DERIVING DECISION RULES FROM ACTIVITY DIARIES

THEO ARENTZE
Eindhoven University of Technology
Urban Planning Group
P.O. Box 513, 5600 MB Eindhoven
The Netherlands

FRANK HOFMAN
Ministry of Transport, Public Works and Water Management
P.O. Box 1031, 3000 BA Rotterdam
The Netherlands

HARRY TIMMERMANS
Eindhoven University of Technology
Urban Planning Group
P.O. Box 513, 5600 MB Eindhoven
The Netherlands

Abstract

This paper presents a system for deriving scheduling decision rules from activity diary data. The proposed system represents constraints as well as preferences in determining the sequence of a given set of activities. Rules are organised in a hierarchy and a systematic search procedure is used to optimise the rule hierarchy. A string alignment technique is used to measure the goodness-of-fit of the model in terms of an aggregate distance between observed and predicted schedules. The proposed system is tested based on a large-scale activity diary data set. The results suggest that the optimised rule-set achieves a considerable reduction of aggregate distance. This rule-based approach is suggested as an alternative to existing simultaneous choice models. Potentially, the system is more flexible in adapting schedules to physical or institutional changes in the environment.

INTRODUCTION

Activity-based models of trip generation and distribution have regained considerable popularity recently. Such models are generally viewed to potentially better capture the complex interrelationships between activity choice, choice of transport mode, destination choice and timing decisions. Many different approaches have been suggested to model these activity patterns. An overview is given in Ettema and Timmermans (1997).

One such model, currently under development by EIRASS/the Urban Planning Group and commissioned by the Dutch Ministry of Transportation and Public Works, is *ALBATROSS*. In this system, activity patterns are the outcome of a sequential decision making process in which individuals (i) select from a household-specific long-term calendar the activities that are to be conducted, (ii) allocate tasks among members of the household, (iii) determine for each individual a schedule in terms of the sequence in which activities are conducted (iv) determine the location of each activity and specify trips in terms of transport mode and route choice and (v) determine the exact timing of the activities. The outcome of this stage is an initial schedule. Next, schedule decisions can be revised either in the scheduling or implementation stage to optimise the schedule, adapt the schedule to unforeseen events and to allow for unplanned activities. This paper focuses on the third step. Given an activity program for a particular individual and a particular day, the (sub)system proposed in this study determines the sequence in which the activities are conducted.

Although it has been argued that individuals consider activity sequencing as an explicit decision (Gärling *et al*, 1989), it has not received much attention as such in current operational activity-based models. In the dominant approach, generating an activity pattern is modelled as a discrete choice between a given set of optional alternative patterns. The number of possible activity patterns soon becomes intractably large, however, so that simultaneous choice models, such as the ones proposed by Ben-Akiva and Bowman (1995), Ettema *et al.* (1997) and Algers *et al.* (1997), have to restrict the choice-set to only a small number of typical patterns. Such a reduction means, however, that these models are limited in describing and predicting in a sufficient degree of detail the variety of patterns one can observe across individuals and physical and institutional settings. The system proposed by Ettema *et al.* (1993, 1994) overcomes this problem by assuming that individuals evaluate decisions to add, delete and reschedule activities sequentially. However, Magic as this system is called was developed primarily to analyse cognitive processes and cannot be readily used for prediction purposes.

The system proposed in this paper represents an alternative approach aimed at overcoming shortcomings of simultaneous choice models. This new model is designed as a rule-based system. That is, activity patterns are explained and represented in terms of a set of heuristics. The use of rule-based systems in activity-analysis is rare, with notable exceptions of Gärling *et al.* (1989) and Vausse (1997), although to our knowledge their work never materialised in an operational model. Rule-based systems, however, have been developed for various other domains. Typically, the rules are derived from either expert knowledge or from protocols or similar knowledge-acquisition methods. We believe, however, that such rules are preferably derived from empirical data.

This paper therefore describes a modelling approach in which rules are derived from observed activity patterns. The remainder of this paper is structured as follows. Section 2 first discusses the proposed scheduling system and formats of the rules. Then, section 3 considers principles of the estimation method used to derive rules from activity diary data. Next, section 4 discusses the results

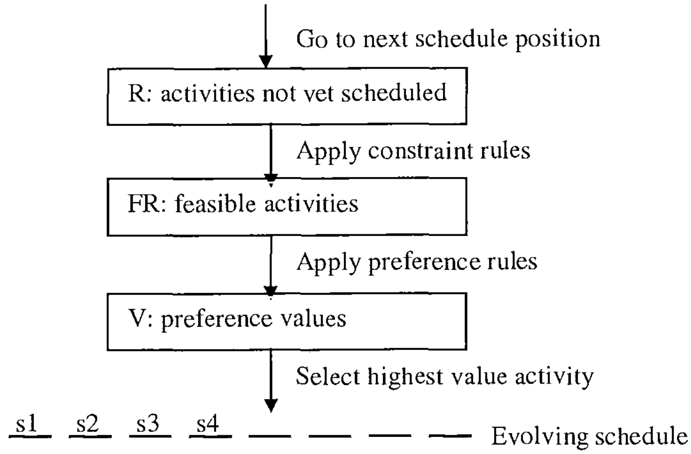

Figure 1. Scheduling algorithm

of an empirical study conducted to test the system. Finally, in the last section we discuss the potential of our approach and ways for further research.

