An Agent-Based System for Simulating Dynamic Choice-Sets

Qi Han, Theo Arentze and Harry Timmermans, (Eindhoven University of Technology)

Davy Janssens and Geert Wets, (Hasselt University)

***Abstract* — Contributing to the recent interest in the dynamics of activity-travel patterns, this paper discusses a framework for the dynamic formation of location choice-sets. It is based on the concepts of aspiration, activation and expected utility. Based on principles of Bayesian learning, re-enforcement learning and social comparison theories, the framework specifies functions for experience-based learning, extended and integrated with social learning.**

I. Context

SO-CALLED activity-based models have rapidly gained interest in the transportation research community. These models predicts and simulate in a coherent fashion multiple facets of activity-travel behavior including which activities are conducted, when, where, for how long, with whom, and the transport mode involved. To the extent that these models have been actively implemented, some type of simulation is used to simulate the implementation of predicted activity-travel in time and space. The majority of these simulations is based on Monte Carlo simulations, other such as Albatross [1-3] and Aurora (e.g., [4]). An overview of these developments is given in [5]. In addition to these comprehensive models, agent-based simulations have also been suggested for particular facets of activity-travel choice (e.g., [6-10]).

As part of the FEATHERS model ([11, 12]), an extension and elaboration of Aurora (e.g., [4]), an agent-based system which incorporates different types of dynamics and learning discussed in [13], this paper discusses the conceptual framework that is suggested. Because it is quite different from commonly used framework on the AI community, it may be useful to introduce it to this research community.

II. The framework

The basic assumption is that an agent acts based on behavioral principles and mechanisms. (S)He holds beliefs (knowledge) about the environment during a certain life course, has preferences and basic needs, leading to plans, agendas and schedules. (S)He carries out those plans, agendas and schedules in time and space. When a deviation exists between his/her expectation and aspiration an agent may start exploring his/her environment for new alternatives. Thus, (s)he learns about the environment and the consequences of his/her actions, in this case the choice of activity locations, and is able to adapt to changing circumstances and improve less effective behavior. Based on experiences, an agent forms habits, reinforces memory traces, updates beliefs about attributes of locations and routes, discovers the conditions under which certain states of the environment are more likely than others, and in so doing makes sense of the world around him/her. Moreover, through social contacts agents exchange information and adjust aspirations, which may trigger actions to explore new alternatives. Thus, for an agent, the composition of the location choice-set for a specific activity under certain conditions is dynamic. The alternatives within the choice-set will be expanded with newly discovered alternatives and reduced with old ones that are discarded or no longer retrievable from memory.

An *aspiration level* is an agent's goal for the outcome of the decision [14, 15]. In theory, aspirations could be defined either at the level of choice alternatives (a bundle of attributes) or individual attributes. We assume that it is more plausible to define aspiration at the level of attributes as it is on that level that an agent may determine goals that give direction to exploration processes (e.g., find alternative stores with a lower price level rather than find stores that have higher utility for my purposes). Defined for an attribute, an aspiration serves as a subjective reference point, which determines what qualifies as a satisfactory outcome for that attribute. An aspiration level is agent and, in case of a dynamic attribute, context-specific and, in the context of this paper, associated with location attributes. The outcome of a comparison between aspiration and actual or expected outcome given current knowledge provides a measure of an agent's satisfaction and willingness to explore new alternatives. A possible discrepancy between the expected outcomes derived from the alternatives within the current choice-set and the agent's aspiration levels may trigger the agent to switch from habitual behavior to a conscious choice mode.

Generally, aspiration levels are context dependent. For example, satisfaction or tolerance about the crowdedness encountered at shopping locations may vary by day-of-the-

Manuscript received April 25, 2007. This paper has been written in the context of the FEATHERS-project, funded by IWT, Vlaanderen, Belgium.

Q. Han, T. Arentze and H. Timmermans are with Eindhoven University of Technology, P.O.Box 513 Eindhoven, 5600 MB Eindhoven, The Netherlands (corresponding author to provide phone: 0031-40-2475403; fax: 0031-40-2478488; e-mail: q.han@bwk.tue.nl; T.A.Arentze@tue.nl; H.J.P.Timmermans@tue.nl).

