Biologically Guided Driver Modeling: the Stop Behavior of Human Car Drivers

Mauro Da Lio, *Member, IEEE*, Alessandro Mazzalai, Kevin Gurney, and Andrea Saroldi

***Abstract*— This paper presents a principled approach to the modeling of human drivers—applied to stop behavior—by uniting recent ideas in cognitive science and optimal control. With respect to the former, we invoke the affordance competition hypothesis, according to which human behavior is produced by resolving the competition between action affordances that are simultaneously instantiated in response to the environment. From the theory of optimal control, we deploy motor primitives based on minimum jerk as the potential suite of actions. Furthermore, we invoke a layered control architecture, which carries out action priming and action selection sequentially, to model the biological affordance competition process. Motor output may be directed to distinct motor channels, which may be partially inhibited, e.g., to model gas pedal release saturation. Within this architecture, two types of motor units—"deceleration" acting on a gas pedal channel and "brake" acting on a brake pedal channel—are sufficient to model, with remarkable accuracy, the various phases that can be observed in human maneuvers in stopping a car, namely: gas release, gas chocked, brake, and final brake release at stop. The model is validated using experimental data collected in 16 different stop locations, from roundabouts to traffic lights. We also carry out a comparison with the well-known Intelligent Driver Model, discuss the scaling of this framework to more general driving scenarios and finally give an example application where the driver model is used, within a mirroring process, to infer the human driver intentions.**

***Index Terms*— Driver modeling, intelligent vehicles, layered control architectures, adaptive behavior, cognitive robotics, motor primitives.**

I. Introduction

MODELING drivers has always been an important part of vehicular and transportation studies. Successful models for many application fields have been developed with many approaches, such as, e.g., psychological, analytical, machine learning, [1], [2], experimental [3], etc..

Concerning the genesis of these models, a generally common trait is that they are developed from *plausible* assumptions regarding the driver decision making process or what the driver is supposed seeking to do. However, incomplete understanding of the principles underlying human behaviors makes it difficult to generalize the assumptions and scale these models – for example, to produce a human model suited for man-machine interactions.

In the domain of robotics and cognitive science a parallel research thread has evolved, that seeks to explain human sensorimotor behaviors in terms of models of the underlying brain architectures and processing. This theoretical framework can be, in principle, used to engineer new driver models grounded into more general cognition theories.

In particular, following the work of people like Cisek [4], this paper argues that human driving can be modelled following the principles of competition between *affordances* – which are possible actions 'invited' by objects and situations in the immediate environment and 'affordable' by the motoric capabilities of the agent (e.g. a bifurcation affords taking either lane). We also assume that these actions are implemented with optimal goal-directed motor strategies.

With these principles, we show how to develop a driver model focusing on the stop behavior of human car drivers, which we analyze theoretically and with real driving data. We show that the stop behavior is surprisingly complex when examined in detail. Nonetheless with the introduced principled approach we generate a model that predicts phenomena such as 'coasting' (no gas and no brake), and the release of the brake to reduce the jerk pulse just before the stop. The switching between different phases preceding the stop (gas release, braking, brake release etc.) is explained with an environmentally-responsive action-selection mechanism, rather than a pre-programmed sequence. Experimental data are used to produce evidence of the different types of affordances, of their optimality and to parametrize their instantiations representing an average driver.

Beyond the interest represented by the detailed knowledge of the stop behavior in itself, we believe that this paper contributes to supporting the thesis that human driver behavior can be modeled with layered control architectures and optimal control motor primitives. Thus, in the conclusions we discuss the prospective application of these studies to human-vehicle interactions, to be achieved by providing the vehicle with a human model that attempts, *in silico*, to 'mirror' the dynamics of the internal cognitive states enjoyed by the human agent *in vivo* [5]–[15].

II. Related Studies

A. Models for Stop Behavior in Transportation Studies

A large number of models can be found in the literature concerning car following (CF) scenarios [16]; many of these include stop.

Manuscript received March 3, 2017; revised August 19, 2017; accepted September 2, 2017. This work was supported by the European Commission under Grant FP7 246587 (interactIVe), Grant FP7 610391 (NoTremor), and Grant H2020 731593 (Dreams4Cars). The Associate Editor for this paper was P. Ye. *(Corresponding author: Mauro Da Lio.)*

M. Da Lio and A. Mazzalai are with the Department of Industrial Engineering, University of Trento, 38123 Povo, Italy (e-mail: mauro.dalio@unitn.it; alessandro.mazzalai@unitn.it).

K. Gurney is with the Department of Psychology, The University of Sheffield, Sheffield S1 2LT, U.K. (e-mail: k.gurney@sheffield.ac.uk).

A. Saroldi is with CRF, the research center of Fiat Chrysler Automobiles, 10043 Orbassano, Italy (e-mail: andrea.saroldi@crf.it).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TITS.2017.2751526

The Intelligent Driver Model (IDM) [17] is among the most famous ones. It is of particular interest for us as it has been used to infer the intentions of drivers at intersections among four alternative hypotheses [18]. Thus, the IDM generates the longitudinal plans for the four hypotheses, which are compared with observations of the actual driver behavior to find the most probable one. In Cognitive Science terms, this method uses the 'mirroring' principle, where another agent (in this case the IDM) *generates* behaviors – like a human – for possible actions and then *selects* the one/ones which best fits the observed behavior of the mirrored agent (the human driver). For the success of this process it is mandatory that the mirroring agent possesses actions similar to those of the mirrored agent [15]. Then, the environment can offer similar affordances to both agents which, in turn, allows (at least in principle) for the mirroring agent to enjoy similar internal mechanisms to its mirrored counterpart. Only in this case, fitting the observed behavior means that the 'intentions' (internal states corresponding to the active action/affordance) of the observed agent are captured correctly. Thus, given that accurate *modelling* of driver behavior is a prerequisite for mirroring, a close comparison between the stop behavior predicted by the IDM and that predicted by this paper's models will be carried out in sections VI and VII.

