

Sharing Information in Teams: Giving Up Privacy or Compromising on Team Performance?

Maaïke Harbers¹, Reyhan Aydođan¹, Catholijn M. Jonker¹, and Mark A. Neerincx^{1,2}

¹Interactive Intelligence Group, Delft University of Technology, The Netherlands

²TNO Human Factors, Soesterberg, The Netherlands

{M.Harbers,R.Aydogan,C.M.Jonker,M.A.Neerincx}@tudelft.nl

ABSTRACT

Human teamwork can be supported by agent technology by providing each human team member with an agent that monitors, supports and advises the human. The agent can, for example, monitor the human's workload, and share that information with (agents of) other team members so that work can be distributed effectively. However, though sharing information can lead to a higher team performance, it may violate the individual team members' privacy. This raises the question what type of and how often information should be shared between team members. This paper addresses this question by studying the trade-off between privacy loss and team performance in the train traffic control domain. We provide a conceptual domain analysis, introduce a formal model of train traffic control teams and their dynamics, and describe an agent-based simulation experiment that investigates the effects of sharing different types and amounts of information on privacy loss and team performance. The results give insight in the extent to which different information types cause privacy loss and contribute to team performance. This work enables the design of privacy-sensitive support agents for teamwork.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Multiagent Systems*

Keywords

Agent-Based Modeling, Information Sharing, Privacy, Team Performance, Teamwork.

1. INTRODUCTION

As teams have become essential in many organizations, the development of high performance teams has been an important research goal [24, 5, 15]. One of the problems faced by human teams is that workload is not always distributed evenly over the team members, in particular in complex and dynamic task environments [23]. Agent technology can offer a solution to this problem by providing each human team member with a personal assistant agent [27]. The

agent can monitor the human's workload (see e.g. [18, 3]), and share this information with (the assistant agents of) other team members. Team members with a low or average workload can then decide to take over tasks of team members with high levels of workload, and thus, workload is distributed more effectively. Multiple studies showed that sharing information in teams has a positive effect on team performance [17].

A potential problem with sharing workload information, however, is that it may violate the team members' privacy [22, 21]. Workload information contains personal information including someone's interaction with the system or physiological measures such as heart rate and galvanic skin response [19]. Sharing this information with other team members may lead to privacy loss. For example, people with a mediate workload may dislike admitting that they have difficulty coping, and someone's stress level may be influenced, besides work, by personal factors.

The above findings imply a trade-off with regard to sharing personal information in teams. On the one hand, knowledge of personal information can have a positive effect on team performance, but on the other hand, it may cause a loss in privacy for the team members. Therefore, to develop personal assistant agents that enhance team performance, but yet, are sensitive to their user's privacy, insight is needed in the effects of sharing different types and amounts of information on privacy loss and team performance.

Earlier work has also stressed that developers need to consider possible privacy implications when developing agent-based applications [2]. In a recent overview on privacy and multi-agent systems, however, Such *et al.* [26] claim that, to their knowledge, privacy is seldom considered in the multi-agent systems research field. A number of researchers have proposed algorithms to protect the privacy of users that interact with other entities by representing the situation as a distributed constraint optimization problem (DCSP), or a distributed constraint satisfaction problem (DisCSP), e.g. [25, 10]. Furthermore, measures to quantify privacy loss have been proposed, e.g. [14]. The key issue in these approaches is to preserve the anonymity of the user as much as possible. In the context of human teamwork, however, preserving anonymity is not the main aim, and team members may want to sacrifice small amounts of privacy if that yields large benefits for team performance. To find appropriate solutions for such situations, the specifics of the team at hand need to be analyzed and taken into account [21].

In this paper we study the trade-off between privacy loss and team performance by analyzing a real world case in the

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

Copyright © 2014, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

domain of train traffic control. The analyzed train traffic control team does currently not share workload information, and has to deal with uneven workload distributions. Domain experts in this field acknowledge this problem, but are also concerned about the negative effects of sharing workload information on the team members' privacy. We make use of agent-based modeling to study this trade-off [7]. We developed an agent-based model that represents the train traffic control team and its dynamics, and performed simulation experiments to determine the effects of sharing different types and amounts of information on privacy loss and team performance. This enables the development of personal assistant agents that enhance team performance while minimizing privacy loss. Furthermore, it may support the specification of working agreements for situated team operations on sharing of sensitive information, i.e., an adaptive information sharing mechanism [4].

The rest of this paper is organized as follows. Section 2 provides a conceptual analysis of teamwork in the train traffic control domain. Section 3 introduces a formal model of train traffic control teams and their dynamics. Section 4 describes the simulation experiments in which we investigated the effects of sharing personal information on privacy loss and team performance. Section 5 provides the discussion and conclusion.

2. DOMAIN ANALYSIS

For the research presented in this paper, we studied train traffic control teams and their dynamics in the Dutch railway system. In this section, we introduce the domain, and describe the development of a personal assistant agent for that domain.

