Lane-Changing and Other Discrete-Choice Situations

Introduction

Simulating any nontrivial traffic situation requires describing not only acceleration and braking but also lane changes. When modeling traffic flow on entire road networks, additional discrete-choice situations arise such as deciding if it is safe to enter a priority road, or if cruising or stopping is the appropriate driver's reaction when approaching a traffic light that is about to change to a red. This lecture presents a unified utility-based modeling framework for such decisions at the most basic operative level.

From the driver's point of view, there are three main actions that directly influence traffic flow dynamics: Accelerating, braking, and steering. The dynamics of steering are part of the vehicle dynamics and therefore the domain of sub-microscopic models.

Traffic flow dynamics describe the dynamics one level higher by directly modeling lane-changing decisions and the associated actions. At this level, the set of possible actions is discrete, i.e., performing a lane change, or not. Details of the lane-changing maneuver such as duration or lateral accelerations are not resolved, and the lane-changing itself is assumed to take place instantaneously.

Discrete decisions and actions can also pertain to the longitudinal dynamics, in parallel to the continuous actions modeled by the acceleration function $a_{mic}(s, v, v_l)$: When approaching a yellow traffic light that is about to turn red, the driver has to decide whether it is safe to pass this traffic light without changing speed, or if it is necessary to stop.

Furthermore, lane changes generally influence the longitudinal acceleration of the decision maker (e.g., preparing for a lane change) or that of the other affected drivers (e.g., cooperatively making a gap to enable a change, or restoring the safety gap afterward). Discrete choices in the traffic-flow context involve several levels:

1. The strategic level (destination choice, mode choice, and route choice) is modeled within the domain of transportation planning.
2. The tactical level includes anticipatory measures to enable or facilitate operative actions such as changing lanes or entering a priority road. This includes cooperative behavior such as allowing another vehicle to merge at a point of lane closure (mode merging). Modeling the tactical level is notoriously difficult and is only attempted in the most elaborate commercial simulators.
3. On the operative level, the actual decision is made.
4. Finally, in the post-decision phase, the actions pertaining to this decision are simulated, e.g. performing the lane change or keeping to one's lane, waiting or entering a priority road, or cruising versus stopping at the traffic light.

In this lecture, we restrict the description to the operative level and the post-decision phase. We model the different discrete-choice situations consistently in terms of maximizing utility functions

associated with each alternative. The utility of a given alternative increases with the (hypothetical) longitudinal acceleration that would be possible once this alternative had been adopted. Using accelerations as a utility ensures the compatibility between the acceleration and discrete-choice models. Furthermore, this is parsimonious since it minimizes the number of parameters and assumptions.
For example, when the acceleration model is parameterized to simulate aggressive drivers, the lane-changing style of these drivers becomes aggressive as well, without introducing further parameters. Generally, any aspect considered in the longitudinal model carries over to the decision model. Specifically, the lane-changing considerations take into account speed differences, brake lights, or anticipative elements if, and only if, these exogenous factors are included in the acceleration model.

General Decision Model
We assume that, at a given moment, the driver can choose from a discrete set K of alternatives k. In the context of lane changes, the alternatives would be the (active) decision to change to the left or right, and the (passive) decision not to change.
When about to enter a priority road, the alternatives would be to initiate the merging, or stop and wait for a sufficient gap between the main-road vehicles. We assume that the drivers are aware of the consequences of their decisions, i.e., they can anticipate, for each alternative, the speeds and gaps of all involved vehicles.
This allows us to calculate all relevant accelerations (i.e., the utilities) using the normal acceleration functions of these vehicles. If the acceleration model is formulated as an iterated map or cellular automaton. In the decision process, the driver maximizes his or her utility (incentive criterion) subject to the condition that the action is safe (safety criterion). Both criteria are based on the acceleration function as follows:

Safety criterion. None of the drivers β affected by the consequences of opting for alternative k (including the decision maker α) should be forced to perform a critical maneuver as a consequence of a decision for alternative k. A maneuver is deemed to be critical if it entails braking decelerations exceeding the *safe deceleration* b_{safe}:

$$a_{mic}^{(\beta,k)} > -b_{safe} \quad (1)$$

The value of the model parameter b_{safe} (of the order of 2 m/s^2) is comparable to the comfortable deceleration b of the IDM or Gipps' model. The safe deceleration can be inherited from these models (b_{safe} = b), if applicable.