D. Janssens and G. Wets are with Hasselt University, Hasselt, Belgium. (e-mail: davy.janssens@uhasselt.be; geert.wets@uhasselt.be).

week and shopping location's category type. Aspiration levels can relate to both (quasi)-static attributes and dynamic attributes (which may fluctuate as a function of the behavior of all agents in the system). Formally, we denote the set of current aspiration values as $A = \{A_k\}$, where $A_k = (c_k, e_k)$, $e_k = (e_{1k}, e_{2k} \ldots, e_{jk})$, e_{1k} represents the aspiration value of the first attribute under the k-th condition, and $c_k = (c_{1k}, c_{2k}, \ldots, c_{lk})$ defines the k-th condition as a set of states of the condition variables considered.

Agents within similar social demographic classes or belonging to the same social network may have similar aspiration levels since they adapt their aspirations based on social comparison. Agents also judge what make up a satisfactory outcome, and have the ability to memorize situations and outcomes (i.e., events). In part, this is context dependent, that is, certain contextual conditions automatically activate particular memory traces that lead to particular levels of awareness. The *activation level* of a location alternative is the indicator of the strength of such a memory trace, and hence reflects the ease with which it can be retrieved from memory. As such, an activation level is associated with each alternative in the current choice-set for each specific contextual condition, for example, defined in terms of type of activity (i.e. purpose of the trip), the previous activity location (i.e. origin location of the trip), day-of-the-week and time-of-the-day.

By repeatedly performing certain behavior under same situational conditions, agents develop habits. By forming and following habits, agents can reduce mental effort involved in constantly evaluating choice alternatives and making choices. By saving cognitive resources for the operation, habits help agents conserve mental resources and time, and free them for other tasks. Habits have been described as learned and scripted behaviors and are capable of being automatically activated by the situational conditions that normally precede the behavior. As such, the activation level of a location represents the degree of an agent's habit of choosing that location under certain contextual conditions. In our framework, habitual behavior involves that agents consistently select from a choice-set the alternative with the highest activation level under the given condition at the moment a choice is to be made. In turn, we define the choice set in a given choice situation as the locations that are retrievable from memory in that situation (i.e., condition). Formally, let q be the number of relevant condition variables, $W_i^t(z_m)$ be the activation level of an alternative i under condition m, where $z_m = (z_{1m}, z_{2m} \ldots, z_{qm})$ represents the states of the q condition variables under condition m, ω be a minimum activation level for memory retrieval ability. Then, the choice-set is defined as $\Phi^t(z_m) = \{L_i \,|\, W_i^t(z_m) \geq \omega\}$. Note that, as implied by this equation, the definition of a choice-set may vary between situations.

The attractiveness of a location is in general influenced by values of its attributes. Depending on the targeted need underlying the activity, the attributes that should be evaluated may be different. For example in case of shopping, the variety of stores is important for entertainment and purchase purposes, while a social need requires some familiarity with the location. Furthermore, the intention of resting attracts attention to spatial layout, while economic considerations emphasize quality and price. Thus, the impact of location may be diverse, that is, the combination of location and activity needs determines the utility of a location. Moreover, some of the attributes are (quasi)-static, while others are dynamic. The (quasi)-static attributes reflect characteristics of the location that are in short term constant, for example the size category, price level, parking space, and presence of stores for certain goods in a shopping centre. We assume that an agent will learn all the (quasi)-static attributes of a location simply through observing them after implementing an activity at that location. This knowledge will keep constant, and only change when the physical conditions are changed externally, for example, after a renovation of the shopping centre.

Dynamic attributes, such as crowdedness and travel time, are subjective and uncertain, and may be dependent on contextual variables. We assume that for each dynamic attribute, X_j, the agent uses some classification, denoted as $Xj = \{x_{j1}, x_{j2}, \ldots, x_{jN}\}$, where x_{j1}-x_{jN} represent possible states of X_j, and specifies his/her beliefs regarding location i based on his/her current knowledge as a probability distribution across X_j denoted as $P_i(X_j)$, which sums up to 1. The degree of uncertainty is given by the degree of uniformity of $P_i(X_j)$. The more evenly the probabilities are spread across possible states, the larger the uncertainty is, and vice versa. For example, consider again the crowdedness of a shopping location. This is a dynamic attribute of a location and therefore may involve uncertain knowledge. An agent could choose four states for crowdedness as {no, little, medium, very}, and specifies his/her beliefs regarding each shopping location i as a probability distribution across these four states.