Most stop maneuvers occur at intersections, and so previous studies have focused on this scenario. Further, they have approached the problem by simply *classifying* different driver behaviors. Thus, Aoude *et al.* [19] used two machine-learning approaches (Support Vector Machines SVM and Hidden Markov Models HMM) to categorize the vehicles approaching an intersection into two classes: compliant or non-compliant. The speed and distance of the approaching vehicles is observed till a given time before the intersection. Then the classification is carried out and, if necessary, warnings may be issued to other vehicles. The time when the classification is carried out, is a tradeoff between classification accuracy and timely/usefulness of warnings: the later the classification is carried out the better; but the lesser time remains for reactions. If the classification is carried out 2 seconds before the intersection, the reported true positive rate is about 65%, with 5% false positive rate.

Classification at intersections is also studied in [20], extending the case to yield scenarios and classification among four possible intentions (no action, stop, creep and go). Three machine-learning approaches are studied and observations of the second vehicle are included to account for the mutual influence between vehicles in the yield scenario.

However, in view of our objectives (*modelling* drivers), a limitation of classification approaches is that they do not provide an explicit explanation for why and how the driver behaves.

B. Biological Plausible Human Behavioral Models

There are two main ideas that can be useful for the development of rational models of human drivers, at least at mid-low levels of sensorimotor control: *1*) the notion of affordance competition, implemented in layered control architectures and *2*) the notion of efficient motor control, that can be emulated by optimal control.

1) Affordance Competition and Layered Control Architectures: The basic idea in layered control is to use multiple, parallel, processing streams, each with perceptual and motor abilities, and of ever increasing processing complexity.

Layered control is a recurrent theme in many disciplines in cognitive science (for a cross-disciplinary review see [21]). For example, in neuroscience it has a long and rich history going back to Jackson in the 1880s [22]. Indeed, the idea of layered systems instantiated in different structures along the neuraxis (from brain stem to frontal lobes) is now part of mainstream opinion.

More latterly, in robotics, the 'subsumption' architecture of Brooks [23] has enjoyed considerable attention. While not intended as a biological model, this architecture showed how simple behaviors could run autonomously without extensive processing but could, if needed, be replaced or 'subsumed' by those requiring more complex control.

Across all layered schemes, a key question is how the different layers of control can be orchestrated to act in a coherent way when confronted with the problem of action selection – determining which of the many possible actions or behaviors should take control of the agent at any one time. In the subsumption scheme in robotics, this is not always clear and is solved in a variety of ways. In biology, however, there is a possibility of a single unified approach – that of a central selection mechanism based on an evolutionary old brain system – the basal ganglia [24]. In this proposal, the brain can simultaneously prime many 'potential actions' or 'action requests' which may originate in any part of a layered architecture. These requests are transmitted to the basal ganglia which, via internal competition, allow only the most urgent or salient actions to prevail. A similar idea has been articulated independently by Cisek under the banner of the *affordance competition hypothesis* [4]. This emphasizes the notion that perceptions for possible actions – or *affordances* – generate action requests that are continually being evaluated for selection. Indeed, both action priming and action selection occur continuously, which permits continuous adaptation of potential solutions to the current tasks in hand, and thereby, selection of the most appropriate behavior.

In the driving domain, with respect to layered control, many driver models have indeed long recognized the hierarchical organization of the driving task, spanning from long-term complex goals at higher levels to short-term simple control actions at lowest levels [25]–[27]. With respect to the affordance competition hypothesis, human decision-making at intersections has been in particular studied in [28], concluding that it looks to be the choice between affordances, where 'drivers take into account both the crossing possibilities and the stopping possibilities'.

2) Optimal Goal-Directed Actions and Motor Primitives: Many studies on human motor control reveal the existence of highly optimized strategies, e.g., [29]–[38], which hold not only for body movements but also for the manipulation and control of objects [39], [40]. Models of such

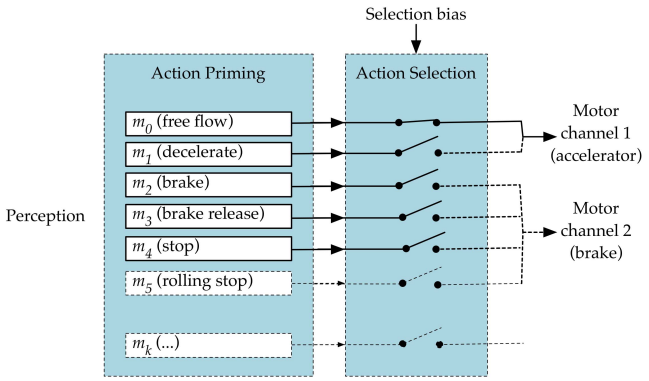

Fig. 1. Layered control architecture for the driver model. Perception-action loop occurs in two phases: 1) based on the perceived environment potential actions $(m_0, \ldots, m_5)$ are primed in parallel; 2) among these, the action that best suites higher-level biasing directives is selected and gated to the motor system for execution. If selected, (m_0, m_1) are directed to motor channel 1 (the accelerator pedal); the remaining actions to motor channel 2 (brake pedal). Only one action is gated at any time. This architecture is scalable with the addition of additional actions m_k, which participate to the same selection procedure. Algorithms that instantiate this architecture for 3 models of increasing complexity will be given in section V.

optimized control include those governed by the minimum jerk and minimum variance principles [35], [36].

The notion of motor primitives, as atomic motor units forming the bottommost layer of a layered architecture, is also present in the literature [41], [42].

In the driving domain, optimal control has been used to *describe* driver behaviors (from a functional point of view) both for long-term tactical plans [43]–[45] and for producing atomic motor units (motor primitives) for layered control architectures [46]–[48]. However, the connection with studies in human optimal sensory-motor control have been rarely recognized.

III. Modeling Human Stop Behavior

A. Layered Control

Following section II.B.1, our artificially engineered layered control architecture aims at reproducing the mechanism of affordance competition in biology [4], [21], [24].

In the schematic representation of Fig.1, action requests related to many plausible goals are generated simultaneously (the action priming block). These 'compete' according to given 'fitness' criterion, and the winner is then gated to the motor system, thereby taking control of the plant (the action selection block). However, the competition process can be steered by biasing the 'fitness' criterion from higher levels, resulting in a layered architecture of the kind described in section II.B.1.

Note how this architecture reverses the traditional logic processing order that sees situation assessment preceding action planning. Here the opposite happens: action planning – in the form of nascent behaviors – occurs before assessment of which is best, and each plan influences the selection process.

