Kennedy, W.G. and Trafton, J.G. (2011) How long is a moment: The perception and reality of task-related absences. International Journal of Social Robotics 3(3), pp 243-252. DOI: 10.1007/s12369-011-0098-7

How long is a moment: The perception and reality of task-related absences

William G. Kennedy and J. Gregory Trafton

Naval Research Laboratory

4555 Overlook Ave. SW, Washington, DC 20585

phone: 1-703-993-9291

Fax: 1-703-993-9290

wkennedy@gmu.edu

Note: The first and corresponding author is now at the Krasnow Institute for Advanced Study at George Mason University.

Abstract

We have investigated actual and perceived human performance associated with a simple task involving walking and applied the developed knowledge to a human-robot interaction. Based on experiments involving walking at a “purposeful and comfortable” pace, parameters were determined for a trapezoidal model of walking: starting from standing still, accelerating to a constant pace, walking at a constant pace, and decelerating to a stop. We also collected data on humans’ evaluation of the accomplishment of a simple task involving walking: determining the transitions from having taken too short a period of time to an appropriate time and from having taken an appropriate time to having taken too long. People were found to be accurate in estimating the task duration for short tasks, but to underestimate the duration of longer tasks. This information was applied to a human-robot interaction involving a human leaving for a “moment” and the robot knows how long the task should take and how time is evaluated by a human.

Keywords: *time estimation; time perception; human-robot interaction; walking*

1 Introduction

How long is a moment? Consider: you have given your elderly mother a personal assistive robot and your mother leaves the room for a moment to go take her evening pills. She comes back in 10 seconds. Was that too short? What if she is not back in 10 minutes? Should the robot go and make sure she is OK? When does the robot’s behavior change from being patient, respectful, but prepared to assist if appropriate to being impatient, distrustful, and overprotective? At least in some instances, it seems to rely on the concept of how long a moment is.

A “moment” is normally not a specific measurement of time. It is one of those commonsense concepts that may be difficult for an assistive social robot to deal with because its definition is situation dependent. Our focus is on the innumerable daily personal and social activities for which an assistive robot may become involved. Typical tasks in the home include making a meal, using the restroom, getting an item from another room, and checking on a sleeping child. To investigate these momentary tasks, we needed a typical task, data on its actual performance, and then data on people's perception of how long the task takes to complete. This paper describes a typical task, experiments on its actual performance and on participants' perception of others performing the task, and an application to a robot in a social situation.

We chose the common, "I'll be back in a moment" kind of task where the task is well understood and involves being out of sight for much of the task's performance. In such a task, the travel time is commonly the major contributor to how long the task takes. Fortunately, the travel time to walk a specific distance can be easily calculated using a simple model and the starting and ending points for multiple distances can be clearly delineated: the departure from view and the return.

Second, we needed to know whether the simple walking model was correct: how long does it actually take to perform such a task. We therefore collected data on how long it took people to walk various distances and return. From that data, we developed parameters for a simple model of walking from point A to point B and returning to point A.

Third, we needed to know what people's perception of the time to perform the task is, i.e., how long do we think it takes to perform such a task. Our focus was not on the time to complete the task per se, but on the perception itself. For example, we wanted to know had the performer taken too long, or too short an amount of time for a given distance. We conducted a second experiment to explore how long people thought a specific walking task **should** take. Clearly, there should be a relationship between how long a task takes and people's perception of how long a task should take.

The first experiment involved participants walking a short distance, performing a simple task, and returning to the starting point. The second experiment involved participants watching videos of a person departing and returning from performing the same task as in the first experiment. Participants were asked to judge the appropriateness of the time interval shown for the task. Finally, we report on a social robotic system that uses the information we collected, both the true times it took to perform the task and human perceptions of the time it took to perform the task. Because walking is the major contributor to the duration of the task, we start there.

2 Background on human walking

The process of tracking humans and modeling their movements is applicable to a wide variety of fields: medical, biomechanical, traffic engineering, computer graphics, animation, and mobile robotics. Much research has been done on the dynamics of the human body while in motion, particularly walking [1, 2]. However, we are less concerned with the details of limb and joint movement than we are with the overall movement of the body while walking. Constant walking speeds are well documented [3, 4], but this is only part of the solution because acceleration and variable speeds may be significant contributors in short trips.

Research supporting reasonable computer-generated visualizations of walking is a source of information, since virtual people in video games and computer animations are expected to act in a realistic manner and their motions are often generated by algorithms [5]. The task of calculating a person's progress along a known path from a known or estimated starting state consists primarily of estimating their speed and acceleration, which is increasingly important as the length of the path decreases. Possible models with increasing accuracy and complexity are a single speed with no acceleration, linear acceleration with a constant cruising speed, and a more accurate model using cubic polynomial to match observed velocities as was done by Brogan and Johnson [5].