In addition, the agent may discover that probabilities of states are conditional upon certain contextual variables. For example, the agent may discover that probabilities of crowdedness of a shopping location depend on day-of-the-week and time-of-the-day (e.g., peak hours and non-peak hours). Learning that some variables have an impact on outcome-states means extending unconditional probabilities $P^t(X_j)$ to obtain conditional probabilities $P^t(X_j \,|\, C)$, where C stands for one or more variables.

A utility function allows the agent to evaluate each location alternative given his/her current beliefs about the attributes of the location (including travel time) and his/her preferences. Using probabilities of the types $P^t(X_j|C)$ to describe the knowledge of the agent, the *expected utility* equation can be expressed as below:

$$EU_i^t(c_k) = EU_i^{static} + EU_i^{dynamic}(c_k) \qquad (1)$$

$$EU_i^{static} = \beta X_i \qquad (2)$$

$$EU_i^{dynamic}(c_k) = \sum_j \sum_n \beta_j x_{jn} P_{ij}^t(x_{jn}|c_k) \qquad (3)$$

Where EU_i^t is the expected utility of location *i* at time *t*, βX_i is the expected partial utility of location *i* for static attributes and preference, and $\beta_j x_{jn} P_{ij}^t(x_{jn}|c_k)$ is the expected partial utility of location *i* under possible states x_{jn} with probabilities $P_{ij}^t(x_{jn}|c_k)$ and preference β_j regarding dynamic attribute *j*. $c_k = (c_{1k}, c_{2k}, \ldots, c_{lk})$ represents the values of relevant condition variables under the *k*-th condition. Thus, expected utility takes into account current beliefs regarding state probabilities as well as an agent's preferences. Of course, static attributes could also be dealt with by these equations, namely as the special case where the believed state has a probability of 1.

In the assumed choice making process, agents go through a mental process to arrive at a choice. They start with implementing their habitual behavior that requires least mental effort, and carry on with conscious choice that asks for more effort only if the habitual choice is not satisfactory, until they find a choice that is satisfactory. Figure 1 schematically shows the main steps of the decision making process by which the model arrives at a location choice. As the aspiration levels are the standards for determining whether an outcome is acceptable, they will try to find the alternative that meets the requirements within a tolerance

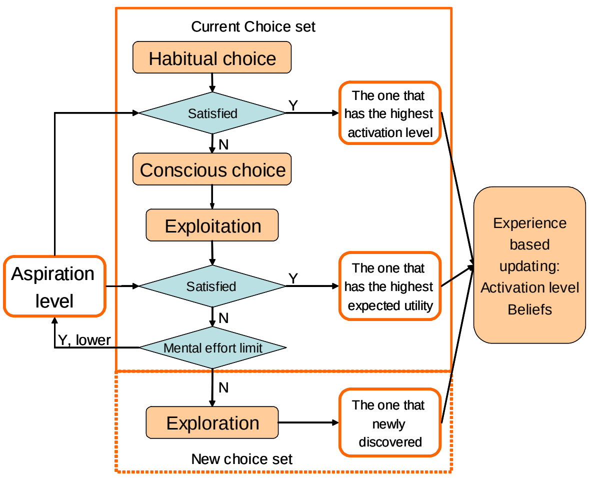

Fig. 1. The decision process of the model

threshold. The tolerance threshold is a predefined and agent specific parameter that reflects a characteristic of the agent. A large tolerance threshold indicates the agent strongly dislikes the mental effort involved to make better actions and is sooner happy with the current situation. Vice versa, a small threshold implies that on the one hand the agent is stricter in what is found acceptable, and on the other hand the agent may have a higher propensity to explore. In general, the larger an agent's threshold is, the higher the probability will be that the agent is satisfied with the expected performance of the current choice-set. Being satisfied with the current situation means less desire to take a risk, invest effort, and change behavior, and consequently, also that it is less likely to explore and possibly make better choices in the future.

