

Information Security, Privacy, and Trust in Social Robotic Assistants for Older Adults

Thomas Given-Wilson, Axel Legay, Sean Sedwards

► **To cite this version:**

Thomas Given-Wilson, Axel Legay, Sean Sedwards. Information Security, Privacy, and Trust in Social Robotic Assistants for Older Adults. Theo Tryfonas. HAS 2017 - International Conference on Human Aspects of Information Security, Privacy, and Trust, Jul 2017, Vancouver, Canada. Springer, 10292, pp.90-109, LNCS. <10.1007/978-3-319-58460-7_7>. <hal-01629094>

HAL Id: hal-01629094

<https://hal.inria.fr/hal-01629094>

Submitted on 6 Nov 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Information Security, Privacy, and Trust in Social Robotic Assistants for Older Adults

Thomas Given-Wilson, Axel Legay, Sean Sedwards

Inria

Abstract. People with impaired physical and mental ability often find it challenging to negotiate crowded or unfamiliar environments, leading to a vicious cycle of deteriorating mobility and sociability. To address this issue the ACANTO project is developing a robotic assistant that allows its users to engage in therapeutic group social activities, building on work done in the DALi project. Key components of the ACANTO technology are social networking and group motion planning, both of which entail the sharing and broadcasting of information. Given that the system may also make use of medical records, it is clear that the issues of security, privacy, and trust are of supreme importance to ACANTO.

1 Introduction

People with impaired physical and mental ability often find it challenging to negotiate crowded or unfamiliar environments, leading to a vicious cycle of deteriorating mobility. This also severely impacts sociability, and increases isolation, that in turn provides an additional cycle of deteriorating health and well-being. To address these issues the ACANTO project¹ is developing a robotic assistant (called a *FriWalk*) that supports its users by encouraging and supporting them to engage in therapeutic group social activities.

The key components of the ACANTO project that act to counteract these vicious cycles are social networking and group motion planning. Both of these entail the sharing and communicating of information about the users. In the social networking setting, to encourage communication between users and organise groups who share common interests and locations. In the group motion planning setting, to coordinate groups of users participating in shared activities while maintaining both group and individual safety and comfort of users (even in different groups) in a shared environment. The goal of the social networking aspect of ACANTO is to support and encourage group activities, while the group activity and related motion planning are the key challenges of the project. Thus, the rest of this work focuses on this setting.

A significant aspect of the ACANTO project is the inclusion of medical professionals and medical information. Users may be prescribed therapeutic activities by a medical professional for maintaining health and well-being or to recover from mobility affecting injury. Thus, the information used in ACANTO about

¹ www.ict-acanto.eu

users may derive from medical records. This yields a particular challenge to information sharing aspects of ACANTO, since such user information must be handled with great care, to respect user security, privacy, and trust.

More generally, the ACANTO project by definition exploits user information to assist in developing social networks for group activities, and to aid in group motion planning. Both of these require the sharing of user information to function effectively. Since the target users of the ACANTO project are likely to be particularly vulnerable (physically or mentally impaired, recovering from injury, etc.), issues of information security, privacy, and trust are major challenges for the ACANTO project.

This paper considers challenges for the implementation of group activities and group motion planning in the ACANTO project. These can be divided into four broad challenges.

The first challenge is that the requirement to plan group activities must account for all the users in the group, thus leaking user information to other members of the group. Since the group activity must not violate the constraints of any user (such as not straying too far from a bathroom), this may be observable to other members of the group, and so members could conceivably infer sensitive private information about other group members.

The second challenge is that the sharing of information during navigation yields information about user location; to other group members and to other ACANTO users (even from different groups). The navigation and reactive planning assistance of the FriWalk exploit information from other ACANTO users and environmental sensors. This can lead to information about the location of other users being inferred, even across different groups.

The third challenge is that using medical information as part of group planning and social networking place a very high burden on the security and privacy of this information. Further, using medical information often has legal requirements that must be met. Thus, information derived from medical records must be handled with particular care.

The fourth is a different kind of challenge; trust in the ACANTO project by users. The users must feel that the ACANTO social network and FriWalk are trustworthy and that they will look after users. Since, without this, users will not use ACANTO and gain the benefits.

With this overview of the kinds of challenges the ACANTO project must address, in the sequel we give more information about the ACANTO project, the technologies exploited, the challenges, and possible solutions in the context of security, privacy, and trust.

The rest of the paper is structured as follows. Section 2 presents a more detailed view of the ACANTO project itself and sets the scope of the project. Section 3 discusses the technological choices made in ACANTO so far. Section 4 considers the challenges of the ACANTO project in more detail, including the limitations and requirements placed upon them by the project and technological choices. Section 5 sketches possible solutions to the challenges that can address the challenge within the bounds of the choices already made. Section 6 concludes.

2 Project Scope

The ACANTO project builds on work done in the DALi project². DALi created a robotic motion planning assistant, based on a standard wheeled walker, to aid those with reduced physical and mildly reduced mental ability to negotiate complex and potentially crowded environments, such as shopping malls or museums. The ACANTO project extends upon this in three directions. (1) the single user scenario of DALi is now generalised to include many users, including users who perform activities and navigate in groups. (2) a social network is created that helps ACANTO users find others to undertake group activities with and to maintain social support. (3) clinical versions of the walker can assist medical professionals with diagnostics and therapeutic activities for users.

These three extensions in ACANTO require various extensions to the work of DALi, and add new aspects that yield new challenges. The rest of this section overviews the ACANTO project and its requirements as they pertain to challenges for human aspects of information security, privacy, and trust.

2.1 FriWalk

The walker in the ACANTO project is called a *FriWalk* and extends upon the DALi walker. Despite these extensions to support group activities and medical assistance, many aspects of the FriWalk are carried over from the DALi walker.

