Recommending messages to users in online participatory media settings

Noel Sardana Robin Cohen
Cheriton School of Computer Science, University of Waterloo, Waterloo, ON

UNIVERSITY OF WATERLOO

Overview/Summary 1

- Participatory media setting (e.g., Reddit, self-help health network, Massive Open Online Courses (MOOCs), etc.)
- Problem: information overload
  - Which messages do users want to see?
  - How can we help users sift through the raft of information?
- Not all peers are credible; not all peers have similar interests

Background 2

Learning Object Annotation Recommendations (LOAR)

- Developed by Champaign et al. [1] to recommend annotations on learning objects (viz., video lectures, quizzes, etc.) to students in an online learning scenario
- Students who experience learning objects leave annotations, which are subsequently rated (good or bad) by other students
- Annotation authors acquire a "global" reputation that depends on the ratings attached to each authored annotation
- Annotations themselves develop a "local" reputation that depend on the ratings it receives from students who experience it
  - Local reputation is customized for each new student by weighting ratings according to the rater's similarity to the given student

Personalized Trust Model (PTM)

- Trust model designed by Zhang et al. [4] to determine the trustworthiness of sellers in an e-marketplace
- Amalgamates evidence from advisors by updating a Beta distribution prior
- Updates are weighted by a buyer's *local* history with an advisor (or, if no suitable history exists, uses an advisor's global reputation)
- Combines third party advice a seller with the buyer's own local knowledge about the given seller (from past interactions) to predict seller benefit (trustworthiness)

Bayesian Credibility Model (BCM)

- Proposed by Seth et al. [3] to derive the credibility of a message in a social network for use in recommender systems
- BayesNet approach: latent variables integrate four facets of credibility: cluster, public, experience, and role-based
- Derived message benefit can be used to impute a ranking to each message and ultimately recommend messages

Motivation and Example 3

- Initial model motivated by *folklore*: the potential spread of false information (inspired by the online student learning environment)
- Idea: adapt LOAR by weighting ratings by credibility as well as similarity

Example: Health forum

- Suppose a health issue is discussed by members of a health forum
- Members can rate messages positively (1) or negatively (0), and can poll other members for advice about messages
- Many similar, inexperienced peers/patients ($p1$,$p2$,$p4$) recommend a message ($m6$) that contains false information about diagnosing/treating a medical condition
- Experienced peer ($p3$, a doctor/nurse) does not recommend the message

Table : User message ratings

	m_1	m_2	m_3	m_4	m_5	m_6
p_1	0	1	1	1	0	1
p_2	1	1	1	1	0	1
p_3	0	0	0	1	0	0
p_4	0	1	1	0	1	1
s	1	1	1	1	1	?

Table : Peer similarities to student s

	p_1	p_2	p_3	p_4
s	0.2	0.6	-0.6	0.2

- Due to overwhelming evidence/advice from inexperienced, non-credible members, $m6$ will be (erroneously) predicted to have a high benefit for user s
- **Conclusion:** peer credibility is important to model

Algorithm (Cred-Trust) 4

Input: The current user, u, his set of peers, P, their credibility scores, $c_p \in [0,1]$, and their corresponding ratings for the annotation in focus, $r_p \in [0,1]$
Output: Beta(α^*,β^*) which encodes the desirability of the current message

$\alpha^* = \beta^* = 1$ `// At the start, user has uniform expectation about message benefit`
foreach $p \in P$ **do**
  $hr_{up} \longleftarrow$ *computeHammingRatio*(u,p) `// A measure of rating similarity`
  if $r_p == 0$ **then**
    `// Adjust the similarity weight by credibility:`
    $\alpha^* += h_{up}(1 - c_p)$
    $\beta^* += 1 - h_{up} \cdot (1 - c_p)$
  else
    `// Else simply compute a credibility-dampened trust score`
    $\alpha^* += c_p \cdot (1 - h_{up})$
    $\beta^* += c_p \cdot h_{up}$
  end
end

