Collaborative Redundant Agents: Modeling the dependences in the diversity of the agents' errors

Laura Zavala[1], Michael Huhns[2], and Angélica García-Vega[3]

[1]Computer Science, University of Maryland Baltimore County, Baltimore, MD, USA
rzavala@umbc.edu
[2]Computer Science, University of South Carolina, Columbia, SC, USA
huhns@sc.edu
[3]Facultad de Física e Inteligencia Artificial, Universidad Veracruzana, Xalapa, Ver., México
angegarcia@uv.mx

Abstract. As computing becomes pervasive, there are increasing opportunities for building collaborative multiagent systems that make use of multiple sources of knowledge and functionality for validation and reliability improvement purposes. However, there is no established method to combine the agents' contributions synergistically. Independence is usually assumed when integrating contributions from different sources. In this paper, we present a domain-independent model for representing dependences among agents. We discuss the influence that dependence-based confidence determination might have on the results provided by a group of collaborative agents. We show that it is theoretically possible to obtain higher accuracy than that obtained under the assumption of independence among the agents. We empirically evaluate the effectiveness of a collaborative multiagent system in the presence of dependences among the agents, and to analyze the effects of incorrect confidence integration assumptions.

Keywords: Multiagent Systems, Collaboration, Data Fusion

1 Introduction

With the increasing availability of sources of information and services, and the involvement of people and organizations as contributors, there are also increased needs and opportunities for the use of collaborating multiagent systems that make use of multiple sources of knowledge for validation and improvement of accuracy and reliability. Potential application examples are situational aware systems [18], sensor networks [12], and semantic sensor web [15]. A common task in these systems is the integration of contributions, which is not straightforward. Common techniques for integrating contributions are based on agent reputation models [7] and the assumption of independence among the agents, such as majority voting and weighted majority voting techniques [11,14]. The dependences among agents are usually not considered despite their crucial role. Consider, for example, a group of three weather expert agents who have the task of collectively producing a weather forecast. Each agent has a reputation (probability of being accurate) of 0.7 which has been asserted by looking at their past performance. Suppose that the particular task is to predict whether or not it is going to rain. How can the contributions of these agents be combined to achieve the most accurate possible result and outperform any expert individual performance? Clearly, this is possible only if the agents compensate for the mistakes, limitations, and vulnerabilities of each other. Suppose, in our example, we get a positive result ("it is going to rain") from the three agents, what confidence should we give to the integrated result ("it is going to rain")? Common approaches used when combining contributions are based on the assumption of independence among the contributions and either use an ad hoc technique that makes use of the agent's reputations to determine the confidence, or do not determine a confidence for the integrated contributions at all (select the "best" contribution by some rule like majority). Assuming independence of contributions, the best way to maximize the probability of an accurate team prediction is to use Bayes rule to determine the confidence of each contribution. In our example, there is only one contribution: "it is going to rain", and the confidence assigned using Bayes rule would be 0.927. However, if all the agents employed the same program to determine rain likelihood, they would fail or succeed identically and the confidence in the integrated result should remain at 0.7. On the other hand, if we have a guarantee that the agents never fail together, the confidence in the integrated result should be 1. There is a confidence uncertainty inherent when integrating the agents' contributions, which, in our example, sets an interval from 0.7 to 0.927 for the confidence in the result, with any additional knowledge of the dependence relation among the agents reducing this confidence interval.

We have proposed[22] the use of multiagent-based redundancy as a basis for robust software and Web services development. In this paper, we focus on the issue of integrating agents' contributions. Particularly, we study the influence that dependence-based confidence determination might have on the results provided by a group of collaborative agents. In order to achieve that, we have developed a model that allows the representation of the full spectrum of potential dependences between pairs of collaborative agents by considering the space of their coincident errors which can be totally coincident, independent and non-coincident. Because of the impossibility to manage a model representing all potential dependences, the model allows to capture a simplified version using a linear structure where only adjacent nodes are directly dependent on each other and conditional independence is assumed between all nonadjacent nodes—a reasonable simplification that at worst, understates the potential error of assuming full contribution independence. Using the model, we can empirically evaluate the effectiveness of a collaborative multiagent system in the presence of known dependences among the agents, as well as analyze the effects of incorrect confidence integration assumptions.

In the next section we discuss related work. In section 3 we present our formal model for the representation of dependences among agents. Section 4 show the experiments conducted, which show how the model can be used to measure the sensitivity of complex collaborative multiagent systems to incorrect assumptions of independence among the agents. In section 5 we present our conclusions and discuss our plans for the integration of the work presented in this paper with trust and reputation networks.