The scheme of Fig.1 has direct analogues with ([4], Fig.1). There, action affordances are generated in the so-called 'dorsal processing stream' in cortex which generate candidates for action selection (action requests) in basal ganglia. The 'action selection' block in Fig.1, indeed, mimics the use of such a dedicated central mechanism in the vertebrate brain. At any time, only one action is gated (only one switch is closed).

Finally, Fig.1 shows two different motor *channels*, defined as of the operation of the accelerator/gas pedal (channel 1), or the brake (channel 2). Actions aimed at controlling the speed in free flow or releasing the gas pedal in early phases of stop are directed to motor channel 1; actions aimed at changing the speed via the brake (including brake release), act on motor channel 2. This will allow modelling the different characteristics of each channel, and in particular, for this paper, the fact that deceleration saturates when the gas pedal is completely released.

1) Enacting Actions, Motor Primitives: This section deals with how actions (such as, e.g., $m_0, \ldots m_4$ in Fig.1) can be enacted. Following section II.B.2, modelling a generic action 'm' is carried out within an Optimal Control framework as follows.

Initial Conditions: Let us consider a goal-directed action m to connect the present state[1] of the vehicle $s(0), v(0), a(0)$ to an aimed final state $s(T), v(T), a(T)$; m may be embodied by a function $s(t), t \in [0, T]$, representing the travelled distance as function of time, and, when necessary, by its derivatives: $v(t)$ velocity, $a(t)$ acceleration, etc. Let $t = 0$ represent the present time and $t = T$ the time at which the final state will be reached. In general, the movement time T is not given beforehand, but is part of the problem.

Let the current state be defined by the following initial conditions:

$$s(0) = 0, \quad v(0) = v_0, \ a(0) = a_0 \tag{1}$$

where a_0 and v_0 are the current acceleration and velocity, and a curvilinear reference frame with origin on the present position and longitudinal abscissa following the lane centerline is assumed (hence $s\,(0) = 0$).

Final Conditions (for Stop): Let a future stop state be defined by the following final conditions:

$$s(T) = s_f, \quad v(T) = 0, \ a(T) = a_f \tag{2}$$

where s_f is the position of the stop line (or any other desired stop point) in the current curvilinear coordinates system, and the final velocity $v(T)$ must be zero. Note that, like T, the final acceleration a_f is not known yet and will be determined by imposing the optimality conditions (this is to model the residual acceleration left at stop by human divers).

Final Conditions (for Free Flow): The final condition (2) is suited for modeling actions that aim at a final rest state (stop), such as m_1, m_2, m_3, m_4. However, the sensorimotor system of Fig.1 includes two actions, free-flow m_0 and rolling stop m_5 which do not aim at stop, but rather at a finite value, v_d, for the final velocity. Hence, let us define an *alternative* set of final conditions to be used for modeling actions such as m_0:

$$s(T) = free, \quad v(T) = v_d, \ a(T) = 0 \tag{2'}$$

[1]We consider the acceleration to be a state variable because longitudinal forces do not change instantaneously.

Plant Model: Let the plant be modelled with a simple kinematic model (3). This choice is motivated in part by simplicity and in part because there is some evidence that human motor planning occurs at the kinematic level in the Cartesian space, e.g., [37], [38].

$$\dot{s}(t) = v(t), \quad \dot{v}(t) = a(t), \ \dot{a}(t) = j(t) \tag{3}$$

where $j(t)$ is the control input, which corresponds to the longitudinal jerk.

Optimality Criterion: Finally, let us formulate the following optimality criterion:

$$J = w_A a_f^2 + \int_0^T w_T + j(t)^2 dt \tag{4}$$

In equation (4), the part in the integral is a trade-off between smoothness (if $w_T = 0$ the integral reduces to the minimum square jerk criterion), and speed (if $w_T \to \infty$ the integral reduces to the minimum time criterion). Defining the optimality criterion as a trade-off between the two clearly serves to model the different speed-accuracy combinations that humans may use. Thus, with large values for w_T, primitives that quickly reach the aimed final state are produced, but they will also use large jerk, which in turn results in large command noise (because human neural noise is proportional to the control signal [35], [38]), and thus in less accurate movements. Vice-versa, maximally smooth maneuvers are produced with $w_T \approx 0$, but they require more time to reach the desired final state.

Thus, for what concerns the integral part, formulation (4) is identical to previously used ones [46]–[48]. The difference in this paper is the new term $w_A a_f^2$; a terminal cost which is introduced to reduce the final acceleration, a_f, making it loosely bound to zero, instead of exactly zero. Indeed, if $w_A = 0$ the final acceleration minimizing (4) would be free to vary and would result quite large; if $w_A \to \infty$ the final acceleration would be forced to zero; but, to obtain this, larger $j(t)$ would be necessary. As the experimental data in Section IV show, human drivers adopt a tradeoff, tolerating some residual acceleration at the stop time, (which results in a slightly uncomfortable jerk pulse). This is here modeled by a finite value for w_A, which will be estimated in section V for different motor primitives.

a) Stop motor primitives: Equations (1), (2), (3), (4) define an optimal control problem, which aims at a final rest state. This problem has an analytical solution, which is given in details in Appendix I. The resulting trajectory is a 5$^{\text{th}}$ order polynomial:

$$s(t) = c_1 t + \frac{1}{2}c_2 t^2 + \frac{1}{6}c_3 t^3 + \frac{1}{24}c_4 t^4 + \frac{1}{120}c_5 t^5, \quad t \in [0, T] \tag{5}$$

where coefficients c_i of the polynomial and the movement time T, are functions of the initial conditions a_0, v_0, of the final condition s_f and of the weights w_T, w_A that set the trade-off among smoothness, time and residual acceleration.

Let us indicate the solution of the optimal control problem with the following concise writing:

$$m = Stop\left(a_0, v_0, s_f, w_T, w_A\right) \tag{6}$$

meaning that m is the particular polynomial (5) with c_i and T computed as shown in Appendix I given a_0, v_0, s_f, w_T, w_A. Alternatively –which is better suited for a computer– m may be a vector collecting the coefficients c_i and T: $m = \{c_1, c_2, c_3, c_4, c_5, T\}$. Either way we have a *representation* of the action.