Simulation Setup 5

- Multiagent environment: 20 agents
- Agents both generate and rate messages according to individual inherent credibility
- In the mean credibility trials, all agent credibilities are generated according to a mean credibility parameter, and each agent rates each message
- In the sparse ratings trials, agents tend to only rate messages they like
- We report the MCC, a measure of the accuracy of recommendations made by each algorithm based on the number of correctly classified messages (higher is better)

Simulation Results: Mean Credibility Trials 6

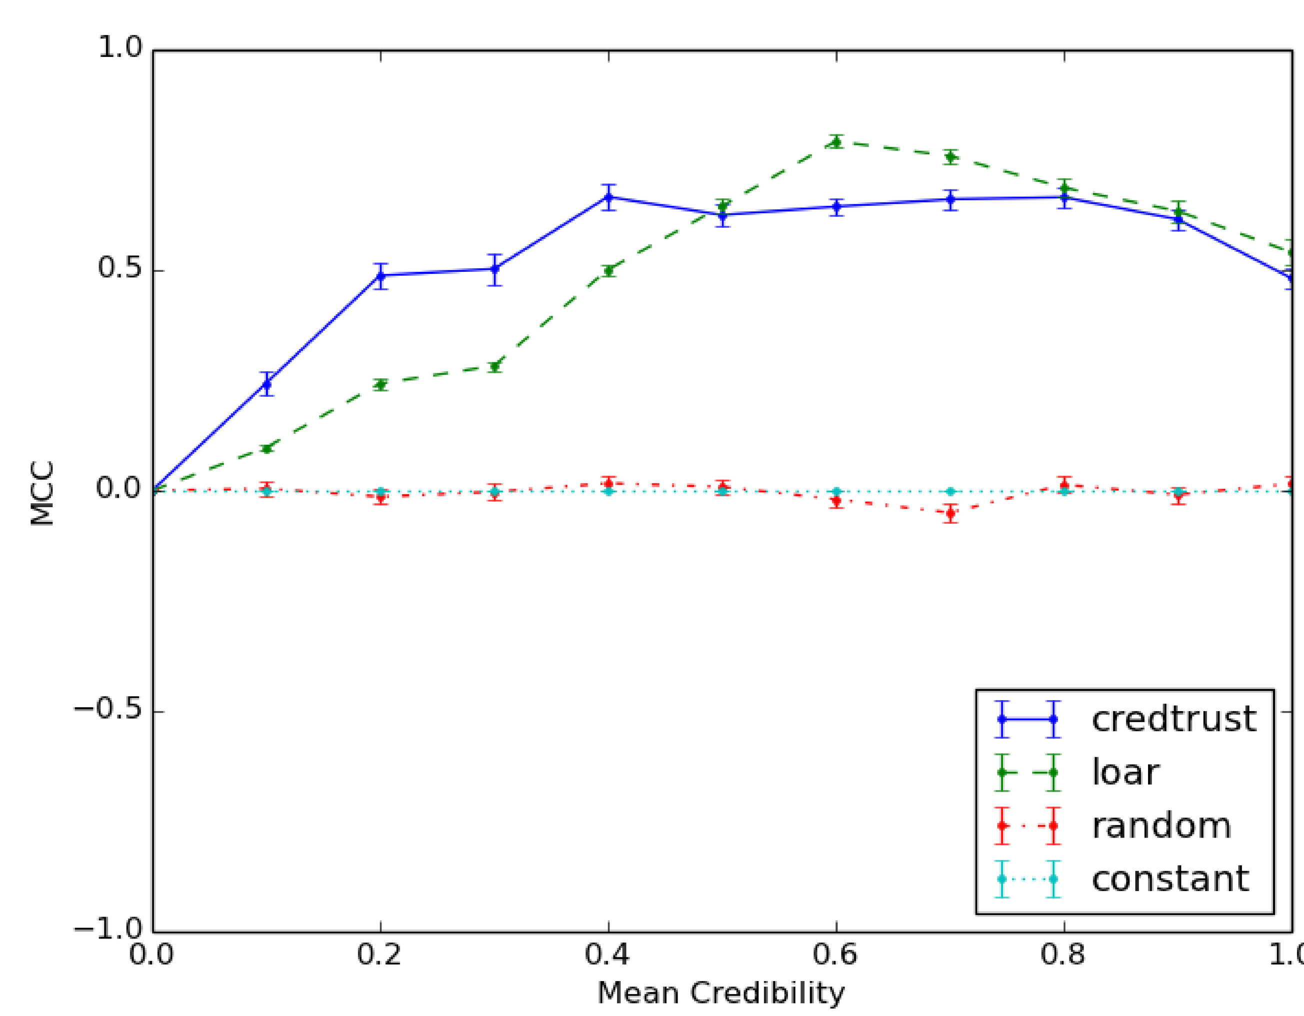

Discussion 7

- Cred-Trust outperforms LOAR when there are a larger number of low credibility agents
- LOAR outperforms Cred-Trust when most agents are moderately credible
- The two algorithms converge as mean credibility in the system approaches 1
- However, this occurs when ratings are "dense"; the LOAR scheme simply discounts peer advice according to similarity
- This suggests that we examine the behaviour of both algorithms in a case where agents rate with some bias and where ratings are not dense
- The sparse ratings case demonstrates how Cred-Trust is able to account for and correct non-credible peer advice when agents are biased with their feedback