A single speed model with no acceleration consists of simply dividing the distance by an average speed, resulting in the time it will take to travel a distance. This type of model could be adjusted with additional speeds for different portions of a path (uphill vs. downhill for example) or the single speed could be set to the average expected speed over a path (this approach would tend to be path dependent). This type of model is the simplest computationally, but obviously inaccurate for short distances where acceleration and deceleration times are significant contributors to the overall time.

Pedestrian traffic literature [3, 6] tends to report velocity as a single value, ignoring acceleration completely, but the similar speeds reported for crossing different size roads indicates that acceleration is not a significant factor (at least in road crossing averaged over a population). This research suggests that people accelerate quickly relative to the speeds and distances involved.

However, humans do not move a constant speed due to their bipedal locomotion. Figure 1 shows a three-phase model of human trunk speeds including oscillations with each step. An average speed during the rhythmic phase is drawn and similar straight lines could be used for the acceleration and deceleration phases. This would result in a trapezoid model [7, 8].

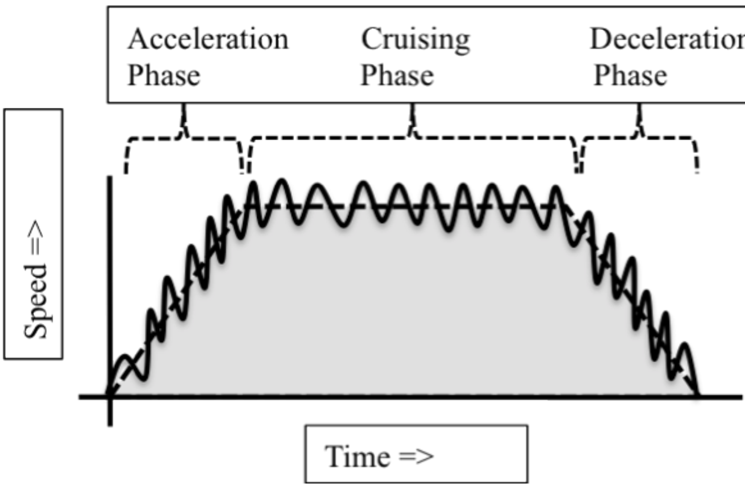

Figure 1. Phases of human walking

Simplifying this figure into a trapezoid model consists of a linear acceleration to the cruising speed (specified over a set time or at a set rate) followed by a period at a constant cruising speed and finally a linear slowdown to stop. For our discussion, we will focus on a straight path without obstacles, hazards, or significant walking surface changes. The constant cruising speed is also affected by the motivation and capabilities of an individual. For higher accuracy, the speeds used in the model could be tailored to the specific person of interest.

Clearly human motion does not consist of piece-wise continuous linear functions, but such a model can be used to reasonably approximate the speed at which a person is moving, given knowledge of their progress through a given path. Brogan and Johnson [5] explicitly measured the velocity profile of the participants while changing speed and found it to be similar to a cubic polynomial. The polynomial was modeled as a lie segment (linear acceleration over a fixed distance). They used an acceleration distance of 1.82 m (5.9 feet) and a deceleration distance of 1.63 (5.33 feet) because stopping was found to occur faster than starting approximately 10 percent). The computational advantages of linear over cubic velocity ramping are obvious.

A cubic acceleration profile could be simplified to a trapezoidal model. The main difficulty with such a model is the computation of velocity and position from the curve. The cubic model should be more accurate during acceleration and deceleration than the trapezoid, but with a well-chosen linear velocity acceleration value, the difference between the cubic and trapezoid models should be insignificant for distances greater than about 4 m (12 feet).

One obvious deficiency of the trapezoid model is the selection of the cruising speed. The problem is that speed tends to be specific to the individual [4] and the circumstances [3]. It should be noted that [3] found the mean street crossing speed of pedestrians under the age of 65 to be 1.51m/s (4.95 ft/sec) while [4] found the mean comfortable speed for adults in their 20s to 70s to be 1.45 m/s (4.79 ft/sec) or less. These results may be effectively (and statistically) the same. Bohannon [4] indicates that walking speeds (both comfortable and maximum) depend on height, weight, and leg strength; with the difference between comfortable speeds being as high as 0.19m/s (0.62 ft/sec) (the difference in cruising speeds was much greater. This leads to the conclusion that a value for cruising speed should consider individual differences and motivation (or situation) for high accuracy.

Ideally, a system for modeling human walking speed/progress would have knowledge about the walker (height, weight, and leg strength), the path (distance and obstacle information), and the situation (motivation of the walker relating to speed). Given this information, a trapezoid model could provide highly accurate walking time predictions. With this background, we conducted an experiment to determine parameter values for the trapezoidal model of human walking performing a specific, but typical, “back in a moment” task.

3 Performing a simple task involving walking

The first experiment was a simple task focused on walking. The goal was to establish the parameters for a trapezoidal model of walking.

3.1 Method

3.1.1 Participants

Twenty-eight students attending George Mason University participated for class credit. There were 9 men and 19 women ranging in ages from 18 to 40 with a mean of 21. Their heights ranged from 1.58m to 1.91m (62 to 75 inches) with a mean of 1.71m (67 inches).