Further, we note that the optimal control problem is somehow 'atomic' in the sense that it cannot be further decomposed. Hence the polynomial solution (5) is called '*motor primitive*' in broad analogy with the biological ones.

b) Free-flow motor primitives: Similarly, if we replace the final condition (2) with (2') we get an optimal control problem aiming at achieving a uniform speed v_d. The solution is another 5$^{\text{th}}$ order polynomial, with coefficient given in Appendix II. Let us call this second type of motor primitive as follows:

$$m = FreeFlow\left(a_0, v_0, v_d, w_T\right) \tag{7}$$

2) Implementing Action Selection: For longitudinal control in a stop scenario, a suitable selection criterion may be based on the initial value of the longitudinal jerk $j(0) = c_3$, i.e., by choosing the action m_i with the smallest initial jerk $j_i(0)$. This also means choosing the motor plan that is initially the slowest, and if this process is continuously updated, the control output from the scheme of Fig.1 is the one which complies with the most severe longitudinal limitation ahead, on a moment-to-moment basis.

However, one should note that this simple selection criterion works only for the stop scenario of this paper. More complex situations call for more complex selection mechanisms [49], requiring the representation of actions in a metric space, and the use of inhibition mechanisms. A discussion of action selection mechanisms for such more general situations in longitudinal control is given in [48], section V.B.

3) Receding Horizon Control, Minimum Intervention Principle: Let m_i be the selected motor primitive. The corresponding optimal control is obtained from (5):

$$j_i(t) = c_{3,i} + \frac{1}{2}c_{4,i} t + \frac{1}{6}c_{5,i} t^2, \quad t \in [0, T] \tag{8}$$

It could be used to specify the requested acceleration rate, which, if tracked, would produce the planned motion. However, if the action priming/selection is continuously updated, e.g. like in a receding horizon framework, we are going to use only the initial part of (8) – ideally only $j_i(0)$ if the agent of Fig.1 runs continuously.

Note that in the re-planning hypothesis, any deviation that might occur, gives birth to an updated plan that still aims at achieving the same final goal, but from the deviated state. Thus, there is no return to the original trajectory and only the deviation components that affects the achievement of the final goal have effect. This is known as the 'minimum intervention principle' [50], [51], which is here an emergent feature of the driver model.

4) Modeling Pedal Saturation: The requested acceleration rate $j_i(0)$ can be produced by either acting on the gas or brake pedal. Since the pedal absolute positions approximately map onto engine torque or brake force, the rate

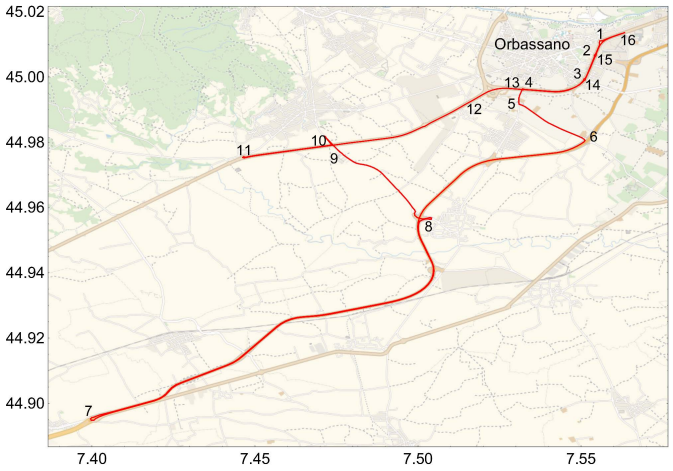

Fig. 2. Map of the test circuit. Numeric labels indicate 16 selected locations.

of pedal stroke approximately maps onto the longitudinal *intentional* jerk (albeit with different gains for gas and brake). Hence, the output of the model in Fig.1 can be interpreted as determining the rate of pedal strokes. However, both pedals saturate, in particular when they are fully released. This can be modeled with thresholding $j_i(0)$: namely, if the gas pedal is completely released negative values for $j_i(0)$ are set to zero (the pedal cannot be further released). The same happens for the brake. In this way actions directed to the gas/brake pedal become non-effectives if they correspond to a further release. As shown in the following, this is enough to reproduce the phenomenon of 'coasting', where a driver releases the gas pedal, but does not yet brake (colloquially speaking, we could say that in the coasting phase the driver 'would like' to further release the gas pedal –which is not effective– but is not ready yet to brake).

IV. Experimental Data

We introduce experimental data that form the basis for both qualitative and quantitative parameterization of the models and model validation. These data were collected during the user tests of the EU FP7 'interactIVe' project [52].

For the project, data of many sensors were logged at 50 ms rate. Signals relevant for this paper, are: longitudinal velocity, longitudinal acceleration, gear, gas pedal position, brake master cylinder pressure and distance of the vehicle ahead.

Velocity and acceleration signals have been improved with a combination of two Kalman filters. One designed for speed above 0.5 m/s (where the odometer is accurate), to remove, in particular, the acceleration bias caused by variation of the road slope and thermal drift. Another designed for the short periods with speed below 0.5 m/s (at which the odometer is less accurate, but slope variation and drift can be neglected) to compensate the loss of accuracy of the odometer at low speed, and consequently to improve the estimation of actual stop positions.

Fig.2 shows the test circuit, which is close to Orbassano, west of Turin, in Italy. The circuit was driven 50 times by 24 users and 1 professional test driver, each driving twice (see [47], [53] for other details).

In the map, 16 locations where stop events have been observed are labeled. Points 1, 2, 3, 4, 9, 11, 13, 14 and 15 are yield locations at the entrance of small or medium-sized roundabouts. Points 6, 7, 8 are yield locations in large roundabouts with complex geometry. Point 5 is a left turn (with yield sign). Point 10 is a ramp (with yield sign). Point 12 is a traffic light. Point 16 is another traffic light (and end of route). Together they represent different conditions that may require a vehicle to stop (today roundabouts are the most frequent situations in EU urban and extra urban roads).

Fig.3 gives a global representation of the speed profiles observed at the 16 locations, as functions of the longitudinal abscissa. Not all drivers stopped at the same position: in some cases, they had to stop behind other vehicles (e.g.: see the traffic lights queues in points 12 and 16).

Vehicle trajectories, over all 16 stopping locations, have been divided into three types.

