable observed in the experiments

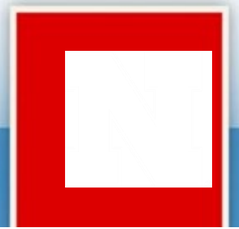

Experiments

- The experimental goal of the authors was to evaluate the effectiveness of the TIDY framework in estimating environmental states
- Specifically, they examined two instances against a change in the number of candidates sampled (referred to as a “budget”)
  - Sources that provided misleading, but independent, information
  - Sources that colluded, i.e. copied each other
- For each experimental instance, they evaluated the TIDY (Diversity-based sampling) framework against three other methods of sampling:
  - Observation-based (OBS)
  - Majority-based (MBS)
  - Random-based (RBS)

Experiments

- Each instance was simulated 10 times, with each simulation run having 100 samples
- The authors used the following parameters for each experiment
  - 100 sources of information in the population
  - 0.1 probability of population change
  - 30 learning interval
  - 0.4 diversity threshold
  - 0.1 report agreement threshold
  - 0.1 report reliability threshold

Results - Small Budget, Malicious Sources

- With a small budget and low proportion of malicious sources, diversity-based sampling (DBS) was on par with all other methods
- With a small budget and medium proportion of malicious sources, there is again no statistically significant difference between sampling methods (although the graph does show a slight advantage for DBS)
- With a small budget and high proportion of malicious sources DBS showed a marked improvement over the other three sampling methods

Results - Large Budget, Malicious Sources

- With a large budget and low proportion of malicious sources,MBS outperformed DBS, with DBS showing no real difference between OBS and RBS

- With a large budget and medium proportion of malicious sources, there is no statistically significant difference between sampling methods (although, once again, the graph does show a slight advantage for DBS)

- With a large budget and high proportion of malicious sources DBS performed significantly better than the other three sampling methods