2 Related Work

Voting algorithms have been used to provide an error masking capability in a wide range of commercial and research applications. There is a large compendium of software systems implementations and experiments reporting some degree of reliability improvement [11,14,22,19,13]. However, all these works usually make use of simple voting strategies and do not address the issue of dependence, relying strongly on the assumption of independence among the versions. A functional classification is used in [9] to provide a taxonomy of the voting algorithms reported in the literature.

In [1] the authors address the issue of dependence and provide a model for analyzing the effect and propagation of potentially incorrect confidence-integration assumptions in a complex MAS with both concurrent and sequential processing in binary scenarios (two possible outputs). Our model is based on [1] but it does not include the sequential case (chains of decisions) and thus, the results are not affected by previous computations. By concentrating in one-time integration cases we are able to observe better the effect of independence assumptions. Furthermore, we have extended the confidence range to include the cases in which it is possible to obtain higher accuracy than that obtained under the assumption of independence among the agents. Specifically, we have extended the confidence range to include non-coincidence. In [1] the confidence, say of a system of 3 agents with $[.8, .8, .8]$ reliabilities is anywhere between $.8$ (when they fail identically) and $.985$ (when their errors are independent). Our model includes one more case for when their errors are mutually exclusive (non-coincident), so that makes the confidence range (in some cases) wider, since the upper bound can reach 1. Finally, our model allows the representation of cases with more than two possible outputs.

There have been several works [16,17,2] in the field of multiagent systems that deal modeling dependencies among agents. Specifically interdependencies among different agents' goals and actions where an agent is said to be dependent on another if the latter can help / prevent him to achieve one of his goals. This type of dependency relations allow an agent to know which of his goals are achievable and which of his plans are feasible (or not) at any moment. In this way, an agent may dynamically choose a goal to pursuit and a plan to achieve it, being sure that every skill needed to accomplish the selected plan is available in the system. Our work, however, is focused on another type of dependencies among agents, namely, dependences in the context of the diversity in the agents' errors. Furthermore, the target application of our approach is to collaborative multiagent systems where agents combine their contributions to achieve the most accurate possible result and outperform any expert individual performance.

3 A Model of Dependences for Collaborative Agents

Our model can be applied to collaborative multiagent systems in which all the agents in the system have functionally equivalent capabilities and always do their best to provide accurate contributions for a particular task. They cooperate in order to increase the overall outcome of the system and are not concerned with their personal payoffs. In other words, all of the agents share the same goals or are trying to maximize a social utility instead of an individual, personal utility. The goal is for the system to exhibit better average performance than that of any single agent in the system. More specifically, we make the following assumptions:

- the agents provide equivalent functionality

- the agents' contributions are not always correct
- the probability of an agent's contributions being accurate is known
- agents can be trusted to always do their best to provide accurate contributions
- the agents' contributions have identical semantics, and therefore can be integrated trivially

To analyze the spectrum of dependence between a pair of collaborative agents, we consider the space of their coincident errors. The errors can be: totally coincident (maximally correlated), in which case the agents always fail together if their accuracies are the same, or every time the more accurate agent fails if their accuracies vary; independent, where the probability of coincident errors is given by the product of the agent's individual probabilities of error; non-coincident, in which case the agents never fail together. Let $P(A_i)$ denote the accuracy of agent A_i and $1 - P(A_i)$ its probability of error. There are four different possible outcomes for each pair of contributions provided by different agents. Assume, for example, that we have two agents A_1 and A_2 with $P(A_1) = .9$ and $P(A_2) = .6$. Table 1 shows the four different possible outcomes and provides the number of occurrences out of 100 tests for the three different cases of dependence among A_1 and A_2.

$P(A_1) = .90$	$P(A_2) = .60$	C	I	NC
fails	fails	10	4	0
fails	succeeds	0	6	10
succeeds	fails	30	36	40
succeeds	succeeds	60	54	50

Table 1: The four different possible outcomes for each pair of contributions provided by agents A_1 and A_2, and the number of occurrences, out of 100, for each outcome for the three different cases of dependence among the agents. C=Coincident, I=Independent, and NC=Non-Coincident

Figure 1 shows this idea in a graphical representation. We can see what happens as the individual accuracies and the difference between them change. As agent accuracies increase, the contributions naturally become more correlated and the difference between the coincident, independent, and non-coincident extremes becomes smaller. When $P(A_i) = 1.0$ there is no difference. Therefore, if our agents are close to perfect, incorrect dependences assumptions will have less impact in the final result.