As implied by the definition of activation level, the alternative that has the highest activation level in the choice-set is the one that is most easily retrieved from memory and requires the smallest amount of mental effort from an agent. In order to determine the level of satisfaction with the habitual choice, the location with the highest activation level is compared to aspiration levels. We assume that if dissatisfaction (i.e., the difference between aspiration and expected level) regarding at least one attribute exceeds the tolerance threshold, an agent will switch to another mode of behavior and start searching consciously for better alternatives. On the other hand, if this threshold is not exceeded, we assume that no active search will take place and that the agent will exhibit habitual behavior that leads to executing the choice that has the highest activation level.

We make a distinction between exploitation and exploration as alternative non-habitual modes of choice making. We assume that when acting in a conscious mode, an agent will first be engaged in exploitation and search within his/her current choice-set (i.e. retrieve alternatives from his/her memory that have a lower awareness) for a better alternative under current conditions. With exploitation, the agent calculates the expected utilities (equation 3) of all the alternatives within the choice-set given current knowledge of the environment and the given conditions, and compares the attributes of the one that has the highest expected utility with aspiration levels. When for none of the attributes dissatisfaction exceeds the threshold, we assume that no active exploration of new alternatives will happen and the agent will choose the location that has the highest expected utility. If for at least one attribute there is a mismatch that exceeds the tolerance threshold, the agent will start to explore new alternatives that might solve the mismatch. We call this exploration. Thus, search is not random, but rather directed. The attribute causing dissatisfaction will guide the agent in what to search for.

Exploration is a process by which new alternatives can enter the choice-set. The probability of a location to be discovered is modeled as a function of attractiveness of the location regarding the attributes that are not satisfied by the alternatives within the current choice-set. Because agents are uncertain in this situation due to limited information, we propose to use the Boltzmann model (see [16]) to calculate discover probabilities across the universal choice set of locations and simulate outcomes of search processes:

$$P(L_i^t) = \frac{\exp(V_i^t / \tau)}{\sum_i \exp(V_i^t / \tau)} \tag{4}$$

Where V_i^t is a utility measure of location i and τ is a parameter determining the degree of uncertainty in the selection of new locations. The larger the value of the τ parameter is the more evenly probabilities are distributed across alternatives and, hence, the higher the uncertainty is, and vice versa. The parameter can be interpreted as the general (lack of) quality of information sources available to the agent, such as social network, public and local media and own observations during travel. V_i^t is a utility calculated based on true levels of attributes of locations. Note that the utility depends on the objective of the search: by including only those attributes that are dissatisfactory in the current best choice, V_i^t reflects the focus of the search. Furthermore, a disutility of travel distance is included in the function for V_i^t for two reasons: (1) the longer the travel distance is, the less likely information about the location is available and, (2) the longer the travel distance is, the less likely the location will be considered by the agent because of the higher generalized travel costs.

Having defined the discover probability distribution across locations across the universal choice set, Monte Carlo simulation will be used to select a new location that will be tried and may be added to the choice-set. Once tried, the new location receives an activation level reflecting memory trace strength and is subject to the same updating and learning process as other alternatives in the choice-set, as will be explained later.

In addition, a mental effort counter is included to prevent an agent from getting trapped in continuous and endless exploration. We assume that the agent will keep a record of how many consecutive times (s)he already tried exploring a new location under the same contextual conditions. Every time a choice is made through exploration, it will add 1 unit of mental effort. A habitual choice or an exploitation choice will break the chain of incrementing the score and restore it back to 0. We assume that when the mental effort involved in search for a better alternative is built up and exceeds a predefined threshold, instead of continuing exploring, the agent will avoid further frustration by lowering the aspiration level (realizing that the current aspiration level is not realistic). Therefore, in the choice process, before engaging in exploration, the system will check whether the accumulated mental effort exceeds this threshold. If this threshold is not exceeded, the agent will continue exploring. When it is exceeded, the agent will replace the current aspiration levels with the attributes levels of the alternative that currently has the highest expected utility, to assure a relatively optimal outcome and maintain high aspiration levels for future choices. As a consequence of choosing it, the activation level of this alternative will be increased.

As a consequence of the above mechanisms, an agent arrives at a selection of a single alternative location each time an activity is to be carried out. Depending on aspiration levels, this alternative could be the one that has the highest activation level (habitual choice), the one that has the highest expected utility (conscious exploitation choice), or the one that was newly discovered (conscious exploration choice).