THE SCHEDULING SYSTEM

The scheduling algorithm

The actual process of scheduling activities is conceptualised as a process in which an individual attempts to realise particular goals, given a variety of constraints that limit the number of feasible activity patterns. The rules that we use to model this process operates within a pre-defined scheduling algorithm. As schematically shown in Figure 1, the algorithm starts with an empty schedule and successively adds activities from a given list of activities called the activity program. More formally, let: AP is the activity program to be implemented on a given day by a particular individual; $R_k \subseteq AP$ is the subset of activities not yet scheduled in the k-th step of the scheduling process and $S_k \subseteq AP$ is the complementary subset of activities that has already been scheduled, i.e. $\forall k\ R_k \cap S_k = \varnothing$ and $R_k \cup S_k = AP$. Initially, $R_k = AP$ and $S_k = \varnothing$. A scheduling decision sd_k consists of adding an activity $a \in R_k$ to the end of list S_k. Formally, a scheduling decision defines the transformation:

$$S_k = \{ s_1, s_2, \ldots, s_k \} \bullet sd_k \Longrightarrow S_{k+1} = \{ s_1, s_2, \ldots, s_k, a \} \quad (1)$$

The selection of the next activity to add is controlled by a set of rules which represents scheduling constraints and preferences. The rules are organised in a hierarchy. From high to low order the rules successively reduce the set of activities until the most preferred activity for the next position is uniquely identified or until all rules have been applied. In the case of indifference, the selection of an activity is random. Formally: the rule at the j-th level of the hierarchy determines for the k-th scheduling step the subset $RS_{k,j} \subseteq RS_{k,j-1}$, if $RS_{k,j-1}$ contains two or more elements. The selected set is initialised at each scheduling step with the activities not yet scheduled, $RS_{k,0} = R_k$.

The present algorithm focuses on generating an initial schedule, which can be revised in later stages. Furthermore, we stress that the objective of the algorithm is merely to put a given set of activities into a sequence. We assume that decisions regarding the selection of activities (which?) and task-allocation (who?) have been made in earlier stages.

Constraint rules

At the highest level, the system uses rules representing constraints that determine the feasibility of a schedule. These include temporal and sequence constraints. Temporal constraints follow from given time-windows for each activity as well as the schedule as a whole. At present, we assume that the duration of each activity is given. This simplifying assumption will be relaxed in later versions of the system. Time-windows are defined in terms of the following variables:

t_a^{es} earliest start time of activity a
t_a^{ls} latest start time
t_a^{ee} earliest end time
t_a^{le} latest end time

The time-windows may reflect institutional constraints (e.g., opening hours of facilities), logical constraints (e.g., not having breakfast late in the evening), household constraints (bringing children to school at given times), social or coupling constraints (e.g., commitments with others to conduct activities jointly) etc.. As opening hours of facilities may vary across locations, the timing constraints of those activities that can be conducted at different locations are unknown in the sequencing stage. In such cases, the system derives a time-window based on a given location choice-set as the minimum and maximum end time across the optional locations.

Logical consistency requires that the following relationships hold:

$$t_a^{ls} \geq t_a^{es} \tag{2}$$
$$t_a^{le} \geq t_a^{ee} \tag{3}$$
$$t_a^{le} \leq t_a^{ls} + t_a^{d} \tag{4}$$
$$t_a^{ee} \geq t_a^{es} + t_a^{d} \tag{5}$$

where t_a^d is the duration of activity a. The 24-hour constraint is specified in terms of a time-window for the schedule as a whole. The start time of the schedule is denoted by t^s and the end time is given by t^e. Then, the 24-hour constraint requires:

$$t^e = t^s + 24 \tag{6}$$

assuming that time is measured in hours. The overall time-window does not only constrain the total duration of the schedule, but also the earliest begin and latest end times of each activity, i.e.:

$$t_a^{es} \geq t^s \tag{7}$$
$$t_a^{le} \leq t^e \tag{8}$$

The proposed rule at the first level deals with these temporal constraints and simulates adding each candidate activity in the current position. The rule rejects the activity if the earliest end time of the resultant chain exceeds the latest begin time of any other activity not yet scheduled or if the earliest end time of the chain plus the total duration of remaining activities is smaller than or equal to the schedule end time. To determine the earliest end time of a given chain this rule accounts for the time window and duration of each activity as well as the minimally required travel time between activities. However, the rule does not account for travel times and time windows of remaining activities, as these depend on how these activities are put into a sequence, which depends on later scheduling decisions. This means that the heuristic tends to overestimate feasibility of candidates.

