TIDY: A Trust-Based Approach to Information Fusion through Diversity

Anthony Etuk[1], Timothy J. Norman[1], Murat Şensoy[1,3], Chatschik Bisdikian[2], and Mudhakar Srivatsa[2]

[1]Computing Science Department, University of Aberdeen, UK
[2]IBM T. J. Watson Research Center, NY, US
[3]Computer Science, Ozyegin University, Istanbul, Turkey
{aaetuk,t.j.norman,m.sensoy}@abdn.ac.uk,{bisdik,msrivats}@us.ibm.com

September 2013

Outline

- Introduction
- Problem
- Diversity and Information Fusion
- Evaluation
- Discussion & Future Work

Introduction

Trust and reputation are significant components in many environments for making informed decisions

- Selecting (reliable) interaction partners
- Mitigating risks in potential transactions

Assessment of trust in information typically relies on reports from multiple sources

- More evidence $\sim$ better assessments
- Minimise the risk of biased opinions

A common approach

- Query as many sources as possible
- Use well-known statistical models to make reliable assessments

UNIVERSITY OF ABERDEEN

But...

Querying all possible sources... needed for existing models?

...is not always realistic!

- Costly, especially in resource constrained environments
  - e.g., sensor networks, emergency response, in terms of time and bandwidth
- Risk of double-counting evidence (fact vs. rumour)

Example Scenario

Conflict Management

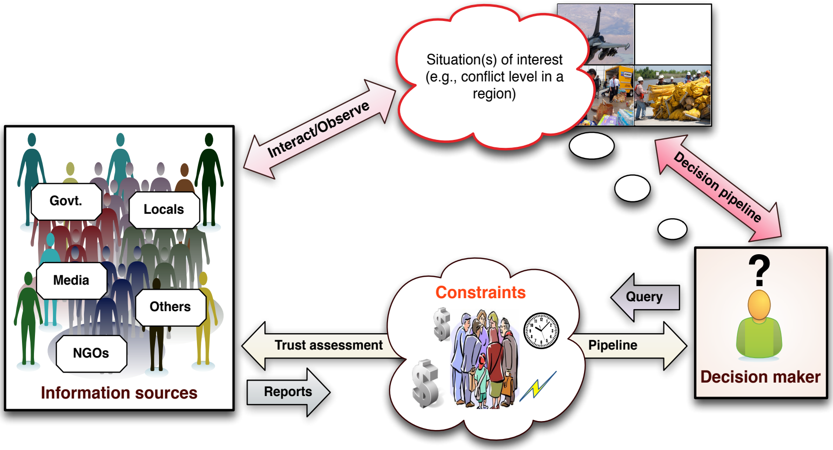

Research Questions

RQ1

How can reliable decisions be reached using evidence from small groups of individuals?

Research Questions

RQ2

How can opinions from diverse sources be taken into consideration, without the risk of double-counting evidence?

Research Questions

Can we intelligently sample from the crowd?

Research Questions

Can we intelligently sample from the crowd?

Challenges

- Information sources may not always provide reliable evidence
  - Malicious, noisy or inaccurate
  - Coordinated (deceptive) actions – collusion
  - Uncertainty in the environment
- Sources may be from different organisations
  - Different motivations/interests/agendas
  - e.g., sensors owned by different organisations
- Trusted partners may leave the system at some point
  - . . . and be replaced by unknown and (possibly) unreliable ones
- Limited capacity to query for evidence
  - e.g. time, bandwidth, information cost

The TIDY Framework

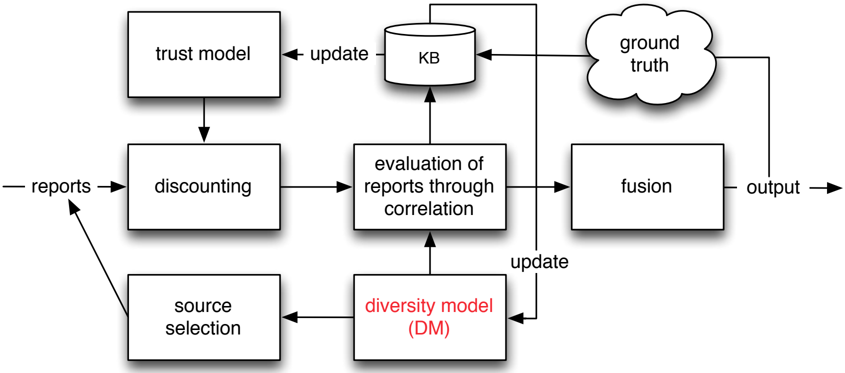

The Diversity Model

Hypothesis

Diversity among information sources may be exploited in order to select a small number of candidates to query for evidence

Aim: Group homogenous sources together in order to reduce the number of queries

Wise/Unwise Crowds

Wise Crowd

- Diversity of Opinion
- Independence
- Decentralization
- Aggregation

+

UnWise Crowd

"...the members of the crowd were too conscious of the opinions of others and began to emulate each other and conform rather than think differently."