3.1.2 Task Design and Procedure

Participants were asked to walk at a "comfortable and purposeful" pace, in an empty indoor hallway. For each trial, the participant started from a stationary standing position on a line taped to the floor, walked a specific distance down an empty, indoor hallway to a doorbell against the wall (see Figure 2), pressed the button for the doorbell, and walked back to the starting line and stopped standing on the taped line. The distances were 15.2m, 30.4m, and 45.6 m (50, 100, and 150 feet) one way. Data was collected for five consecutive trials at each distance with the distance for the block of trials randomly ordered. Longer distance under the same conditions, i.e. straight indoor hallway, were not available nearby.

3.1.3 Measures

Two times were recorded for each trial using a digital stopwatch. The first was the time from the start of motion to the pressing of the doorbell. The second was the time from the start of motion to stopping upon return to the starting line.

Figure 2. Doorbell and button to be pushed at the turning point in the task

3.2 Results

The mean times for each subject to perform the complete task for the three distances were about 1 second more than 20, 40, and 60 seconds, respectively, as shown in Table 1.

Table 1. Mean performance times (with standard deviation in parentheses) for a simple task involving walking.

Distance (one way)	Time to ring Doorbell in seconds	Time to complete task in seconds
15.2m (50 feet)	11.3 (1.35)	21.8 (2.78)
30.4m (100 feet)	21.3 (2.54)	41.7 (5.18)
45.6m (150 feet)	31.2 (3.66)	61.1 (7.04)

To examine the relationship between walking speed and gender, age, and height, we took each participant's average walking speed for each trial and performed a correlation between speed and each auxiliary variable. There was no relationship between walking speed and gender, $r(26) = -0.15$, $p = 0.5$, walking speed and age, $r(26) = -0.24$, $p = 0.22$, or walking speed and height, $r(26) = 0.02$, $p = 0.9$.

3.3 Analysis

The goal of this experiment was to establish the parameters for a trapezoidal model of walking. We assumed that the acceleration and deceleration periods could be treated as equal given their small contribution to the overall task. We also assumed participants walked at the same cruising speed independent of the direction and distance involved for these relatively short distances, i.e., there was no significant fatigue and their speeds were constant for the duration in both directions. At the turning point in the task, we assumed that the time to push the doorbell and turn around were insignificant and included with the walking portions of the task such that we could ignore the time to perform that small part of the task. Overall, the only parts of the task that contributed to the data collected were the acceleration/ deceleration times and the cruising speeds. With these simplifications, the two data points taken for each task could be transformed into the two parameters for the trapezoidal model of walking.

This resulted in:

cruising speed = 1.56 m/s (min: 1.27, max: 1.92, SD = 0.18) or

5.12 ft/sec (min: 4.16, max: 6.29, SD = 0.57)

and time to achieve cruising speed or to slow down and stop

= 0.64 sec (min: 0.21, max: 1.15, SD: 0.22).

3.4 Model

From analysis of walking data and the two parameters, the predicted times for participants to walk to the doorbell and press the button ("time out") and the time from the start until they complete the task ("time back") are:

time out = 2 * (time to achieve cruising speed) + (distance/cruising speed)

= 1.28 sec + (distance) / (1.56m/s or 5.12 ft/sec) seconds, and

time back = 4 * (time to achieve cruising speed) + 2 * (distance/cruising speed)

= 2.56 + (distance) / (3.12m/s or 2.56 ft/s) seconds.

The model can be evaluated in two different ways: compared to mean performance and compared to individual performance. As Figure 3 suggests, the model captures mean performance quite well; all 3 model points are within 95 percent confidence bars of the empirical walking distances and RMSD of the model is 0.11.

A second way to evaluate the model is to compare the model to each individual participant's average for each distance walked. In this case, the RMSD was 2.7, 5.1, and 6.9 for 50, 100, and 150 feet, respectively. These RMSD values are a consistent 12 percent error. Additionally, we ran a correlation between each individual participant's average for each distance walked and the model prediction; r^2=.90, $p < 0.05$. This model thus accounts for 90 percent of the variance at the individual level.