A set of decisions made by agents is concurrent if the decisions are made without knowledge of any of the other decisions in the set [1]. Figure 2a illustrates the concurrent decisions of n agents related to node W, which is the state of the world that the agents are trying to match. In our model, W can take on any number of values, though for simplicity we will illustrate it for two values, 1 and 0. In our initial weather agents example, this abstraction could represent $W = \{rain, no - rain\}$ Given a set of n agents as in Figure 2a, there are 2^n subsets or groupings that can be made out of the n agents, and therefore, potential dependences between them. Consequently, a full model representing all the potential dependences between the agents, becomes unmanageable as the number of agents grows. To keep the number of dependences manageable, we

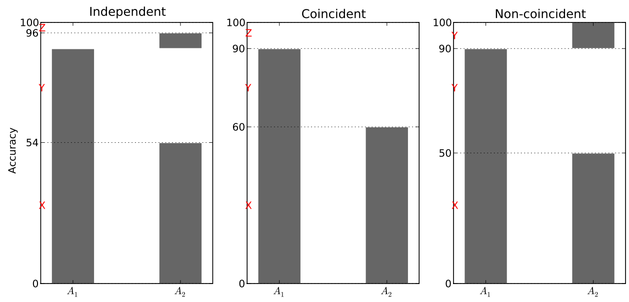

Fig. 1: Dependence relations of two agents. The black bars indicate cases where the agent's result is correct. The percentage of black versus white is the accuracy of the agent. Region X represents where both agents' results are correct; region Y is where their results disagree, and region Z represents where both agents' results are wrong.

simplify the model to the linear structure assumed in [1], where only adjacent nodes are directly dependent on each other and conditional independence is assumed between all nonadjacent nodes—a reasonable simplification that at worst, understates the potential error of assuming full contribution independence. Figure 2b shows the subset of dependences represented in our model. For each pair of nodes with a dependence relation (adjacent nodes) we introduce two nodes, a C node and a NC node, to represent coincident and non-coincident dependence information, respectively, between A_i and A_{i-1}. For simplicity of notation, we will use C_i to refer to the node $C_{i-1 \rightarrow\ i}$ and NC_i to refer to the node $NC_{i-1 \rightarrow\ i}$. When both nodes are $false$ ($C_i = false$ and $NC_i = false$), the errors in A_i and A_{i-1} are independent. In other words, A_i is conditionally independent of A_{i-1} given W, $P(A_i|W) = P(A_i|W, A_{i-1})$.

When $C_i = true$, the errors between A_i and A_{i-1} are maximally correlated with one another (as in the coincident case in Figure 1). When $NC_i = true$, the errors between A_i and A_{i-1} are totally exclusive (as in the non-coincident case in Figure 1).

We can assign C_i and NC_i a value between 0 and 1 to represent the range of dependence from fully coincident ($P(C_i = true) = 1$), to conditionally independent ($P(C_i = true) = 0$) and ($P(NC_i = true) = 0$), to non-coincident ($P(NC_i = true) = 1$).

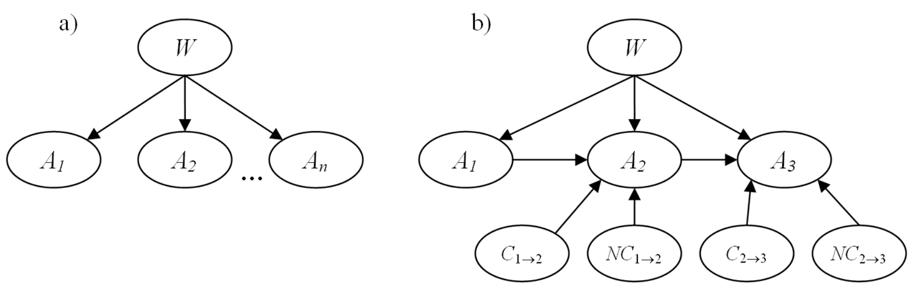

Fig. 2: a) Concurrent decisions of n agents related to node W; b) The subset of dependences represented in our model

Tables 2 to 4 show the conditional probability tables (CPTs) for agent A_2 in the fully coincident, conditionally independent, and non-coincident cases respectively. For intermediate values of $P(C_i = true)$ or $P(NC_i = true)$ (only one can have an intermediate value, while the other must have a 0 value), $P(A_i|W, A_{i-1})$ is obtained by combining the mixed distributions defined by either one or the other. That is: $P(A_i|W, A_{i-1}) = P(C_i = true)P(A_i|W, A_{i-1}, C_i = true) + P(C_i = true)P(A_i|W, A_{i-1})$, when $0 < P(C_i = true) < 1$ or $P(A_i|W, A_{i-1}) = P(NC_i = true)P(A_i|W, A_{i-1}, NC_i = true) + NC_i = true)P(A_i|W, A_{i-1})$, when $0 < P(NC_i = true) < 1$.