The second rule in the hierarchy eliminates activities which do not meet certain sequence constraints. Sequence constraints are defined in terms of a set of Boolean variables, $S_{aa'}^{s}$, denoting whether an activity of type a can be performed before an activity of type a'. Generic rules for defining sequence constraints do not exist, as these constraints depend on the chosen typology of activities. In discussing the empirical study, we will describe how these constraints were specified in that particular application.

Preference rules

The rules at the subsequent levels of the hierarchy are different in the sense that they represent preferences which may vary between individuals. The specification and place in the hierarchy of these rules are not a-priori known, but rather are to be derived or 'estimated' from schedule data. As will be explained in the section, the estimation procedure relies on a system for exhaustively generating candidate rules. This section discusses the proposed rule formats and a method of deriving candidate rules. We use a classification of preferences on three major dimensions.

Time-of-day preferences
First, individuals may display time-of-day preferences for certain activities. For example, individuals may try to perform a necessary shopping activity at off-peak hours, to avoid congestion on the road or inside centres or stores. Preferences of this type can be readily represented in general form as:

> Assign a negative/positive value to candidate activity a for the current position, if a meets property P and the begin time falls in time range T.

A negative value represents a preference for delaying the activity, and a positive value indicates a preference for selecting the activity. P and T are parameters of the rules which are systematically varied to generate instances of the rule. In this rule and the rules below, we suggest that P is systematically varied based on classifications of activities. Activity dimensions such as location (e.g., in-home/out-home), activity type (e.g., leisure, mandatory), transport mode (e.g., car, bike), accompanying person (e.g., with others, alone), duration (e.g., long or small) and timing constraints (flexible or strict) provide a suitable basis for this.

Combination preferences
Second, positive or negative interactions between activities, such as carrying-over effects, may create negative or positive preferences for certain combinations. For example, individuals may wish to avoid a social visit directly after a sports activity to allow for a refreshment or redressing activity. Tendencies to cluster (positive interactions) or space (negative interactions) certain activities can be represented by a general rule of the form:

> Assign a negative/positive value to candidate activity a for the current position, if a meets property P and the last scheduled activity meets property Q

P and Q are parameters which are systematically varied to create possible instances of this rule. A positive value indicates a wish to cluster activities, whereas a negative value represents a desire to separate the activities. Dependent on parameter settings, this rule could for instance lead to clustering out-home activities. A more generalised version of this rule type is:

> Assign a negative/positive value to candidate activity a for the current position, if a meets property P and the *current activity chain* meets property Q

Substituting 'last scheduled activity' by 'current activity chain' in this rule reflects the idea that properties of entire chains can create negative or positive conditions for a next activity. For example, the current length of a trip chain can give rise to a preference for an in-home activity as the next activity (i.e., limiting the number of stops in a trip chain). As another example, the total time of engagement in mandatory activity can create a desire to schedule a leisure activity next (e.g., having a break after a work block of two hours).

Sequence Preferences
Rules of this type represent preferences for certain sequences per se, independent of time-of-day and combination preferences. Rules of this kind do not consider the current position of the schedule, but they determine the priority of candidate activities using the following general format:

> assign priority value z_1 if candidate activity a is an instance of class c_1
> assign priority value z_2 if candidate activity a is an instance of class c_2
> .....
> assign priority value z_n if candidate activity a is an instance of class c_n

where $z_1..z_n$ are preference values and $c_1..c_n$ denote mutually exclusive classes (on some dimension). The classification scheme $c_1..c_n$ is considered the parameter to be varied here. For example, assigning higher priority to in-home activities would reflect a preference for completing in-home activities before engaging in out-home activities. Or as another example, assigning higher priority to non-leisure activities would lead to a tendency to delay leisure activities until all non-leisure activities in the activity program have been scheduled.

Complex scheduling behaviour

Although these rule-formats are simple, when arranged in a rule-hierarchy they are capable of complex behaviour. At each step of the scheduling process, the highest level rules eliminate activities that do not fit in the current position given timing and sequence constraints. The preference rules, then, successively narrow down the set of remaining activities. For example, the system may display a preference for completing non-leisure activities such as for example shopping or household tasks before leisure activities such as for example watching tv or reading. At the same time, a more specific rule, at a higher level, may overrule this general preference and select a leisure activity after a block of non-leisure activities, and so on.