Results - Malicious Sources

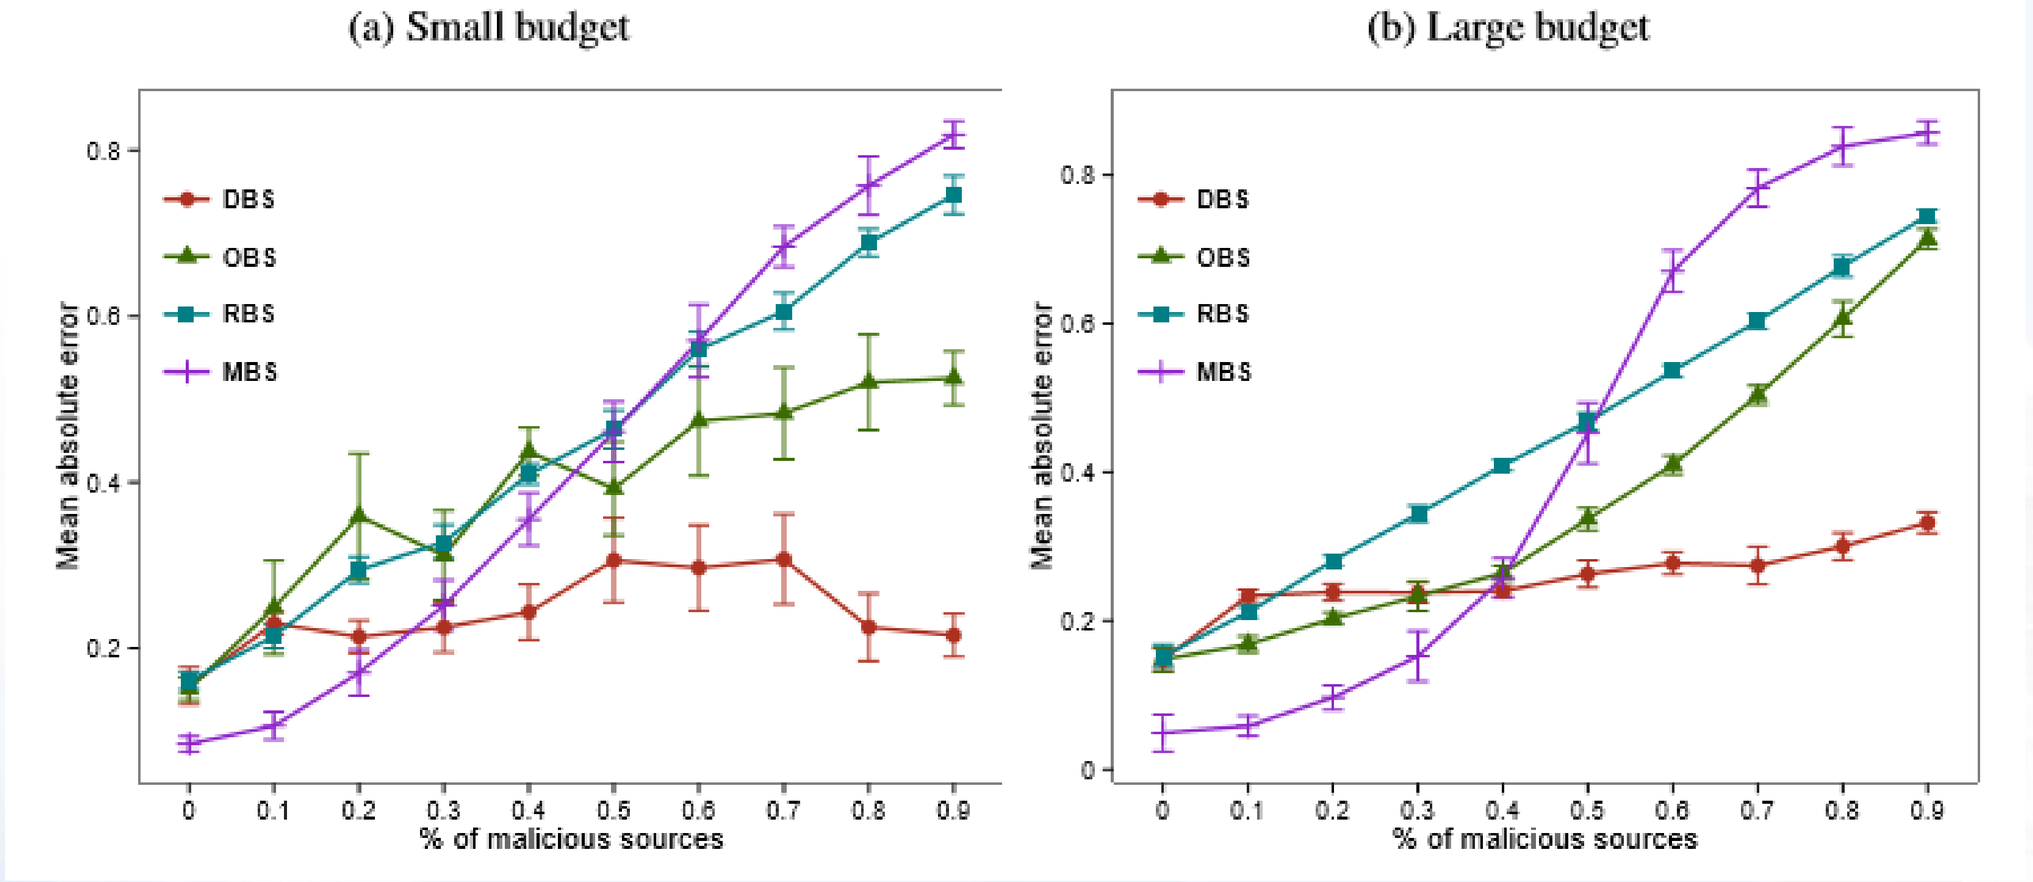

Images: Etuk, A., T. J. Norman, M. Sensoy, and M. Srivatsa (2017), Journal of Autonomous Agents and Multiagent Systems, 31:531-560.