Incentive criterion. Choosing among all safe alternatives $\acute{k}$, the driver α selects the option of maximum utility U:

$$k_{selected} = arg \max_{\acute{k}} U^{(\alpha,\acute{k})} \quad (2)$$

As in most other discrete-choice models, the incentive criterion is based on a rational decision maker (also called homo oeconomicus) who maximizes his or her utility. In the simplest case, the utility is directly given by the acceleration function:

$$U^{(\alpha,k)} = a_{mic}^{(\alpha,k)} \tag{3}$$

In contrast to the standard framework for discrete decisions (multinomial Logit and Probit models and their variants), we do not assume explicit stochastic utilities unless the acceleration model itself contains stochastic terms

For some discrete-choice situations such as discretionary lane changes, one needs an additional threshold preventing all active decisions (e.g., a decision to change lanes rather than to stay put) when the associated advantage is only marginal. Such a threshold prevents unrealistically frequent withdrawals of an active decision taken in the last time step which could, for example, lead to frantic lane-changing actions.

Traffic rules (such as a "keep right" directive) may also enter the utility. Finally, one can include all consequences of a decision to other drivers by introducing a politeness factor.

Lane Changes

Figure 1 depicts the general situation. The vehicle α of the decision maker (speed v_α) is located in the center. There are three alternatives:

- Change to the right,
- change to the left, and
- No change.

Without loss of generality, we compare only the last two alternatives. Here, and in the following, we denote the vehicle of the decision maker with α, the leading vehicle with l, and the following vehicle with f. All accelerations, gaps, or vehicle indices with a hat refer to the new situation after the lane change has been completed while quantities without a hat denote the old situation.

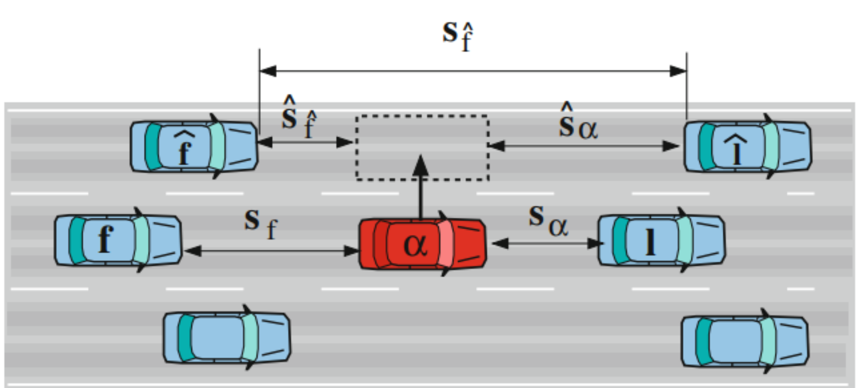

Fig. 1 Notation for a lane change of the center vehicle α to the left. All quantities with a hat pertain to the new situation after the (possibly hypothetical) lane change.

Safety Criterion

Assuming that the present situation (i.e., the alternative "no change") is safe, the safety criterion Equ. (1) Refers to the acceleration $\hat{a}_{\hat{f}}$ of the new follower ($\beta = \hat{f}$) after a possible change, and also to the new acceleration $\hat{a}_{\alpha}$ of the decision maker him or herself ($\beta = \alpha$). For the follower, this criterion becomes:

$$\hat{a}_{\hat{f}} = a_{mic}\left(\hat{s}_{\hat{f}}, v_{\hat{f}}, v_{\alpha}\right) > -b_{safe} \qquad safety\ criterion \tag{4}$$

In order that this condition also prevents lane changes whenever there are following vehicles on the target lane at nearly the same longitudinal position (the gap $\hat{s}_{\hat{f}}$ is negative, i.e., a change would result in an immediate accident), the acceleration function $a_{mic}(s, v, v_l)$ should return prohibitively negative values if $s < 0$.