∼ Surowiecki, 2004

Measuring Diversity

$\Delta : 2^{\mathcal{S}} \to G$

Challenge

Different similarity metrics may define different subgroups in a population

Proposed solution

Exploit domain knowledge and historical evidence to attempt to disambiguate what metrics lead to better stratification

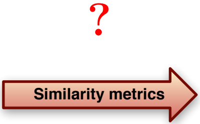

Working assumption

Correlation between features and reports of sources

Learning Diversity

M5 model tree learning

Input: Sources' Features + Reports

- Collection of training instances
- Features $(f_1, f_2, \ldots, f_n)$ e.g., country, location, expertise
- Each f_i coded by numeric values
- Reports $(\mathcal{R}^t_{s,\rho})$ e.g., no. of casualties in a war

Output: Linear regression models (LM)

- Used to predict similarity between sources

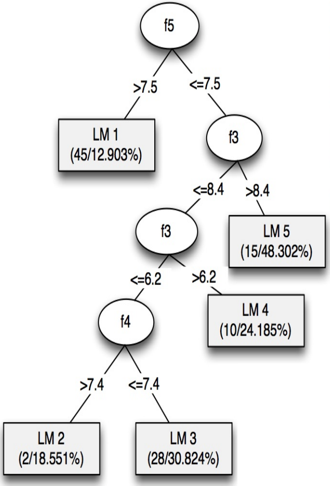


```
f5 > 7.5 : LM1 (45/12.903%)
f5 <= 7.5 :
|  f3 <= 8.4 :
|  |  f3 <= 6.2 :
|  |  |  f4 > 7.4 : LM2 (2/18.551%)
|  |  |  f4 <= 7.4 : LM3 (28/30.824%)
|  |  f3 > 6.2 : LM4 (10/24.185%)
|  f3 > 8.4 : LM5 (15/48.302%)

LM num: 1
class =
      -0.0055 * f1
      + 0.0013 * f2
      + 0.0029 * f3
      + 0.8007 * f4
      + 0.4015 * f6
      + 11.0971
```

Source Stratification

Hierarchical clustering

Procedure

- Uses linear regression models constructed by the M5 algorithm
- Takes the feature vector of any source pair as input
- Obtains a similarity score ($\mathcal{M}$), specifying degree of 'closeness' of the source pair
- The $\mathcal{M}$ measure is used to cluster the sources into groups
- Process terminates when a predefined stoppage condition (diversity threshold) is met

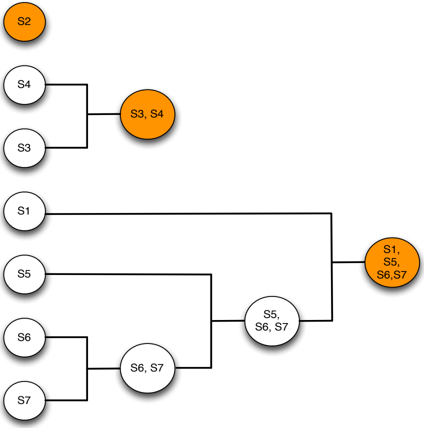

Group Trust

- Trust score is maintained for each group
- Used as an expected reliability of members encountered in a group
- Individual trust is computed using subjective logic (~ Jøsang, 2013)
- Trust of a group is computed as mean trust of group members
- In the absence of any evidence, good or bad outcomes are considered equally likely

Exploiting Diversity for Fusion

Sampling and Fusion of Reports

- Given sampling budget (Φ)
  - expressed in terms of number of sources
- Maintain a competitive exploratory advantage at a reduced cost

Case 1: $\Phi \geq |G|$

- $budget(g_i) = |g_i| \times (\Phi/|\mathcal{S}|)$
- Representative candidates are then randomly selected from g_i according to $budget(g_i)$

Case 2: $\Phi < |G|$

- Each group ($g_i \in G$) is ranked in order of trustworthiness
- $budget(g_i) = 1$, if group rank $> \Phi$; $budget(g_i) = 0$, if group rank $\leq \Phi$

What is the conflict level at region xyz?

