Rustin Haase
David Snavely

How to trust a few among many

Anthony Etuk, Timothy J. Norman, Murat Şensoy, Mudhakar Srivatsa

Introduction

- Problem Statement
  - Deciding on which information sources to trust has become challenging due to hardware constraints and increasing supply. The presence of numerous unreliable and misleading reports adds more complexity. An efficient model for selecting accurate sources is necessary.

- Proposed Solution
  - This paper introduces a diversity based approach
    - Goals:
      - Minimize redundancy
      - Reduce bias
      - Make decisions more effective
  - Method
    - Separate sources into subgroups and sample from them.

Trust in Information through DiversitY(TIDY) Framework

- Definitions:
  - Task - the derivation of an environmental state at some time interval
  - Information Source - a tuple containing a unique identifier and vector for the source's features
  - Report - a message from an information source containing the result of a task
  - Report History - The set of reports from a given source
  - Feature - An observable attribute of an information source
  - Similarity Metric - A function that gives a score representing the degree of similarity between two sources (revised to focus on dissimilarities between reports instead of agreements)
  - Diversity Structure - A division of the set of sources into exhaustive, disjoint groups
  - Report Agreement - A function that evaluates the degree of agreement between two reports
  - Source Agreement - An aggregation of a sequence of agreements between reports from two sources at the same time
  - Report Assessment - An assessment of how truthful a report is
  - Source Trustworthiness - An assessment of a source based on the sequence of reports received during a given time period
  - Information quality - A function to denote how good a source's estimate of its environment's state.
  - Sampling cost - The cost of environmental sampling. Assumed to be stable over time.
  - Fusion - a function that computes an estimate of the environmental state given a set of reports.