The parameter b_{safe} indicates the maximum deceleration imposed on the new follower which is considered to be safe. If one simulates heterogeneous traffic with individual acceleration functions, the acceleration function $\hat{a}_{\hat{f}}$ of the new follower $\hat{f}$ is calculated with the function and parameters of this driver-vehicle unit.

Regarding the safety of the decision maker him- or herself ($\beta = \alpha$), condition (4) prevents changes if the new gap $\hat{s}_{\hat{\alpha}}$ is dangerously low such that $\hat{a}_{\alpha} = a_{mic}(\hat{s}_{\hat{\alpha}}, v_{\alpha}, v_{\hat{l}}) > -b_{safe}$.

In any case, the condition on the acceleration function to return prohibitively negative values for negative gaps guarantees that changes are prohibited if the leader on the target lane is essentially at the same longitudinal position ($\hat{s}_{\alpha} < 0$) which would result in an immediate crash.

Incentive Criterion for Egoistic Drivers

Most lane-changing models formulate the incentive criterion exclusively from the perspective of the decision-maker ignoring the advantages and disadvantages to the other drivers. Furthermore, lane-changing behavior depends on the legislative regulations of the considered countries. For example, a right-overtaking ban is in effect on most European highways. Here, we will restrict to the simpler situations of lane changes on highways, or more generally to lane changes in city traffic, where lane usage is only mildly asymmetric. Then, the incentive criterion for the egoistic driver reads:

$$\hat{a}_{\alpha} - a_{\alpha} > \Delta a + a_{bias} \tag{5}$$

$$a_{\alpha} = a_{mic}(s_{\alpha}, v_{\alpha}, v_l) \quad \hat{a}_{\alpha} = a_{mic}(\hat{s}_{\alpha}, v_{\alpha}, v_{\hat{l}}) \tag{6}$$

The lane-changing threshold Δa prevents lane changes when the associated advantage is only marginal (Table 1). Furthermore, the constant weight a_{bias} introduce a simple form of asymmetric behavior. If a keep-right directive is to be modeled, a_{bias} would be positive for changes to the left and reverses its sign for changes to the right. This contribution should be relatively small ($|a_{bias}| \ll b_{safe}$) but greater than Δa. Otherwise, vehicles would not change to the right lanes if the highway was essentially empty (see Table 1).

Table 1: Parameters of the lane-changing models.

Parameter	Typical value
Limit for safe deceleration b_{safe}	2 m/s^2
Changing threshold Δa	0.1 m/s^2
Asymmetry term (keep-right directive) a_{bias}	0.3 m/s^2
Politeness factor p (MOBIL lane-changing model)	0.0–1.0

The parameters b_{safe} and Δa apply to any changing model, $a_{\text{bias}} \neq 0$ only if asymmetric driving rules are to be modeled, and $p \neq 0$ if the drivers are not purely egoistic

Lane Changes with Courtesy: MOBIL Model

The changing conditions (4) and (5) characterize purely egoistic drivers who consider other drivers only via the safety criterion. If the lane change is mandatory as in lane-closure or merging situations, this behavior is plausible (and, additionally, the changing threshold Δp = 0). On the other hand, if the lane change is not necessary (also termed a discretionary lane change), most drivers refrain from changing lanes if their own advantage is disproportionally small compared to the disadvantage imposed on others, even if the safety criterion is satisfied. This can be modeled by augmenting the balance of the incentive criterion with the utilities of the affected drivers, weighted with a politeness factor p,

$$\hat{a}_\alpha - a_\alpha + p\left(\hat{a}_{\hat{f}} - a_{\hat{f}} + \hat{a}_f - a_f\right) > \Delta a + a_{bias} \qquad MOBIL\ incentive \qquad 7$$

For the special case when politeness p = 1 (corresponding to a rather altruistic driver), no bias (a_{bias} = 0), and negligible threshold (Δa = 0), a lane change takes place if the sum of the accelerations of all affected vehicles increases by this maneuver. Hence the acronym for this model:

MOBIL—minimizing overall braking deceleration induced by lane changes.

