

Information Merging with Trust

Jonathan Ben-Naim and Emil Weydert*
Computer Science and Communications
University of Luxembourg

The goal of the present paper is to propose a new framework for information merging based on trust networks. This is a special instance of a more general paradigm called deep merging by Weydert. The idea is to see merging not just as an operation on – even prioritized – multi-sets of domain assertions, but as a function exploiting in addition more sophisticated background information about the sources producing these assertions. More specifically, we are going to consider background assumptions about the trust relationships between agents.

Let us proceed step by step. First, imagine several agents, each making assertions about a certain topic. In the real world, their claims may be incomplete as well as incorrect, so that conflicts between the statements of different agents may arise. A major problem is now to merge such conflicting views.

There have been a number of proposals on how to do this in a purely qualitative context, but without additional information about the sources, these procedures do not seem to inspire much confidence. An obvious extension would be to exploit an appropriate pre-order or valuation over the agents, e.g. reflecting their degrees of reliability. Corresponding approaches have for instance been discussed in [BDL⁺98, BDPW99, BDKP02, DDL06].

The question is how we arrive at these preferences or degrees, and whether they include all the information required for reasonable merging decisions. The basic assumption is that there is some higher authority attributing these priorities. But there is an alternative, bottom-up perspective, which may be more realistic. In fact, we can try to collect and combine the local reliability judgments of the agents themselves. That is, the idea is to determine the actual trust network and to use it to control the merging process. Thus, whereas traditional prioritized merging is based on what we may interpret as the global trust attributions of some oracle agent, we want to consider approaches aimed at local, actually existing trust attributions among the agents. These attitudes could be obtained directly from the agents or estimated from their past behaviour.

When we are talking about trust, we mean a kind of generic subjective reliability judgment of one agent about another agent, or himself. Here the relation between reliability and trust somehow mirrors – although from a less stringent perspective – the one between frequentist and personal probability. A trust

*jonathan.ben-naim,emil.veydert@uni.lu

network at a specific time point is defined by a trust function mapping pairs of agents to trust types characterizing directed trust relationships. More formally, we consider trust functions $\mathcal{T} : S^2 \rightarrow V$, where S is a collection of agents, and V is the domain of a trust value structure $\mathcal{V} = (V, \dots)$.

In practice, \mathcal{V} is often of the form $(D \times Tp, \prec)$, where D collects the degrees of trust, ordered by \prec , and Tp the topics focusing trust. For instance, $\mathcal{T}(a, b) = (d, \tau)$ tells us that a trusts b on topic τ with strength d . Of course, \mathcal{V} may also carry operations indicating how to combine trust values under certain conditions. In this context we emphasize that – even for a given topic – \prec is not necessarily a total order. In fact, there seems to be a difference between the absence of any trust judgment, and the explicit withholding of trust, i.e. the attribution of the minimal trust value. V should therefore include a distinguished and incomparable element \perp expressing full ignorance, or a missing link in the network representation. In what follows, to keep things simple, we will assume a single topic and restrict ourselves to trust functions from S^2 to $V = [0, \infty] \cup \{\perp\}$ (i.e. labelled networks).

The question is how we can exploit such a trust network to improve the fusion of information from different agents. A straightforward approach is to make use of the tools developed in the field of trust and reputation and to derive a ranking of the agents from the network, which brings us back to prioritized merging. Because different trust networks can easily lead to the same agent prioritization, maybe, we are losing valuable information here which could be relevant for the fusion process. For instance, the fact that there are no undirected paths linking two agents may be evidence that they constitute independent information sources. Consequently, we are going to propose another type of merging strategies which try to exploit the network structure more thoroughly.

One central idea is to proceed through local prioritized merging, deriving the priorities from the agent’s local trust valuation, and then replacing its original assertions by the result of the local prioritized fusion process. What does this mean? Suppose agent a trusts b (to some degree), although their statements conflict. This can easily happen because trust relationships are generic, i.e. allow for exceptions, or the agents have not yet taken into account the trust information. Intuitively, from the perspective of the merging agent, the information given by b would then – by default – take precedence over that given by a . So, a simple strategy would be to merge for each agent a the information from the trusted agents with its own one, giving higher priority to the former. This seems intuitively justified. It also tends to decrease the number of conflicts between the agents and therefore to make the fusion more informative. Let us illustrate the approach by the following example, where the merging operators are taken from [BDP95].

Let us assume that the information given by an agent is represented by a set of propositional formulas. The plain merging of several sets of formulas is obtained by first taking their union and then removing every element involved in some conflict. More formally, the merging operator \star is defined as fol-

lows: $\star(\Gamma_1, \dots, \Gamma_n)$ is the set of those elements of $\Gamma_1 \cup \dots \cup \Gamma_n$ that belong to no minimal inconsistent subset of $\Gamma_1 \cup \dots \cup \Gamma_n$. Note that [BDP95] write $Free(\Gamma_1 \cup \dots \cup \Gamma_n)$.

