Strategies for Truth Discovery under Resource Constraints

(Extended Abstract)

Anthony Etuk
Timothy J. Norman
Nir Oren
Dept. of Computing Science
University of Aberdeen, UK

Murat Şensoy
Dept. of Computer Science
Özyeğin University, Turkey

ABSTRACT

We present a decision-theoretic approach for sampling information sources in resource-constrained environments, where there is uncertainty regarding source trustworthiness. We exploit diversity among sources to stratify the population into homogeneous subgroups to both minimise redundant sampling and mitigate the effect of source collusion. We show through empirical evaluation that our model is as effective as existing truth discovery approaches with respect to accuracy, while significantly reducing sampling cost.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous

Keywords

Trust, Truth Discovery, Diversity, Sampling

1. INTRODUCTION

The tremendous surge in the number and variety of information sources available to support decision-making calls for efficient methods of harnessing their potential. Information sources may be unreliable, biases may be introduced due to source collusion, and misleading reports can affect decisions. Truth discovery mechanisms [1, 3] typically rely on reports from as many sources as possible. In many real-world contexts, however, capturing and distributing evidence would require utilising scarce resources such as bandwidth or energy. Furthermore, relying on responses from all possible sources may adversely affect utility, for example when the cost of sampling some sources outweighs the value derived.

In this research, we combine a form of stratification and trust assessment in order to optimise the selection of sources, an approach that allows us to effectively manage truth discovery under resource constraints. We exploit diversity among sources to form groups, made up of those likely to provide similar reports. We then use reinforcement learning to identify effective sampling strategies across groups. In this way, we relax important limiting assumptions underlying truth discovery and trust approaches: (i) the greater the number of reports acquired, the better the estimate; (ii) reports are independent; and (iii) acquiring reports is cost-free. We show empirically that our model, DRIL (Diversity modelling and ReInforcement Learning), performs as well as classical trust approaches in estimating ground truth, but by sampling significantly fewer sources.

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.*

2. TRUTH DISCOVERY STRATEGIES

We consider an agent that has the task of monitoring an environmental state, θ (e.g., the weather condition at a location). For a particular query, Θ represents the set of possible values of θ. The value of $\theta \in \Theta$ can change over time, and the agent must, therefore, repeatedly update its estimate at time t, $\hat{\theta}^t$, of the environmental state, θ^t.

To acquire a value of $\hat{\theta}^t$, sources of varying trustworthiness may be queried. Let $\mathcal{N}$ denote the set of sources known to the agent. A report received from a source, $i \in \mathcal{N}$ at time t regarding θ^t is denoted o_i^t. Querying a subset of the sources, $N \subseteq \mathcal{N}$ incurs a cost to the agent. We define this *sampling cost* as a function: $cost : 2^N \rightarrow \mathbb{R}$. Our goal is to optimise the selection of sources by learning subsets of $\mathcal{N}$ (a diversity structure), and then deciding how to query those subsets.

2.1 Modelling Source Diversity

A *diversity structure*, $\mathcal{DS} = \{G_1, \ldots, G_K\}$, is a stratification of the source population into K homogeneous and non-overlapping subgroups according to some criteria. We assume that sources have observable features; for example, the number and types of followers it has in a social network. We learn these criteria by looking for correlations between reports received over time and source features. In this way, we can cluster sources that are likely to provide similar reports. The diversity structure is then computed using source features, not reports from individual sources, which means the accuracy of allocating a source to a group does not depend on the number of reports from that source [2].

2.2 Sampling Decision-Making

Our sampling problem is formulated to exploit source similarity as captured in a diversity structure, $\mathcal{DS}$. Reinforcement learning provides a means for an agent to identify good allocations across diverse groups of sources.

States: For $G_k \in \mathcal{DS}$, let τ_k and σ_k denote the trust and similarity parameters of G_k respectively. A sampling state, s_t is a tuple $\langle \mathcal{T}, \Sigma \rangle$, where $\mathcal{T} = <\tau_1, \ldots, \tau_K>$ and $\Sigma = <\sigma_1, \ldots, \sigma_K>$ are vectors corresponding to the trust and similarity levels of groups $G_k \in \mathcal{DS}$, $1 \leq k \leq K$.

Actions: Actions are defined in terms of sampling allocations to the different groups in $\mathcal{DS}$. Let $\mathcal{G}_k = \{0, 1, \ldots, |G_k|\}$

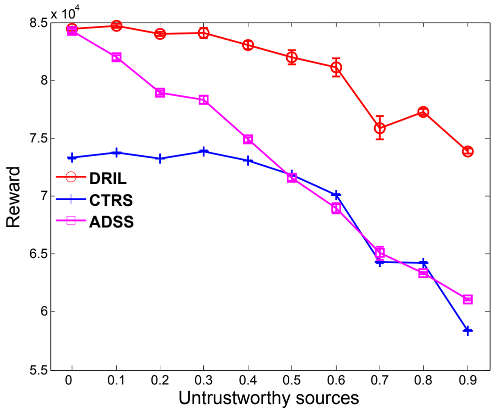

Figure 1: Total reward, $\lambda = 0.8$

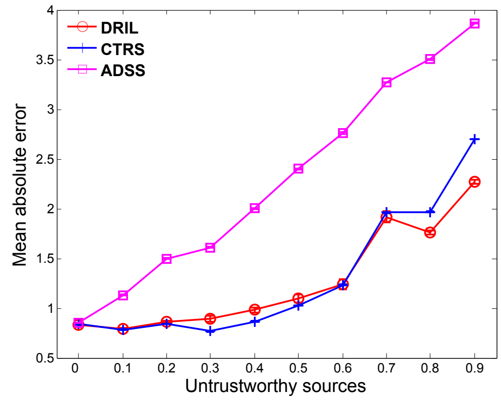

Figure 2: Estimation accuracy, $\lambda = 0.8$

be the set of possible allocations to G_k. A sampling action, a_t is a tuple $\langle g_1, \ldots, g_K \rangle$ such that $g_k \in \mathcal{G}_k$, $1 \leq k \leq K$.