Simulation Results: Sparse Ratings Trials 8

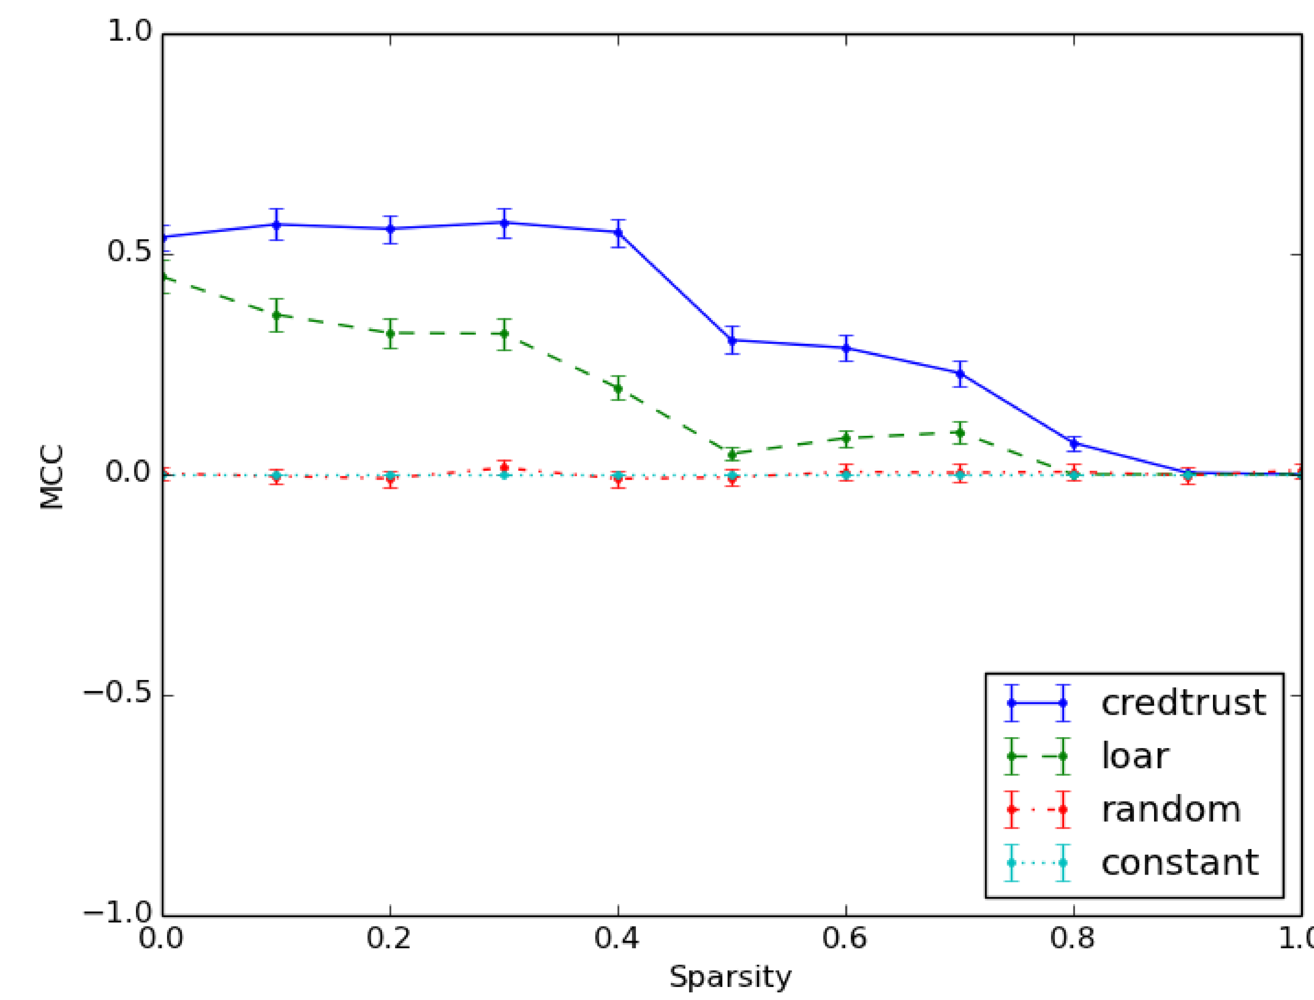

Toward Robust Classification and Decision Making 9

- Most recommendation techniques compute a benefit metric, and then use it to classify items of interest according to some acceptability threshold
- Idea: migrate entire process into a Partially Observable Markov Decision Process and directly decide whether or not to show a given message
- A decision making agent chooses whether to poll advisors for information (tuples that include a message rating, similarity, credibility, etc.), or whether to recommend a message
- Polling for advice allows the decision making agent to update its belief about a message according to a probabilistic observation function
- Inspired by trust models like BLADE [2], this novel approach uses POMDPs to make classification decisions; this allows us to integrate learning (about observation functions) and user utilities (in the form of rewards) to better-classify messages
- Based on beliefs about the underlying message state, the agent develops and follows a decision making procedure for users (show a message or not)
- Currently validating this new approach in simulation and also analyzing using real-world data from Reddit.com and Epinions.com

Conclusions and Future Work 10