TIDY Framework

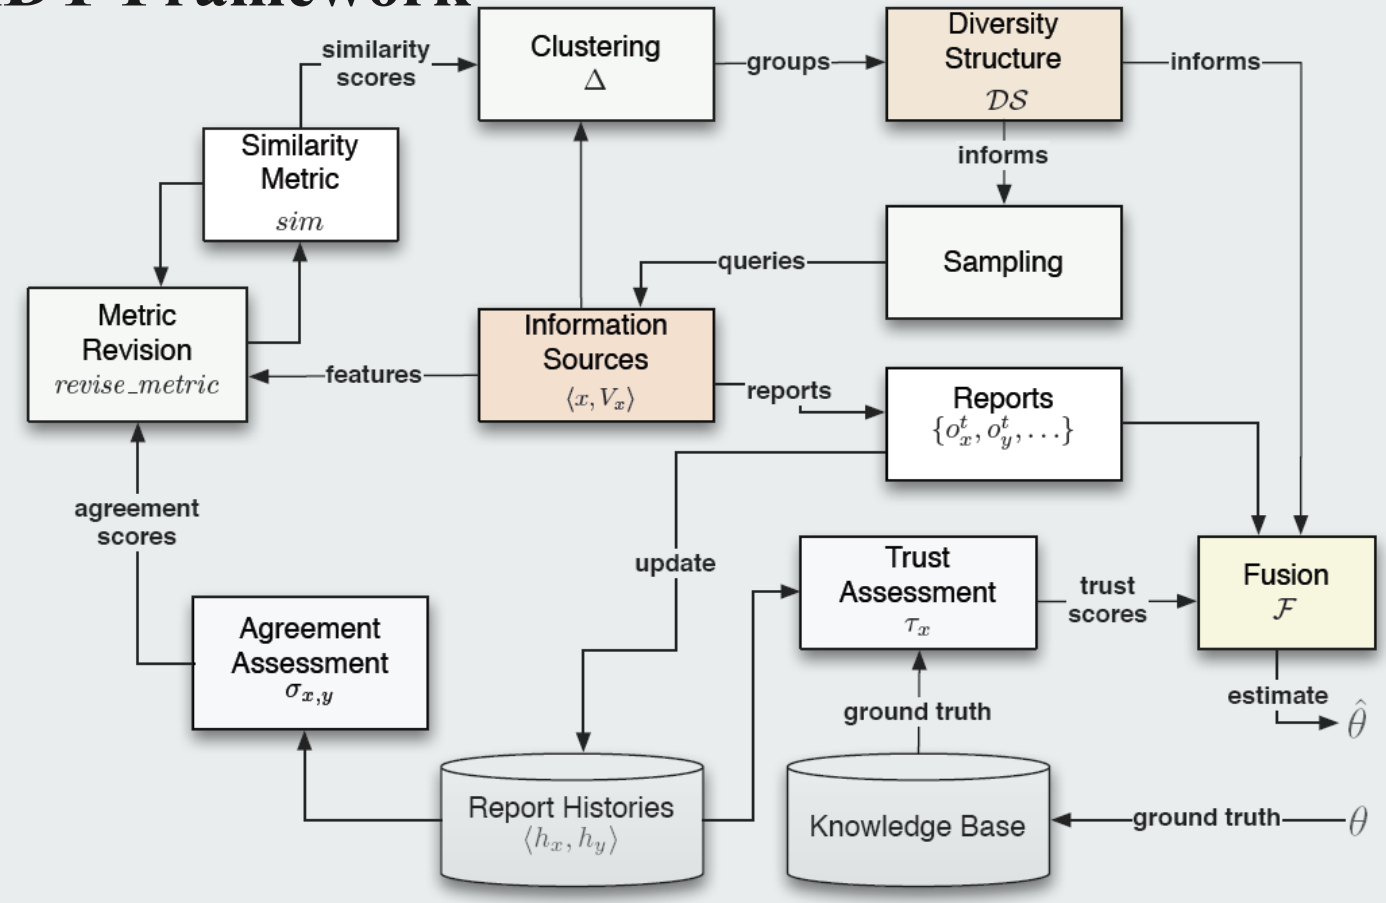

TIDY Framework explanation

- Uses Diversity of Information
- Uses histories of reports from sources with certain features to form a similarity metric
- Uses that metric to cluster the sources into a diversity structure
- The Diversity structure informs a sampling strategy over the set of sources to create an environmental estimate.
- Example given: Weather station

A set of Reports => Conglomerate of information => Estimate of truthfulness

Weather Station Example

- Task: Give periodic weather reports
- Information Source: Sensors (thermometers, rain gauges)
- Reports: Sensor readings sent to the station
- Features: Sensor’s owner, cost, battery life, etc
- Diversity Structure: Usage of sensors across different locations and owned by different companies
- Information Quality: The accuracy of a given sensor
- Fusion: The actual weather estimate

Updates to TIDY($TIDY_0$)

- Defined a method to quantify both source agreement and trustworthiness by counting instances of agreement/disagreement between either two sources or a source and the environment, respectively, and used those counts in a Beta distribution
- Use of model tree learning to create a similarity metric based on Euclidean distance between features in a pair of sources.
- Use of hierarchical clustering to create the diversity structure
- Use of a learning interval for model revision
- Two strategies for sampling
    - Budget for a group is dependent on its size, and individual sources within the group are chosen randomly
    - If budget is insufficient to cover all groups, then select a single source from the most trustworthy group, then the second most, and so on until the budget is exhausted
- Fusion occurs by averaging the sampled sources within groups, and creating a weighted average of the group reports based on the trustworthiness of those groups

Evaluation

- Two sets of experiments
  - 1st set's independent variables:
    - Sampling budget and proportion of malicious sources (independent but misleading)
  - 2nd set's independent variables:
    - Sampling budget and proportion of colluding sources (likely to copy each other's reports)
- Question: How effective is diversity-based sampling as budget constraints and source trustworthiness vary
  - To answer, compare the following methods
    - Diversity-based sampling ($TIDY_0$)
    - Observation-based sampling - uses assessments of individual sources to guide sampling
    - Majority-based sampling - filters out reports more than one standard deviation from the mean report
    - Random sampling(no filtering or weighting)

Table 3 Experimental Parameter

Parameter	Value	Description
$\mathcal{N}$	100	No. of sources in popl.
P_l	0.1	Popl. change probability
ψ	0.4	Diversity threshold
L	30	Learning interval
δ_{agr}	0.1	Report agreement threshold
δ_{tru}	0.1	Report reliability threshold

Experiments

Table 4 Source profiles

ID	f_1	f_2	f_3
p_1	x		x
p_2		x	x
p_3	x	x	x
p_4			x
p_5	x	x	

- Each source is assigned a profile that determines its reporting behavior
- Each profile has three features, with randomly drawn values
- Each feature has its own distinct Gaussian distribution to draw values from
- Informative profile features have a small standard deviation, while uninformative profile feature follow a uniform distribution
- Each profile has a conformity parameter to declare how correlated source reports in a profile are (based on probability)
    - If a source conforms, it randomly selects an opinion held by a fellow profile member, otherwise, it keeps its opinion
    - Adds additional challenge to model, due to added noise
    - By default, the conformity probability is .8
- Sources may freely (probabilistically) leave/join the system at any time.
    - When a source leaves, it is replaced by a new source to keep the number of sources constant
- Each source has a reliability parameter that determines whether it reports honestly/maliciously