Probably, constraint rules apply in every case, as these rules represent basic scheduling logic. On the other hand, preference rules are possibly conditional upon characteristics of households, individuals, the environment in which the activities are conducted as well as the day of the week and contents of the activity program. For example, a preference to cluster out-home activities, e.g. scheduling shopping directly after work, may be selected only on days when the time pressure on the schedule is high, in the sense that the available time for leisure activities is limited. To model *conditional* application of preferences, we use the decision table. The advantage of this technique is that it allows one to verify and validate the completeness and consistency of the model (see Vanthienen *et*

C1	household type	I		II	
C2	time pressure on schedule	Low	High	Low	High
A1	in-home activities first	X	-	-	-
A2	cluster out-home activities	-	X	-	-
A3	mandatory activities first	-	-	-	X
A4	non-mandatory activities first	-	-	X	-

Figure 2. Arbitrary example of a decision table

al., 1994; Lucardie 1994). Figure 2 gives an example, whereby a preference rule is selected dependent on household type and time pressure. In general, the table structure allows one to define for which cases which preference rule applies.

THE ESTIMATION METHOD

The data used for estimation consists of the activity programs and related schedules of a group of individuals. Let AP_j^o and S_j^o be the observed activity program and schedule of the j-th individual, S_j^g be the generated schedule and $z(S_j^g, S_j^o)$ is a measure of dissimilarity (distance) between two sequences. Then, the objective of estimation is to find the subset of preference rules that minimises:

$$Z = \sum_j z\left(S_j^g, S_j^o\right) \tag{9}$$

The proposed estimation method assumes that it is possible to define an exhaustive set of possible preference rules (as well as the conditions under which they apply). In this section, we briefly discuss the principles of the proposed algorithm and similarity measure.

Estimation algorithm

As we have shown in an earlier study, an inductive learning algorithm could be used to identify the rules that best reproduces observed schedules (Arentze, Hoffman and Timmermans, 1997). Inductive learning involves incrementally adjusting weights of rules using the observed schedules as training examples. Each time when supplied with an activity program, the system selects a rule for each position in the rule hierarchy based on an error value associated with each rule. A probabilistic selection function is used such that rules with a lower error value have a higher chance of being selected. Each time a rule is used its error value is updated dependent on the degree of dissimilarity between the observed and generated schedule. The weight of earlier cases is a parameter of the error value update function. As learning proceeds, the error value of each rule will reflect a weighted average error score across all the cases in which the rule was applied. On a simulated data set it was shown that, when the rule selection and error update function was properly specified, the system was able to identify the rules that were used to generate the data. For more details we refer to the mentioned study.

The similarity measure

To measure the degree of similarity, z, between generated and observed activity sequences, we use a string alignment technique. The used method, referred to as String Alignment Method or SAM, was first introduced in biology for comparing DNA strings. Wilson (1996) proposed the use of SAM for comparing activity patterns. Recently, Joh *et al.* (1997) proposed an extension of SAM to deal with multi-dimensional activity patterns. However, for the present purpose we use the original uni-dimensional SAM as our problem at this stage is only one-dimensional.

SAM determines the degree of similarity between two strings of symbols in terms of the amount of effort required to make the two strings identical using insertion, deletion and substitution operations. SAM allows one to assign different weights to these operations. The weighted sum of operations required for alignment is known as Lehvenstein distance. Then, the measure is defined more specifically as the minimum Lehvenstein distance across optional paths to the solution.

Traditional Euclidean measures use a position-based comparison of elements and are, therefore, highly sensitive to mismatches caused by one arbitrary element. For example, inserting one element somewhere in a string may have a big impact in terms of Euclidean distance (as all positions after the insertion point are shifted to the right). In contrast, as the SAM is sensitive to sequential information only, the impact of an arbitrary element on Lehvenstein distance is never bigger than the effort associated with a single deletion operation. Furthermore, SAM can handle strings with different lengths using the same distance criterion.

The weights associated with insertion, deletion and substitution operations can be set dependent on the application. For the present application, there is no reason to use different weights. Therefore, we use a weight of one unit for both insertion and deletion and a weight of two units for substitution (being a combination of a deletion and insertion). Given these settings, the measure has a clear interpretation. For example, a distance value of 10 units means that alignment requires 10 times the equivalent of an insertion/deletion operation or 5 times the equivalent of a substitution.

EMPIRICAL APPLICATION

As a first step, only a limited set of initial rules was tested. In this stage, we did not test the inductive learning algorithm. In stead, we focused on a more fundamental question, namely whether the system is sensitive enough to identify rules for reproducing real schedules.