Central to our dynamic process is the notion that choices are contingent upon the outcome of previous choices. By repeatedly making decisions, an agent acquires knowledge (learns) about the environment and thereby forms expectations about attributes of the environment. It should be noted that adaptation and learning processes involve two operations. One concerns updating an agent's perception of the environment. Through repeated experience, agents will update their expectation of attributes of locations (and routes), which are considered relevant for making choices, and discover conditions having an influence on outcomes. The other operation concerns the formation of habits to avoid the needless repetition of effortful memory retrieval and evaluation tasks. In this section, we will consider these two processes in turn, starting with habit formation.

A mechanism similar to reinforcement learning will be used for updating activation levels to simulate memory process. In line with evidence in cognitive psychology [17], the basic assumptions are that an alternative that has higher utility stays longer in memory, and that memory is reinforced when an alternative is chosen and memory decays if it is not chosen. Every time a location is chosen, the activation level of that location will be incremented to simulate the strengthening of a memory trace. The reinforcement rate is an increasing function of the experienced utility of the chosen location which in turn is a function of the location's attributes (as before). Limited memory retention capacity is simulated in the system by a parameter that determines rate of decay over time. If one alternative has not been chosen for some time, its activation level will decrease. When its activation level drops below some predefined threshold, it will be removed from the current choice-set to reflect the limited human ability of memory retrieval.

Formally, the strength of a memory trace of a particular activity location i in the choice-set is modeled as follows:

$$W_i^{t+1}(z_m) = \begin{cases} W_i^t(z_m) + \gamma U_i^t(z_m) & if\ I_i^t = 1 \\ \lambda W_i^t(z_m) & otherwise \end{cases} \tag{5}$$

Where $W_i^t(z_m)$ is the strength of the memory trace (awareness) of location i at time t under a configuration of conditions z_m and $I_i^t = 1$, if the location was chosen at time t, and $I_i^t = 0$, otherwise, $0 \le \gamma \le 1$ is a parameter representing a recency weight, which is relevant only when the location is chosen; and $0 \le \lambda \le 1$ is a parameter representing the retention rate. $U_i^t(z_m)$ is the experienced utility attributed to location i that is calculated based on experienced states of the attributes of location i, including both (quasi)-static and dynamic variables. The calculation (based on a utility function similar to the one represented by equation (1)), uses observed states of the dynamic attributes, such as crowdedness and travel time. Thus, at each time step the strength is reinforced or decays depending on whether the location has been chosen in the last time step. The coefficients γ and λ determine the size of reinforcement and memory retention respectively and are parameters of the system. Based on the current value of memory strength, the system determines whether or not the item is included in the choice-set in the next time step based on the simple rule stating that it is included if it exceeds a threshold level and is not included, otherwise. This rule can be written in the following general form:

$$\Phi^t(z_m) = \{L_i \,|\, W_i^t(z_m) \ge \omega\} \tag{6}$$

We assume that agents make personal observations and update their beliefs of their environment based on these observations in order to be able to make better predictions about what can be expected in the next time step. Each time a location is chosen when an activity is implemented, the agent updates beliefs $P^t(X_j|C)$, where C is the condition or, if multiple condition variables are involved, the condition configuration experienced. Learning implies two processes: parameter learning and structural learning. The first process involves incrementally updating the conditional belief distributions across the possible states for each observed attribute of the location after experiencing the actual states. The second process is aimed at discovering the conditions that have an influence on the likelihood of states of the system. Thus, the second process determines the form of the conditional probabilities that are kept up to date through the first process. This is done by periodically reconsidering splitting or merging condition states based on condition variables to update a tree structure that better predicts states based on observed outcomes. In the field of Bayesian Networks, the two processes are generally known as parameter and structural learning respectively.