The central component of the MOBIL criterion is the politeness factor indicating the degree of consideration of other drivers if there are no safety restraints. Since a degree of consideration amounting to p = 1 is rare (which would correspond to “Love thy neighbor as thyself”), sensible values are of the order 0.2.

Application to Car-Following Models

The general lane-changing criteria presented above return explicit rules only when combined with a longitudinal acceleration model. In principle, the safety criterion (4) and the incentive criteria (5) or (7) are compatible with any longitudinal model providing the acceleration function a_{mic} either directly (time-continuous car-following models) or indirectly via Eq. (11) (time-discrete iterated coupled maps) that explained in the previous lecture.

When applying the safety criterion (4) to any acceleration model satisfying the general plausibility conditions discussed, we obtain a minimum condition for the lag gap $\hat{s}_{\hat{f}}$ of the new follower behind the changing vehicle on the new lane,

$\hat{s}_{\hat{f}} > s_{safe}(v_{\hat{f}}, v_{\alpha})$ 8

The safe gap function $s_{safe}\ (v_f, v)$ is obtained by solving the equation defining the marginal safety of the follower,

$a_{mic}(s_{safe}, v_f, v) = -b_{safe}$ 9

For the gap s_{safe}. Notice that a unique solution s_{safe} exists by virtue of the plausibility condition (2) in the previous lecture stating that, in the interaction range, the function a_{mic} increases strictly monotonically with respect to s. This means the safety criterion allows changes if the following gap (lag gap) on the target lane is greater than some minimum value depending on the speeds of the changing vehicle and the new follower $\hat{f}$, i.e., the safety criterion becomes a generalized gap-acceptance rule for the lag gap.
Similarly, the general incentive criterion (5) of egoistic drivers can be written as a generalized gap-acceptance rule for the lead gap of the changing vehicle on the new lane,

$\hat{s}_{lead} = \hat{s}_{\alpha} > s_{adv}(s_{\alpha}, v_{\alpha}, v_l, v_{\hat{l}})$ 10

The advantageous gap function $s_{adv}(s_{\alpha}, v_{\alpha}, v_l, v_{\hat{l}})$ is obtained by solving the equation:

$a_{mic}(s_{adv}, v, v_{\hat{l}}) - a_{mic}(s, v, v_l) = \Delta a + a_{bias}$ 11

Defining a marginal change of utility, for s_{adv}. Again, condition (2) in the previous lecture ensures that a unique solution s_{adv} exists if
$a_{mic}(s, v, v_l) + \Delta a + a_{bias} < a_{free}$ (v)
Where:
a_{free} (v) = a_{mic} (∞,v, v_l) is the free-flow acceleration function.
In contrast to the safety condition, however, this is not always satisfactory. Then, s_{adv} is not unique or even does not exist. Obviously, this corresponds to an infinite advantageous gap reflecting the fact that there is no need to change lanes because one can either drive freely on the old lane, or there is an obstruction but it is so small that the finite threshold $= \Delta a + a_{bias}$ prevents lane changing for marginal utility improvements, even if the target lane is free. In the following, we discuss the application of three specific longitudinal models.