A_1	W	C_2	1
1	1	1	$\min(\frac{P(A_2)}{P(A_1)}, 1)$
0	0	1	$1 - \min(\frac{P(A_2)}{P(A_1)}, 1)$
1	0	1	$\min(\frac{1-P(A_1)}{1-P(A_2)}, 1)$
0	1	1	$1 - \min(\frac{1-P(A_1)}{1-P(A_2)}, 1)$

Table 2: Conditional probability table for agent A_2 under the fully coincident model. The errors between A_2 and A_1 are maximally correlated with one another (as in the coincident case in Figure 1).

A_1	W	0	1
-	0	$P(A_2)$	$1 - P(A_2)$
-	1	$1 - P(A_2)$	$P(A_2)$

Table 3: Conditional probability table for agent A_2 under the independent model. The errors in A_2 and A_1 are independent. In other words, A_2 is conditionally independent of A_1 given W.

A_1	W	NC_2	1
1	1	1	$\frac{P(A_2)-(1-P(A_1))}{P(A_1)}$
0	0	1	$1 - \frac{P(A_2)-(1-P(A_1))}{P(A_1)}$
1	0	1	0
0	1	1	1

Table 4: Conditional probability table for agent A_2 under the non-coincident model. The errors between A_2 and A_1 are totally exclusive (as in the non-coincident case in Figure 1).

4 Experiments and Analyses

Our model can be used to measure the sensitivity of complex collaborative multiagent systems to incorrect assumptions of independence among the agents when integrating their contributions. In this section we demonstrate the use of our model with homogeneous and heterogeneous multiagent systems of different sizes and under different scenarios given by real versus assumed dependence relations between the agents.

To determine the accuracy of a particular system configuration, we determine the confidence of the integrated result for each combination of agents' contributions A_i in

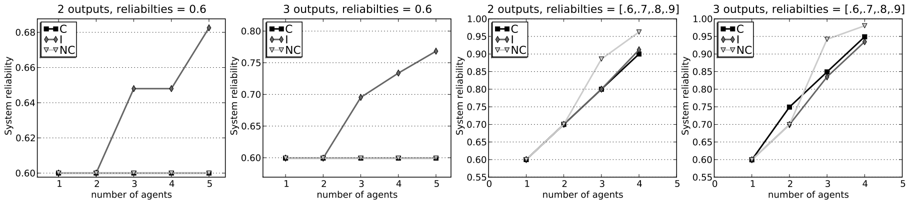

Fig. 3: Reliability of the system for fully coincident (C), independent (I), and non-coincident (NC) agents—homogeneous and heterogeneous cases with 2 and 3 outputs each.

the model. Assuming the system answer is the value of W that maximizes $P(W|A)$, the confidence of an integrated result is then $max_W P(W|A)$. Since we want to evaluate the sensitivity of a system to particular incorrectly assumed dependence relation, we measure the expected belief that our answer is correct by taking a weighted average over all possible instances of the observed data, $\sum_A P(A) \times max_W P(W|A)$ and compare it to the expected belief obtained under the assumed dependence relations, $\sum_A P(A) \times max_W P(W|A, P(C_i), P(NC_i))$.

For simplicity of notation refer to $P(C = true)$ as $(P(C))$ and $P(NC = true)$ as $(P(NC)$. Figure 3 shows the accuracy of homogeneous and heterogeneous systems for fully coincident, independent, and non-coincident agents with 2 and 3 outputs. For homogeneous systems, there is no difference, in terms of accuracy, between fully coincident and non-coincident. The difference in these cases becomes clear when looking at heterogeneous systems.