We will adopt the approach proposed by Arentze and Timmermans [13]. In their approach, a method of parameter learning is used that is derived from Bayesian principles. Moreover, for structural learning, the proposed approach assumes a process of incrementally splitting and merging conditions based on events experienced in the past and stored in memory using some split criterion. In specific, the problem can be defined as a well-known problem considered by decision tree induction methods, namely as the problem of finding the most efficient way of splitting a set of known observations on predictor variables into partitions c_k that are as homogeneous as possible in terms of a response variable. For example in case of estimating the crowdedness of a location, the state of crowdedness is the response variable and time-of-the-day and day-of-the week serve as predictor variables. Then, the problem is to split the sample of observations on the condition variables such that observations within partitions are as homogeneous as possible in terms of crowdedness. Different criteria for finding the best splits, such as Chi-square or expected information gain can be used for this problem. Condition variables that are not significant in the current time step may become so at some next moment in time when more observations have been stored. Therefore, splitting and merging operations are periodically reconsidered. For more detail, readers are referred to [13]. The result of a structural learning step, generally, is that subsequent parameter learning is based on a new belief structure. The new conditional probabilities can be derived from the event base in a straight-forward way.

Agents are not isolated from each other, but participate in social networks. Participation in social networks may lead to adaptation of aspirations and diffusion of knowledge, which in turn may trigger changes in activity-travel choice behavior. Modeling the dynamic formation of social contacts between agents based on social links is beyond the scope of this paper (for a possible model of these processes, see [18]). In this section, we consider social links as given and focus on the impacts of social interactions on agents' aspiration levels and knowledge about activity locations.

According to social comparison theory, people often obtain information about their performance by comparing themselves to others [19]. Social comparison theory posits that people are generally motivated to evaluate their opinions and abilities and that one way to satisfy this need for self-evaluation is to compare themselves to others. Information gathered from these social comparisons can then be used to provide insights into one's capacities and limitations, which may motivate them to achieve higher goals since people are motivated to maintain or increase positive self-evaluation.

Following this theory, we assume that when two agents P_1 and P_2 meet, agent P_1 will evaluate and update his/her aspiration levels based on the best performances of agent P_2, if P_2 belongs to the reference group of P_1. More specifically, for each contextual condition of which agent P_1 has defined aspiration levels, P_1 will ask P_2's best performance. Agent P_2 will provide as feedback the attribute information of the alternative that has the highest expected utility within his/her choice-set under the corresponding conditions, since this alternative reflects his/her highest possible achievement given his/her current knowledge.

After receiving the information from agent P_2, agent P_1 first makes a decision on whether or not (s)he will change his/her aspiration levels. For this, P_1 compares the expected utility that is calculated using attributes values from agent P_2's answer and his/her own preferences with the expected utility that is derived from his/her current aspiration levels. We assume that only if a positive discrepancy between the two expected utilities exist (i.e., $U(P_2) > U(P_1)$) which exceeds a tolerance threshold of P_1, then P_1 is willing to update his/her aspiration levels; we say the agent is in an updating mode. If the discrepancy is not positive or the threshold is not exceeded, we assume that no adjustment will take place implying that P_1 will leave his/her aspiration levels unchanged. We assume that when in an updating mode, P_1 will upgrade the aspiration levels on those attributes on which the alternative conveyed by P_2 has the better value. Note that, updating aspiration levels may lead to a switch from a habitual to a conscious choice mode, which in turn may lead to exploration of new alternatives and, hence, adaptation of the person's choice-set.

Besides social comparison, when two agents P_1 and P_2 meet, P_1 will also update his/her knowledge by integrating the new information provided by P_2. In the system, P_2 presents a list of all the locations (s)he knows to P_1. After receiving the list from P_2, P_1 checks the list with his/her knowledge to find out if the list of P_2 includes alternatives that are new to him/her. Each location alternative that is unknown to P_1 activates P_2 to provide further information about the attributes levels of the location. Then, P_1 checks whether there are constraints (e.g., opening times, travel time) that limit the use of the new alternative, and add the new known alternative to condition-dependent choice-sets, if any, for which the new alternative is appropriate. When added to a choice-set, the new alternative is specified according to the attribute information conveyed by P_2 and the activation level is initialized. Once added, the new location is subject to the same selecting, updating and learning processes as other alternatives within the choice-set.

III. Discussion

This paper has outlined the conceptual framework that will be used to model the dynamic process of agents' activity location choice in a micro-simulation system. The framework considered dynamic formation of the choice-set. It integrates cognitive learning and social leaning. In the proposed approach, cognitive leaning focuses on updating beliefs about a non-stationary environment that will impact the expected utility of alternatives and habit formation, while social learning emphasizes on deriving and updating aspirations that may trigger re-evaluating currently known alternatives (exploitation) or search for new alternatives (exploration). As such, it provides a modeling approach for distinguishing habitual choice, exploitation choice and exploration choice.