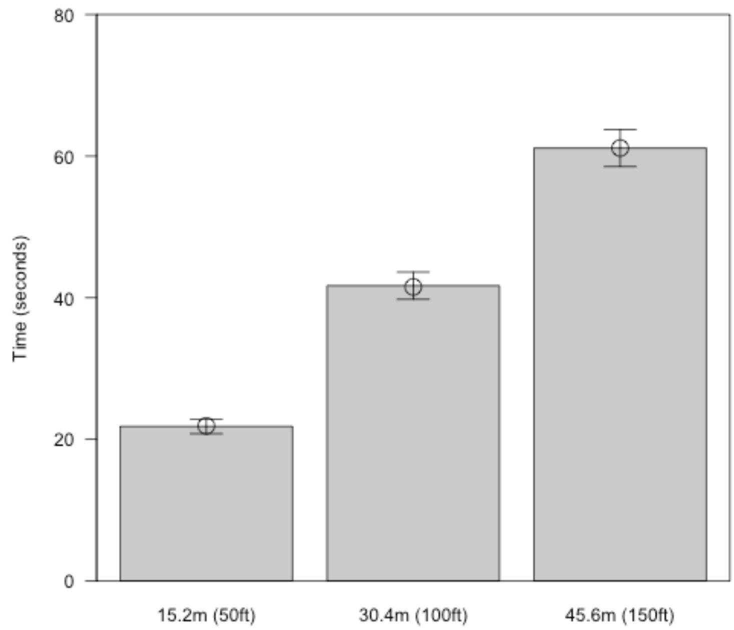

Figure 3. Reality and trapezoidal model of performing a simple task involving walking (the bars are mean value of 28 participants, error bars are 95 percent confidence intervals, and circles are model predictions)

3.5 Discussion

We found that the trapezoidal model of walking is quite good at estimating how long it takes people to walk distances of about 15, 30, and 45 meters (50, 100, and 150 feet) and return under the conditions of our experiment: indoors, unobstructed, and motivated to walk in a "purposeful and comfortable" pace. The small size of the standard deviations of the model and the experimental measurements for the overall task (as shown in Figure 3) and for the half task (up to pushing the door bell, not shown) support this.

We expect this model would be useful for predicting human walking for arbitrary but relatively short distances under similar conditions, such as the many simple daily tasks that take "just a moment" around a person's home or group home. Unlike more general population walking data [3, 4, 6, 7], we found no effect of the walkers' sex, height, or age. Of course, our participants were university students who perform a lot of purposeful walking and may be unusually practiced at "purposeful and comfortable" walking at the same speed as others. Had we obtained participants at a shopping mall or elder care center, we expect we would have had a wider variation in observed walking speeds.

With knowledge of the actual performance of a simple task involving walking, the next question is how people evaluate others performing such a task. That is the subject of the second experiment.

4 Evaluation of the duration of a simple task

What do you do when someone you are working with needs to leave for a moment but does not come back for a while? There are several different approaches to this everyday event. You or a personal assistive robot could mentally simulate the task step by step or could recall previously recorded performance times for similar tasks or could numerically calculate an appropriate expectation. There is evidence that people are good at estimating short time intervals, under a minute, but not for longer intervals [9-11] and there is evidence that scientists use mental simulation only when they have to [12]. We expect humans would expect the same capabilities they have to be in their assistive robots [13, 14]. Our second experiment was intended to discover how good people are at evaluating a moment, i.e., the duration of the performance of a simple, but unobserved, task.

We looked into how accurate people are at evaluating the time to perform a simple task that involves walking. The typical scenario is having a companion who goes away for a moment; say she goes to another room in the house, performs a simple task like retrieving her evening pills, and returns. Further, presume our companion is out of sight for most of the time, i.e., we cannot watch the task being performed. The issue is then not how good or poor we are at estimating distances, walking speeds, and calculating time intervals numerically, but how good we are at predicting task completion events. We want to know how well people do at this evaluative task so that a robot can appropriately provide assistance such as questioning whether the task was actually performed (or was it forgotten) or could investigate or call for help if there is something seriously wrong based on taking too long. Our hypothesis is that a system that represents and reasons like people do will be better at dealing with people respectfully and appropriately than one that does not [15, 16]. In accordance with our hypothesis, our second experiment investigates whether there is a difference between actual performance and our reasonable expectations.

4.1 Method

4.1.1 Participants

Twenty-five George Mason University students participated for course credit. There were 11 men and 14 women. Their ages ranged from 18 to 35 years with a mean of 20.4. None of these students had participated in the previous experiment.

4.1.2 Task Design and Procedure

Each participant evaluated videos of performance of the first experiment at the same three distances, 15.2, 30.4, and 45.6m (50, 100, and 150 feet) (one way distances down an empty, indoor hallway) and classified the video as the person having taken “too long”, “about right” or “too short” to perform the task. A computer program presented a series of videos for one distance and participants provided their evaluations until the transition points between their assessments of “too long”, “about right”, and “too short” were identified. They then worked on another distance until they had identified transition points for all three distances. The ordering of the three distances to be evaluated was randomized for each participant.

Each video started by showing the distance to be walked and then the image of a person leaving the doorway to perform the task and eventually returning. The distance to be walked was presented at the beginning of every video by showing a walker’s view of the distance and panning the camera from looking down at the starting position on the floor, up to horizontal showing the doorbell in the distance (see Figure 4), and then back down to the starting position. In this way, the participants were shown but not told the distance involved. Following this display of the walking distance, the camera position was shifted to inside a room off the hall and looking out a doorway into the hall to see just the feet of a person, the walker. It was not clear whether the person walking was male or female. The video then showed the walker leave the doorway area and, after a variable period of time with the screen showing just the empty doorway, the walker returned into view and stopped at the starting position. Participants were then asked to judge whether the person in the video took “too long,” “too short,” or “about right” time to perform the task. Based on their response, the computer system selected the next video. The videos were constructed in one-second intervals from 5 seconds to 2 minutes of walking time.