2.1 Train traffic control teams

The Dutch railway network is used by multiple passenger and cargo transporters, and ProRail is the organization that is responsible for controlling this train traffic. ProRail has thirteen regional control centers and one national control center from where it controls train traffic in the Netherlands. Each regional control center is occupied by a group of people with different roles. In this research, we focus on the team of operators (*treindienstleiders*) that control the signals and switches on the rails of a limited area, e.g. a station, and communicate with train personal in case of disruptions. In the remainder of this paper we will use the term operator to refer to this role.

Under normal circumstances, train traffic is automatically regulated according to fixed schedules that describe the train traffic flow. The operators mainly monitor the situation, and their workload is low. In case of a disruption, the operators actively have to regulate the train traffic by manually entering changes to the controls and switches on the rails, and informing train drivers about the changes. In case of a simple disruption, e.g. a train that is delayed, there exist procedures that prescribe how the problems can be resolved. For the operators this is routine work, and it usually yields an intermediate level of workload. In case of a larger or multiple simultaneous disruptions, there usually is no single procedure to solve the problem, and adhoc solutions have to be created. This falls out of the operators' routine and can cause high levels of workload.

Each operator is responsible for a particular section of the railway network, and the division of work over opera-

tors is purely based on who is responsible for which railway section. The assignment of railway sections to operators is determined beforehand and does normally not change during a shift. It is possible for operators to take over tasks of others, but in practice, this does not happen a lot because operators tend to try to solve disruptions themselves as much as possible. Therefore, when there is a disruption that only affects certain areas of the railway network, it may happen that some of the operators in a control center experience a rather high and others a relatively low workload. This disparity is undesired when the workload of some of the operators becomes so high that it leads to a decrease in their performance.

2.2 Introducing a personal assistant agent

In a research project in collaboration with ProRail, we are currently developing a personal assistant agent that can support operators in (re)distributing workload effectively over the team members. The aim is to minimize the individual operator's decrease in performance due to a high workload, and thus maximize performance of the team.

To redistribute workload effectively, it is necessary to have insight in the operators' workload. In the context of our project, we developed a component of the personal assistant agent that can predict an operator's performance by assessing his cognitive and affective load. This monitor component is based on Neerincx's model about cognitive task load [18, 20], which describes the effects of task allocations on performance of operators working in dynamic, critical and high-demand task environments. To measure affective load, the monitor component implements an adaptation of Mehrabian's pleasure-arousal-dominance model [16, 19]. In this paper it is not possible to provide a detailed description of the working of the monitor component due to space limitations, but it is not necessary for following the ideas presented in this paper.

In a demonstration session at ProRail, the monitor component of the personal assistant agent was demonstrated to a group of four domain experts employed by ProRail. The experts acknowledged the value of a personal assistant agent with regard to the distribution of workload among operators. The participants, however, also raised the concern that operators may be unwilling to share information about their cognitive and affective state with colleagues because it violates their privacy. In the demonstration session, the domain experts were not able to provide exact indications about the extent to which the personal assistant agent would benefit team performance and harm privacy.

To further develop the personal assistant agent, we need to develop a mechanism and an interface for sharing the information collected by the monitor component to the operators in the team. The information sharing mechanism determines what types of information are shared, and how often and under what conditions it is shared. The interface design determines how the information is presented to the operators. A possibility is to provide operators with a display that continuously shows the cognitive and/or affective load of their team members. This way of providing information resembles information provision through awareness panels [1], a common solution in operational environments. An alternative to continuous information provision is to let the assistant agents send updates to the operators every now and then. An advantage of this solution is that personal

information is shared less often, and may thus cause less privacy loss.

To conclude, the results of the demonstration session at ProRail are in line with other findings in literature on technology, information sharing and privacy [22, 21]. Thus, in order to design an information sharing mechanism and interface for the personal assistant agent for train traffic operators, more insight is needed in the effects of different types and amounts of information sharing on privacy loss and team performance.

3. FORMAL MODEL OF THE TRAIN TRAFFIC CONTROL TEAM

In this section we provide a formal model of the train traffic control team and its dynamics described in the previous section. As mentioned earlier, the railway network is divided into sections, and each operator is responsible for the train traffic on a particular section. Therefore, tasks are initially always assigned the operator responsible for that railway section. We formalize this as follows.