References

[1] T.A. Arentze, and H.J.P. Timmermans, *Albatross: A Learning Based Transportation Oriented Simulation System*, EIRASS, Eindhoven, The Netherlands, 2001.

[2] T.A. Arentze, and H.J.P. Timmermans, "measuring impacts of condition variables in rule-based models of space-time choice behavior: method and empirical illustration," *Geo. An.*, 35, pp. 24-45, 2003.

[3] T.A. Arentze, and H.J.P. Timmermans, "A learning-based transportation oriented simulation system," *Transpn. Res.B*, 38, pp. 613-633, 2004.

[4] Joh, C.-H., Arentze, T.A. and H.J.P. Timmermans, "Measuring and predicting adaptation behavior in multi-dimensional activity-travel patterns," *Transportmetrica*, 2, pp. 153-173, 2006.

[5] H.J.P. Timmermans, T.A. Arentze and C.H. Joh, "Analyzing space-time behavior: new approaches to old problems," *Progress in Human Geography*, 26(2), pp. 175-190, 2002.

[6] D. Charypar and K. Nagel, "Generating complete all-day activity plans with genetic algorithms," *Proc. 10th Int. Conf. Travel Beh. Res.*, Lucerne, Switzerland (CD-Rom), 2003.

[7] M. Balmer, K. Nagel and B. Raney, "Large scale multi-agent simulations for transportation applications," *Proc. Beh. Responses to ITC*, Eindhoven (CD-Rom), 2004.

[8] R.J.F. Rosetti and R. Liu, "A dynamic network simulation model based on multi-agent systems," *Proc. 3rd Works. Agents in Traff. And Transpn.*, pp88-93, AAMAS, New York, 2004.

[9] G. Hertkort and P. Wagner, "Adaptation of time use patterns to simulated travel times in a travel demand model," in *Progress in activity-based analysis*, H. Timmermans, Ed. Amsterdam, pp. 161-174, 2005.

[10] G. Rindsfuser and F. Klugl, "The scheduling agent using sesam to implement a generator of activity programs," in *Progress in activity-based analysis*, H. Timmermans, Ed. Amsterdam, pp. 115-137, 2005.

[11] T.A. Arentze, H.J.P. Timmermans, D. Janssens and G. Wets, "Modeling short-term dynamics in activity-travel patterns: from Aurora to Feathers," Paper presented at the *innovations in travel modeling conference*, May 21-23, 2006, Austin, Texas.

[12] D. Janssens, G. Wets, H.J.P. Timmermans and T.A. Arentze, "Modeling short-term dynamics in activity-travel patterns: the Feathers model," Paper prepared for presentation at the *WCTR conference*, June 24-28, 2007, Berkeley, California.

[13] T.A. Arentze, and H.J.P. Timmermans, "Modeling learning and adaptation processes in activity-travel choice," *Transportation*, 30, pp. 37-62, 2003.

[14] J.W. Payne, D.J. Laughunn and R. Crum, "Translation of gambles and aspiration level effects on risky choice behavior," *Management Science*, 26(10), pp. 1039-1060, 1980.

[15] P. West and S. Broniarczyk, "Integrating multiple opinions: the role of aspiration level on consumer response to critic consensus," *Journal of Consumer Research*, vol. 25(1), pp. 38-51, 1998.

[16] R.S. Sutton and A.G. Barto, *Reinforcement Learning: An Introduction*. London, MIT Press, 1998.

[17] J.R. Anderson, *The Architecture of Cognition*. Cambridge, Harvard University Press, 1983.

[18] T.A. Arentze, and H.J.P. Timmermans, "Social networks, social interactions and activity-travel behavior: a framework for micro-simulation," In: *Proceedings of the 85th Annual Meeting of the Transportation Research Board*, Washington, D.C. (CD-ROM: 18 pp.). To be appeared, 2006.

[19] L. Festinger, "A theory of social comparison processes," *Human Relations*, 7, pp. 117-140, 1954.