1) Type 1 (the light blue curves in the online version of this paper) is formed by vehicles that reached a real stop condition (i.e., zero speed).
2) Types 2 and 3 have been formed by dividing the non-stopping vehicles in two classes: 'slow' (yellow curves) and 'fast' (violet curves). The 'fast' type represents vehicles that either do not reduce speed at all (e.g., at the green traffic lights at location 12) or reduce the speed only for what is strictly necessary to comply with the safe speed in the curves ahead. The 'slow' type represents vehicles that initially reduce speed in a way that is compatible with stopping, until a point where they resume speed, when they have enough information to take the decision to go.

A close examination of Fig.3 shows that, except for the traffic lights (locations 12 and 16) the speed profiles of the three types become distinctly separated only when very close to the possible stop point.

A. Evidence of Hierarchically Structured Behaviors

Fig.4 gives a first example of a maneuver ending with complete stop at the first roundabout (point 1).

The recordings begin from the precedent peak of the speed, which occurred about 21 seconds earlier, and which marks the beginning of the approach to the stop point. From that moment, the maneuver proceeds with different phases: 1) an initial partial use of the gas pedal (*b*-*c*), which corresponds to very slight deceleration; 2) a subsequent phase with complete choking of the gas (*c*-*d*); 3) a phase where the brake pressure is gradually increased, causing increasing deceleration (*d*-*e*); 4) a final phase where brake pressure is reduced, to suppress the uncomfortable jerk pulse that would otherwise happen at the stop (*e*-*f*). In Fig.4, brake and gas signals, which represent the intentional driver control, originally sensed as percent of the gas stroke or percent of the maximum brake cylinder pressure, have been scaled to show the correspondence with the longitudinal acceleration.

The composite nature of this example is a first indication that the stop behavior is actually made by sub-behaviors.

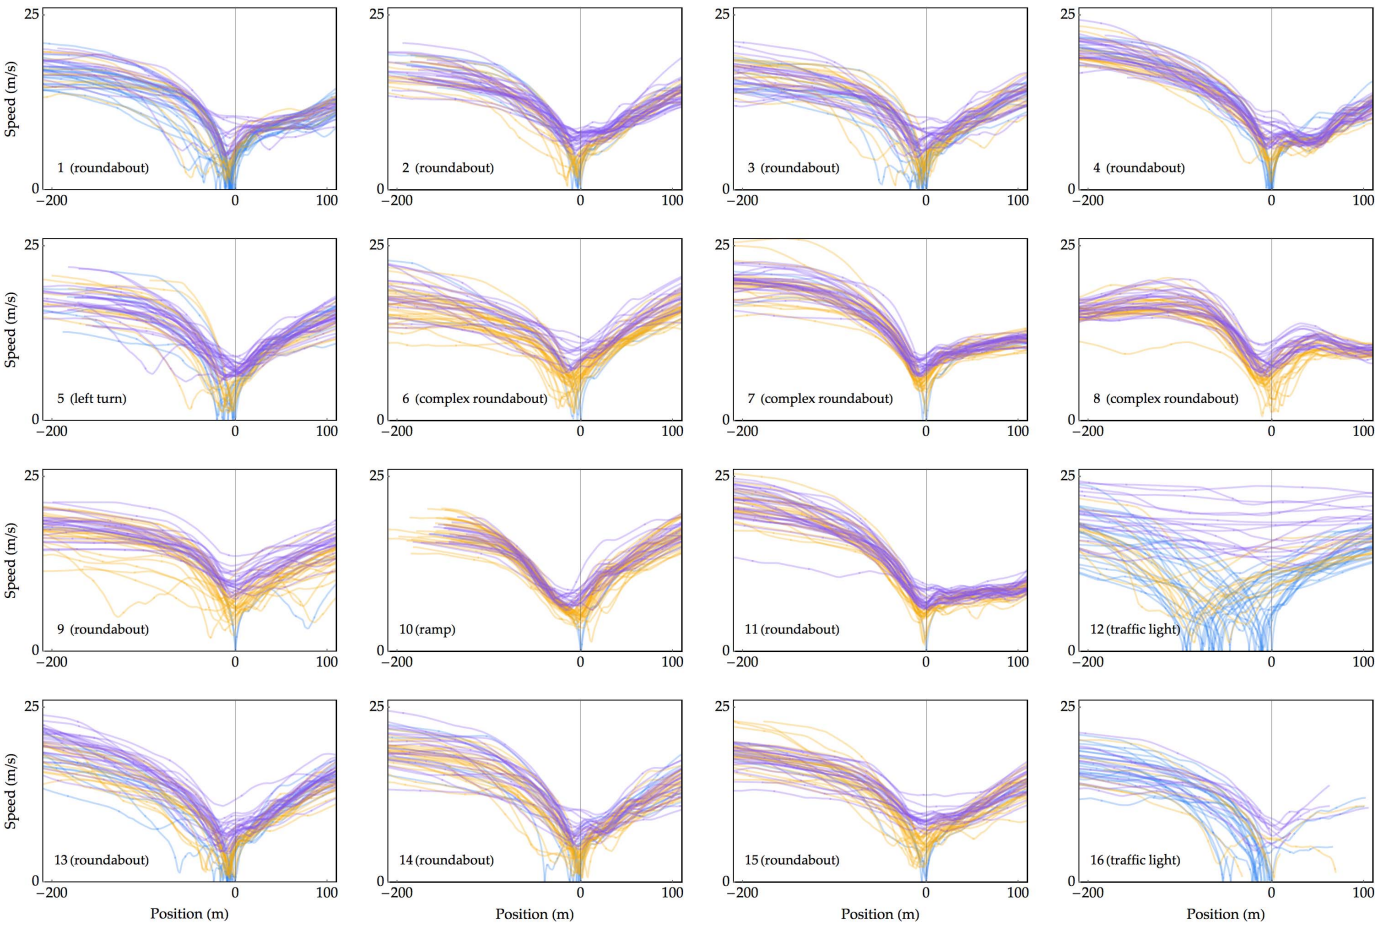

Fig. 3. Speed of vehicles in the 16 selected locations. Light blue tracks represent stopping vehicles; yellow and violet are two clusters that represent vehicles slowing down/not slowing down for possible stop (see text).