- O is a team of operators, where $o \in O$ is an operator that belongs to team O
- K is a set of tasks, where $k = (t, d, l, s) \in K$ is a task generated at time t , with duration d , a level of information processing l required to perform the task (where $l \in [low, medium, high]$), and a severity level s of the disruption due to which the task was generated (where $s \in [low, medium, high]$)
- $K_o(t, f)$ is the set of tasks that an operator o performed in time frame f at a certain time t , where f lasts from $t-f$ to t
- A is a set of task assignments, where $a_o \in A$ ($a_o \in 2^K$) denotes the set of tasks assigned to operator o
- M is a set of messages, where $m = (s, r, i, t) \in M$ is a message m sent by sender s , to receiver r , with information type i (where $i \in [cognitive\ load, affective\ load, total\ load, offer, accept]$), at time t

In the model, we assume that an operator can only perform one task at a time and that tasks are always achieved when executed. The duration of task performance is fixed, but operators can complete tasks with varying degrees of quality. Quality of task performance is determined by the operator's cognitive and affective load during task performance, i.e. the higher the loads, the lower the quality of task performance. Once a task has been achieved it is removed from the respective operator's set of task assignments. In our model, the levels low, medium and high for information processing and severity correspond to the values 1, 2 and 3, respectively.

3.1 Cognitive load

The notion of cognitive load in our formal model is based on Neerinx's model of cognitive load [18]. According to this model, an operator's cognitive load (CL) in a specific time frame f is determined by three factors: level of information processing, time occupied, and task set switching. We will explain each of these factors, and show how we adapted the definitions from Colin *et al.* [3] for this study.

The first factor, *level of information processing* (LIP), gives an estimation of the complexity of the mental activity operator o performed in time frame f , and is calculated by the following formula.

$$LIP(K_o(t, f)) = \frac{\sum_{i=1}^n l_i d_i}{\max(l) f}$$

The n is the number of tasks in set K_o . This formula differs from Colin *et al.* in the sense that they divide by time frame f , whereas we also divide by $\max(l)$, the maximum value of l (3 in our model). Dividing by $\max(l)$ normalizes the outcome of the formula to a value between 0 and 1. Furthermore, we added a time t to this (and the following two) formula to indicate the time at which LIP is determined, i.e. the ending of time frame f .

The second factor, *time occupied* (TO), indicates the proportion of time during which an operator is performing tasks.

$$TO(K_o(t, f)) = \frac{\sum_{i=1}^n d_i}{f}$$

Again, n is the total number of tasks. We took this formula directly taken from Colin *et al.* since it already returns a value between 0 and 1.

The third factor, *task set switching* (TSS), considers the workload associated to switching attention from one task to another.

$$TSS(K_o(t, f)) = \frac{\sum_{i=1}^{n-1} |K(t, f)|}{f}$$

The n is the number of tasks. The TSS formula is a simplification of Colin *et al.*'s proposal. The original formula involves the concept of an information domain, through which it is taken into account that switching between similar tasks is considered less demanding than switching between different tasks. In our formula, in contrast, any task switch has the same effect on TSS. Another adaptation we made to the TSS formula is that we divide by f to normalize the outcome of the formula.

With these three factors, LIP, TO and TSS, someone's *cognitive load* (CL) can be calculated. For that, the three factors are plotted along three axes in a 3-dimensional space. Then, the cognitive load of an operator o in time frame f ending at time t is calculated as follows.

$$CL(K_o(t, f)) = \frac{d_{origin} - d_{diagonal}}{\sqrt{3}}$$

In this formula, d_{origin} is the Euclidean distance of the point (LIP, TO, TSS) to the origin, and $d_{diagonal}$ is the shortest distance from the point (LIP, TO, TSS) to the line where LIP=TO=TSS. Cognitive load is always a value between 0 and 1. We refer to the paper of Colin *et al.* for more details about these metrics.

3.2 Affective load

Affective load is usually assessed by physiological measures such as heart rate, skin conductance, and facial expressions [19]. The current model, however, does not include the physiological states of agents, and we therefore use an alternative measure for assessing an operator's emotional response. In our model, affective load is determined by two factors: expected cognitive load and severity level of the current task.

The first factor, *expected cognitive load* (ECL), is deter-

mined by the set of tasks that are currently assigned to an operator. In other words, ECL depends on the operator's upcoming work, and the assumption is that more upcoming work yields more stress.

$$\begin{aligned} ECL(t) &= CL(K_o(t, f_{ECL})), \text{ if } \sum_{i=1}^n d_i < \vartheta_{ECL} \\ ECL(t) &= 1, \text{ otherwise} \end{aligned}$$

The first part of the formula (n is the number of tasks) states that the expected cognitive load of operator o at time t is the cognitive load of the set of all tasks K_o that are currently assigned to operator o . Cognitive load is calculated for a time frame between t - f and the current time t . Expected cognitive load, in contrast, is calculated over all tasks assigned to the operator, denoted by f_{ECL} . If the total duration of tasks, however, is larger than the threshold ϑ_{ECL} , expected cognitive load is 1. This models the phenomenon that a lot of upcoming work increases stress, but that after a certain point, more upcoming tasks does no longer increase stress. Note that for the calculation of expected cognitive load, the cognitive load of the task the operator is currently working on is included.