- We proposed a method for combining similarity and credibility when considering messages to recommend to users
- This method outperforms LOAR, the model predecessor, in scenarios where there is a dichotomy between peer credibilities and when agents are biased when providing feedback about messages
- Future: comparisons with BLADE using real-world data

References 11

[1] John Champaign, Jie Zhang, and Robin Cohen, *Coping with poor advice from peers in peer-based intelligent tutoring: The case of avoiding bad annotations of learning objects*, Proceedings of User Modeling, Adaptation and Personalization (UMAP), 2011, pp. 38–49.

[2] Kevin Regan, Pascal Poupart, and Robin Cohen, *Bayesian reputation modeling in e-marketplaces sensitive to subjectivity, deception and change*, Proceedings of the 21st National Conference on Artificial Intelligence (AAAI), 2006, pp. 1206–1212.

[3] A. Seth, J. Zhang, and R. Cohen, *Bayesian credibility modeling for personalized recommendation in participatory media*, Proceedings of the International Conference on User Modeling, Adaptation and Personalization (UMAP), 2010, pp. 279–290.

[4] Jie Zhang and Robin Cohen, *Evaluating the trustworthiness of advice about seller agents in e-marketplaces: A personalized approach*, Electronic Commerce Research and Applications (2008), 330–340.

nisardan @ uwaterloo.ca