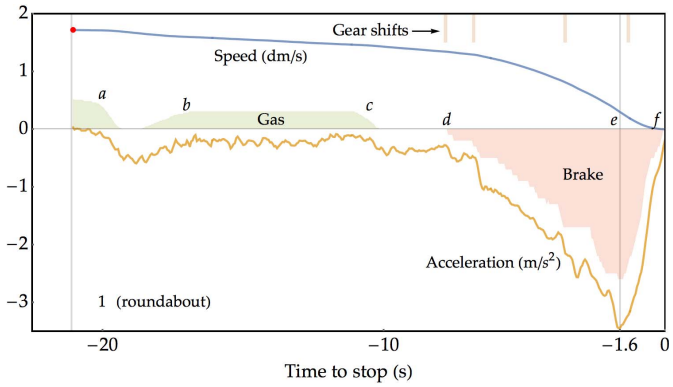

Fig. 4. Example of an unperturbed stop maneuver shows it is made of sub-behaviors (partial gas, chocked gas, brake, brake release). The y-axis scale serves for both speed (blue curve) and acceleration (orange curve), the units of which are labeled next to each curve (dm/s for speed and m/s^2 for acceleration). Brake and gas pedal signals, originally in percent of stroke or max pressure, have been scaled to show the quasi proportionality that relates pedal commands to acceleration. Brake and gas represent the intentional actions of the driver, which result in a proportional acceleration control. The coding of this figure has been used also for Fig.5 and 6.

Here we noted: partial gas, chocked gas, brake and brake release, at least.

The situation of Fig.4 still is a comparatively simple case, in particular because not perturbed by traffic, where the stop behavior is achieved by executing the sub-behaviors in ordered sequence. However, depending on the context (in particular traffic, changes in visibility and available information, etc.) the driver may switch back and forth between sub-behaviors.

Fig.5, shows one such example for a left turn (point 5): the driver brakes a first time between d' and f' because of front vehicles, as the violet 'car following' vertical bands mark (as detected by the frontal laser-scanner and confirmed by visual inspection of front camera recordings). When the driver returns in the 'free flow' condition, he/she returns to using the gas in a'', until the final stop maneuver is carried out in *d-e-f*. A video showing these events is included in the supplement materials. Note that, in this case, the final brake release is incomplete, ending with a finite acceleration at stop.

Fig.5 is given as an example of the dynamic adaptation that layered control architectures can produce: while the long-term goal of the driver was stopping, the continuous selection of the most appropriate behavior let him/her suspend the stop behavior and brake at d'-f' in response to a temporary condition caused by front vehicles.

To model car following the layered architecture should be extended with new specific actions (say in place of m_k); in particular using a new type of final conditions in place of (2), (2'); such as used in e.g., [46]. This is not done here because out of scope. However, for the parameterization of

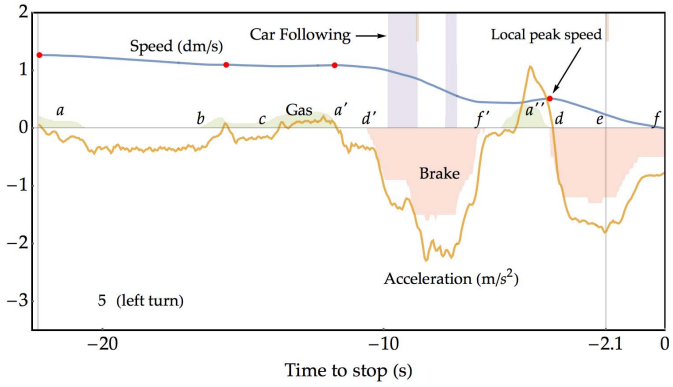

Fig. 5. Example of alternation of sub-behaviors induced by traffic. Note the acceleration adjustment after traffic clears at a''. The coding is the same of Fig.4: the y-axis scale serves for both speed (dm/s) and acceleration (m/s^2). Brake and gas are scaled to show proportionality with acceleration.

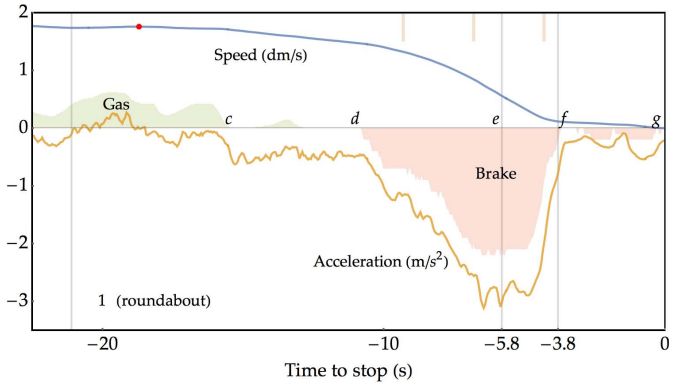

Fig. 6. Example 'rolling' stop. Close to stop the driver modulates the brake pressure to slowly creep across the yield location.

actions we will have to ignore maneuvers that are conditioned by traffic.

Fig.6 gives a final example, which is similar to Fig.4, except that, close to stop, the driver modulates the brake in order to creep across the yield line, as shown in the new phase $f-g$. This kind of behavior is referred as 'rolling stop'.

V. Model Instantiations and Parametrization

A. Model Instantiation

We study three different models, derived from the architecture given in Fig.1 and made of increasing number of motor primitives, i.e., with increasingly richer motor repertoire (we do not include rolling stop though).

1) Model 1: The simplest possible instantiation switches between only two motor primitives: m_0, to travel at constant desired speed v_d, and m_2 to stop by braking (Algorithm 1).

Algorithm 1

Require: $a_0, v_0, v_d, w_{T,0}, s_f, w_{T,}, w_A$
$m_0 \leftarrow FreeFlow(a_0, v_0, v_d, w_{T,0})$
$m_2 \leftarrow Stop(a_0, v_0, s_f, w_{T,}, w_A)$
$i \leftarrow argmin(j_0(0), j_2(0))$
Return m_i

2) Model 2: Model 2 switches among free-flow m_0 and two different ways of reducing the speed: m_1, by releasing the gas pedal, and m_2 by braking (Algorithm 2). It is assumed that m_1, makes no attempt at reducing the final acceleration to zero, i.e., $w_{A,1} = 0$. The gas $g \in [0, 1]$ and brake $b \in [0, 1]$ pedal positions are required and they are exclusively non-zero (when one is non-zero, the other must be zero).