The FriWalk assists in motion planning at two levels; a *long term planner* [4] and a *reactive (short term) planner* [3]. The architecture of the ACANTO motion planner is similar, but we wish to provide robotic guidance to a number of people with reduced ability who are taking part in a group activity within the same type of environments. The typical goal of such an activity is to facilitate social interaction and provide therapeutic exercise in an enjoyable way. Thus both the long term and reactive planners need to consider the notion of a group. To facilitate this group notion and maintain group cohesion, communication between FriWalks is considered a vital part of ACANTO.

An activity will be defined in advance, considering many user preferences. This in turn has requirements about how information about users is shared and handled outside the FriWalk and planning. For further detail see below.

During the activity, the ACANTO system comprises a fixed server and mobile client applications on the FriWalks, which interact via radio communication. The server collects, processes and distributes information gathered by sensors attached to the client devices. This information comprises the pose and location of the user and any other agents visible to his sensors. (These agents are any other humans in the environment, and are not assumed to be recognised by the sensors as being other FriWalk users or not.)

We assume that the high level goals of an activity are (eventually) translated into a *global plan* (a path for the group to follow through an environment), based on pre-existing knowledge of the layout of the environment. Hence we do

² www.ict-dali.eu

not consider the activity explicitly and simply assume the existence of a global plan. Each FriWalk attempts to follow the global plan, using a *reactive planner* that makes planning decisions according to the dynamic local conditions (e.g., the position of other agents) and the constraints and requirements imposed by the individuals and the activity. Dynamic conditions may require that the global plan is modified (e.g., an encountered obstruction not present in the plan of the environment), but in this work we focus on the challenges related to information security, privacy, and trust when performing distributed group planning and communication. Note that changes to the activity plan may also result from feedback during the activity, and so some care has to be taken with information leakage through (re)planning as well.

We want to guide the users to follow the global plan as a group, while allowing them to move around within the group. We want to achieve this by efficiently distributing the problem among the FriWalks and server, while ensuring that the loss of an individual will not cause the activity to fail. We must accommodate the possibilities that people have different physical abilities and that some members of the group will not be cooperative and may decide to temporarily or permanently quit the group.

2.2 Social Network

The creation of activities is largely driven by the social network of the ACANTO project. Users of the social network can propose activities, join groups to find or be recommended activities, or add their preferences and be recommended activities by the social network. Once an activity has been created, an activity plan is generated taking into account the goals of the activity, the preferences of the users participating, and the safety requirements assured by the ACANTO system. This requires gathering significant information about the users, and balancing potentially conflicting or competing requirements.

For example, an activity may include visiting various locations in a shopping mall. Taking into account user preferences, the activity global plan may need to ensure the following. The global plan does not travel too far from any bathroom. The distance traveled is within a lower and upper bound. The projected time taken is within lower and upper bounds. The global plan does not include any flights of stairs. There is at least one trained medical professional included among the users. Observe that these requirements will likely be synthesised from information gathered regarding the users subscribed to the activity. (Note that in theory it may be impossible to meet all requirements, or the solution may be unlikely to be achieved in practice, these concerns are not considered here.)

Observe that to generate such activity plans, it is necessary for user information and preferences to be considered. It would be impossible to create such an activity plan without user information, and unsafe to proceed with some plans if user information is withheld. For example, when a user requires a medical professional to be nearby at all times, and the activity plan may not include a medical professional in the list of users, and this would be unsafe.

2.3 Therapeutic Rôle

The ACANTO project also includes a therapeutic aspect where FriWalks are used in diagnosis and treatment. In particular, FriWalks may be used to gather, compare, and exploit medical information in collaboration with medical professionals. This allows for therapeutic care and support for those recovering from mobility impairments. Some FriWalks are planned to be certified as medical devices equipped with more sensors and capable of (assisting in) diagnosing or monitoring patient health. Since FriWalks may thus yield medical information that will be considered and exploited in therapeutic activities, the handling of medical information requires some care. However, these diagnostic rôles of FriWalks shall not be considered directly here, as this paper focuses on aspects related to the group activities and group motion planning.

As part of therapeutic care and other support, medical professionals may recommend a regime of activities to patients. Thus, medical information (both direct and inferred) may be used by the ACANTO system. This is most obvious in the creation of activity plans; since these must consider the users and their preferences, which in turn may be information directly by medical records, diagnoses, and treatment plans.

3 Technology

This section overviews the technology choices made in the ACANTO project. The focus here is on the technological choices made from the group motion planning—the generation of a global plan is straightforward, and the aspects of social networks and medical records are already well known from other contexts. Thus, the rest of this section presents key points of the technology used in ACANTO for group motion planning.

Activity planning generates an a priori global plan of therapeutic and social activities defined by the activity generator. Reactive planning refers to local motion planning that copes with the actual conditions encountered by the users, given the activity plan. In addition to accounting for unforeseen changes to the environment and other pedestrians who are not part of the activity, the reactive planner also accounts for the random, potentially uncooperative behaviour of users of the system. Activity monitoring is performed in real time and ensures that the concrete suggestions offered to the users will achieve the goals of the activity with high probability.

Although the ACANTO system will have powerful centralised infrastructure, communication latency and potential interruption require that reactive motion planning is both autonomous and cooperative. The algorithms must therefore be efficient because the motion planning problem is complex and the algorithms will be executed on low powered embedded hardware. In general, we require the system to be robust and able to take advantage of increased computational power and additional information as these become available.

This section presents key details about the chosen hierarchical framework to analyse the local environment and classify the behaviour of moving agents

or groups of agents. The framework takes as input a series of instantaneous snapshots of behaviour observed by the sensors. From these it constructs traces that evolve over time. In future, following developments of the ACANTO sensor technology, the framework will take traces or partial traces as input.