The second factor, *severity level of current task* (SLT), captures someone's affective response to disruptions and is determined by the nature of the disruption due to which the operator has to perform that task. Tasks caused by minor disruptions, such as a delayed train, have a low severity level, and tasks caused by severe disruptions, e.g. a collision, yield a high severity level. SLT is determined by dividing the severity level of a task with the highest severity level possible (3 in our model).

$$SLT(t) = \frac{s_{currenttask}}{max(s)}$$

This formula implements the idea that a task, e.g. changing a switch, yields more affective load when it is related to a severe disruption than when it is related to a minor disruption. Yet, all tasks are considered as routine tasks of the operator.

With the above two formulas, affective load can be calculated as follows.

$$AL(t) = \alpha ECL(t) + (1 - \alpha) SLT(t)$$

The α determines the weight of both factors contributing to affective load. In the simulation experiments described in this paper, we used $\alpha = 0.5$.

3.3 Sharing information and task distribution

Sharing information about workload helps operators to predict each other's performance, and based on that, to decide when to take over other's tasks. In our model, information is shared by sending messages of type $m=(s,r,i,t)$ such as defined above. There are three types of information i that contain information about workload: cognitive load, affective load, and total load. *Total load* (TL) is a combination of cognitive and affective load and can be determined as follows.

$$TL(t, f) = \beta CL(t, f) + (1 - \beta) AL(t)$$

The value of β determines the respective weights of cognitive and affective load on total load. For the simulation exper-

iments, we used $\beta = 0.5$ to model an equal impact of both factors.

To mimic the costs of processing a message containing workload information, the probability that an operator processes an incoming message depends on his cognitive and affective load, where higher load values yield a lower probability to process a message. The following linear function indicates the probability that a message with information about workload will be processed.

$$P(m \text{ processed}) = 1 - TL(t, f)$$

If an operator processes a message that contains information about cognitive, affective or total load, it compares the load value to his own respective load. If the other operator's load is higher than the threshold $\vartheta_{sendoffer}$, he will offer that operator to take over one of his tasks by sending a message with information type $i=offer$. This threshold is introduced to avoid that tasks are reassigned in (almost) every time step. Messages of information type *offer* are always processed. If the operator to which the offer is sent has not accepted any offers in that time step, it will accept the offer by sending a message with information type $i=accept$. When an offer to take over a task is accepted, the first task in the accepting operator's task queue is reassigned to the operator that sent the offer. A confirmation before reassigning the task is needed to avoid that two operators take over the same task of a third operator.

3.4 Performance and privacy measures

A scenario consists of a set of tasks K and task assignments A . Given such a scenario, team performance is determined by two factors: *speed* of completion of all tasks in the scenario and *quality* of task performance. Speed is determined by the time that is needed to complete all tasks in the scenario.

To determine quality of task performance of the team, we first need to determine the quality of task performance of the individual team members. *Quality of individual task performance* (QITP) of an operator o at a certain time is determined as follows.

$$QITP(o, t, f) = 1 - (\gamma CL(t, f) + (1 - \gamma) AL(t))$$

The formula implements the idea that high cognitive and affective loads have a negative impact on performance. This formula is a simplification of reality because, though high cognitive load has a negative impact on performance, extremely low values of cognitive load can cause boredom and also have a negative effect on performance. The value γ determines the respective effects of cognitive and affective load on quality of task performance. In the simulation experiments we took $\gamma = 0.5$.

Quality of team task performance (QTTP) is an aggregation of the quality of task performances of the individual team members. There are multiple ways to aggregate values, for instance, by taking the average, minimum, maximum, first, or last of the values. In the domain of train traffic control, the performances of all operators matter for team performance, which makes the minimum and average aggregate functions suitable for this domain. For the simulation experiments described in this paper we use the average, which can be calculated as follows.

$$QTTP(O, t, f) = \frac{\sum_{i=1}^{|O|} QTTP_i(t, f)}{|O|}$$

In this formula, the set O is used to denote all operators in the team.

To assess the quality of team performance of a whole scenario, i.e. a finite set of tasks, we introduce a formula to determine the *quality of team task performance of a scenario* (QTTPS). Quality of team performance of a scenario is determined by taking the average quality of team task performance of all time steps in the scenario, as follows.

$$QTTPS(O) = \frac{\sum_{i=1}^{t_{end}} QTTP(O, i, f)}{t_{end}}$$

In the formula, t_{end} denotes the end time of the scenario.

Sharing information about cognitive, affective or total load causes privacy loss. The amount of privacy loss depends on the kind of information that is shared. An operator’s cognitive load is determined by his interaction with the system, and his affective load by physiological measures. The domain experts in our workshop considered the latter more obtrusive than the former. Therefore, in our model, sharing information about cognitive (system interaction), affective (physiological measures) and total (system interaction and physiological measures) load cause a low, medium and high privacy loss, respectively. Messages of the type ‘offer’ and ‘accept’ do not involve any loss of privacy. To express the amount of privacy loss, the information types cognitive load, affective load, total load, offer and confirm correspond to the values 1, 2, 3, 0 and 0, respectively. The privacy loss of an operator o can now be calculated over the set of messages for which it holds that $s = o$ (the operator is the sender) as follows.

$$PL(M) = \sum_{j=1}^{|M|} i$$

Note that for privacy it does not matter whether a message is processed or not. The sender of the message, for whom privacy loss is calculated, only knows that the information may have been processed.