Reward: We define the optimisation objective or reward, r_{t+1}, as a function of information quality, or the deviation of the estimate, $\hat{\theta}^t$ from ground truth, θ^t (*qual*) and the cost of sampling the selected sources (*cost*):

$$r_{t+1} = \lambda r_{t+1_{qual}} + (1 - \lambda) r_{t+1_{cost}}$$

$r_{t+1_{qual}}$ and $r_{t+1_{cost}}$ represent rewards for the quality and the cost criteria respectively: $r_{t+1_{qual}} = 1 - \frac{|\hat{\theta}^t - \theta^t|}{\theta^t}$, $\theta^t \neq 0$, and $r_{t+1_{cost}} = \frac{cost(\mathcal{N}) - cost(N)}{cost(\mathcal{N})}$. The parameter, $\lambda \in [0, 1]$, is a coefficient that controls the trade-off between the different optimisation criteria considered (i.e. quality versus cost).

We asssume here that we have access to ground truth, θ^t, at some point, but too late for it to inform decision-making.

3. EVALUATION

In evaluating our approach, we focus on the following metrics: (i) the total reward or utility up to a certain time-point; (ii) accuracy (error) of estimates with respect to ground truth. We compare our technique (DRIL) to the following strategies for truth discovery: "ADSS" (ADaptive Stratified Sampling), which selects a strategy that leads to the least variance in the estimate, and "CTRS" (Classical Trust and Reputation Sampling), which selects all sources and uses trust assessments to discount reports received.

Design. The system consists of 100 sources, each randomly assigned to one of three profiles that determines its reporting pattern. Each profile has three features. Feature values for each profile are drawn from Gaussian distributions. Each profile has a correlation parameter, P_c that specifies the degree to which reports of sources in a profile tend to be correlated. With probability P_c a source will provide a similar report to other profile members, and independent reports with probability $1 - P_c$. The P_c parameter is set at 0.8 for all profiles. Each source has a reliability parameter, P_r that determines the type of reports it provides (i.e., reliable or malicious reports). Reliable reports (provided by sources with high P_r) are distributed according to $N(\theta^t, 0.01)$. Malicious reports (provided by sources with low P_r) are distributed according to $N(\theta^t + \varepsilon, 0.01), \varepsilon \in [0, 5]$. The reinforcement learning model was instantiated with the following parameter values: the learning rate, $\eta = 1.0$; the discount factor, $\gamma = 0.1$; and the temperature, T_{mp} fixed at 0.1.

Results. Figures 1 and 2 show reward and estimation error when there is a high preference for information quality ($\lambda = 0.8$): a more demanding case for DRIL. Our model, DRIL, performs better than both CTRS and ADSS in terms of reward (Figure 1). A one-way analysis of variance showed a highly significant effect of truth discovery strategy on reward ($p < 8 \times 10^{-4}$). Multiple comparisons (Tukey's HSD, α-level $= 0.05$) indicated that the only significant two-way contrasts showed that DRIL out-performed both CTRS and ADSS. Similar analyses showed that DRIL and CTRS outperform ADSS in estimation accuracy, but that there is no significant difference between DRIL and CTRS (Figure 2).

4. CONCLUSIONS

In this research we have studied the problem of how to optimally sample a population of sources to estimate ground truth. This is an important problem to solve for real-world applications, such as sensor networks and crowd-sourced sensing, where working within resource constraints is critical. The DRIL model combines source diversification and reinforcement learning to drive sampling strategy. We argue that this technique is robust to source collusion, and have demonstrated empirically that this approach performs as well as classical trust approaches in estimating ground truth. Further, by sampling fewer sources DRIL significantly reduces the cost of generating a good estimate.

Acknowledgments

This research was supported by the Petroleum Technology Development Fund (PTDF) Overseas Scholarship Scheme (OSS) Nigeria, and the U.S. Army Research Laboratory and the U.K. Ministry of Defence accomplished under Agreement Number W911NF-06-3-0001.

REFERENCES

[1] X. Dong, B. Saha, and D. Srivastava. Less is more: Selecting sources wisely for integration. *Proc. of the VLDB Endowment*, 6(2):37–48, 2012.

[2] A. Etuk, T. J. Norman, M. Şensoy, C. Bisdikian, and M. Srivatsa. TIDY: A trust-based approach to information fusion through diversity. In *Proc. 16th Int. Conf. on Information Fusion*, pages 1188–1195, 2013.

[3] X. Yin, J. Han, and P. S. Yu. Truth discovery with multiple conflicting information providers on the web. *IEEE Trans. on Knowledge and Data Engineering*, 20(6):796–808, 2008.