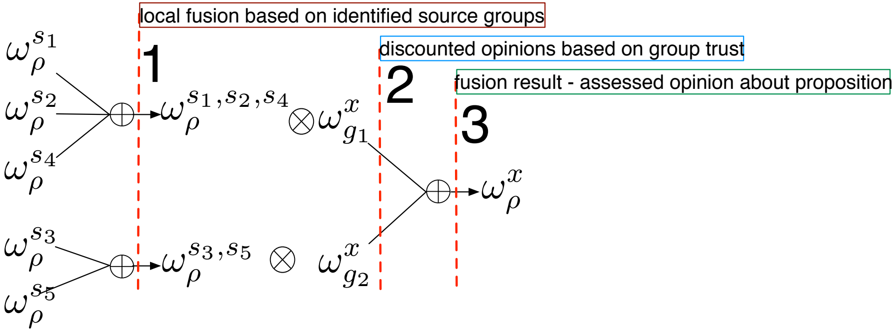

Evaluation

Experimental Conditions

- Experiments based on a simulation test bed
- Measures the effectiveness of TIDY in making accurate assessments
  - Experts (malicious/honest) with knowledge of ground truth
  - Non-experts reporting inconsistently on the ground truth
- Explores the effect of correlated behaviour in the population
  - Subjectivity due to conditioning factors e.g., organisation policy, collusion
- Effect of different budget constraints on trust assessments
- Compares technique to popular trust approaches in literature
  - Observation-based sampling (OBS) e.g., Teacy et al., 2006
  - Majority-based sampling (MBS) e.g., Zhang et al., 2006

Evaluation

Experimental Parameters

- Report types
  - Honest report: closer to the ground truth, small gaussian noise $N(0, 0.01)$
  - Malicious report: significantly deviated from the ground truth, $N(1, 0.01)$
- Source population: 100
- Number of profiles: 3 $(p1, p2, p3)$
- Profiles reliability probability (P_r): $p1 = 0.2$, $p2 = 0.8$, $p3 = 0.9$
- Profiles conformity probability (P_c): $p1 = 0.8$, $p2 = 0.8$, $p3 = 0.8$
- Population change probability (Pl): 0.1
- Diversity threshold (δ): 0.4
- Report similarity threshold (η): 0.01

Results

Robustness to deception with experts (Small budget)

- Small budget: 5 sources sampled in each sampling round
- Increasing percentage of malicious sources
- Graph shows estimation accuracy of different approaches

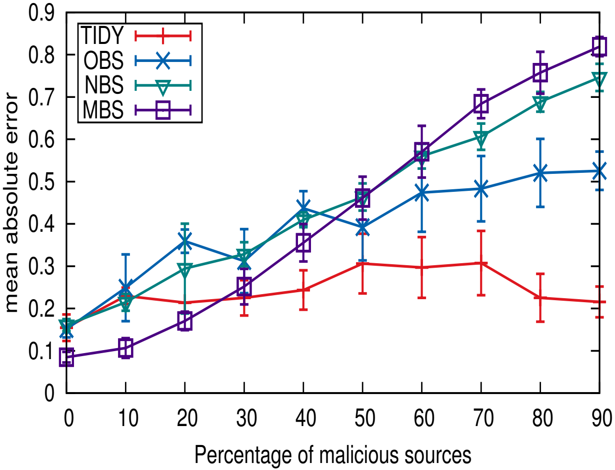

Results

Robustness to deception with experts (Large budget)

- Large budget: 75 sources sampled in each sampling round
- Increasing percentage of malicious sources
- Graph shows estimation accuracy of different approaches

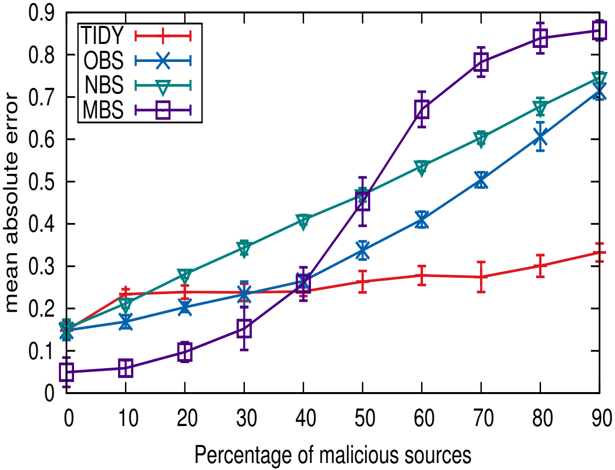

- Small budget: 5 sources sampled in each sampling round
- Increasing percentage of source dependence
- Graph shows estimation accuracy of different approaches