When priorities are considered, we just need to treat the case of two sets. The way we merge a set Γ with low priority and a set Δ with high priority is as follows: we fully take Δ and then add those elements of Γ which are involved in no conflicts. More precisely, we choose the operator \star_p given by $\star_p(\Gamma, \Delta) := \Delta \cup \star(\Gamma, \Delta)$. Clearly, $\star_p(\Gamma, \Delta)$ is coherent.

At present, consider the following simple situation, where three agents A_1 , A_2 , and A_3 provide respectively the pieces of information: $\Gamma_1 := \{\alpha, \beta\}$, $\Gamma_2 := \{\neg\alpha, \beta\}$, and $\Gamma_3 := \{\neg\alpha, \neg\beta\}$. In addition, A_1 trust A_2 and A_3 . Let us exploit the trust network according to our previous strategy, i.e. we replace Γ_1 by $\Gamma'_1 = \star_p(\Gamma_1, \star(\Gamma_2, \Gamma_3)) = \star_p(\Gamma_1, \{\neg\alpha\}) = \{\neg\alpha\} \cup \star(\Gamma_1, \{\neg\alpha\}) = \{\neg\alpha\} \cup \{\beta\} = \{\neg\alpha, \beta\}$. It remains to merge the revised views to get the final result: $\star(\Gamma'_1, \Gamma_2, \Gamma_3) = \{\neg\alpha\}$, which is indeed more informative (in a justified way) than the merging of the original views: $\star(\Gamma_1, \Gamma_2, \Gamma_3) = \emptyset$.

The previous method is however just a simple instance of a much more general approach. Let F be a trust evaluation mechanism which maps a global prior trust valuation $\nu_0 : S \rightarrow V$ and a trust function $\mathcal{T} : S^2 \rightarrow V$ to a global posterior trust valuation $F(\nu_0, \mathcal{T}) : S \rightarrow V$. Here ν_0 may be interpreted as the local default trust attribution of the merging agent $m \notin S$. It is meant to specify the global trust value for those a which have no incoming arcs in the trust network given by \mathcal{T} (i.e. where $\mathcal{T}(x, a) = \perp$ for all $x \in S$). Furthermore, let $\Phi : S \rightarrow 2_{fin}^L$ be a source base, associating with each agent a finite set of statements in a language L , and \sqcup be a prioritized merging strategy. Then the full trust-based merging procedure is a function, which – parametrized by a trust evaluation mechanism F and a prioritized merging strategy \sqcup – maps each triple consisting of a source base Φ , a trust function \mathcal{T} , and a prior global trust valuation ν_0 , to a finite set of statements $\sqcup[F, \sqcup](\nu_0, \mathcal{T}, \Phi)$.

In the simple example discussed above, $\sqcup[F, \sqcup](\nu_0, \mathcal{T}, \Phi)$ is $\star(\Phi')$, where Φ' is the revised version of Φ obtained by \star_p and \mathcal{T} , \sqcup is basically the composition of \star and \star_p , F is the trivial evaluation mechanism, and ν_0 is irrelevant.

To summarize, we propose to govern merging by a trust network. One way to realize this is to replace the information given by an agent a by the result of some processing of the information given by those agents which a trusts. But there are many other ways to take advantage of trust networks, which we plan to explore in future work.

References

- [BDKP02] S. Benferhat, D. Dubois, S. Kaci, and H. Prade. Possibilistic merging and distance-based fusion of propositional information. *Annals of Mathematics and Artificial Intelligence*, 34(1-3):217252, 2002.

- [BDL⁺98] S. Benferhat, D. Dubois, J. Lang, H. Prade, A. Saffiotti, and P. Smets. A general approach for inconsistency handling and merging information in prioritized knowledge bases. In *6th Int. Conf. Principles of Knowledge Representation and Reasoning*, pages 466–477. Morgan Kaufmann, 1998.
- [BDP95] S. Benferhat, D. Dubois, and H. Prade. How to infer from inconsistent beliefs without revising? In *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI'95)*, pages 1449–1455, 1995.
- [BDPW99] S. Benferhat, D. Dubois, H. Prade, and M. Williams. A practical approach to fusing and revising prioritized belief bases. In *EPIA '99, LNAI 1695*, page 222236. Springer Verlag, 1999.
- [DDL06] J. P. Delgrande, D. Dubois, and J. Lang. Iterated Revision as Prioritized Merging. In *10th International Conference on Principles of Knowledge Representation and Reasoning (KR'06)*, pages 210–220, 2006.