EMS-Vision: Mission Performance on Road Networks

R. Gregor
and *E. D. Dickmanns*

Institut für Systemdynamik und Flugmechanik,
Universität der Bundeswehr München (UBM),
D-85577 Neubiberg, Germany
Tel +49 89 6004 4149 - Fax +49 89 6004 2082

Rudolf.Gregor@unibw-muenchen.de

Abstract

The paper gives a survey on the state of the art of complex mission performance at UBM. It shows the interaction between the upper decision and planning modules and the experts specialized for locomotion and perception. The scenario of a complex mission on road networks including a turn-off maneuver is chosen to illustrate the overall control and information flow.

Keywords perceptual architecture, autonomous vehicles, mission performance, landmark navigation, knowledge bases

1 Introduction

Since nearly two decades autonomous systems are a topic of intense research all around the world. Countless approaches have been developed and abandoned over the years. As one of the pioneers of the very first hour, UBM has concentrated its efforts on the development of vision systems for the interpretation of video streams in real-time (25Hz). The 4-D approach to dynamic machine vision developed in the mid 80ies has proved its superior performance over the years in various applications. With the latest generation implementation, the Expectation-based Multifocal Saccadic Vision (`EMS-Vision`) system, the basis for a new level of performance has been set. In contrary to former implementations, mostly working with static configurations and optimized for specific tasks or domains, `EMS-Vision` is a flexible system, which is able to configure itself dynamically during operation depending on the actual situation. The explicit representation of the system's capabilities allows the direct activation of specific capabilities just in time when needed. Special decision units responsible for different functions assign tasks to specialized experts and supervise their actions. While in [1] a survey of the EMS-Vision system architecture is given, in this paper the control and information flow between the decision units and the experts for perception and locomotion is shown in detail.

2 Static Background Knowledge

In order to achieve goal-oriented and responsible behavior in complex environments, an expectation based autonomous vehicle needs background knowledge about this environment. Knowledge can be divided into a static part, constant during at least one mission, and the dynamic part acquired during mission performance. The dynamic knowledge accumulated by the EMS-Vision system has been presented by Gregor et al. [1]. In the following section the general structure of the static background knowledge used for mission planning is described. The Mission Expert for autonomous road vehicle guidance developed at UBM uses three static knowledge bases:

1. digital maps for road networks and
2. for static objects beside the road used as landmarks as well as
3. a statistical knowledge base about the performance of the own capabilities.

The road map consists of three layers:

The base layer describes the topology of the road network. The road network itself is divided into node point elements (like intersections, entrances and exits) and the connecting roads. This layer also contains a vectorized description of the roads and the position of the road elements in a WGS84 coordinate system.

The medium layer represents the locomotion-specific data content of the knowledge base. Besides specifying the road type (domain), additional information like the number of lanes, driving restrictions or maximal velocity are stored here.

The top layer represents the perception-specific data content of the knowledge base. It serves as the link between the road map and the road models in the object model database of EMS-Vision used for perception (see [2]). In combination with data from the lower layers it provides information for the instantiation and initialization of hypotheses for road objects. Additionally, references to objects contained in the landmark map visible from this road element are given.

The landmark map contains data about static objects in the world. A generic model is used to describe object geometry. The map entry for an object contains a reference to the 4D model used for perceiving this object; for the initialization of this model a set of geometric parameters, photometric properties and the position of the object in the WGS84 coordinate system are given.

The third static knowledge base contains information about the own vehicle. On the one side, static information about the vehicle's geometry and performance limitations are stored. On the other side, it contains statistical data about the probability of success for complex capabilities implemented. Due to their varying complexity not all capabilities can be performed with equal efficiency and safety. For safety reasons it might be necessary to exclude the use of not always successful capabilities during critical parts of missions. At the moment, the content of this database is limited to locomotion capabilities.

3 Decision units

One of the most significant differences between the EMS-Vision system design and its predecessors is that a hierarchy of decision units (`DUs`), each specialized for a field of functions, is responsible for task assignment to specialized experts and for supervision and control of expert behaviors. For optimal resource exploitation, experts are started by the `DUs` through the `System Control` module just in time when needed and may be terminated after task completion. Three `DUs` have been designed and partially implemented at the moment, but the control structures used by the `DUs` have already been completed:

Behavior Decision for Gaze and Attention (`BDGA`) on the one side is responsible for task assignment to specialized Perception Experts (`PEs`). On the other side, it calculates optimal gaze strategies to keep the most important objects in the field of view (for detailed information see [3]).