As a final note, the experimental design, with its alternate testing of the two thresholds, was intended to reduce the potential for an anchoring bias by subjects. Subjects were shown interleaved videos altering between narrowing in on the threshold between “too short” and “about right” and narrowing in on the threshold between “about right” and “too long”. By this method, subjects were not able to use the previous video as a basis for evaluating the next video.

Figure 4. Doorbell 30.4m (100 feet) down the empty hall

4.1.3 Measures

The duration of the video showing the walker's task performance was varied based on the evaluation provided by the participant and which transition was the current focus. To explain the data collection, consider that the subject's evaluation of the previous video was that it was "too long" and the current focus was the transition between "too long" and "about right". Then the next video would be longer than the previous video that had been judged to be "about right" and less than the most recent video that was judged to be "too long". This way the sequence of videos got closer and closer to the subject's transition point. When the gap between videos with evaluations on either side of the transition was down to one second, the transition had been found. This process was also followed for the other transition, i.e., between "about right" and "too short". To minimize subject biases, the system alternated working on the two transitions so that the subject would not know which threshold was the current focus and just saw a series of videos with no obvious pattern of video lengths. This search process quickly and reliably found the two transitions for each subject.

4.2 Results

The times at which the participants' evaluations changed from "too early" to "about right" and from "about right" to "too late" are shown in Table 2.

Table 2. Evaluations of the duration of a simple task involving walking

Task Distance (one way)	Range of "about right" evaluations	Transition from "too early" to "about right" Mean time (SD)	Transition from "about right" to "too late" Mean time (SD)
15.2 m (50 feet)	15 – 32 sec.	16.8 sec. (4.4)	25.5 sec. (4.13)
30.4m (100 feet)	20.5 - 60.5 sec.	33.0 sec. (11.1)	47.7 sec. (11.8)
45.6m (150 feet)	34 - 73.5 sec.	40.0 sec. (9.14)	63.4 sec. (12.5)

The transition points bound the range of times the participants considered to be "about right" for each distance. Figure 5 shows the middle of these ranges, i.e., mean times between the transition points of the task with a 95 percent Confidence Intervals.

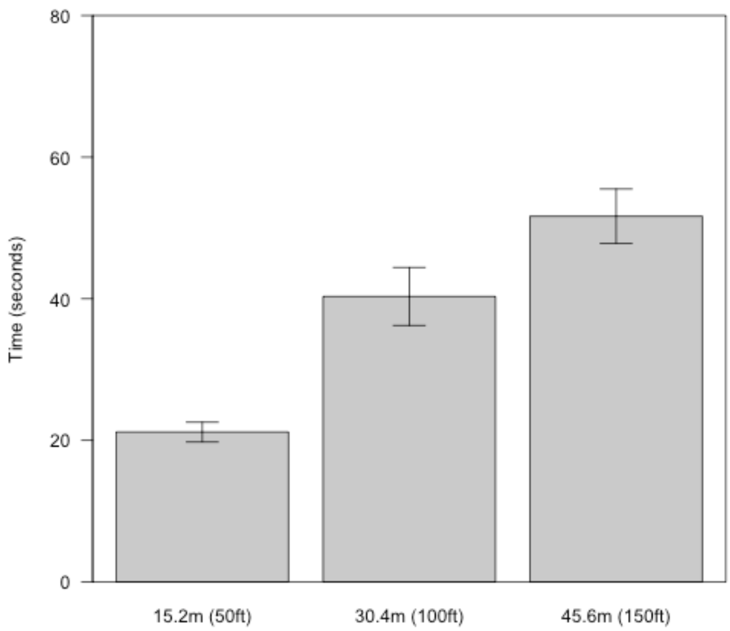

Figure 5. Mean of transitions between "too short" and "about right" and between "about right" and "too long" with 95 percent Confidence Intervals.

4.3 Analysis and Model

In addition to providing assessments of when is too early and when is too late, the range between these transitions provides an assessment of when the walker is expected to return. The mean between these transitions for each distance compares well with the measured task completion times from experiment 1 as shown in Table 3 and Figure 6.