Data collection

The activity diary data used were collected through a paper-and-pencil survey among a sample of 1223 households in the Netherlands. Of each household, one or more individuals older than 18 years was asked to fill out a diary for two days. An open format was used, whereby individuals could specify activities in terms of activity type, start and end time, travel time, location, transport mode used to travel to the location and accompanying persons. The start time of each schedule was set to 3 AM and the end time to 3 AM the next day. This time frame was chosen to make sure that in almost all cases the first and last activity concerned a sleep-episode. Individuals could specify activity type in terms of a pre-defined list of 49 activity categories. Travel was not considered a separate activity, but rather a characteristic of an activity, namely the means to reach the activity location. Inconsistencies, errors and missing values were as much as possible corrected by using a set of deterministic rules. The rules were designed to improve the quality of the data with respect to time notation (24 or 12 hour basis), activity start time (including / excluding travel), reporting the return trip, overlapping time intervals due to simultaneous performance of different activities (e.g., eat and watch tv) etc..

The cleaned data-set was further prepared for analysis by (i) classifying activities into a smaller set of activity categories and (ii) specifying time constraints in terms of a time-window for each activity (Table 1). In-home activities were generally considered to be unconstrained. The only exceptions are eating and childcare for which limited flexibility around observed times was assumed. There were no individual-specific data available on social commitments prior to scheduling. Therefore, no constraints were specified for social visits nor for the presence of accompanying persons in other activities. Constraints of facility-based activities such as shopping, personal business and recreational out-home activities were set based on region-specific opening hours by facility type and by day of the week. As individual-specific data on constraints of out-home work activities were not available, the time-window of these activities were set to a small range centred around the observed start and end time (assuming some flexibility in timing). Other activities including medical visits, bring/get persons and the morning and evening sleep episodes were completely fixed on observed start and end times (assuming no flexibility).

Table 1 - Used activity classification

Fully constrained	partly constrained	Unconstrained
- work out-home	- grocery / service	- work /study in-home
- primary work / school	- grocery shopping	- household tasks
- voluntary work etc.	- postal, financial service	- leisure in-home
- union-based activities	- other personal business	- social visits
- bring / get persons	- non-grocery shopping	- tour (walk, bike, car)
- primary sleep	- sports and fitness	
	- other recreational activities	
	- secondary sleep	
	- eat (breakfast, lunch, diner)	

Activity durations were set to observed values. However, leisure in-home (e.g., watching tv, reading) and household tasks were treated different in this respect. Long duration episodes of these activities were split up into smaller units of maximally 30 minutes. This reflects the notion that these activities can often be interrupted by other activities, so that larger blocks reflect an explicit scheduling decision to cluster smaller units.

Results

The data set contained 2919 schedules and was subdivided into two sets of approximately equal size. One subset was used for estimation and the other one for validation. Thus, the estimation results discussed in this section are based on one halve of the data set only. For each observation, the activities in the schedule were put in a random order and rule-based models were tested on their ability to reproduce the sequence of the activities.

A relatively simple rule hierarchy was tested. The used hierarchy consisted of the following levels:

1. Rules evaluating temporal constraints;
2. Rules for determining the timing of eat episodes;
3. Rules representing sequence preferences based on activity type;
4. Rules representing sequence preferences based on activity duration.

The second level consists of a rule for timing of eat-episodes. These rules attempt to realise a preferred timing of breakfast, lunch, diner or in-between tea/coffee breaks (if any). Specifically, the rules assign a low priority to (i) an eat activity when the current time slot is too early and (ii) a non-eating activity when adopting that activity would imply a too long delay of eating. The third and fourth level consists of rules representing sequence preferences. Third level rules determined priority based on activity type. The fourth level was added to solve remaining ties, if any. The rule that was tested assigns a higher priority to longer duration activities.

With respect to timing, different rules were tested. Preferred eat times were defined either relative to waking up time (e.g., having lunch 4 hours after waking up) or in terms of absolute times (e.g., having lunch at 12 o' clock irrespective waking up time). Furthermore, different assumptions regarding the length of the acceptable range around an ideal time were tested. As it turned out, rules based on absolute timing and a 30 minute time interval gave the best goodness-of-fit. Significant differences between groups could not be identified. Sequence preference rules at the third level deserves more attention, as these rules determine the overall preference structure of schedules.

Activities were classified on two dimensions: mandatory versus leisure and out-home versus in-home. This resulted in four classes:

1. Mandatory out-home (MO: grocery/service, non-grocery shopping and sports);
2. Mandatory in home (MI: household tasks, touring);
3. Leisure out-home (LO: social visits and recreational activities);
4. Leisure in-home activities (LI).