E1 Results

- 10 repetitions, with 100 sampling rounds
- Small budget
  - Low/medium malice - no statistically significant difference
  - High malice - 24% higher accuracy than OBS, 39% higher than RBS, 44% over MBS
- Medium budget
  - Low/medium malice - no statistically significant difference
  - High malice - 26% higher than OBS, 38% higher than RBS, 50% higher than MBS
- Large budget
  - Low malice - 13% lower than MBS, no other statistically significant difference
  - Medium malice - no statistically significant difference
  - High malice - 26% over OBS, 34% over RBS, 49% over MBS
- In general, DBS performs very well in situations with high levels of malice among the sources, but performances diminish with low malice, but no worse than others

E1 Graphs

(a) Small budget

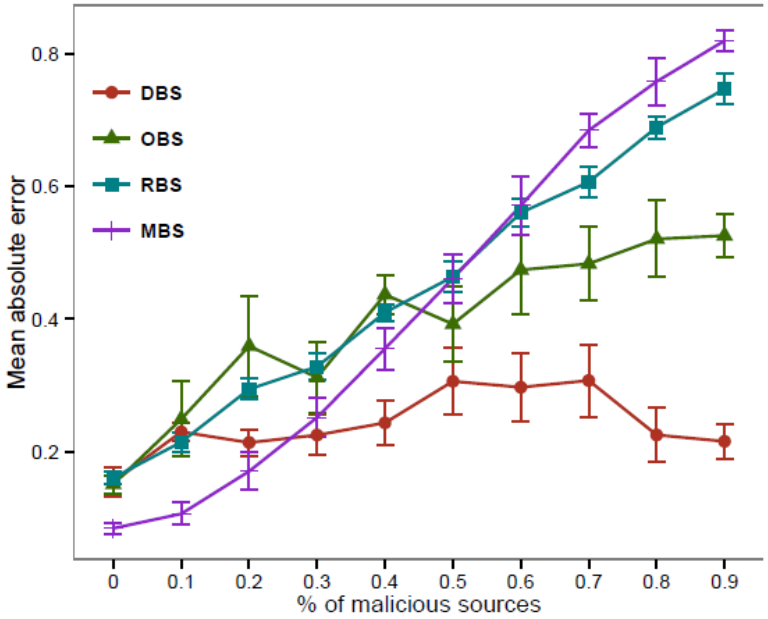

(b) Large budget

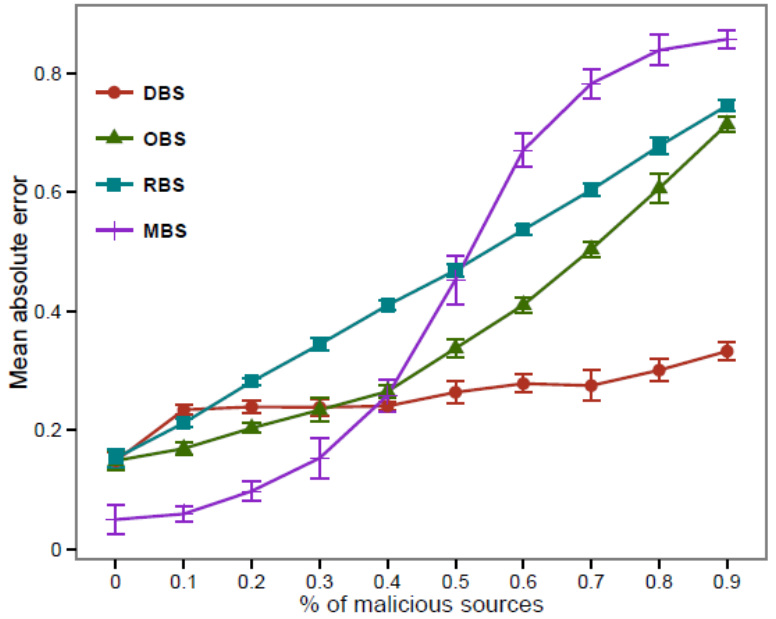

E2 Results