4. SIMULATION EXPERIMENTS

This section describes the simulation experiments we performed to get insight in the effects of sharing personal information on privacy loss and team performance. We adopted the formal model described in Section 3 and implemented our simulation tool in Java.

To get reliable insights, we developed an automatic scenario generator that allows us to generate a diverse set of scenarios in a systematical way. As mentioned before, a scenario consists of a set of tasks and their assignments to the operators. The generator firstly generates a given number of tasks (60 in the experiments), and randomly generates the values of the task parameters within the following ranges: generation time (1-100), duration (2-8), and information processing level (1-3). Note that though the maximum generation time of a task is at time step 100, completing the scenario may take longer than 100 time steps. We used teams of 3 operators, where each operator works on a different disruption with a randomly assigned severity level (1-3). Tasks are randomly assigned to the operators, and get the severity level of the disruption their operator is working

on. In the experiments, we studied 16 different conditions, and for each condition, we considered 1000 randomly generated scenarios. We used the same set of scenarios for all conditions. The parameters used in the experiments were chosen after testing a number of values. In total, we thus performed 16000 simulation runs.

4.1 Study 1: Information type

The first study investigates the effects of sharing different information types on privacy loss and performance. We distinguish three types of information that can be shared in a team: cognitive, affective and total load. That results in the following four experimental conditions: 1) sharing no personal information, 2) sharing information about cognitive load, 3) sharing information about affective load, and 4) sharing information about total load (a combination of cognitive and affective load).

In the experiment, all operators share the same type of information (or no information) at each time step, and their team members may check the information, depending on their own load. Sharing information at every time step imitates information sharing through an awareness panel as described in Section 2. For each simulation run, we logged privacy loss, scenario completion time, quality of team task performance, quality of worst performing team member, number of messages sent containing personal load information, and number of tasks that were reassigned to another operator. Table 2 displays the averages and standard deviations.

The results show that, as expected, privacy is not violated when no information is shared (0), least when only information about cognitive load is shared (691), more when information about affective load is shared (1294), and most when information about total load is shared (1965). With regard to performance, the results show that for completion time sharing information leads to a faster completion of the scenario, and thus has a positive effect on team performance. Furthermore, sharing information leads to a higher quality of team performance and higher quality of lowest individual performance. Affective and total load have greater impact on the outcome of these three measures than cognitive load.

None of the information types scores well on both minimizing privacy loss and maximizing team performance. According to these results though, sharing information about affective load is more favorable than total load because there is considerable less privacy loss, while performance is equal or even slightly better. An interesting difference between affective and total load is that in the affective load condition almost twice as many tasks are reassigned to another operator. That means that in the simulation, sharing information about affective load most likely leads to a task reassignment.

The three performance measures we used (completion time, quality of team task performance, and minimum individual performance) yielded similar results. A potential disadvantage of the completion time measure is that it may depend too strongly on the generation time of the last task in the scenario, in particular when the scenario involves few tasks. The results of total team performance are more diverse than the results of minimum individual performance. Therefore, in our next experiments we will use quality of team task performance as a measure for team performance.

4.2 Study 2: Information amount

In the previous study, information sharing yielded rather

Information type	Privacy loss		Completion time		Total quality of performance		Min quality of performance		Nr. of sent messages		Nr. of task reassignments	
	mean	σ	mean	σ	mean	σ	mean	σ	mean	σ	mean	σ
No information	0	0	131	11.7	1.39	0.14	0.55	0.046	0	0	0	0
Cognitive load	691	48	125	8.0	1.47	0.11	0.56	0.046	691	48	15	9
Affective load	1294	68	118	5.6	1.57	0.10	0.59	0.038	647	34	55	11
Total load	1965	108	119	6.0	1.55	0.09	0.59	0.037	655	36	28	11

Table 1: Effects of sharing different types of personal information (n=1000).

large amounts of privacy loss. A way to reduce privacy loss is to share load information less frequently. In the real world, this could correspond for example to a pop-up message that appears every now and then. To investigate the effects of less frequent sharing of information on privacy loss and team performance, we will compare the following three conditions to each other: sharing information 1) in every time step, 2) every 5th time step, and 3) every 10th time step. Figure 1 and Figure 2 show the effects of different amounts of information sharing for all three information types (cognitive, affective, and total load) on privacy loss and (quality of) team performance, respectively.