Conceptually, sports activities belong to the group of leisure activities. However, a first analysis suggested that, in terms of priority ranking, sports activities are stronger associated with the mandatory out-home group. For the same reason, touring was considered a household task, because in many cases it seemed to involve letting out the dog(s) of the household (which may have an obligatory character). In total 8 rules were formulated. The rules differ in the choice of priorities on these dimension as well as the dimension having the highest weight. For example, the rule denoted as 'MI, MO, LI, LO' assigns higher priority to both mandatory and in-home activities and gives higher weight to the M/L dimension.

Table 2 - Performance of base-models by group: Average Lehvenstein distance across schedules

	FW, NC	FW, C<12	FW, C>12	MW	AW	NW, NC	NW, C<12	NW, C>12	WE, NC	WE, C<12	WE, C>12
1. MI, MO, LI, LO	3.454	3.185	3.543	5.905	7.487	6.557	9.155	7.116	5.868	7.155	6.182
2. MO, MI, LO, LI	3.258	3.052	3.333	5.333	6.974	6.244	8.969	7.163	5.576	6.602	5.849
3. LI, LO, MI, MO	4.722	4.385	5.162	8.000	9.180	9.954	12.957	12.674	7.840	9.512	8.515
4. LO, LI, MO, MI	4.650	4.282	5.029	7.619	9.333	9.718	12.671	12.558	7.500	9.236	8.273
5. MO, LO, MI, LI	3.268	2.993	3.219	5.357	7.231	6.466	9.081	7.372	5.821	6.634	5.879
6. LO, MO, LI, MI	4.175	4.119	4.667	6.905	9.026	9.099	12.422	11.814	7.208	8.764	7.970
7. MI, LI, MO, LO	3.650	3.304	3.829	6.286	8.000	6.985	9.615	7.837	6.132	7.545	6.606
8. LI, MI, LO, MO	4.639	4.489	5.143	8.024	9.128	9.916	12.857	12.535	7.689	9.480	8.303

LI: Leisure in-home activities;
LO: Leisure out-home activities;
MI: Mandatory in-home activities;
MO: Mandatory out-home activities.

NC: no children;
C < 12: youngest child is younger than 12 years;
C > 12: youngest child is 12 years or older;
FW: Full time Work: more than 6 hours this day;
MW: Morning Work: 3-6 hours dominantly in the morning;
AW: Afternoon Work: 3-6 hours dominantly in the afternoon;
NW: No Work: less than 3 hours;
WE: Weekend.

Schedules were grouped into 11 categories based on (i) household composition; (ii) total time engaged in out-home work, and (iii) weekdays/weekend. Goodness-of-fit of the rules were determined for each of the groups, to find out whether group differences exist. The results are shown in Table 2. The figures represent the average difference between observed and generated schedule across cases, given optimal specifications of the constraint and eat-timing rules.

Although scores vary considerably within groups dependent on length of schedules and time-windows of the activities, the relative performance of rules is remarkably stable across groups. For all groups, rules 3, 4 and 8 belong to the three worst performing rules, rules 1, 2 and 5 belong to the three best-performing rules and rules 6 and 7 take in a middle position. The rank seems to correlate strongly with the priority assigned to in-home leisure activities. Rules assuming a high priority perform poorly and rules assigning a low priority perform well. In-home leisure activities constitute the most frequently reported group and tend to be postponed until other activities have been completed. Rule 2 is the best fitting rule for nearly all groups and second-best otherwise (NW-NC, NW-C>12). In these latter groups, the difference with the best-fitting rule is not significant. This finding suggests that the same rule is suitable for reproducing the main structure of schedules, irrespective household composition, work characteristics and day of the week. This rule prioritises activities in the order of 'MO, MI, LO, LI'. This suggests that individuals generally try to complete mandatory activities before leisure activities and, at a second level, out-home activities before in-home activities.

This rule was taken as a starting point for further refinement. More specific rules were tested describing different ways of ranking activities within the MO and LO categories. The variant maximising overall fit prioritises MO activities in the order 1. grocery/service, 2. non-grocery shopping and 3. sports and LO activities in the order 1. social visits and 2. recreational activities. Taking all groups together (N = 1467), the fit of this model in terms of aggregate distance across generated and observed schedules equals 8160 (average m = 5.576 and standard deviation s = 4.061).