- Small Budget - no signifigant differences between DBS and RBS
  - Low dependency - 2.4% better performance than OBS, 8.2% better than MBS
  - Medium dependency - 6.3% better than OBS, 12.8% over MBS
  - High dependency - 12% over OBS, 20% over MBS
- Medium Budget
  - Low dependency - 5% over MBS, no other significant differences
  - Medium dependency - 9% over OBS, 13% over MBS
  - High dependency - 16% better than OBS, 13% over RBS, and 22% better than MBS
- Large Budget
  - Low dependency - no statistically significant differences
  - Medium dependency - 11% better than OBS, 10% over RBS, 14% over MBS
  - High dependency - 15% better than OBS, 14% over RBS, 20% better than MBS

E2 Graph

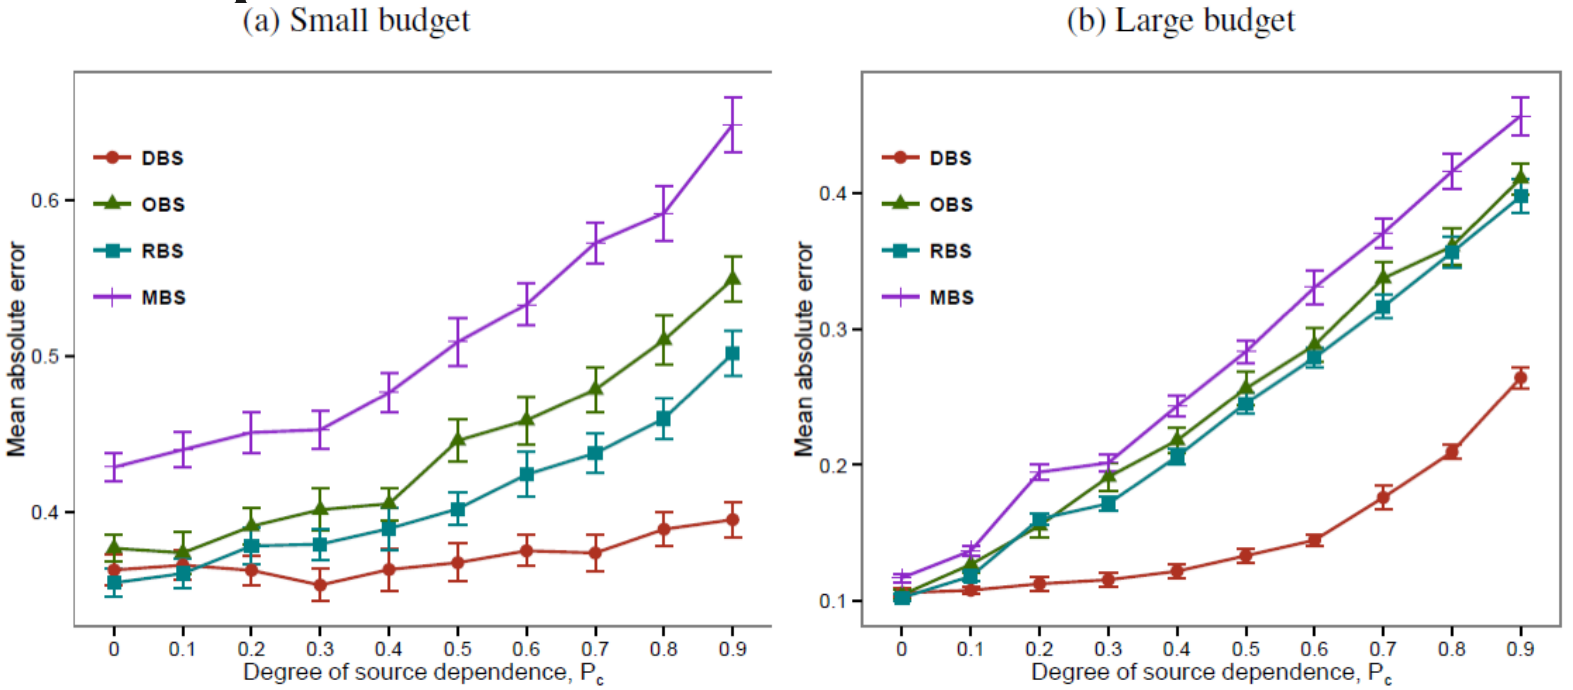
(a) Small budget
(b) Large budget
DBS
OBS
RBS
MBS
Mean absolute error
Degree of source dependence, P_c
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9