Table 3. Comparison of mean (and SD) between transition and actual task completion times

Task Distance (one way)	Mean between evaluation transitions, or the "about right" time (SD)	Mean time to actually perform the task (SD)
15.2m (50 feet)	21.2 seconds (3.56)	21.8 seconds (2.78)
30.4m (100 feet)	40.3 seconds (10.4)	41.7 seconds (5.18)
45.6m (150 feet)	51.7 seconds (9.82)	61.1 seconds (7.04)

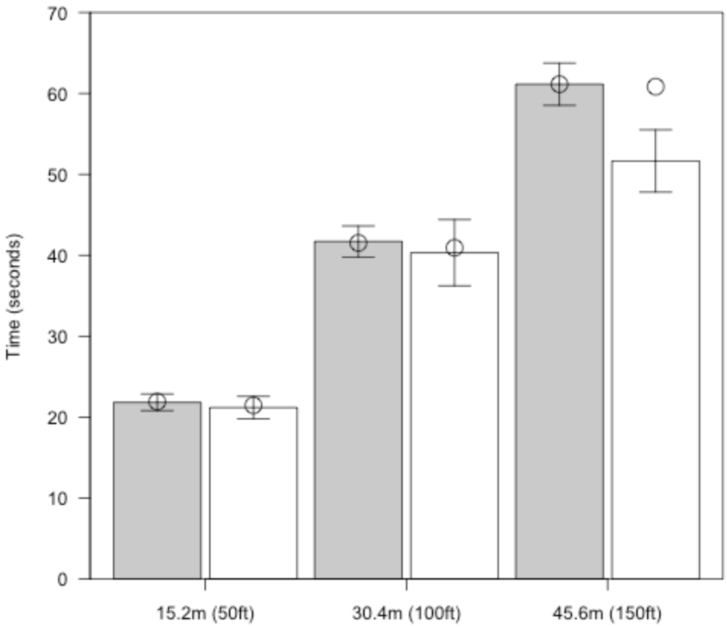

Figure 6. Reality, trapezoidal model of reality, and perception in human walking tasks (gray bars are round trip walking data, white bars are evaluations (perception), and circles are modeled data; error bars are 95 percent Confidence Intervals).

4.4 Discussion

The trapezoidal model of the walking is accurate in estimating the time to actually walk each of the three distances. However, means of people's estimates are accurate as well but only for the first two distances (within the 95 percent Confidence Intervals, see Figure 6) and less than the actual time for tasks involving taking approximately one minute (clearly not within the 95 percent Confidence Intervals). This result is consistent with the literature that we are accurate in time estimates up to 45-60 seconds and poor above.

It is also interesting to note that the variation of the participants' estimates for the "about right" performance did not grow linearly with the length of time involved.

In summary, we have shown that people are reasonably good at estimating how long someone should be gone if it is a short distance and the "moment" is less than a minute, but tend to under estimate on longer distances and "moments".

With this information, we can apply this information to have a personal assistive robot behave appropriately when someone is gone to perform a task that takes just a moment.

5 Application to human-robot interactions

We applied this developed knowledge of actual and perceived performance of a simple task to a mobile robot to demonstrate appropriate human-robot interactions in the course of human absences from the immediate area of human-robot social interactions in accordance with our representational hypothesis.

5.1 Robotic System

The robot we used is an iRobot B21r. It is a human-scale robotic platform with a zero-turn-radius and designed for indoor environments. The robot has a set of sensors and effectors associated with movement and a flat-panel display with an animated face [17, 18]. The raw inputs from the sensors are processed by onboard software and converted into symbolic, feature representations for use by our cognitive model in real time. Cognitively driven requests for movement of the robot are passed from cognitive model to the robot's motion control subsystem [19]. Speech output requests are sent to a commercial speech generation system, Cepstral. The animated face is synchronized with the speech output and turns to face the appropriate direction to indicate a change in visual attention.

5.2 Cognitive Model

We built a cognitive model for the robot that noticed when a human companion left the immediate area, noted where the human said he or she was going, and acted appropriately if they did not return as expected.

The cognitive architecture we used started with ACT-R [20]. We modified and embodied it in a robot as ACT-R/E [21]. The basic system is a hybrid symbolic/sub-symbolic production-based system. It has modules that are intended to represent specific cognitive functions such as visual and auditory perception, declarative (fact-based) and procedural (rule-based) memory, manipulation, vocalization, and time perception. These modules each have anatomical correspondences with recent fMRI data [20]. Our additions provide for processing of the robot's visual and auditory sensors, localization of the robot, and its movement to arbitrary locations.

The ACT-R/E system includes a time estimation module based on the work of Taatgen, Rijn, and Anderson [22]. This module provides a logarithmic accumulator for the passage of time and includes variations in its measurements to match human performance in a range of prospective time estimation tasks.

A cognitive model implemented in the ACT-R/E architecture is primarily an initial set of declarative and procedural memories that determine behavior. The architecture repeatedly matches the conditions of all productions against the current state of the buffers associated with its internal modules, selects one production to fire serially, and through the firing of that production modifies buffers or makes functional requests of modules which result in updates to buffers, and the cycle repeats. Module request cause actions by the robot involving the environment, such as visual sensing and movement of the robot to a specified location.

Our model is initialized with two times to complete the task at each of several distances. The first time is based on using the trapezoidal model and parameters derived from the first experiment. It is the robot's estimate of the accurate time for a human to complete the task. The second time is the mean time of the evaluated transition from an "about right" duration to having taken too long. This data came from the second experiment.