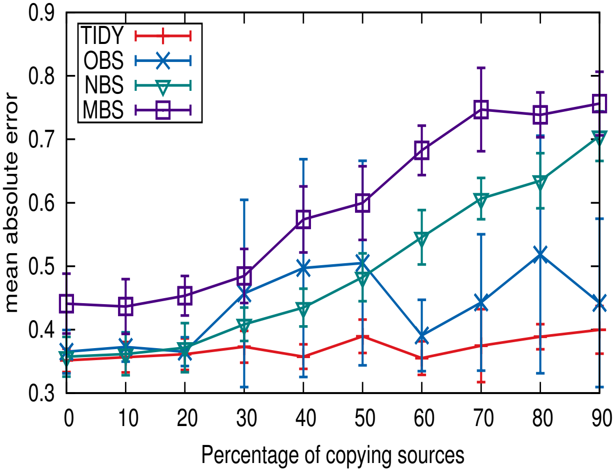

UNIVERSITY OF ABERDEEN

Results

Robustness to deception with non-experts (Large budget)

- Large budget: 75 sources sampled in each sampling round
- Increasing percentage of source dependence
- Graph shows estimation accuracy of different approaches

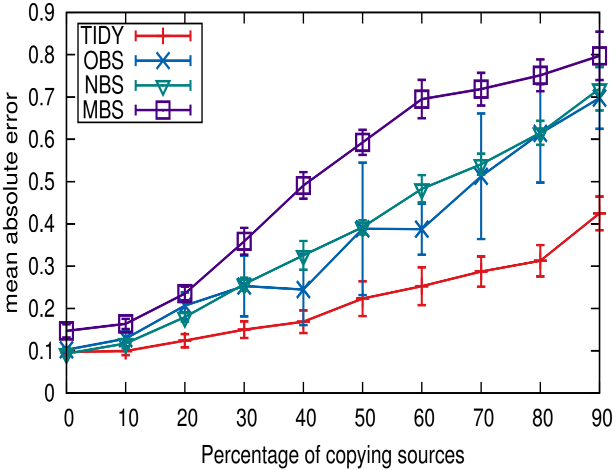

Conclusion

- Existing approaches to information fusion exploiting trust and reputation could be problematic
  - Not always realistic to query many sources for evidence due to costs e.g., time, bandwidth
  - Reports from multiple sources expose one to the risk of double-counting evidence

- Where hidden networks or patterns defining group behaviour exist in the population
  - Relevant features and evidence from past reports of sources can be exploited to stratify the source population
  - Resulting models can be exploited to sample a small number of sources and to protect against biases

What Next?

- Robust and principled decision-theoretic framework to handle complex source selection strategies
- Address more dynamic settings involving streaming information from multiple sources
- Apply model to real-life applications like crowdsourcing and sensor networks

Questions

Thank you!

Appendix I

Subjective Logic (SL) ~ Jøsang, 2013

- A type of probabilistic logic that explicitly takes uncertainty and belief ownership into account
- Arguments in SL are subjective opinions about propositions
- A binomial opinion of an agent x about the truth of a proposition ρ is represented by the quadruple $\omega_\rho^x = (b, d, u, a)$, where: b is the belief that ρ is true; d is the belief that ρ is false; u is the uncertainty about ρ; and a is the base rate
- $b + d + u = 1$ and $b, d, u, a \in [0, 1]$
- Opinions are formed on the basis of positive and negative evidence
- The variables p and q, represent the number positive and negative observations about ρ respectively, and can be used by x to obtain an opinion about ρ as:

$$b = \frac{p}{p+q+2}, d = \frac{q}{p+q+2}, u = \frac{2}{p+q+2}.$$

The probability expectation value of an opinion is defined as:

$$E(\omega_\rho^x) = b + u \times a.$$

UNIVERSITY OF ABERDEEN

Appendix II

Trust and Report similarity computation

	s_1	s_2	s_3	gT
q_1	1	0	1	1
q_2	0	1	0	1
q_3	1	1	1	1
q_4	0	1	0	1
q_5	1	0	1	1
τ	0.57	0.57	0.57	

(a) Trust relationship

	(s_1,s_2)	(s_1,s_3)	(s_2,s_3)
q_1	0	1	0
q_2	0	1	0
q_3	1	1	1
q_4	0	1	0
q_5	0	1	0
φ	0.28	0.85	0.28

(b) Similarity relationship

Report matrix

- For source s_1, the number of positive evidence (complying with ground truth gT) p is 3, and the number of negative evidence (conflicting with gT) q is 2
- Positive evidence p represents instances a source pair gives similar reports, and negative evidence q are those instances the pair gives conflicting reports