E2 explanations

- OBS' vs DBS
  - OBS cannot easily exploit sources' models due the sources not being on the extremes of reliability
  - OBS assumes some amount of independence, which is the antithesis of what was tested
- MBS vs DBS
  - No clear experts, so outliers can't be filtered out by MBS
  - Since the reports aren't very diverse, MBS can't compensate for individual errors
- RBS vs DBS
  - Very similar in low budget cases
  - As budget increases and models can sample more, RBS is more vulnerable to correlated biases
- In general, DBS copes better with low budget

Varying Diversity Thresholds

- The diversity threshold allows for the control of group formations
  - bounded between 0 and 1
- Determines the similarity between source categories
  - If similar, combine
  - If dissimilar, keep seperate
- Greater thresholds allow for greater differences in a source category

Varying Diversity Thresholds

(a) Small budget

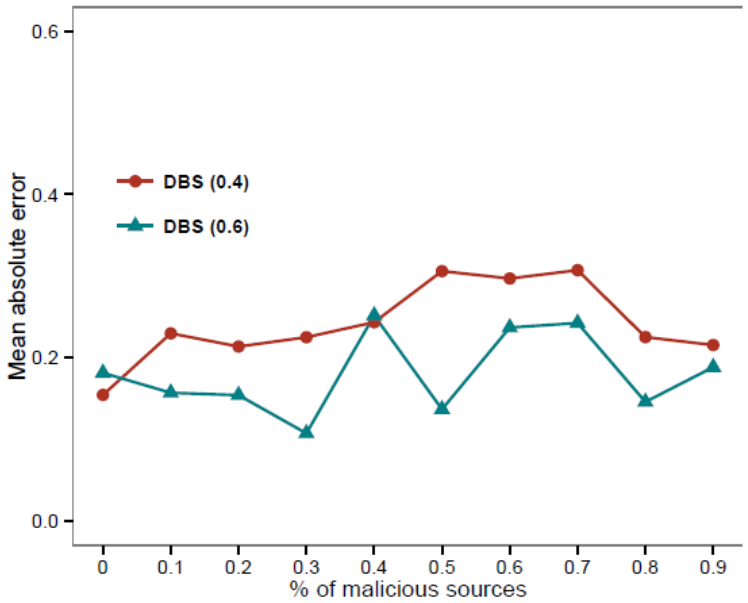

(b) Large budget

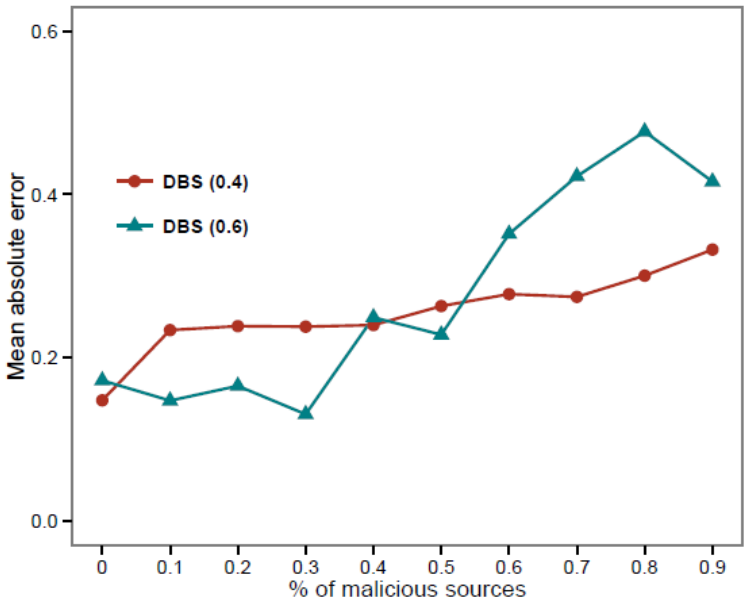

Our Conclusions on the paper

- The paper was pretty long but made for an easy read
- It is clear Diversity Based Systems are superior to RBS, MBS, and OBS

- Relation to Multiagent Systems
  - Agents take in a large amount of information from their environment, and need to organize it effectively to make a decision.
  - Agents need to be able to determine good information from bad information and noise.
  - Agents need to be able to get rid of redundancies in their information to be able to operate efficiently.

Questions?