Behavior Decision for Locomotion (`BDL`) is responsible for vehicle control. Given complex task descriptions, e.g. "Turn right at intersection 1 onto crossroad 2", (`BDL`) realizes this task by sequentially performing several low-level locomotion maneuvers.

Central Decision (`CD`) resides on the top level of the hierarchy with highest authority for decision. If no autonomous driving mission is actually specified by the human operator, `CD` may decide autonomously, whether a stationary observation mission or an off-line data processing for model-refinement and map update shall be performed. During mission performance it supervises the actions of `BDL` and `BDGA`.

For conflict avoidance, tasks are ranked according to their relevance for the system. Minor conflicts within the field of functions of `BDL` and `BDGA` are directly solved there. Conflicts between `BDL` and `BDGA` or conflicts that cannot be solved under given constraints are announced to `CD`. `CD` then may vary the constraints and leave conflict solution to `BDL` and `BDGA`, or determine and command a solution directly.

The following section describes the interaction between the perception experts, the `DUs` and the Mission Expert.

4 Control Flow

In the EMS-Vision system, object hypotheses can be generated by different modules. For performing a turn-off maneuver within a road network, the hypotheses for the according road elements representing intersection and crossroad are instantiated by the Mission Expert. Therefore, instances of the according objects are inserted into the scene tree and initialized. The initialization contains both relative position to the ego vehicle, determined by the Mission Expert, and road geometry parameters contained in the road map. The instantiation of a new object for perception triggers `BDGA` to assign the perception task to one of the perception modules specialized for the according object class. Considering all objects with visible features in the scene tree, the `BDGA` module has to calculate an optimal gaze strategy for the perception of these objects. As a matter of fact, with multiple objects to be perceived in parallel, the requirements concerning gaze direction for all objects may conflict. In order to solve these conflicts, a relevance value is assigned to each object. The relevance of an object itself depends on the tasks that actually are performed or that are planned to be performed. The experts specialized for actions assign the relevance value to the objects they need for task performance, e.g. for BDL the own road is relevant for locomotion, while perception modules only accumulate data about these objects. This already gives a good relation between the relevance of objects for the same kinds of tasks. Additionally, the relevance of different tasks for the complete system has to be considered. Therefore, different relevance ranges are allocated for different kinds of tasks.

Within the system-wide relevance scale all objects can be compared directly. This overall relevance range is divided into several parts, the relevance classes. Each expert is limited to a certain relevance class according to the relevance of the task performed. During operation, the relevance class of an object may change dynamically, if another expert needs this object for task performance. An example may explain the mechanism in detail: before a turn-off maneuver, multiple objects may be perceived in parallel. While approaching an intersection with multiple branches, at least the own road and one crossroad have to be perceived. However, additional crossroads may be useful as landmarks and therefore recognized for navigation purposes. The relevance classes of roads used for locomotion are set by BDL, those of the landmarks by the Mission Expert. Near by the intersection, only one crossroad can be focussed with the appropriate camera. As locomotion is more important than localization, the relevance range of BDL is ranked higher, so that BDGA can determine easily appropriate gaze maneuvers to focus on the most important crossroad. If the recognition of the crossroad onto which the turn-off maneuver shall be performed fails, an alternative behavior has to be generated and performed by BDL. Besides a stop, a turn-off maneuver to another crossroad, previously used as landmark only, may be chosen. BDL increases the relevance class of this crossroad and hence it will be attention-focussed by BDGA.

5 Mission planning

The first step of a complex driving mission is the planning phase. CD generates a planning order to the Mission Planning expert (MP) containing travel destination and temporal aspects like desired arrival time. The Mission Planning expert, evaluating the static background knowledge bases, computes several routes to the destination point using different criteria for route optimization. After the planning phase, the relevant information on the planned routes, containing duration, start- and final-time, route length and the required capabilities, are sent back to CD, which then may choose one of these or restart the planning with changed parameters. The chosen route is converted into a task list, the mission plan.

6 Mission plan

The mission plan contains all planned tasks for perception and locomotion during one mission. In general, it is organized as a sequential list of complex tasks, the mission elements (misel). The definition of the mission elements is directly connected to the appropriate maneuvers (see [4]) for locomotion within the mission element. The appropriate maneuvers and their parameters are constant during one mission element. Optionally, mission elements contain a task list of variable size for perception, including references to object data items within the static knowledge bases (landmarks and roads).