Table 3 - Model fit on the estimation set (N = 1467, m = 14.77, s = 5.60)

Model	Z (total dist.)	m_i (average)	s_i (st.dev.)	m_i-m_{i+1} (reduction)	m_1=100 (% reduction)	t_i-value
Random sequence	17608	12.0027	5.283			
Constraints	12216	8.3272	4.922	3.676	69.4	19.496
Timing eat	11382	7.7587	4.844	0.568	64.6	3.153
Type sequence	8160	5.5624	4.0261	2.196	46.3	13.355
Duration sequence	8006	5.4574	3.9651	0.105	45.5	0.712

System performance and validation

There are several ways to analyse and evaluate the performance of the optimised rule-set. As an indicator of relative performance, Table 3 shows reductions in aggregate and average distance value, when layers of rules are sequentially added to the system. The zero case is defined by using no rules, i.e. schedules are generated by arranging activities in a random order. In this case, the aggregate distance equals 17608 units (m = 12.003). Adding constraint rules leads to a reduction to 69.4 % of this base level (m = 8.327). The t-value of difference of difference in means indicate that the improvement in fit attributable to constraint rules is highly statistically significant (t_1 = 19.496 against $t_{\cdot=0.95}$ = 1.645). It should be noted that this improvement is not achieved independent of observation. Specifically, the time-constraints of medical visits, work out, bring/get activities were set dependent on observed start and end times. Although it cannot be fully interpreted in terms of performance, this figure does indicate, however, that under present specifications there is still much freedom of choice left for preference rules at lower levels.

Adding rules for timing eat-episodes from this point results in a further reduction to 64 % of the base level (t_2 = 3.153). Next, adding the optimal sequence preference rule at the third level further reduces the distance value to 46.3 % (t_3 = 13.355). Finally, the complete system, which results when ties are solved by ranking activities according to decreasing duration, achieves a reduction to 45.5 % (t_4 = 0.712).

Table 4 shows the results when the same set of rules is applied to the validation set. As these figures show, the distance reductions related to each layer are of the same magnitude as for the estimation set. This suggests that the rules are transferable to a different set of observations.

Finally, to give a visual impression of the system's ability to predict schedules, Table 5 shows observed and generated schedules for a representative case.

Table 4 - Model fit on the validation set (N = 1452, m = 15.18, s = 5.66)

model	Z (total dist.)	m_i (average)	s_i (st.dev.)	m_i-m_{i+1} (reduction)	m_1=100 (% reduction)	t_i-value
random sequence	17934	12.3512	5.327			
constraints	12758	8.7865	5.112	3.565	71.1	18.399
timing eat	11566	7.9656	4.842	0.821	64.5	4.443
type sequence	8546	5.8857	4.242	2.080	47.7	12.312
duration sequence	8342	5.7452	4.169	0.141	46.5	0.900

Discussion of results

The results suggest that compared to a null-model a strong reduction in aggregate error can be achieved even by a limited set of rules. In this case, a large part of the reduction is attributable to preference rules. Furthermore, the same set of rules seems to apply to all groups. This is not to say that there are no differences between the activity patterns, but that such differences are given by the activity programs rather than sequencing rules.

Table 5 - Typical example of a generated and observed schedules (n = 13, distance = 4)

	begin time constraints	end time constraints	begin time	end time	duration (minutes)
1. sleep	3.00- 3.00	7.00- 7.00	3.00	7.00	240
2. eat or drink	6.00- 8.00	6.30- 8.30	7.00	7.30	30
3. work, out	3.00-26.00	4.00-27.00	8.00	9.00	60
4. work, out	3.00-26.00	4.00-27.00	9.00	10.00	60
5. social visit, out	3.00-24.30	5.30-27.00	10.30	13.00	150
6. work, out	12.45-13.45	16.30-17.30	13.15	17.00	225
7. household tasks	3.00-26.00	4.00-27.00	18.30	19.30	60
8. leisure, in	3.00-26.30	3.30-27.00	20.00	20.30	30
9. leisure, in	3.00-26.30	3.30-27.00	20.30	21.00	30
10. leisure, in	3.00-26.30	3.30-27.00	21.00	21.30	30
11. leisure, in	3.00-26.30	3.30-27.00	21.30	22.00	30
12. leisure, in	3.00-26.30	3.30-27.00	22.00	22.30	30
13. sleep	22.30-22.30	27.00-27.00	22.30	27.00	270

Observed string:
SLE-EAT-WORKO-WORKO-SVIS-WORKO-HHT-LEII-LEII-LEII-LEII-LEII-SLE
Predicted string:
SLE-EAT-WORKO-WORKO-HHT-WORKO-SVIS-LEII-LEII-LEII-LEII-LEII-SLE

The present four-level system provides a basis for elaboration through adding levels to the hierarchy. A follow-up study will be concerned with elaborating the system by adding interactions and time-of-day-preference rules to the system.

CONCLUSIONS AND DISCUSSION

This paper described a rule-based model for predicting activity sequences given an activity program for a particular day. In contrast to existing activity-based models, the proposed system assumes a sequential decision making process and uses rules for representing scheduling constraints as well preferences. Where constraint rules are a-priori specified, preference rules are derived from observations. An empirical case-study suggests that optimising only a limited set of preference rules already leads to a considerable reduction in aggregate error compared to a null model. This result could be repeated on a subset of observations that was not used for estimation.