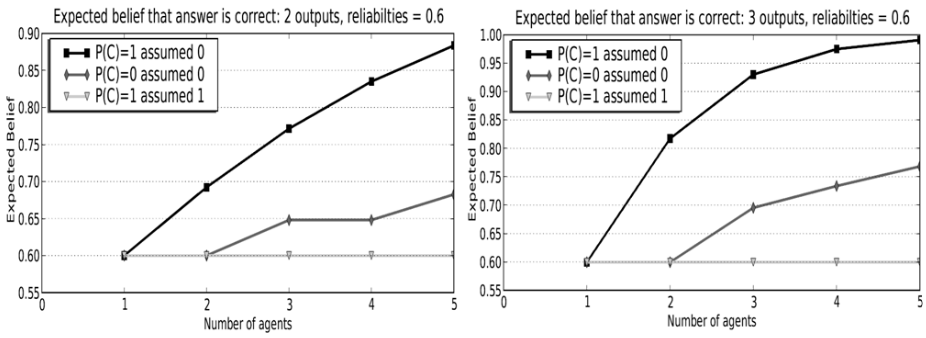

Fig. 4: The effect of assuming agents are independent when they are fully coincident (maximally correlated).

The sensitivity of multiagent systems to incorrect assumptions of independence among the agents is shown in figures 4 to 7. Figure 4 shows the effect of assuming agents are independent when they are fully coincident, for homogeneous systems with two and three outputs. $P(C) = 0, assumed0$ shows the expected belief for the case when the agents are conditionally independent which is the usual assumption when applying Bayes rule. But if we make that assumption and the agents are not conditionally independent, the expected belief is overestimated, at worst, as shown by $P(C) = 1, assumed0$, which is the case when the agents are fully coincident.

Figures 5 to 7 show the sensitivity of multiagent systems to incorrect assumptions of independence among the agents for heterogeneous systems with two and three outputs. As Figure 2b shows, the model implies a direction of the dependence relations. If the system is homogeneous (agents have the same reliabilities), the direction does not matter. But that is not the case for heterogeneous systems. Therefore, we include two cases in figures 5 to 7. The first case is when more reliable agents depend on less reliable agents. The second case is when less reliable agents depend on more reliable agents.

Figure 5 shows the effect of different dependence assumptions when agents are fully coincident. Specifically, assuming agents are either independent or non-coincident when they are fully coincident. The plots on the top row show the case when more reliable agents depend on less reliable agents. The plot for two outputs shows how when agents are fully coincident, adding more reliable agents improves the expected belief, but it can never be greater than that of the most reliable agent at any particular system configuration (number of agents). The plots on the second row show the case when less reliable agents depend on more reliable agents. The expected belief does not improve adding less reliable agents, and again, it can never be greater than that of the most reliable agent. Finally, the expected belief is underestimated when assuming independence and non-coincidence.

Figure 6 shows the effect of different dependence assumptions when agents are independent. Specifically, assuming agents are either coincident or non-coincident when they are independent. The plots on the top row show the case when more reliable agents depend on less reliable agents. When agents are independent, adding more reliable agents improves the expected belief. Contrary to the case when agents are dependent (Figure 5), the expected belief becomes greater than that of the most reliable agent at a certain system configuration (number of agents). The improvement is bigger for systems with a larger number of outputs. Finally, the expected belief is overestimated when assuming coincidence and underestimated when assuming non-coincidence.

Figure 7 shows the effect of different dependence assumptions when agents are non-coincident. Specifically, assuming agents are either coincident or independent when they are non-coincident. The plots on the top row show the case when more reliable agents depend on less reliable agents. When agents are non-coincident, adding more reliable agents improves the expected belief. As in the independent case (Figure 6), the expected belief becomes greater than that of the most reliable agent at a certain system configuration (number of agents), but the improvement is even greater. The improvement is also bigger for systems with a larger number of outputs. Finally, the expected belief is overestimated when assuming independence and coincidence.

5 Conclusions and future work

We presented a domain independent model for the study of the effect that incorrect assumptions about the dependences among collaborative agents have on the overall system accuracy estimation. We showed results of experiments using our model with homogeneous and heterogeneous multiagent systems of different sizes and under different scenarios given by real versus assumed dependences relation between the agents.