The more complex behaviour evident in the traces is clustered to infer grouping and other metrics that also evolve over time. The interpretation of these dynamic metrics allows ever more complex patterns of behaviour to be classified. To improve efficiency, we propose a group-based model abstraction that takes advantage of the fact that the motion of people walking together is strongly correlated. We thus motion plan at the level of groups, while incorporating a sliding level of abstraction that allows groups to consist of a single pedestrian.

3.1 Reactive Planning

Figure 1 gives a diagrammatic overview of the ACANTO reactive planner, the key elements of which are summarised below.

The global objectives, comprising the specification of the chosen activity and the preferences of the users, is provided as input a priori. During the course of the ensuing activity, sensors locate the users and other pedestrians with respect to the fixed objects in the environment. This information is used to parametrise a predictive stochastic model of human motion based on the social force model (SFM [9, 2]). The SFM is overviewed in Section 3.4. This model is used to simulate multiple future trajectories with respect to alternative immediate behaviour of the users. The sets of simulated trajectories corresponding to each alternative immediate behaviour are validated against the global objectives using statistical model checking (SMC). The basic notions of SMC are described in Section 3.3. The immediate behaviour that maximises the probability of achieving the global objectives is recommended to the users.

The measurements of the sensors contain an element of noise, but for the purposes of the SFM are treated as deterministic best approximations. To account for their potential inaccuracy and the fact that the model is necessarily an incomplete representation of reality, we include a random “noise” term that allows the SFM to explore non-smooth behaviour and behaviour that is not explicitly modelled by forces (see also Section 3.4). Simulations of the model are therefore samples of a random variable and it is




Fig. 1. Overview of reactive planning.

for this reason that we then use SMC to estimate the probability of “successful” trajectories.

Our reactive planner is in fact a combination of reactive and predictive planning. A purely reactive approach might be adequate if we could guarantee perfect sensing with no latency. On the other hand, with a perfect predictive model we would have minimal need to sense the environment. Since neither of these are feasible, we adopt a “predictor-corrector” approach. We make a recommendation to the user based on a prediction with an efficient human motion model (i.e., the SFM), then correct our recommendation with updated predictions as the user progresses. Using this approach we significantly improve on the performance of the SFM [2] and can accommodate unpredictable eventualities that would be difficult to include in any reasonable model.

3.2 Group Motion Planning

The notion of groups of pedestrians and their interaction is key to ACANTO. While the reactive planning approach described in Section 3.1 has been shown to be efficient on embedded hardware for the case of a single user [2], using the same ideas with a group of users leads to a potential exponential explosion of hypothesised initial behaviour. Given that each user may go left, right or straight with respect to their current direction, there are a minimum of $3^{\#users}$ alternatives to try. In practice there are also different degrees of left and right choices, so the number of alternatives is much higher. Our approach is to construct the group behaviour compositionally and to first hypothesise alternative initial behaviour of groups as a whole. Suggestions to individuals within the group will aim to respect the group motion and maintain the group cohesion. Similar ideas for crowd simulation have been explored in [12].

One of our principal concerns is therefore the detection of groups. Although the group of users involved in the activity may be known a priori, this notional grouping may not adequately reflect the actual grouping for the purposes of motion planning. For example, it is known that when part of a large group people often prefer to walk together in smaller groups of between two and four to facilitate conversation [10]. In addition, for technical reasons (e.g., temporary loss of network connection), it may not always be possible for an individual FriWalk to recognise all of the other members of a group. Moreover, most people in the environment are likely not part of the the therapeutic activity, but may nevertheless be moving in ad hoc groups [10]. It is therefore necessary to infer grouping directly from the trajectories of pedestrians.

3.3 Statistical Model Checking

Statistical model checking (SMC) is a variety of probabilistic model checking that avoids an explicit representation or traversal of the state space and estimates the probability of a property from an empirical distribution built by verifying a property ϕ against multiple independent executions (simulations) of the system.

Given N independent simulation traces ω_i and a function $z(\omega_i) \in \{0, 1\}$ that indicates whether $\omega_i \models \phi$ (read “ ω_i satisfies ϕ ”), the probability γ that, in general, $\omega \models \phi$ can be estimated using the unbiased estimator $\tilde{\gamma} = 1/N \sum_{i=1}^N z(\omega_i)$. The confidence of the estimate can be guaranteed by standard statistical bounds, allowing SMC to trade certainty for reduced confidence plus tractability. For example, the sequential probability ratio test [13, 14] efficiently evaluates the truth of an hypothesis without needing to calculate the actual probability, while the Okamoto bound [11] asserts a level of confidence for a given number of simulations N , expressed as $\Pr[|\gamma - \tilde{\gamma}| > \epsilon] < 1 - \delta$. In words, this formula reads that the probability that the absolute error of the estimate is greater than ϵ is less than $1 - \delta$, where δ is a function of N and ϵ . In comparison to the ‘certain’ varieties of model checking, SMC does not require a finite or even tractable state space. This makes SMC particularly suitable for the present application that considers continuous time and space.

3.4 The Social Force Model

The social force model (SFM) [9, 8, 7, 6] combines real and psychological “forces” to predict the behaviour of pedestrians in crowds under normal and panic situations. The model recognises that pedestrians are constrained by the physical laws of motion and also by social rules that can be modelled as physical forces. The model considers an environment comprising fixed objects (walls) and moving agents (pedestrians) that respond to attractive and repulsive forces that arise from social and physical interactions.