Algorithm 2

Req.: $g, b, a_0, v_0, v_d, w_{T,0}, s_f, w_{T,1,}, w_{T,2,}, w_{A,2}$
$w_{A,1} \leftarrow 0$
$m_0 \leftarrow FreeFlow(a_0, v_0, v_d, w_{T,0})$
$m_1 \leftarrow Stop(a_0, v_0, s_f, w_{T,1,}, w_{A,1})$
$m_2 \leftarrow Stop(a_0, v_0, s_f, w_{T,2,}, w_{A,2})$
$i \leftarrow argmin(j_0(0), j_1(0))$
If $g > 0$ **Return** m_i // gas pedal pressed
If $b > 0$ **Return** m_2 // brake pedal pressed
$h \leftarrow argmin(j_0(0), j_1(0), j_2(0))$
Return m_h

If the gas pedal is pressed ($g > 0$) the algorithm selects between m_0 and m_1. If the brake is pressed ($b > 0$) the algorithm selects m_2. If both pedals are released the algorithm choses the action with minimum $j_h(0)$. If the latter acts on a saturated channel (e.g., $j_1(0) < 0$), it remains ineffective until superseded (e.g., by $j_2(0)$).

3) Model 3: Model 3 switches among free-flow m_0, two different ways of reducing the speed: m_1, and m_2 and an additional brake primitive m_3 specifically aimed at modelling the brake release phase e-f.

Algorithm 3

R: $g, b, a_0, v_0, v_d, w_{T,0}, s_f, w_{T,1,}, w_{T,2,}, w_{A,2}, w_{T,3,}, w_{A,3}$
$w_{A,1} \leftarrow 0$
$m_0 \leftarrow FreeFlow(a_0, v_0, v_d, w_{T,0})$
$m_1 \leftarrow Stop(a_0, v_0, s_f, w_{T,1,}, w_{A,1})$
$m_2 \leftarrow Stop(a_0, v_0, s_f, w_{T,2,}, w_{A,2})$
$m_3 \leftarrow Stop(a_0, v_0, s_f, w_{T,3,}, w_{A,3})$
$i \leftarrow argmin(j_0(0), j_1(0))$
$l \leftarrow argmin(j_2(0), j_3(0))$
If $g > 0$ **Return** m_i // gas pedal pressed
If $b > 0$ **Return** m_l // brake pedal pressed
$h \leftarrow argmin(j_0(0), j_1(0), j_2(0), j_3(0))$
Return m_h

Like Model 2, this algorithm selects between m_0 and m_1 when the gas pedal is pressed and between m_2 and m_3 when the brake is pressed.

Caveats: Stop motor primitives become singular when $s_f \to 0$ (close to the stop line). That is because when $s_f \to 0$, also $T \to 0$. We prevent the singularity by switching to a special 'stop' maneuver when the speed is below 0.5 m/s which is implemented as a maximally smooth ($w_T = 0$) free-flow primitive aimed at zero speed: $m_4 = FreeFlow(a_0, v_0, 0, 0)$.

In the cognitive architecture of Fig.1, the switch from one of Models 1, 2, 3 to m_4, can be obtained by biasing action-selection to m_4.

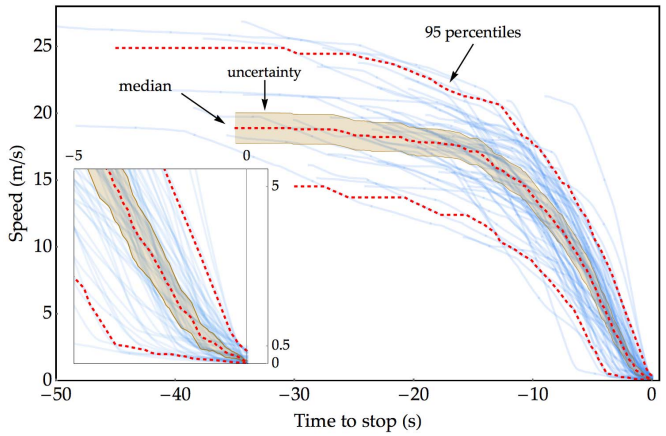

Fig. 7. Speed of 66 stop maneuvers not affected by traffic. The median, 95 and 5 percentile speed and an estimation of the uncertainty band are shown.

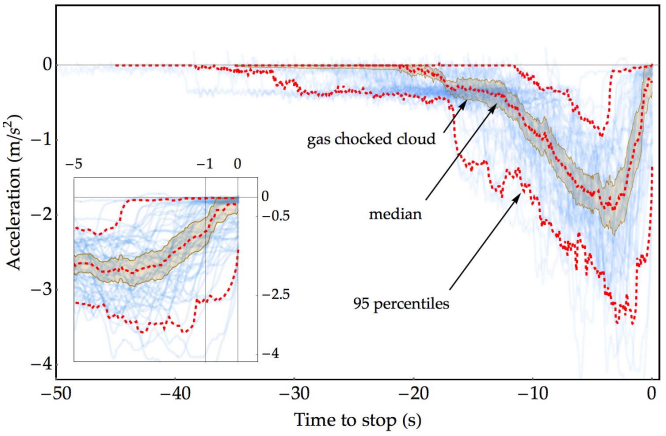

Fig. 8. Acceleration of 66 stop maneuvers not affected by traffic. The median, 95 and 5 percentile speed and an estimation of the uncertainty band are shown. Note a dense cloud at about -0.37 m/s^2 which corresponds to travelling with fully release gas pedal (and brake not pressed), i.e., the coasting phase. Note that a fraction of maneuvers, about a half, end with finite acceleration.

B. Parameterization of Models

For parameterization of the motor primitives of the 3 models we have considered only the maneuvers that were not delayed by the traffic, such as, e.g., Fig.4 and Fig.6, but not Fig.5.

1) All Types of Stop: Considering al stop locations, there are 66 stop maneuvers that were not influenced by traffic. For these, Fig.7 shows the speed profiles, an estimation of the 95 and 5 percentiles, the median speed and an estimation of the uncertainty region (taken as $3\sigma/\sqrt{n}$ where σ is the standard deviation and $n = 66$). In the inset, a magnified view of the last 5 seconds is given, which allows appreciating the effect of rolling stop maneuvers (particularly clear for the 5-percentile curve) and the final residual acceleration.