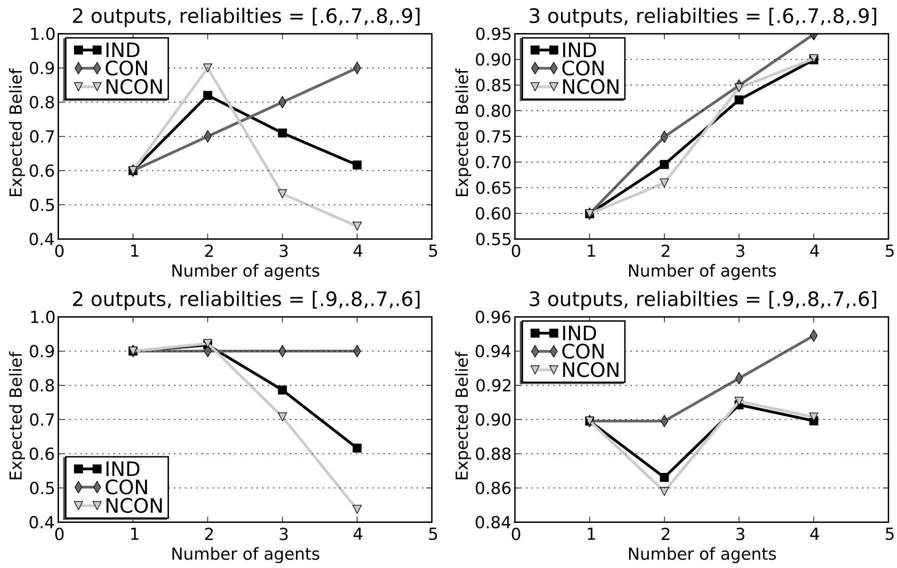

Fig. 5: The effects of assuming agents are either independent (IND) or non-coincident (NCON) when they are fully coincident (CON). When agents are coincident, the expected belief can never be greater than that of the most reliable agent. The expected belief is underestimated when assuming independence and non-coincidence.

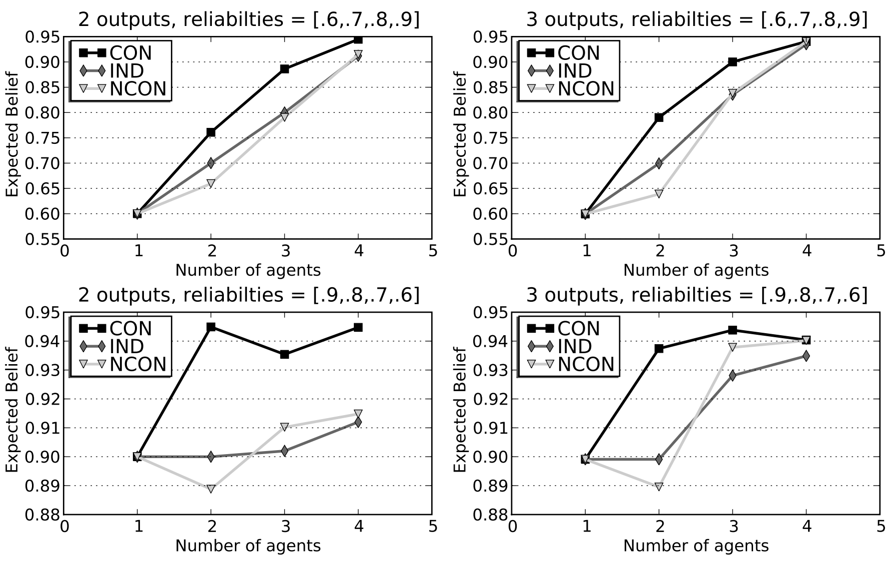

Fig. 6: The effects of assuming agents are either coincident (CON) or non-coincident (NCON) when they are independent (IND). When agents are independent, the expected belief becomes greater than that of the most reliable agent at a certain system configuration (number of agents). The expected belief is overestimated when assuming coincidence and underestimated when assuming non-coincidence.

We showed that the accuracy of a collaborative multiagent system varies according to the diversity of the errors of the agents. We also showed that the lower bound on the accuracy is given by the case where the agents errors are maximally correlated; and that the upper bound is not given by the conditional independence case, but by the case where agents errors are non-coincident (never fail together), usually 1.

We did not specify how to obtain the degree of dependence among agents, which can be a difficult task and varies with the application. In software fault tolerance, for example, testing as well as static programs analyses techniques have been used in trying to determine potential dependences among programs[4,20,10,21]. A challenge, as we mentioned earlier, is that the number of potential dependences grows exponentially with the number of agents. Also, multiagent systems are typically open and dynamic, which makes it difficult to measure dependences among the agents. A future work is the study of dynamic and adaptive strategies to estimate and model dependences among the agents, including alternative formalisms , e.g. Dempster-Shafer theory.