The SFM here considers (groups of) agents that have a mass a their centre, a velocity, and an ellipsoid shape. Additionally, fixed objects in the environment are modelled with solid forms according to their footprint. The SFM then determines the forces that act upon the (groups of) agents due to their desired path, current velocity, and the forces of other elements of the system. Thus the influence of all these forces can be used to predict the (group of) agents’ future path. In addition to these forces, random “noise” is added, that serves to represent fluctuations not accounted for by the model, and to avoid deadlocks.

In general the forces in the SFM can be defined and exploited in many different ways. For example, these may include: repulsive forces away from unknown agents, repulsive forces from fixed objects, attractive forces towards friends or other users, attractive forces towards activity goals, attractive forces to maintain proximity to a perimeter, etc. Thus, many different effects upon the agents can easily be represented and balanced by manipulating these forces in the SFM.

3.5 Our Approach

In ACANTO the reactive planner will be collaborative and cooperative. Sensor information obtained by each user will be shared between other users of the ACANTO platform, effectively giving the planner a much wider view. The planner will explicitly consider group motion and identify pedestrians who are part of

the same social activity as the user. Due to potential communication latency or interruption, planning will nevertheless be local to the user. The significantly increased complexity of the planning task thus necessitates an approach that is efficient. Figure 2 shows a diagrammatic representation of our behaviour classifier.

The classifier takes as input sensor information provided by the sensor board, as shown in Figure 1. This information will at least contain the estimated positions and velocities at a given time point of moving agents in the vicinity of the sensors of a number of users of the ACANTO platform. The current sensor technology (based on Microsoft Kinect) does not recognise the identity of individuals between consecutive readings by the same sensor. Future developments may allow the sensor technology to directly infer traces or partial traces and to identify users of the FriWalk, reducing the computation from sensor input to traces and clusters.

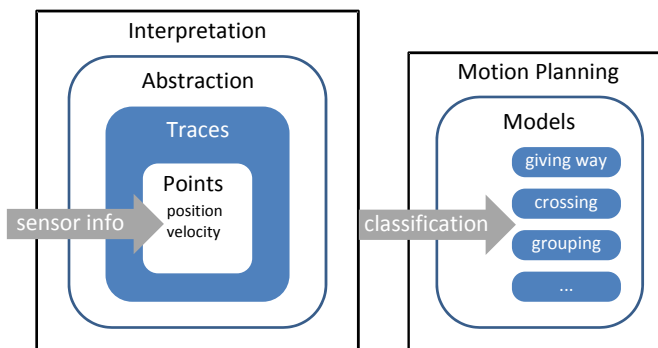


Fig. 2. Hierarchical structure of behaviour classifier.

Combining the output of multiple sensors incurs the additional challenge of identifying pedestrians who leave the view of one sensor and appear in another. Pedestrians may also appear and disappear as a result of sensors being obscured, because of communication unreliability or because users just leave. Our trace inference algorithm therefore makes minimal a priori assumptions about the data, but will take advantage of whatever information is available. If no additional input is available from other users, an individual reactive planner can still function using the information provided by its local sensors, along the lines of the DALi short term planner [3].

Having inferred a set of active traces, the classifier then clusters them into groups of traces with characteristics that imply the pedestrians are or will be moving as a group. Mere proximity is not a sufficient indicator, since two close pedestrians may be trying to get away from each other. The classifier therefore also considers velocity and acceleration (inferred from successive observations of velocity or predicted by the SFM). It is possible for the framework to include higher level information (e.g., we may know that two pedestrians are part of the same activity group), but if pedestrians are close and moving in a similar direction at a similar speed, for the purposes of motion planning they are already moving as a group, regardless of whether they are involved in the same activity. Finally, note that groups are not necessarily disjoint and may overlap.

Identifying de facto groups allows us to plan motion at a more efficient level of abstraction. When hypothesising the alternative directions for a number of users of the platform, we believe that it is a reasonable compromise to only hypothesise the overall motion of the groups to which they belong. We feel it is not necessary to consider all the possible combinations of suggestions to those within the same group given that, by virtue of how we define a group, their motion is strongly correlated. Note that suggestions are nevertheless tailored to the actual position of an individual within the group, in order to maintain its “social” structure. A further advantage of this approach is that we may also identify behavioural templates at the level of groups, rather than at the computationally prohibitive level of individuals. We may also quantify temporal properties over traces of group-related metrics.

Finally, it is important to note that a group may comprise a single pedestrian, so our framework allows us to choose a level of abstraction that is appropriate for the available computational power. In general, our approach is to plan the motion of an individual against an abstraction of the environment that may be as detailed or complex as the available computational capacity allows.

4 Challenges

This section considers information security, privacy, and trust challenges raised within the ACANTO project. Some of these are very strong: requiring careful handling of medical information in a social and collaborative setting. While others are more general, relating to handling of private location information, and indirect information leakage. The goal of this section is to provide an overview of four main areas of challenge, particularly as they relate to the group activities and group motion planning required for the ACANTO project.

4.1 Group Planning Leakage

One main area of challenge is in the leakage of information through the group activity planning. The planning of an activity must take into account all the constraints of all the group members, and so the end result must account for all of these constraints. This in turn implies that *some* potentially secure or private information must be shared and thus could be inferred by other group members based on the chosen group activity and global plan. The rest of this section considers the scope and risks of such information leakage in ACANTO.

To illustrate this leakage, consider the scenario where one group member has a particular medical issue that requires them to always be within five minutes of a bathroom. To satisfy this user, the activity and global plan must keep the group within five minutes of a bathroom at all times. In general this can lead to global plans that will not appear optimal or natural to a user who does not have this constraint.

For example, consider the paths presented in Figure 3. The shortest path is represented in gray, while the path that remains within five minutes travel time of a bathroom is in red.