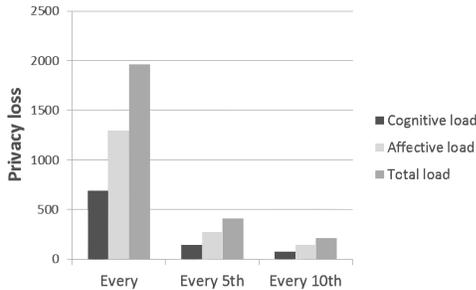


Figure 1: Effects of sharing different amounts of information on privacy loss (n=1000).

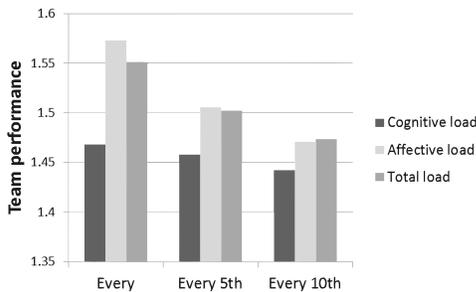


Figure 2: Effects of sharing different amounts of information on team performance (n=1000).

Figure 1 shows that, as expected, privacy loss considerably drops when personal information is only shared at every 5th time step, and even more when it is only shared at every 10th time step. Figure 2 shows that sharing information about affective load leads to the highest quality of performance. When information is shared less frequent, though, quality of team performance drops more when sharing affective and total load than when sharing cognitive load. Sharing cog-

nitive load information thus becomes more attractive when information is shared relatively infrequent. Combining the results of the two figures, the trade-off between privacy loss and team performance is clearly visible, higher team performance co-occurs with higher privacy loss and vice versa.

4.3 Study 3: Sharing policy

The information sharing policies in the previous study were rather simple by only considering time. A more intelligent way to limit the amount of shared information is to only share under certain conditions, e.g. when the sharing operator has a high cognitive, affective or total load, and would benefit from the reassignment of one or more of his tasks. For this third study, we introduce an information sharing policy in which an operator only shares personal information when his cognitive, affective or total load is above a certain threshold. In the experiment, we compare the following conditions to each other: sharing information 1) every time step, 2) every 5th time step, and 3) when above the threshold of 0.5. Table 2 shows the results.

With regard to privacy loss, the threshold sharing policy scores fall between those of sharing information in every time step and every 5th time step. Regarding team performance, however, the threshold policy scores are equal to or higher than those in all other conditions. In particular when information about affective and total load is shared, the threshold policy outperforms the other policies.

In all studies together, we collected results for 4 information types and 4 sharing policies. These results are combined in Figure 3, displaying a scatter plot of privacy loss and team performance. Note that the information type ‘no information’ has equal results for all sharing policies.

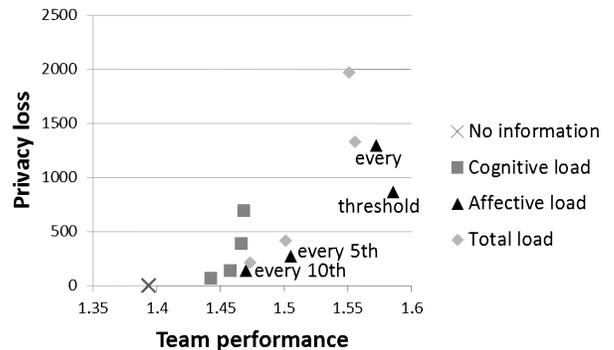


Figure 3: Team performance and privacy loss of all 16 conditions (4 information types x 4 sharing policies) (n=1000).

In Figure 3 it is clearly visible that sharing information

Information type	Every time step				Every 5 th time step				Threshold			
	Privacy loss		Performance		Privacy loss		Performance		Privacy loss		Performance	
	<i>mean</i>	σ	<i>mean</i>	σ	<i>mean</i>	σ	<i>mean</i>	σ	<i>mean</i>	σ	<i>mean</i>	σ
Cognitive load	691	48	1.47	0.11	141	10.1	1.46	0.11	386	43	1.47	0.11
Affective load	1294	68	1.57	0.10	273	16.5	1.51	0.11	866	108	1.59	0.10
Total load	1966	108	1.55	0.09	411	25.2	1.50	0.10	1324	144	1.56	0.09

Table 2: Effects of sharing information through different sharing policies (n=1000).

about affective load (black, labeled triangles) yields relatively high performances with relatively little loss of privacy. The most favorable information sharing policy (see labels) depends on preferences with regard to giving up privacy or compromising on team performance, but the ‘every 5th’ and ‘threshold’ sharing policies are both good candidates.

4.4 Discussion of the results

Sharing information about affective load was the most favorable option with regard to optimizing performance and minimizing privacy loss in most conditions. Interestingly, affective load considers an operator’s current and future tasks (Section 3.2), and cognitive load considers the tasks the operator just performed (Section 3.1) in our model. This may be the reason that affective load better predicts the operators’ performances, and thus better indicates when tasks should be reassigned. The reason that cognitive load yields lower performance scores might be due to the scenarios we used in the simulations. In real scenarios, an operator’s past tasks may be a better predictor of his performance than the randomly generated scenarios used in the experiments.