5.3 Resulting Behavior

When the robot needed to estimate how long a task that is performed by its companion should take, the system retrieves both times. Using the robot's internal time module, when the internal time estimate reached the mean time of the transition of humans' assessment from having taken "about right" amount of time to having taking too long, the robot comments on the fact. This allows the remaining humans in the area to be patient and wait until the actual performance time. When the time exceeds the mean actual time for having taken too long, the robot acts and goes looking for its companion. (We used the mean time to perform the task plus one standard deviation.) While waiting for the companion to return, the robot with such a cognitive model could perform other actions. Video is available at http://www.mllab.com/WalkBot/WalkBot.htm.

5.4 Discussion

Implementing a cognitive model using the information developed on how humans perform a simple task involving walking and how they evaluate the performance of such a task raised an interesting issue. While we found that humans may be quite consistent and predictable in their "purposeful and comfortable" walking task, their evaluations of such performance were not accurate. As a result, a robot has two different answers for how long is too long to perform a simple task involving walking: how long humans think it takes to perform a task and how long it actually does take. For short time periods, up to 45 seconds or one minute, the two times are effectively the same because people are good at such estimations. However, starting with slightly longer times, a minute or longer, humans expect tasks to take less time than they actually do. An issue for our implementation was that a human would consider the companion too late before the robot would. We chose to implement the robot commenting when the human would be erroneously expected to have returned and had the robot act when the companion had actually taken too long plus one standard deviation. An alternative implementation would have been to have the robot act when a remaining human would consider the missing teammate too late, which is prior to when they are actually being late. Our implementation, therefore, integrates both sources of knowledge on human performance of simple tasks involving walking.

6 General Discussion, Limitations, and Future Work

We report the results of measuring university students performing a simple task involving walking, evaluating another person performing the same task, and then we report on implementing these results in a robot. We found that when our participants were asked to walk at a "purposeful and comfortable" rate, their performances were very consistent and their walking could be accurately characterized using a trapezoidal model. When we asked other participants to evaluate the appropriateness of the task performance times shown in videos, their evaluations of too early, about right, and too long were much more inconsistent. However, their evaluations of performance times for tasks of relatively short distances (50 and 100 feet or 15 and 30 meters) was comparable to the actual performance time but participants underestimated the time necessary when the people in the videos walked 150 feet or about 45 meters. Because of this difference, implementation of these results in a robot raised the dilemma that the robot could either act like a person when evaluating the absence of a companion knowing a human underestimates longer tasks or the robot could act on the knowledge of the actual task performances and not act like a person. If your elderly mother had a guest, the assistive robot could calm the guest if the guest would likely underestimate the performance time or, in the opposite case, confirm to the guest that assistance may be needed. In our video demonstration, we had the robot note when a human would evaluate the performance as too long and take action based on when the absence would actually be too long.

The findings of this work have direct implications for human-robot interaction. First, we have determined university students' "purposeful and comfortable" walking parameters for a trapezoidal model. Second, we have confirmed that people are good at estimating the duration of short tasks, those taking under a minute. Third, we have confirmed that people think tasks that take longer than a minute take less time than they actually do. With this information, the human-robot interactions can be informed of the actual and perceived performance of simple tasks involving walking.

There are limitations in this work based on the human subjects used, the methods used, and the generality of the results. These experiments used university students as subjects. While we may believe they are useful for the evaluative aspects of the experiment, they may not be representative of our elderly mothers based primarily on the known effects of aging on walkers [3, 4]. In addition, in the referenced studies, the elderly walkers are walking outside and without the assistance of canes or other devices. Our method attempted to control for different motivations of walkers by asking them to walk at a "comfortable and purposeful pace". In other studies, pedestrians crossing a street are also motivated to move quickly. However, our elderly parents may not be so motivated and, in the privacy of their own homes, may significantly not be so purposeful and this may in turn impact their walking speed and variation. Another limitation is that we studied a task taking about a minute or less. If the time involved is much beyond a minute, such as a visit to a store by car or an elderly person walking from their living room to an outside mailbox and back, then the divergence we report in the evaluations may be greater. Subjects traveling in different environments with different motivations and taking more time may require additional study.

Our work lays the groundwork for studies of the social interactions between humans and personal assistive robots involving performance of small tasks and the perception of a "moment". Whether people are as predictable and consistent in performing the task as our university students were or if the performance is unknown and additional data collection is necessary, we have shown that there are perception differences between actual performance and evaluations of having taken "too long". In a more general way than we have used, a personal assistive robot could collect data on actual performance of their companion and allow for observed noise, say a standard deviation or more, before providing assistance. Further research or personal choice may be the best way to establish the duration of a moment and when the robot would change from being a patient, respectful, but monitoring assistant into the robot being perceived as impatient, distrustful, and overprotective.

Acknowledgments

This work was performed while the first author held a National Research Council Research Associateship Award and was partially supported by the Office of Naval Research under N001406WZ20001 and N0001406WR20156. John Hutchison contributed to the background material, Kristina O'Connell conducted the experiments, and Malcolm McCurry prepared the videos and program to obtain evaluation data. The views and conclusions contained in this document should not be interpreted as necessarily representing official policies, either expressed or implied, of the U.S. Navy or the National Research Council.