Of course, inferring the cause for a particular global plan may not be a trivial exercise, particularly when several competing or complementary constraints are in play. After all, if bathrooms and access ramps are co-located in an environment, the inference may be that a group member is unable to use stairs, rather than a member of the group having issues that require a bathroom.




Fig. 3. Paths with and without bathroom constraints.

Such issues are not always as simple and obvious as the global plan chosen. Groups may have members that have upper or lower bounds on exercise, on time spent, on rest taken, and many other factors. Thus, the activity plan may approximate all of these initially, but require the plan to be recalculated and changed based upon these constraints, that may appear arbitrarily to another user.

For example, if the group has been moving too slowly, a re-plan to ensure sufficient calorie burn may lead to longer (non-intuitive) paths between the next activity locations. Similarly, if too much time has been spent resting, the global plan may alter or even drop activities, which would indirectly yield information both about the status and requirements of the users, but also about the importance of the activities themselves.

4.2 Shared Server Leakage

Another path for information leakage to occur is through the shared server infrastructure. Since information is shared between FriWalks through the server, it is possible for information about one FriWalk user to be leaked to others, even FriWalk users not in the same group. Similarly to the above, the goal of collaborative and group activities in ACANTO requires some communication and thus information leakage between users. Although this could be largely ignored for planning as discussed here, this would lead to significantly less user safety, which is considered a top priority in the ACANTO project.

The underlying goal of the shared server infrastructure is to allow for shared sensor data to be used, improving planning for both the long term planner, and the reactive planner. However, this also allows information to be inferred about where other agents are.

Consider the example illustrated in Figure 4. A user (here in blue with sensor range shown as the blue triangle) approaches a corner that obstructs vision. Another user (green here) around the corner may have shared (via the server) that an agent (in red) is about to turn around the corner from the obscured side. The reactive planner will advise the blue user whose vision is obstructed (both sensor vision, and natural vision) to turn away or stop in order to avoid a collision.

In such a scenario it is easy to see how leakage of user location can occur, since the FriWalk reactive planner will react to things that can only be sensed from another source. Thus, the approximate location of the other source could be inferred.

On the larger scale, the server will also be aware of obstructions and traffic behaviours from different locations due to information from other agents. Thus, a change in global plan could yield information about the prior locations of other users.

For example, the global plan may have originally been to cross a food court in a mall, however a previous user or group attempted to cross the food court (or is in the process of doing so) and reports significant traffic and issues with avoiding collisions. Thus, a user may have their plan updated due to server information that the path ahead will be densely trafficked and may not be suitable for them.

Thus, information about the location of other agents will be implicitly shared with users of the system. Although this is considered an overall benefit to the design of the system, and should improve overall safety and comfort for users, it cannot be ignored that some location information could be inferred about users by other users.

4.3 Medical Information

One particular challenge in the ACANTO project is the use of medical information. The ACANTO project includes medical personnel as part of the social aspects, and in assisting impaired adults with designing and completing preventative and recovery activities. However, this implies the use of medical data and potentially medical records in various aspects of the project. Further, some FriWalks are also equipped with more advanced sensors that can both aid in, and perform, medical diagnostics.

The use of medical records and medical information is more obvious in the planning for (group) activities. As part of rehabilitation and preventative care, activity plans should account for specified actions. These can include: minimum and maximum calories spent in activities; minimum and maximum distances

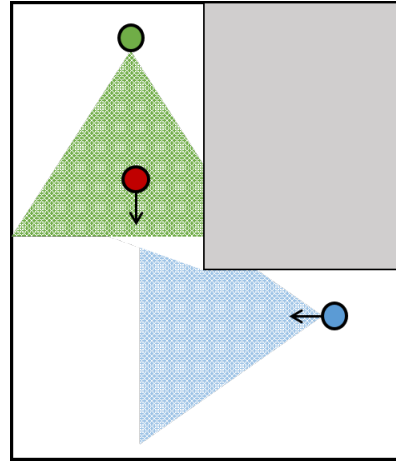


Fig. 4. Inferring User Location Around Corners.

walked; minimum and maximum time spent in a single walk; minimum and maximum sustained pace during walking; requirement to always be accompanied by a trained medical professional, etc. Thus an activity plan must adhere to these constraints that are derived from medical information without the source medical information being made available.

For example, consider when a patient’s medical records indicate that they must always have a trained medical professional in the vicinity, perhaps due to some heart condition that may require intervention. Clearly the global and reactive planners must ensure that this user is not separated from any medical professional who is part of the group. However, in general the group need not remain in a single group and can instead split into sub-groups for parts of the navigation and activities. If this vulnerable user is to be appropriately cared for, this will add new constraints on the planning, and may yield behaviour that makes it clear this user has some medical issue.

It follows that medical data must be used as part of the planning and activities, but this should be handled with great care since even indirect information may yield significant breaches of patient information security and privacy.

Legal Limitations To further complicate medical records use, in most jurisdictions there are many specific laws with respect to the use of medical records and in the general handling of privileged medical information. Although the ACANTO project is being developed in partnership with medical professions and with these in mind, larger outcomes and a general application of the ACANTO system will need to consider these issues with great delicacy.

4.4 FriWalk Trust

A different challenge is for adoption of ACANTO by users, there will also need to be user trust in the FriWalk and ACANTO social network. In addition to ensuring security to a sufficient level for users to feel comfortable with the FriWalks and social network, there must also be trust in the FriWalk itself behaving in a good, reliable, and user-centric manner. That is, FriWalks should rarely conflict with user perceptions and about planning (or only do so in a comprehensible manner), and must establish trust with the user that the FriWalk has their interests foremost.