Results - Small Budget, Dependent Sources

- With a small budget and low proportion of dependent sources, DBS outperformed OBS and MBS, with no real difference compared to RBS
- With a small budget and medium proportion of dependent sources, there is no statistically significant difference between DBS and RBS, but RBS again outperforms OBS and MBS
- With a small budget and high proportion of dependent sources DBS performed better than OBS, significantly better than MBS, and no better or worse than RBS

Results - Large Budget, Dependent Sources

- With a large budget and low proportion of dependent sources, DBS performed no better or worse than any of the other methods
- With a large budget and medium proportion of dependent sources, DBS outperformed all other methods
- With a large budget and high proportion of dependent sources DBS again performed better than all other methods

Results - Dependent Sources

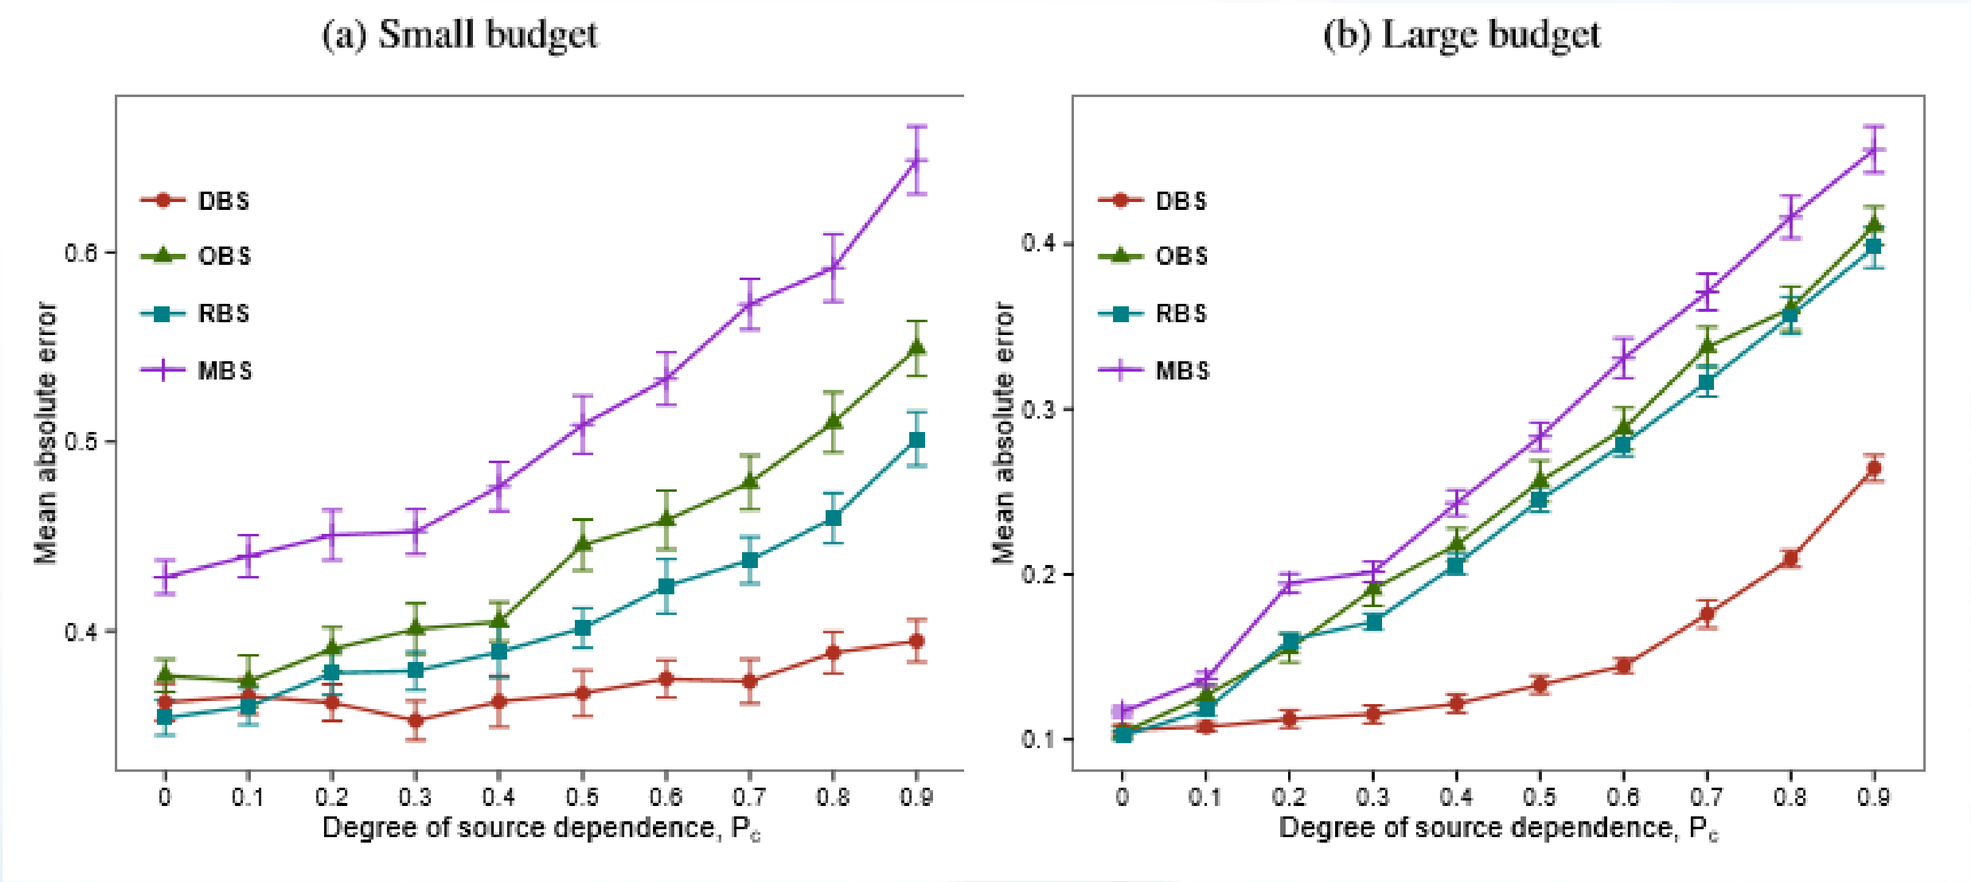

Images: Etuk, A., T. J. Norman, M. Sensoy, and M. Srivatsa (2017), Journal of Autonomous Agents and Multiagent Systems, 31:531-560.

Results - Summary

- Malicious Sources
  - DBS is on par with other methods when there is a small percentage of malicious actors regardless of the budget size, but it shows a noticeable improvement as the percentage of malicious actors increases in all budget sizes

- Dependent Sources
  - DBS starts on par* with other methods in both small and large budgets when there is a lower concentration of dependent actors , but again, it shows improvement over all other methods as the percentage of dependence increases

* Although RBS is actually a little better for very low dependence cases

Author Conclusions

- DBS groups sources based on perceived similarities
- The performance of all approaches is affected by budget
- This method ensures that there is no double counting by ensuring that sources in the population belong to distinct groups
- A selected source in DBS inherits the trust score of the group
- DBS has one limitation which is generality
- Each sampling model has its limitations under different environments but DBS helps the decision-maker intelligently sample for evidence, such that reliable assessments can be made within different budget constraints

Our Conclusions and Discussions

- Crowd-sourcing (majority based sampling) is not always the most accurate, as inaccurate sources can copy each other to appear more reputable than they are
- The TIDY approach accounts for colluding sources, which is a growing issue in the world of social media
- RBS performed surprisingly well in some circumstances

Questions