Approaching a Traffic Light
When approaching a signalized intersection and the traffic light switches from green to yellow, it is necessary to decide whether it is better to cruise over the intersection with unchanged speed, or to stop (Fig. 2). This can be modeled within the general discrete-choice framework situation, the decisions are determined by the safety criterion alone: “Stop if it is safe to do so”. In our general framework, the decision to stop is considered a_{safe}if the anticipated braking deceleration will not exceed the safe deceleration b_{safe} at any time of the braking maneuver. For models with a plausible braking strategy, it is sufficient to consider the braking deceleration for this option at decision time. To calculate this deceleration, we model the traffic light as a standing virtual vehicle

(v_l = 0, Δv = v, desired speed v_0 = 0) of zero extension such that s denotes the distance of the front bumpers to the stopping line. This results in the simple rule,

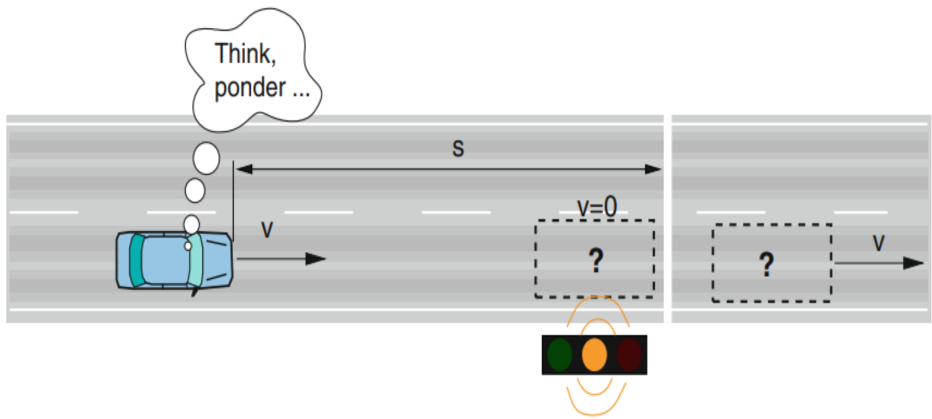

Fig. 2 Illustration of the decision to stop or to cruise at a traffic signal about to go red.

Cruise if $a_{mic}(s, v, v) < -b_{safe} \Leftrightarrow s < s_{crit}(v)$ 12

Stop otherwise.

Obviously, the critical distance s_{crit} where the decision changes is a special case of the safe gap function (8) of the safety criterion,

$$s_{crit}(v) = s_{safe}(v, 0) \quad 13$$

It is particularly instructive to apply this rule to the IDM safe gap for the common situation when the driver approaches the signalized intersection at his or her desired speed, and the IDM parameters satisfy approximately a = b = b_{safe}. In this case, $s_{crit}(v)$= s∗(v, v) is equal to the dynamic desired IDM gap for Δv = v, and condition (24) becomes,

Cruise if $s < s^* = s_0 + v_0 T + \frac{v_0^2}{2b}$ 14
Stop Otherwise

When setting the desired time gap T equal to the driver's reaction time, this means that one stops if, at decision time, the distance to the stopping line is greater than the stopping distance. This is perfectly consistent since this distance (which we have already introduced when formulating Gipps' model) is necessary to stop in a controlled way taking into account reaction time.

Figure 3 shows that also for more general parameter settings, the critical distance increases quadratically with the speed, while the critical time-to-collision (TTC) (here defined as the time to reach the stopping line at an unchanged speed) increases essentially linearly. We emphasize that the critical TTC of 3 s for 50 km/h, and 4 s for 70 km/h is consistent with European legislative regulations for the minimum duration of yellow phases of traffic lights on streets with the respective speed limits.