Finally, we have discussed our plans to use our work in the area of trust networks and online reputation systems. The dependences in reputation systems can be obtained from the rating similarities between pairs of users.

5.1 Modeling and using dependences in trust and reputation networks

We are interested in investigating the relation of our work with trust networks and online reputation systems. In this section we briefly discuss our insights on these tasks which we have planed as a future work.

The benefits of linking our work with trust networks are in both directions, i.e. using trust networks as an instrument in the dependences model, and using the dependences model as an instrument in trust networks. In the first case, a trust network could be used in assigning dependency probabilities among entities (e.g. the same contribution by a group of friends is likely to have some underlying dependences —similarities in the process of generating the contribution). In this case we can say that we use the trust network to navigate between the lower and upper (confidence) bounds generated by the model. In the second case, the dependences model could be used to assist existing trust propagation algorithms [6] in predicting trust between any two nodes in the network.

We plan to use our work on modeling dependences among entities when integrating contributions as a baseline to develop a Bayesian inference-based recommendation approach for online reputation systems[8] such as Flixster and Epinions. The dependences correspond to rating similarities between a pair of users, which can be measured by a set of conditional probabilities derived from their mutual rating history. Integrating contributions corresponds to calculating a rating score for an item based on the ratings (contributions) of the users. Suppose for example, in the electronics domain, a user is considering buying a certain product and wants to know what the general opinion is about it. Other user's rating of the product can be combined as an averaged score or can be better combined using likely dependences among the users that have rated the product to obtain a score that accounts for similarities among users. Further, the dependences of the user who is querying the product with those who have rated the item can also be factored into the result so that the obtained score is biased towards the user's preferences. In the first case we would be calculating an overall (not averaged) score that reflects the reputation of the product in the community, and in the second case we simply give more weight to some users' contributions.

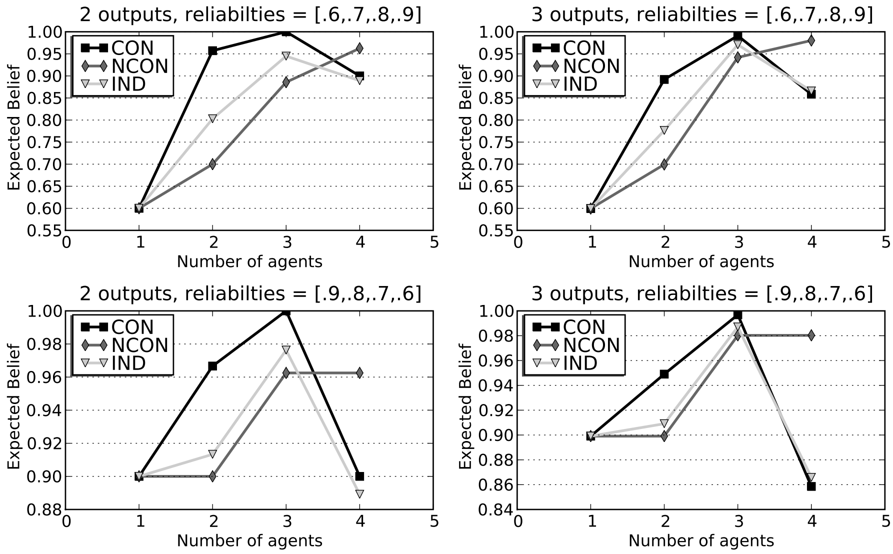

Fig. 7: The effects of assuming agents are either coincident (CON) or independent (IND) when they are non-coincident (NCON).

References

1. Becker, R., Corkill, D.: Determining confidence when integrating contributions from multiple agents. In: Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 07). pp. 449–456. The International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS) (May 2007)
2. Conte, R., Castelfranchi, C.: Simulating multi-agent interdependencies. a two-way approach to the micro-macro link. In: Social Science Microsimulation [Dagstuhl Seminar, May, 1995]. pp. 394–415. Springer-Verlag, London, UK (1996), `http://portal.acm.org/citation.cfm?id=647378.724484`
3. Corkill, D.D.: Collaborating software: Blackboard and multi-agent systems & the future. In: Proceedings of the International Lisp Conference. New York, New York (October 2003), `http://mas.cs.umass.edu/paper/265`
4. DAI, Y., XIE, M., POH, K., NG2, S.: A model for correlated failures in n-version programming. IIE Transactions 36(12), 1183–1192 (December 2004)
5. Grosz, B.J.: Aaai-94 presidential address: Collaborative systems. AI Magazine 17(2), 67–85 (1996)
6. Guha, R., Kumar, R., Raghavan, P., Tomkins, A.: Propagation of trust and distrust. In: Proceedings of the 13th international conference on World Wide Web. pp. 403–412. WWW '04, ACM, New York, NY, USA (2004), `http://doi.acm.org/10.1145/988672.988727`
7. Huynh, T.D., Jennings, N.R., Shadbolt, N.R.: An integrated trust and reputation model for open multi-agent systems. Autonomous Agents and Multi-Agent Systems 13(2), 119–154 (2006)