In our simulations, sharing information about total load causes most privacy loss. These results were expected since total load has a privacy loss value of 3 in our model (Section 3.4), whereas cognitive and affective load have values 1 and 2, respectively. Total load received this high value because it is based on both cognitive and affective load information. One could however argue that, because two information types are combined into one new value, personal information is lost, and that total load should therefore have a lower privacy loss value. If so, sharing total load information would become a more favorable option.

To conclude, as is always the case in agent-based simulations, the results of the experiments heavily depend on model design and parameter settings. Though we carefully made these choices based on literature and expert knowledge, agent-based models always remain an approximation of the real world. However, as we will explicate in the next section, we do not view agent-based models and simulations as an end product, but as a part of the design process. We believe that even an imperfect model can yield valuable insights for the development of our personal assistant agent.

5. DISCUSSION & CONCLUSION

In this paper we studied the trade-off between privacy loss and performance in teams, by analyzing a real world case in the train traffic control domain. We provided a domain analysis, introduced a formal model of a train traffic control teams and its dynamics, and described a series of agent-based simulation experiments investigating the effects of sharing different types of information according to different policies on privacy loss and team performance. The simulation results show that with regard to minimizing pri-

vacuity loss and maximizing team performance, the preferred information type is ‘affective load’ and the preferred sharing policy is ‘threshold’ or ‘every 5th’ time step.

We started the paper by stating that the development of high performance teams is an important research goal. In the field of multi-agent systems, research on optimizing team performance mostly focuses on the direct effects of introduced teamwork mechanisms on the time and costs spent on achieving the team’s goals, and the quality of the team’s output [13]. For teams that only contain software agents (e.g. [28, 11]) this may be an effective way to develop high performance teams. For human-agent teams (e.g. [27, 12] or human teams (e.g. [15, 24]), however, the performance metrics of time, cost and quality may not always be adequate to grasp all team dynamics [24].

In contrast to most software agents, humans esteem values such as privacy, autonomy, and trust [8]. Violating these human values is not only undesirable in itself, it can also have negative effects on team performance in the long run. For example, violating a team member’s privacy over a long period of time may lead to a decrease in motivation, which in turn can lead to lower team performance [5]. It is thus important to account for human values when creating high performance teams that contain humans.

The importance of accounting for human values in the design of technology has received an increasing amount of attention over the last years [9, 8]. Though a range of tools and methods for analyzing human values have been proposed, there is less work on how to incorporate these values in an actual design [6]. Thus, when designing agents for optimizing human team performance, literature provides little guidance on how to take values into account in this process.

The work presented in this paper may be a first step towards a methodology to account for human values in the development of intelligent agents that interact with humans. In this paper, we focused on questions around the value of privacy in order to develop a privacy-sensitive personal assistant agent for teamwork. The process of developing an agent-based model and performing simulation experiments helped us to obtain insight into the effects of different options (which and when to share information) on privacy loss. This will help us to make better design choices for the personal assistant agent. Furthermore, the simulation results may make users and domain experts more aware of their preferences and priorities with regard to the sharing of personal information in teams, which may help them to formulate clear requirements for the assistant agent.

Besides further developing a methodology for developing value-sensitive agents, we suggest three other directions for future work. First, the agent-based model can be extended to obtain more insight in the trade-off between privacy loss and performance, e.g. by modeling the effects of social relationships, organizational roles or location. Second, the re-

sults of the simulations can be used to propose a mechanism and an interface for sharing personal information in teams. In other words, the second option involves the development of (parts of) the personal assistant agent. Third, once a prototype of the personal assistant agent has been developed, user experiments need to be performed that test the value of privacy-sensitive assistant agents for teamwork.

Acknowledgments

This research was conducted within the RAILROAD project and is supported by ProRail and the Netherlands organization for scientific research (NWO) (under grant 438-12-306). Special thanks to the ProRail experts in the demonstration session for their feedback on the monitor, Willy Siegel for sharing his domain knowledge, and Victor Lesser for his valuable input on the paper.