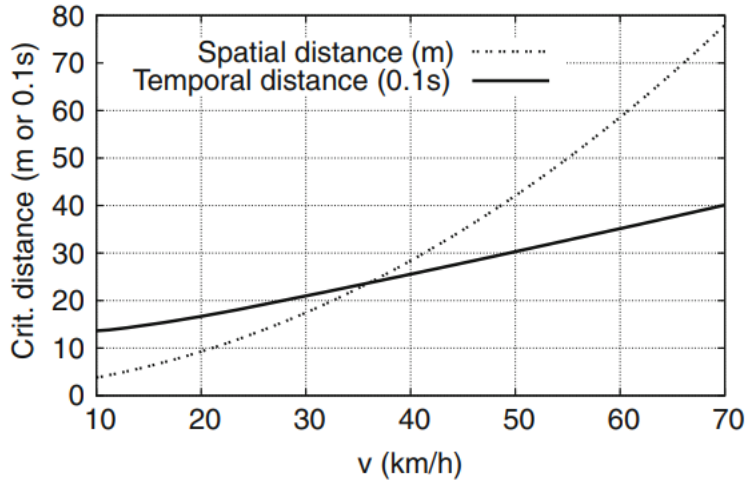

Fig. 3 Critical distance to the stopping line of a traffic light at decision time as a function of the speed for the IDM with b = 2 m/s^2, b_{safe} = 3 m/s^2, and further IDM parameters. Also shown is the associated TTC value.

Entering a Priority Road

This situation can be considered a special case of mandatory lane-changing decisions:

- The (nearest lane of the) main road corresponds to the target lane of the lane-changing situation.
- The merging action to enter the road corresponds to the lane-changing maneuver.
- The speed of the merging/lane-changing vehicle is very low or zero (the latter is true if there are stop signs, or the entering vehicle is already waiting).
- And the incentive criterion is always satisfied.

In contrast to normal (discretionary) lane-changing decisions, entering a priority road implies two safety criteria, one for the new follower and one for the merging vehicle itself. The latter was not necessary for discretionary lane changes since, there, a fulfilled incentive criterion automatically implies safety for the decision maker him or herself. When formulating the criteria, we assume that the driver of the merging vehicle can anticipate his or her speed v_α, the speedsv_f and v_l of the follower and leader, and the corresponding gaps s_f and s_l, respectively, at merging time (Fig.4).

$$merge \quad if\ s_f > s_{safe}(v_f, v_\alpha) AND\ s_{lead} > s_{safe}(v_\alpha, v_l) \tag{15}$$

Stop or wait otherwise.

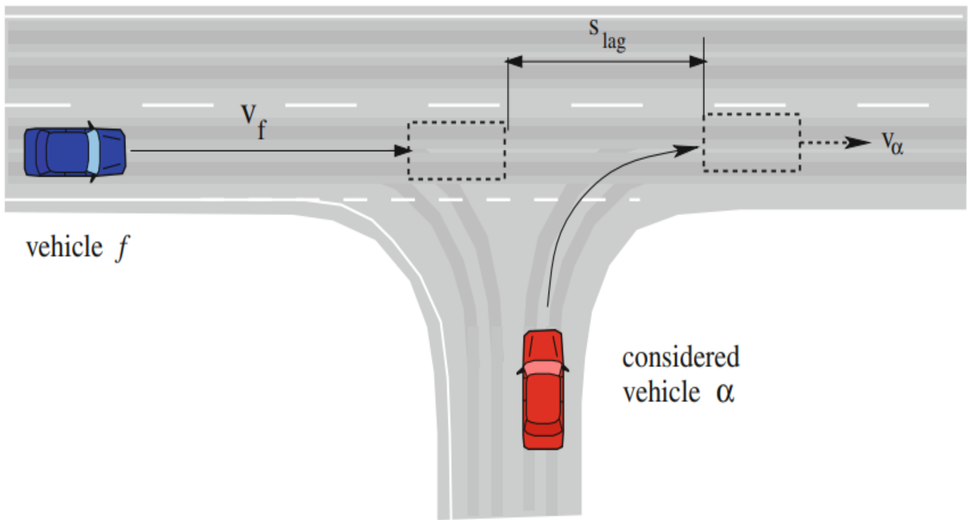

Fig. 4 Illustration of the safety criterion for the decision "stopping/waiting or merging" when entering a priority road.