8. Jøsang, A., Ismail, R., Boyd, C.: A survey of trust and reputation systems for online service provision. Decision Support Systems 43, 618–644 (March 2007), http://portal.acm.org/citation.cfm?id=1225318.1225716
9. Latif-Shabgahi, G., Bass, J.M., Bennett, S.: A taxonomy for software voting algorithms used in safety-critical systems. IEEE Transactions on Reliability 53(3), 319–328 (2004)
10. Littlewood, B., Miller, D.R.: A conceptual model of the effect of diverse methodologies on coincident failures in multi-version software. In: Fehlertolerierende Rechensysteme / Fault-Tolerant Computing Systems, 3. Internationale GI/ITG/GMA-Fachtagung. pp. 263–272. Springer-Verlag, London, UK (1987)
11. Looker, N., Munro, M., Xu, J.: Increasing web service dependability through consensus voting. In: COMPSAC '05: Proceedings of the 29th Annual International Computer Software and Applications Conference (COMPSAC'05) Volume 2. pp. 66–69. IEEE Computer Society, Washington, DC, USA (2005)
12. Mitchell, H.: Multi-Sensor Data Fusion. Springer (2007)
13. Parker, J.R.: Multiple sensors, voting methods, and target value analysis. In: Kadar, I. (ed.) Proc. SPIE Vol. 3720, p. 330-335, Signal Processing, Sensor Fusion, and Target Recognition VIII, Ivan Kadar; Ed. Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference, vol. 3720, pp. 330–335 (Jul 1999)
14. Parker, J.: Voting methods for multiple autonomous agents. In: Intelligent Information Systems, 1995. ANZIIS-95. Proceedings of the Third Australian and New Zealand Conference on. pp. 128–133. IEEE Computer Society, Perth, WA, Australia (1995), citeseer.ist.psu.edu/339155.html
15. Sheth A., Henson C., S.S.: Semantic sensor web. Internet Computing, IEEE 12(4), 78–83 (2008)
16. Sichman, J.S.a., Conte, R.: Multi-agent dependence by dependence graphs. In: Proceedings of the first international joint conference on Autonomous agents and multiagent systems: part 1. pp. 483–490. AAMAS '02, ACM, New York, NY, USA (2002), http://doi.acm.org/10.1145/544741.544855
17. Sichman, J.S.: Depint: Dependence-based coalition formation in an open multi-agent scenario. J. Artificial Societies and Social Simulation 1(2) (1998)
18. Smiley, M.: Carry small, live large: Situational awareness. White paper, Intel Corporation: Intel Developer Forum (August 2008), http://electronics.wesrch.com/Paper/paper_details.php?id=EL11TZQLMTGDE&paper_type=pdf
19. Townend, P., Groth, P., Xu, J.: A provenance-aware weighted fault tolerance scheme for service-based applications. In: ISORC '05: Proceedings of the Eighth IEEE International Symposium on Object-Oriented Real-Time Distributed Computing. pp. 258–266. IEEE Computer Society, Washington, DC, USA (2005)
20. Vouk, M., McAllister, D., Eckhardt, D., Kim, K.: An empirical evaluation of consensus voting and consensus recovery block reliability in the presence of failure correlation. Journal of Computer and Software Engineering 4, 367–388 (1993)
21. Zavala, L., Huhns, M.N.: Analysis of coincident failing ensembles in multi-version systems. In: Proceedings of the 19th IEEE International Symposium on Software Reliability Engineering: Dependable Software Engineering Workshop. IEEE Press, Los Alamitos, CA, USA (November 2008)
22. Zavala Gutierrez, R.L., Huhns, M.N.: On building robust web service-based applications. In: Cavedon, L., Maamar, Z., Martin, D., Benatallah, B. (eds.) Extending Web Services Technologies: The Use of Multi-Agent Approaches, chap. 14, pp. 293–310. Kluwer Academic Publishing, New York (2004)