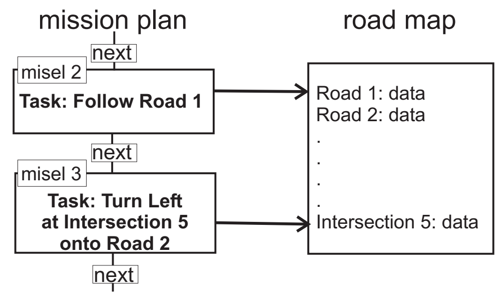

Figure 1: References to maps

Figure 1 shows a section of a mission plan, consisting of two mission elements. Mission element 2 contains one task for locomotion, "Follow Road". The road object used as reference for road-following itself is specified as "Road 1". Perception of "Road 1" is the first planned task for the perception experts. The mission elements refer to the entries in the road map for the expected road objects. As mission element 3 specifies a turn-off maneuver at "Intersection 5" onto "Road 2", the perception of the crossroad "Road 2" already has to be initiated during mission element 2.

During mission performance, the transition between mission elements has to be determined by the Mission Expert. Therefore, several criteria are set for successful or failed transitions. The criteria may be generated by combinations of the following aspects:

1. Temporal aspects: duration of mission element or absolute point in time.
2. Spatial aspects: covered stretch within one mission element or position relative to a static object.
3. Events: perception of an expected object, e.g. a crossroad.

The parameters for the criteria, e.g. the stretch covered, are dynamically calculated and set during mission performance by the Mission Expert, considering static background knowledge and dynamic ego-state data.

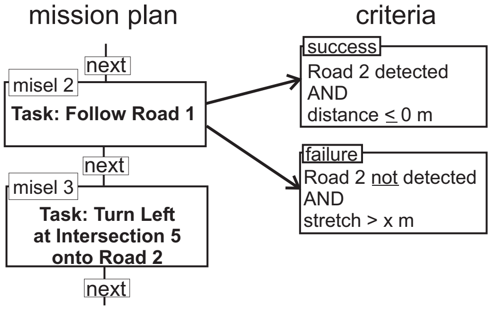

Figure 2: Transition criteria for mission elements

Figure 2 shows as an example, how several aspects are combined to determine the transition between mission elements 2 and 3. The main event, set as criterion for successful completion of mission element 2, is the detection of the expected crossroad; only then the turn off maneuver can be performed and the second transition criterion, that the vehicle enters the intersection (distance $\leq$ 0m), can be checked. The failure criterion is specified by the distance covered by the vehicle while travelling on Road 1. If this stretch exceeds a certain value, the vehicle must already have passed by the intersection without detecting it or the road map may be obsolete. The exact value is determined dynamically during mission performance (see section 7).

If a failure criterion is met, e.g. because the expected crossroad had not been detected, and thus it is not possible to fulfill the original mission plan, this state is announced to `CD`. `CD` then may initiate a replanning. In order to achieve higher robustness and to have background knowledge for at least the time horizon necessary for replanning, alternative mission elements are inserted at node points within the road network, like at intersections.

The following sections describes in detail, how the Mission Expert interacts with the system and evaluates data to determine actual mission progress.

7 Mission performance

At the beginning of a mission, only coarse information about the vehicle's position in the world is available. Using a commercial C/A-code GPS-receiver, the standard deviation σ of the position measurement is $\sigma \approx 100m$. With the verfication of the hypothesis, that the vehicle is located on a road and position information from the road map, the circular region of uncertainty can be reduced to a longitudinal position uncertainty σ_x along the road. During mission performance, the variance of the position estimate may even increase according to the accuracy of the sensors used for dead-reckoning.

A reduction of the variance can only be achieved by localization relative to landmarks. Therefore, the objects contained in the landmark knowledge base or characteristic road elements, like crossroads or sharp curves, may be used. The variance of the actual position is considered by the Mission Expert in two ways:

1. Hypotheses for expected objects must be instantiated in time to avoid that the vehicle has already passed by.
2. If the perception of an object is specified as criterion for transition, the actual variance of position estimation is used for parameterizing the failure condition.

Monitoring the vehicle's actions, the Mission Expert cyclically determines the progress within the complete mission as well as within the actual mission element and reports it to `CD`. As both spatial and temporal aspects are considered as transition criteria, the Mission Expert determines progress in space and time.