References

1. Hirasake E, Moore ST, Raphan T, Cohen B (1999) Effects of walking velocity on vertical head and body movements during locomotion. Exp Brain Res 127: 117-130
2. Imai T, Moore ST, Raphan T, Cohen B (2001) Interaction of the body, head, and eyes during walking and turning. Exp Brain Res 136: 1-18
3. Knoblock RL, Pietrucha MT, Nitzburg M (1996) Field Studies of Pedestrian Walking Speed and Start-Up Time. Transp Res Rec (1538): 27-38
4. Bohannon RW (1997) Comfortable and maximum walking speed of adults aged 20-79 years: reference values and determinants. Age and Ageing 26: 15-19
5. Brogan DC, Johnson NL (2003) Realistic Human Walking Paths. in Sixteenth International Conference on Computer Animation and Social Agents (CASA)
6. Gates TJ, Noyce DA, Bill AR, Van Ee N (1982) Recommended Walking Speeds for Pedestrian Clearance Timing Based on Pedestrian Characteristics. Transp Res Rec: 38-47
7. Rose J, Gamble JG eds (2005) Human Walking. Third ed. Lippincott, Williams and Wilkins: Philadelphia
8. Lettre C, Contini R (1967) Accelerographic Analysis of Pathological Gait, Techinical Report No. 1368.01. NY University School of Engineering and Science: New York
9. Buhusi CV, Meck WH (2006) Internal timing with gaps and distractors: Evaluation of the ambiguity, switch, and time-sharing hypotheses. J Exp Psychol: Anim Behav Process 32: 329-338
10. Zakay D, The evasive art of subjective time measurement: Some methological dilemmas, in Cognitive models of psychological time, RA Block, Ed. Erlbaum: Hillsdale, NJ. 59-84
11. Zakay D, Block RA (1997) Temporal cognition. Curr Dir Psychol Sci 6: 12-16
12. Trickett SB Trafton JG (2007) "What if...": The Use of Conceptual Simulations in Scientific Reasoning. Cogn Sci 31(5): 843-875
13. Breazeal C, (2003) Towards social robots. Robot Auton Syst 42(3-4)

14. Heerink M, Krose B, Evers V, Wielinga B (2010) Assessing Acceptance of Assistive Social Agent Technology by Older Adults: the Almere Model. Int J Soc Robotics 2: 361-375
15. Trafton JG, Schultz AC, Bugajska MD, Mintz FE (2005) Perspective-taking with robots: Experiments and models. I Workshop Robot Human Interactions IEEE Press 580-584
16. Kennedy WG, Bugajska MD, Marge M, Adams W, Fransen BR, Perzanowski D, Schultz AC, Trafton JG (2007) Spatial Representation and Reasoning for Human-Robot Collaboration. in Proceedings of the Twenty-Second Conference on Artificial Intelligence (AAAI2007) AAAI Press: Vancouve 1554-1559
17. Parke F, Waters K (1996) Computer facial animation. Natick, MA: AK Peters, Ltd
18. Simmons R, Adams W, Atrash A, Bugajska M, Coblenz M, MacMahon M, et al (2003) GRACE: an autonomous robot for the AAAI Robot challenge, in AI Magazine 51-72
19. Schultz AC, Adams W, Yamauchi B (1999) Integrating exploration, localization, navigation, and planning with a common representation. Auton Robots 6: 293-308
20. Anderson JR (2007) How Can the Human Mind Occur in the Physical Universe? Oxford: Oxford University Press
21. Kennedy WG, Bugajska MD, Adams W, Schultz AC, Trafton JG (2008) Incorporating Mental Simulation for a More Effective Robotic Teammate, in Proceedings of the Twenty-Third Conference on Artificial Intelligence (AAAI 2008) AAAI Press: Chicago 1300-1305
22. Taatgen NA, van Rijn H, Anderson JR (2007) An Integrated Theory of Prospective Time Interval Estimation: The Role of Cognition, Attention, and Learning. Psychol Rev 114(3): 577-598

William G. Kennedy conducts research and teaches in the fields of cognitive science and computational social science in the Krasnow Institute for Advanced Study at the George Mason University in Virginia, USA. He published previously in this journal, in Cognitive Systems Research journal, and at conferences of the Association for Advancements in Artificial Intelligence (AAAI), Behavior Representation in Modeling and Simulation (BRiMS), International Conference on Cognitive Modeling (ICCM), and several social science conferences. His current research focuses on improving cognitive modeling within social simulations.

J. Gregory Trafton leads the research team in intelligent systems at the Navy Center for Applied Research in Artificial Intelligence located outside Washington, DC. He has published widely in cognitive science, cognitive robotics, and human-robot interactions. His current research focuses on robotic Theory of Mind, object recognition, and person recognition and the cognitive architecture to support these topics.