The behaviour and advice of the FriWalk needs to be reliable and comprehensible to the extent that the user trusts that the directions and suggestions made are sensible. This is particularly pertinent when the directions given by the FriWalk may be in conflict with the expectations of the user.

For example, the user may be familiar with the environment and already know the “best” way to the next location in the activity plan. However, the FriWalk may suggest an alternative path. This could reduce user trust if the user does not see any reason or benefit to the suggested alternative plan.

Another example would be when the reactive planner suggests a direction that seems at odds with the immediate observations by the user. This could

occur when the server has provided information about obscured agents moving in the environment, and the user may not (yet) be aware of their existence (as discussed above in Section 4.2).

These general scenarios bring into consideration how to best support the user while providing good information, even in conflict with the user’s knowledge or observations.

Another dimension of trust is that the behaviour of the FriWalk must act in the best interest of the user. It is possible that a user may feel that the FriWalk is acting to force them into some conformity with the group, rather than taking into account the individual’s needs. This is most likely to appear when activities are designed to meet the requirements of many, and so may be suboptimal for many (or even all) users if considering the activity individually. This can include planning poor paths through the environment, choosing unnatural movement patterns to maintain group cohesion, guiding members to continue when they feel they need a break, or alternatively suggesting breaks or delays when users are keen to continue.

All of these provide a complex interplay of balancing the individual desires of a user, and the group plans and actions. To some degree the FriWalk should incentivise the user towards maintaining group cohesion and following the group activity plan. At the same time, it must be flexible and reactive to the needs of an individual; perhaps splitting the group easily when one user indicates a need to visit the bathroom and adding others to this sub-group to ensure no user is left alone.

5 Proposed Solutions

This section considers possible solutions to the challenges of Section 4. The focus here is upon how to solve the challenges within the framework of the ACANTO project, the technologies chosen, limitations, and in a manner that does not introduce new overhead or concerns. The goal of this section is to consider such possible solutions and their effectiveness. The details of their implementation (and related experiments) are left to future work.

5.1 Group Planning Leakage

In a general sense the issue of some information being leaked in such scenarios is unsolvable; it is not possible to create a global plan that both achieves the constraints required, and does not yield any information about those constraints. That said, it is feasible to mitigate the leakage of information, and address the manner in which it is leaked.

The most obvious “solution” to this issue is in the complexity and conjunction of the constraints themselves. While a constraint such as “must always be close to a bathroom” may appear strict, many other constraints could also lead to the kinds of paths features in Figure 3. As hinted in Section 4.1, it may not be obvious that this is the constraint imposed, since avoiding stairs could coincide.

More generally, this kind of leakage can be mitigated by the conjunction of other constraints. Consider that the global plan may be due to wishing to increase caloric burn, duration, or to avoid traffic and other issues the users are not aware of. All of these other plausible explanations make inferring a particular constraint much more complex, particularly given limited information.

To complicate such inferences further, the global plan is not made evident to the users initially. Thus, the users may not even be aware of the global plan having this initial constraint. Since replanning may occur due to a variety of factors, it is quite conceivable that the path followed by the users was emergent rather than designed.

Indeed, such emergent paths through the environment may dominate any global plan that was initially created. Considering that any local traffic factor or updated information could change the global plan, it is likely that perturbations of the global plan would be normal rather than an exception.

Even further, the above all assume that the users follow the directions and plans without agency. However, one key consideration in ACANTO is that the FriWalk provides guidance, but the user may ignore or alter their behaviour. Thus, even if there was the potential to reasonably infer some secure or private information from observing the plan followed, it would not be clear that this was planned, or simply emergent from the actions of users, reactive planning, and general constraints to maintain group cohesion.

For example, the global plan may have indicated a direct path that did not remain close to bathrooms. However, one user could have opted to ignore the suggested direction (or varied their actions due to reactive planning) and ended up shifting the whole group down an alternate path that was always close to bathrooms. (The reverse is also possible, with users opting to ignore the guidance to remain close to a bathroom. In such a scenario the FriWalk would strongly suggest directions to the users to maintain the constraint of staying close to a bathroom, but this cannot be forced by the FriWalk.) Thus, an observed path cannot be reliably assumed to have been the global plan chosen to satisfy user constraints, and thus inferring information from the global plan is non-trivial, and likely to be highly erroneous in practice.

In instances where it is obvious that a replan has occurred due to violation of constraints (such as when key activities are dropped, or the activity prematurely ended), it is still unclear which possible constraint this could be related to. Consider that a premature end could signal an overrun of: calorie burn, distance, time, scheduling, etc., or even that some emergency has occurred or some updated traffic information made the plan impossible.

Thus, although information leakage is impossible to avoid, the details are suitably obscured to make this a minor issue in the implementation of ACANTO.

5.2 Shared Server Leakage

Like the previous challenge, the leakage of information between users or other sensors in the environment cannot be completely prevented (and indeed would

contradict the choice in ACANTO to exploit this data to improve reactive planning). Again the solution here is to limit the amount of data that is directly evident to a user.

In a general sense the challenge is to prevent the location of another FriWalk user from being easily inferred by exploiting the location of agents that is provided by the server. This can be mitigated in four different ways.

In many environments the infrastructure also includes several fixed sensors, such as fixed cameras. This allows for the environment to be augmented with agent information that does not come from any user or FriWalk. That is, the fixed sensors of the environment can also provide the location (and trajectory information) of agents. Thus, when reactive planning exploits information about agents outside the sensor range of the FriWalk (doing the reactive planning), it is not certain that the agents being considered were observed by another user/FriWalk. Thus, while there *may* be another user in the vicinity, it is not necessary for another user to be in the vicinity to have information about the location of agents outside sensor range. That said, if the location of fixed sensors is known, the shared information through the server can still leak approximate location information about other users/FriWalks.