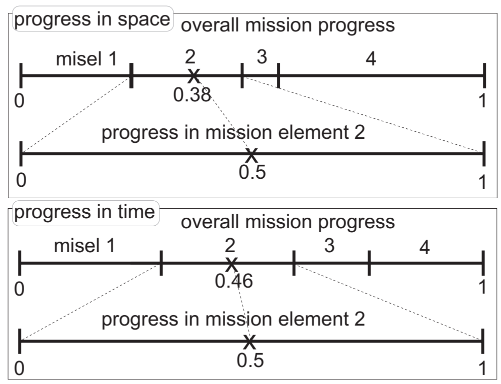

Figure 3: Monitoring of mission progress

The absolute value of the distance the vehicle has to cover during mission is derived from map data. The expected duration is calculated considering route length and desired velocity on the route. Progress is diagrammed on a scale ranging from $[0, 1]$. Figure 3 shows, how actual progress is reported. The two graphs show actual mission progress in space and time for one mission consisting of 4 mission elements. The top scale of each graph shows the overall mission progress. The actual overall mission progress in space of 38%, marked by an x, corresponds to a progress of 50% within mission element 2. The actual mission progress in time is shown on the lower graph. The top scale shows, that actually 46% of the estimated mission duration have gone by, corresponding to 50% of

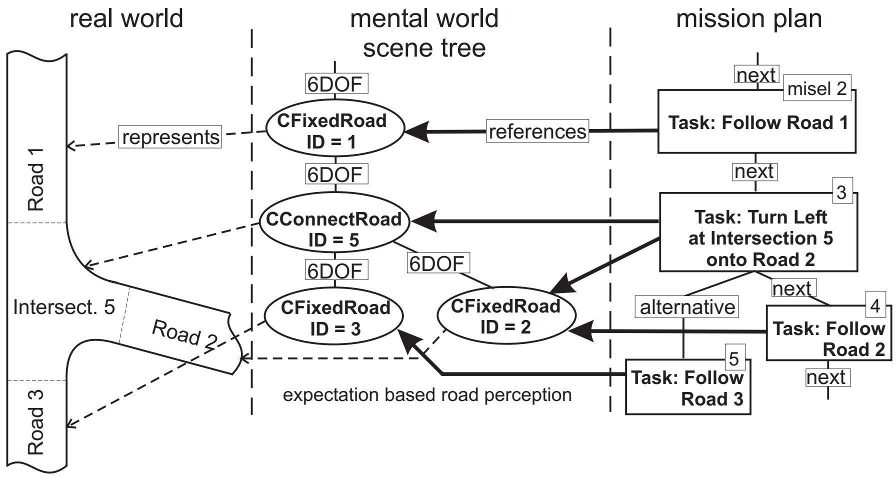

Figure 4: Mission realization using perceptual and behavioral capabilities

the duration of mission element 2. Of course, progress in space and time within one mission element is equal for constant velocity. Due to certain circumstances, the duration of a mission element or the complete mission may exceed the estimated value, e.g. because the vehicle is blocked in a traffic jam. Thus, progress in time may exceed the 100% level as well. This indicates, that external disturbances might jeopardize mission success, and as a consequence `CD` can initiate a replanning or increase travel speed on free stretches.

Figure 4 shows the control flow during one part of a complex driving mission. On the left side, the real world containing physical road objects is shown. Within the system, a mental world (scene tree) containing mental objects (shown as ellipses) is built up to represent the real world. While following a road, the vehicle will approach an intersection with one crossroad to the left. At this intersection, the vehicle shall turn left onto the crossroad. On the right hand side of fig. 4, the corresponding part of the mission plan is shown. While performing the mission, the information in the mission elements and the references to background knowledge are evaluated by the Mission Expert to instantiate object hypotheses for the road elements relevant for locomotion and navigation. Starting with mission element 2, "Follow Road", the Mission Expert inserts an object of class "CFixedRoad" ($ID = 1$) into the scene tree and directs the references for mission element 2 to this object. This triggers `BDGA` to activate a perception expert for this task. `BDL` evaluates estimation data on "Road 1" accumulated by the perception expert for performing road-following. One necessary condition for the transition to mission element 4 is the perception of crossroad 2. Thus, at a certain distance from the intersection, but still within mission element 2, the Mission Expert initiates the perception of crossroad 2 by inserting additional objects for intersection 5 and crossroad 2. In order to prepare the system for the alternative mission element 5, an additional object for the second branch ($ID = 3$) is instantiated, too. `BDL` starts two behavioral capabilities for locomotion in parallel:

1. The capability "turning-off" is started in the active mode. It has the priority to overrule the actual capability "road-following" at any time. While approaching the intersection, it cyclically computes feed-forward control variables for the turning-off maneuver. At the optimal distance to the intersection it takes over vehicle control from the active capability "road-following".
2. The capability "road-following" for the alternative mission element is started with lowest priority in a passive mode. Only if the transition to mission element 4 fails, this capability may be directly activated by `BDL`.

For detailed information about locomotion see [4].

The following section presents experimental results of a turning-off maneuver performed with the test vehicle VaMoRs on a campus road of UBM.

8 Results

Figures 5 and 6 show data generated by the mission expert during one mission over system time. The top graph in fig. 5 shows the actual mission elements: beginning with mission element "Follow Road" the vehicle "Turns left" onto the crossroad, follows the crossroad and finally stops.