We feel that the proposed approach is potentially useful, but that further research is required to improve its value for prediction and analysis purposes. First fundamental research is required to develop ways to improve system's sensitiveness for identifying rules. The current string alignment method (SAM) is sensitive for differences in sequential information. Through combinations of deletion, insertion and substitution operations the source string is made identical to the target string and the total amount of effort is taken as a measure of similarity. When applied to the same element, a combination of a deletion and insertion can be interpreted as replacing that element to another location in the string. The SAM treats deletions and insertions separately and does not take the distance between existing and target location into account. In biological applications, this is reasonable because differences between DNA strings are considered to be the result of independent mutations. In the context of activity patterns, however, one may argue that keeping everything else constant dissimilarity between activity patterns increases with the distance over which elements have to be replaced. An extension of the similarity measure that takes such distance information as well as sequence information into account would provide more informative feed back to the scheduling system. Then, a model would perform better the closer it locates activities to target positions. We expect that such a refinement would significantly improve the ability of the system to identify patterns in data.

Second, the classification of activities deserves attention. The classification used in the present study was based on commonly used definitions in the literature. Yet, some of the activity categories may comprise activities which are treated quite different by individuals. For example, the sports category included activities such as fitness and team sports, while individuals may display different time-of-day or clustering preferences in scheduling these activities. Future research should focus on using statistical techniques to find homogeneous groups of activities from a scheduling point of view. Furthermore, future research will be concerned with developing and testing additional levels of the rule-hierarchy.

REFERENCES

Algers, S., Daly, A. and Widlert, S.(1997) Modelling travel behaviour to support policy making in Stockholm. In P. Stopher and M. Lee-Gosselin (eds.), **Understanding Travel Behaviour in an Era of Change**, pp. 547- 570. Pergamon, Guildford.

Arentze, T.A., Hofman, F. and Timmermans, H.J.P. (1997) Estimating a rule-based system of activity scheduling: A learning algorithm and results of computer experiments. Paper presented at the Informs San Diego Conference, San Diego, US.

Ben-Akiva, M.E. and Bowman, J.L. (1995) Activity-based disaggregate travel demand model system with daily activity schedules. Paper presented at the Workshop on Activity Based Approaches, 25-28 May 1995, Eindhoven, The Netherlands.

Ettema, D.F. and Timmermans, H.J.P. (1997) Theories and models of activity patterns. In D.F. Ettema and H.J.P. Timmermans (eds.), **Activity-Based Approaches to Travel Analysis**, pp. 1-36. Elsevier Science, Oxford.

Ettema, D., Daly, A., De Jong, G. and Kroes, E. (1997) Towards an applied activity based travel demand model. Paper presented at the IATBR-conference, Austin, Texas, USA.

Ettema, D.F., Borgers, A.W.J. and Timmermans, H.J.P. (1993) A simulation model of activity scheduling behaviour. **Transportation Research Record 1413**, 1-11.

Ettema, D.F., Borgers, A.W.J. and Timmermans, H.J.P. (1994) Using interactive computer experiments for identifying scheduling heuristics. In **Proceedings of the 7th International Conference of the Association for Travel behaviour Research**, Santiago, Chile.

Gärling, T., Brännäs, K., Garvill, J., Golledge, R.G., Gopal, S. , Holm, E. and Lindberg, E. (1989) Household activity scheduling. In **Transport Policy, Management and Technology Towards 2001: Selected Proceedings of the Fifth World Conference on Transport Research 4**, 235-248, Western Periodicals, Ventura.

Joh, C.H., Arentze, T.A., Hofman, F., and Timmermans, H.J.P. (1997) Activity pattern similarity: Towards a multidimensional sequence alignment. Paper presented at the IATBR Meetings at Austin, Texas.

Lucardie, G.L. (1994) **Functional Object-Types as a Foundation of Complex Knowledge-Based Systems**, Ph.D.-Dissertation, Eindhoven University of Technology, Eindhoven, The Netherlands.

Vanthienen, J. and Wets, G.(1994) From decision tables to expert system shells. **Data and Knowledge Engineering 13**, 265-282.

Vause, M. (1997) A rule-based model of activity scheduling behaviour. In D.F. Ettema and H.J.P. Timmermans (eds.), **Activity-Based Approaches to Travel Analysis,** pp. 73-88, Elsevier Science, Oxford.

Wilson, C. (1996) Activity pattern analysis using sequence alignment methods. Paper presented at the Conference of International Association of Time Use Research.