Even in this case, it is not clear that the information from the server can be used to infer the location of another agent with high precision. The information sent from the server is an n -dimensional cube (three dimensions: x , y , and maybe z for Euclidean space, and t for time) and so does not contain all the information observed by other users/FriWalks. Thus, it is not in general possible to infer precise location information about the sensor(s) that observed the other agents. Consider the diagram in Figure 5. The blue user may be able to infer that there is another user/FriWalk in the area to be informed about the red agent, but any of the green locations (as well as many others) are potential locations for the other user/FriWalk.

To further complicate the problem of inferring information from agent locations provided by the server, the server and reactive planner may both project the future locations of recently observed agents, even if no longer “visible” to any sensor. This arises because observations are recorded and sent asynchronously by FriWalks, while it may be necessary for both the server and individual FriWalks to predict the current state of the environment from stored data (for the purpose of disambiguating different observations of the same agent and if communication breaks down). Thus, even though an agent’s location was observed in the

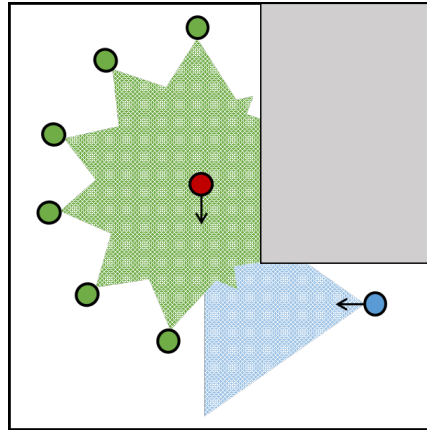


Fig. 5. Possible information about other user/FriWalk from server information.

past and sent to a FriWalk, this does not imply that the agent is still observable. Their location may be projected forward by the reactive planner, thus giving the location of an agent that is not actually observed to be there by any sensor. This makes inferring possible sensor information (and from this other user/FriWalk location information) much more difficult.

Lastly, location information is not perfect and the reactive planning itself is not deterministic. The location information is refined by further observations, so it is possible that locations will be made more accurate or altered when the server aggregates information from multiple devices. This can lead to location information changing (slightly). Further, the reactive planner itself is not deterministic: adding noise to the SFM and using SMC to derive the best outcome. This ensures a non-deterministic outcome, and so means that even with the same input, it is possible for different suggestions to be made to the user. Since the user only observes the direction suggestions of the reactive planner, this can obscure information that might lead to leakage of other user/FriWalk locations.

5.3 Medical Information

The handling of medical data in the ACANTO project raises several complications, particularly related to the shared information aspects of the project. The proposed solutions to prevent issues here fall into three general categories: legal solutions, data obfuscation solutions, and user trust solutions.

Legal Solutions In a sense the legal complications and challenges raised when handling medical data are the easiest to resolve. While the various jurisdictions and variations in legal requirements are on their own a challenge, the solution can be easily applied on the scale of the project implementation. At this stage the ACANTO project is approaching clinical trials overseen by medical professionals and in a research environment. Thus, for now the solution is to abide by the legal requirements for medical equipment trials, and thus comply with all legal requirements in the countries involved so far (UK, Spain). As an EU project, however, ACANTO must eventually comply with all subscribing nations.

In the future the goal of ACANTO is to build a larger scale social network that can support mobility impaired and older adults in various locations. This will involve rolling out the social network to different jurisdictions and with different legal requirements. However, the solution is clear in all cases – to abide by the legal requirements for medical data used in all jurisdictions.

Obfuscation Solutions Another solution to the challenge of handling medical data is to obfuscate or anonymize the data being used. This approach is inspired by traditional approaches to handling medical data in research, and in data obfuscation used for anonymization and as discussed in Sections 5.1 & 5.2.

Various approaches are used to anonymize medical (and other) data that is released to be used in research or as part of the results of a study. These techniques can be used within the ACANTO project to reduce leakage of medical

information being tied to any particular user. Although this is non-trivial to resolve and relies upon the data set and other information, techniques such as differential privacy [5] could be used at the activity and group planning stages to ensure that the data does not leak medical information about any individual.

Further efforts to reduce the medical information leakage can also be taken, such as adding noise to the various results that derive from medical data. This could be considered by extending the bounds of differential privacy to ensure greater distances, or by adding some randomness to the outcomes of derived results that depend upon medical information.

User Trust A further solution to the challenge is to provide the users with the choice of what medical information can be used in social and planning aspects of ACANTO. This could be approached in a similar manner to various current social networks that let the user choose what information applications and other users have access to. Thus, a user of the ACANTO network could choose to share their preferences (derived from medical information) or keep these secret.

For example, a user may be happy to allow the ACANTO social network and activity planning to know that they require rests regularly or have a very low top speed. The user may not be recognised as the source of this information due to other proposed solutions (above), or may be happy to share this and merely note it as part of their recovery. On the other hand, a user may choose to keep such requirements to themselves and merely trust that the activity monitoring and planning aspects of the activity planner and FriWalk will adjust for the slow speed and frequent rests during the activity.

This approach allows users to have greater control over how their medical information may be used by the ACANTO system. This should provide transparency and control to users, allowing them to better understand how the ACANTO systems works, and to gain trust in the ACANTO system accounting for their needs.

5.4 FriWalk Trust

This section considers general approaches and concerns with gaining user trust of the ACANTO system and FriWalk. This area is difficult to approach within the technology and scope of the implementation of the ACANTO project, yet on the other hand is also relatively straight forward to gain useful information from other domains and from user feedback during clinical trials.

The challenge of reliable and comprehensible behaviour for the FriWalk significantly falls back to the FriWalk respecting the requirements and constraints of the user. Thus, the global and reactive planners should clearly behave in the best interest of the user. This is simple to implement for an individual, though slightly more complex when multiple users may have competing constraints.