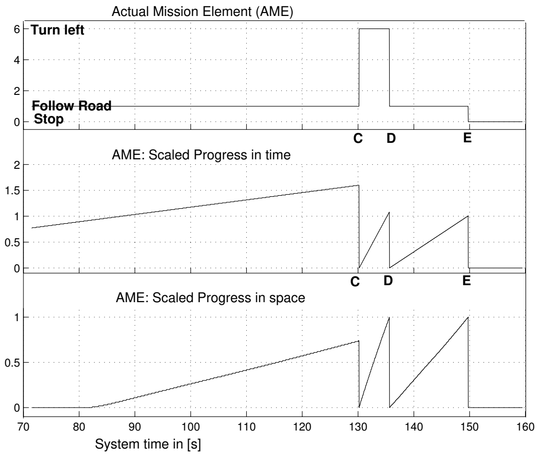

Figure 5: Complex driving mission

The center graph shows the scaled progress in time for each mission element, the bottom graph shows the corresponding progress in space.

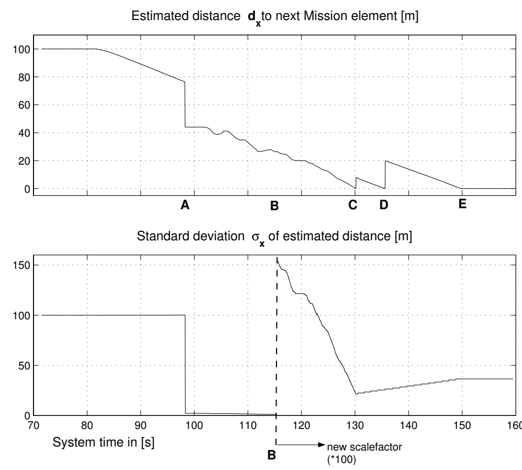

Figure 6: Estimation data

Figure 6 shows the distance d_x to the next mission element and its standard deviation σ_x estimated by the mission expert. In the first mission element "Follow Road", the perception of a crossroad is specified as criterion for transition to mission element "Turn left". While approaching the intersection without perceiving the crossroad, position is estimated using solely odometry as sensor input. Thus, the value for $\sigma_x \approx 100m$ is slightly increasing from its initial one according to the accuracy of the odometric sensor. After detecting the intersection, the road expert delivers estimation data for egostate relative to the intersection. These data are used continuously by the mission expert for a more accurate estimation of d_x and σ_x, indicated by the significant step in both graphs at point A. Please note, that at point B the scale for graph 2 in fig. 6 is increased by a factor of 100 for better resolution. By reaching the intersection ($d_x = 0$), the transition to the next mission element is performed at point C. Due to the high position uncertainty within the first mission element, it is finished at a scaled progress in space of $\sim 70\%$ and, due to a delayed mission start, at a scaled progress in time of $\sim 160\%$. After entering the intersection, again solely the odometric sensor is used for position estimation and thus σ_x increases again. At point D the turn off maneuver has been finished and the next mission element, "Follow Road", is activated. In this mission element, the vehicle has to cover a distance of $20m$. At point E the scaled mission progress in space has reached 100% and the mission is terminated by the transition to mission element "Stop".

9 Conclusions and outlook

EMS-Vision, the latest generation implementation of UBM vision systems based on the 4D approach has been presented. It constitutes a powerful basis for reaching levels of performance for autonomous vehicles unknown up to now. The principle capabilities for turning-off maneuvers, demonstrated first by [5], have been reimplemented and integrated into a complex control flow scheme. The approach has been verified with tests on the UBM testtrack with a modernized version of VaMoRs. Besides the extension of specialized capabilities, future work will concentrate on the expansion of the decision units used to activate the growing number of capabilities implemented.

References

[1] R. Gregor, M. Lützeler, M. Pellkofer, K. H. Siedersberger, and E. D. Dickmanns. EMS-Vision: A perceptual system for autonomous vehicles. In *this volume.*

[2] M. Lützeler and E. D. Dickmanns. EMS-Vision: Recognition of intersections on unmarked road networks. In *this volume.*

[3] M. Pellkofer and E. D. Dickmanns. EMS-Vision: Gaze control in autonomous vehicles. In *this volume.*

[4] K. H. Siedersberger and E. D. Dickmanns. EMS-Vision: Enhanced abilities for locomotion. In *this volume.*

[5] N. Müller and S. Baten. Image processing based navigation with an autonomous car. In *Proc. Int. Conf. on Intelligent Autonomous Systems*, Karlsruhe, March 1995.