6. REFERENCES

- [1] J. M. Carroll, M. B. Rosson, G. Convertino, and C. H. Ganoe. Awareness and teamwork in computer-supported collaborations. *Interacting with Computers*, 18(1):21 – 46, 2006.
- [2] S. Chopra and L. White. Privacy and artificial agents, or, is google reading my email? In *IJCAI*, pages 1245–1250, 2007.
- [3] T. R. Colin, T. Mioch, N. J. Smets, and M. A. Neerincx. Estimating an operator’s cognitive state in real time: A user modeling approach. In *RO-MAN*, pages 627–633. IEEE, 2012.
- [4] T. De Greef, H. Arciszewski, and M. Neerincx. Adaptive automation based on an object-oriented task model: Implementation and evaluation in a realistic c2 environment. *Cognitive Engineering and Decision Making*, 4:152–173, 2010.
- [5] L. A. DeChurch and J. R. Mesmer-Magnus. The cognitive underpinnings of effective teamwork: a meta-analysis. *Journal of Applied Psychology*, 95(1):32, 2010.
- [6] C. Detweiler, K. Hindriks, and C. Jonker. Principles for value-sensitive agent-oriented software engineering. In *AOSE*, pages 1–16, Berlin, Heidelberg, 2011. Springer-Verlag.
- [7] X. Fan and J. Yen. Modeling and simulating human teamwork behaviors using intelligent agents. *Physics of Life Reviews*, 1(3):173–201, 2004.
- [8] M. Flanagan, D. C. Howe, and H. Nissenbaum. *Embodying Values in Technology: Theory and Practice*, pages 322–353. Cambridge University Press, Cambridge, MA, 2008.
- [9] B. Friedman, P. H. Kahn, and A. Borning. Value sensitive design and information systems. In *Human-Computer Interaction and Management Information Systems: Foundations*, pages 348–372, 2006.
- [10] R. Greenstadt, J. P. Pearce, and M. Tambe. Analysis of privacy loss in distributed constraint optimization. In *AAAI*, volume 6, pages 647–653, 2006.
- [11] B. J. Grosz and S. Kraus. The evolution of sharedplans. In *Foundations of rational agency*, pages 227–262. Springer, 1999.
- [12] E. Kamar, Y. K. Gal, and B. J. Grosz. Modeling information exchange opportunities for effective human-computer teamwork. *Artificial Intelligence*, 195:528–550, 2012.
- [13] L. C. Lee, H. S. Nwana, D. T. Ndumu, and P. De Wilde. The stability, scalability and performance of multi-agent systems. *BT Technology Journal*, 16(3):94–103, 1998.
- [14] R. T. Maheswaran, J. P. Pearce, E. Bowring, P. Varakantham, and M. Tambe. Privacy loss in distributed constraint reasoning: A quantitative framework for analysis and its applications. *Autonomous Agents and Multi-Agent Systems*, 13(1):27–60, 2006.
- [15] J. E. Mathieu, T. S. Heffner, G. F. Goodwin, E. Salas, and J. A. Cannon-Bowers. The influence of shared mental models on team process and performance. *Journal of applied psychology*, 85(2):273, 2000.
- [16] A. Mehrabian. Pleasure-arousal-dominance: A general framework for describing and measuring individual differences in temperament. *Current Psychology*, 14(4):261–292, 1996.
- [17] J. R. Mesmer-Magnus and L. A. DeChurch. Information sharing and team performance: a meta-analysis. *Journal of Applied Psychology*, 94(2):535, 2009.
- [18] M. Neerincx. *Cognitive task load design: model, methods and examples*, pages 283–305. Lawrence Erlbaum Associates, Mahwah, NJ, 2003.
- [19] M. A. Neerincx. Modelling cognitive and affective load for the design of human-machine collaboration. In *EPCE*, pages 568–574, Berlin, Heidelberg, 2007. Springer-Verlag.
- [20] M. A. Neerincx, S. Kenndie, M. Grootjen, and F. Grootjen. Modeling the cognitive task load and performance of naval operators. In *Foundations of Augmented Cognition*, pages 260–269. Springer, 2009.
- [21] H. Nissenbaum. *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press, 2009.
- [22] J. S. Olson, J. Grudin, and E. Horvitz. A study of preferences for sharing and privacy. In *CHI*, pages 1985–1988. ACM, 2005.
- [23] C. O. Porter, J. R. Hollenbeck, D. R. Ilgen, A. P. Ellis, B. J. West, and H. Moon. Backing up behaviors in teams: the role of personality and legitimacy of need. *Journal of Applied Psychology*, 88(3):391, 2003.
- [24] E. Salas, N. J. Cooke, and M. A. Rosen. On teams, teamwork, and team performance: Discoveries and developments. *Human Factors*, 50(3):540–547, 2008.
- [25] M.-C. Silaghi and D. Mitra. Distributed constraint satisfaction and optimization with privacy enforcement. In *IAT*, pages 531–535. IEEE, 2004.
- [26] J. M. Such, A. Espinosa, and A. Garcia-Fornes. A survey of privacy in multi-agent systems. *Knowledge Engineering Review*, pages 1–31, 2013.
- [27] K. Sycara and M. Lewis. Integrating intelligent agents into human teams. In *Team cognition: understanding the factors that drive process and performance*, pages 203–231, Washington, DC, 2004. APA.
- [28] M. Tambe. Towards flexible teamwork. *Artificial Intelligence*, 7:83–124, 1997.