Competing constraints and scenarios where there is no “good” solution is a hot topic in the related domain of autonomous vehicles [1]. The proposed resolutions of ethical dilemmas in that field can potentially provide a basis for solutions within the ACANTO project.

To further support such choices, ACANTO also has clinical trials that can be used to gather more information, and also implement different solutions to gain user feedback. This is discussed further below in Section 5.4.

The other aspect of when the user and the FriWalk/FriTab believe separate paths are “better” can be resolved by the FriWalk not being rigid in the choices. By having the global and reactive planners both willing and able to replan as the situation changes, ensures that the user can override what they think is a “poor” choice and have the FriWalk smoothly adjust to this change. This is similar to GPS/navigation in vehicles, whether the user not taking the nominated path causes a replan, rather than the device attempting to override the user or force them back to the original plan.

Another challenge is when the individual and the group may be in competing positions. For example, when an individual needs to use the bathroom or in some other manner wishes to diverge from the group/activity. In such scenarios, the proposed solution is to provide the user of a FriWalk an option to notify their device of this scenario. Thus, the device can react to the requirements of the individual user, and also relay this to the server and other devices. This will allow for the user’s FriWalk to immediately support their needs, overriding any constraints imposed by the group. Further, this can allow other users to react accordingly, such as ensuring a companion or medical professional is aware of the situation and can support the user, or perhaps the group will be split to ensure the user is not left behind.

Clinical Trials Various proposed solutions have been suggested above or in related works and projects. In addition to considering and learning from other projects and results, the ACANTO project will be conducting clinical trials of groups of users and will be able to experiment with different approaches. This can be used to determine which solutions work best, and also which behaviours of the FriWalk and the ACANTO project as a whole are most accepted, and also which cause tension or distrust in users. It is expected that feedback from clinical trials will be used to refine proposed solutions, and be able to test different solutions to see which are most effective and more trusted by users in practice.

6 Conclusions

The ACANTO project’s aim to assist physically and mentally impaired users with therapeutic and social support via social networking and group activities clearly raises several challenges in the context of security, privacy, and trust. Further, the choices of technology to implement the ACANTO project also influences and implies other challenges.

These challenges range across several areas. Information sharing where users of ACANTO join or participate in group activities that must account for the safety and support of all users. Information leakage during motion planning where other users of ACANTO even outside the group may have location information leaked. The handling of medical information of users, both in the legal

and privacy dimensions. Gaining user trust and providing a trustworthy platform for users so they can feel safe in interacting with ACANTO.

Several ways to address these challenges have been presented, and their feasibility to implement within the design and technological choices of ACANTO. In general these provide effective solutions that balance the need to, for example, share information with the need to maintain user security, privacy, and trust.

References

1. J.-F. Bonnefon, A. Shariff, and I. Rahwan. The social dilemma of autonomous vehicles. *Science*, 352(6293):1573–1576, 2016.
2. A. Colombo, D. Fontanelli, D. Gandhi, A. De Angeli, L. Palopoli, S. Sedwards, and A. Legay. Behavioural templates improve robot motion planning with social force model in human environments. In *Emerging Technologies Factory Automation (ETFA), 2013 IEEE 18th Conference on*, pages 1–6, Sept 2013.
3. A. Colombo, D. Fontanelli, A. Legay, L. Palopoli, and S. Sedwards. Motion planning in crowds using statistical model checking to enhance the social force model. In *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*, pages 3602–3608, Dec 2013.
4. A. Colombo, D. Fontanelli, A. Legay, L. Palopoli, and S. Sedwards. Efficient customisable dynamic motion planning for assistive robots in complex human environments. *Journal of ambient intelligence and smart environments*, 7:617–633, 2015.
5. C. Dwork. Differential privacy. In *33rd International Colloquium on Automata, Languages and Programming, part II (ICALP 2006)*, volume 4052, pages 1–12, Venice, Italy, July 2006. Springer Verlag.
6. D. Helbing, I. Farkas, P. Molnár, and T. Vicsek. Simulation of pedestrian crowds in normal and evacuation situations. In M. Schreckenberg and S. D. Sharma, editors, *Pedestrian and Evacuation Dynamics*. Springer, 2002.
7. D. Helbing, I. Farkas, and T. Vicsek. Simulating dynamical features of escape panic. *Nature*, 407:487–490, September 2000.
8. D. Helbing, I. J. Farkas, and T. Vicsek. Freezing by heating in a driven mesoscopic system. *Phys. Rev. Lett.*, 84:1240–1243, 2000.
9. D. Helbing and P. Molnár. Social force model for pedestrian dynamics. *Phys. Rev. E*, 51:4282–4286, May 1995.
10. M. Moussaïd, N. Perozo, S. Garnier, D. Helbing, and G. Theraulaz. The walking behaviour of pedestrian social groups and its impact on crowd dynamics. *PLoS ONE*, 5(4):e10047, April 2010.
11. M. Okamoto. Some inequalities relating to the partial sum of binomial probabilities. *Annals of the Institute of Statistical Mathematics*, 10:29–35, 1959.
12. S. Raupp Musse and D. Thalmann. Hierarchical model for real time simulation of virtual human crowds. *Visualization and Computer Graphics, IEEE Transactions on*, 7(2):152–164, April 2001.
13. A. Wald. Sequential Tests of Statistical Hypotheses. *Annals of Mathematical Statistics*, 16(2):117–186, 1945.
14. H. Younes and R. Simmons. Probabilistic verification of discrete event systems using acceptance sampling. In *CAV*, volume 2404, pages 23–39. Springer, 2002.