Fig.8 is the analogous of Fig.7 for the acceleration. A dense cloud at approximately $a_c = -0.37$ m/s^2, between 25 and 10 s before stop, represents the 'chocked gas' states (e.g., *a-b* Fig.5). Note also that the 5-percentile curve ends with zero acceleration indicating 3-4 seconds rolling stops (inset). The median curve has a very small residual end-acceleration a_f which is instead much larger for the 95-percentile curve.

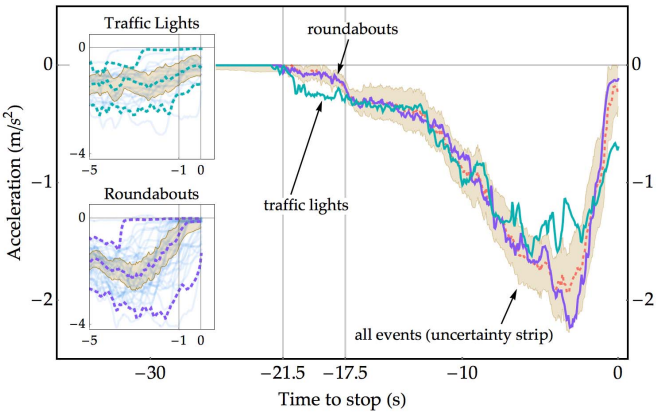

Fig. 9. Comparison of the median acceleration for roundabouts, traffic light and the whole set.

2) Traffic Light and Roundabout Subsets: To investigate the influence of the stop type, the stop maneuvers corresponding to the two traffic lights (locations 12 and 16) have been separated, and the corresponding distributions and medians computed. The same has been done for the 9 simple roundabouts (locations, 1, 2, 3, 4, 9, 11, 13, 14, 15).

Fig.9 compares the median acceleration for traffic lights (21 maneuvers), roundabouts (37 maneuvers) and the whole set in the background (i.e., the same of Fig.8, 66 maneuvers including also locations 5, 6, 7, 8, 10 that are neither traffic lights nor simple roundabouts, Fig.3).

A Wilcoxson-Mann-Whitney test has been used to test whether the roundabout and traffic lights medians are significantly different in the three regions compared below.

1) One first difference may be noticed for the gas release, which for traffic lights tends to occur ∼5 s earlier. This may be explained with the fact that the future states of the traffic light, and thus whether it is necessary to stop, tend to be predictable earlier, as traffic lights are placed in long straights and visible from long distance. Conversely, at roundabouts drivers will know later, and closer to the yield line, whether they can cross. Statistically, this difference is however not significant, probably because the distributions are bimodal, in this region, and not different enough, even if the medians appear different.

2) The maximum deceleration at traffic lights is less than at roundabouts (i.e., overall traffic light stop occur with less deceleration over longer time). This difference is significant (Wilcoxon-Mann-Whitney $p < 0.002$ at $t \in [-4.4s, -2.1s]$).

3) Finally, the residual acceleration at stop tends to be non-zero for traffic lights, but essentially zero at roundabouts. This difference is also significant (Wilcoxon-Mann-Whitney $p \approx 0.006$ at $t \in [-0.4s, 0s]$).

From the insets, it can be seen that at roundabouts the median acceleration is close to zero for about the last second, which represents frequent rolling stop conditions lasting abut one second for the median case, and about 3-4 s for the 5-percentile case. This latter difference can be explained with the fact that at traffic light there is a 'hard' stop condition, whereas at roundabouts there is a yield sign; hence drivers creep into the roundabout till they can decide whether to stop or cross.

TABLE I
COMPARISON OF MODELS

Model	Parameters	Acceleration error rms
Complete data set (all location types)		
Model 1. Layered control with 1stop motor primitive.	$w_T = 0.545\ m^2s^{-6}$ $w_A = 2.088\ s^{-1}$	$0.177\ m/s^2$
Model 2. Layered control with 2 stop motor primitives.	$w_{T,1} = 0.0507\ m^2s^{-6}$ $w_{T,2} = 0.778\ m^2s^{-6}$ $w_{A,2} = 6.891\ s^{-1}$	$0.090\ m/s^2$
Model 3. Layered control with 3 stop motor primitives.	$w_{T,1} = 0.0507\ m^2s^{-6}$ $w_{T,2} = 0.368\ m^2s^{-6}$ $w_{A,2} = 0.1317\ s^{-1}$ $w_{T,3} = 1.689\ m^2s^{-6}$ $w_{A,3} = \infty$	$0.049\ m/s^2$
Intelligent Driver Model	$s_0 = 0.252\ m$ $s_1 = 0$ $\tau = 0.573\ s$ $b = 1.677\ ms^{-2}$	$0.065\ m/s^2$
Traffic light data subset (location 12 and 16)		
Model 1. Layered control with 1stop motor primitive.	$w_T = 0.168\ m^2s^{-6}$ $w_A = 0.299\ s^{-1}$	$0.154\ m/s^2$
Model 2. Layered control with 2 stop motor primitives.	$w_{T,1} = 0.0260\ m^2s^{-6}$ $w_{T,2} = 0.319\ m^2s^{-6}$ $w_{A,2} = 0.663\ s^{-1}$	$0.082\ m/s^2$
Model 3. Layered control with 3 stop motor primitives.	$w_{T,1} = 0.0260\ m^2s^{-6}$ $w_{T,2} = 0.320\ m^2s^{-6}$ $w_{A,2} = 0.607\ s^{-1}$ $w_{T,3} = 0.419\ m^2s^{-6}$ $w_{A,3} = 0.769\ s^{-1}$	$0.082\ m/s^2$
Intelligent Driver Model	$s_0 = 4.96\ m$ $s_1 = 0$ $\tau = 1.078\ s$ $b = 1\,440\ ms^{-2}$	$0.093\ m/s^2$
Roundabout data subset (locations 1, 2, 3, 4, 9, 11, 13, 14, 15).		
Model 1. Layered control with 1stop motor primitive.	$w_T = 0.676\ m^2s^{-6}$ $w_A = 3.997\ s^{-1}$	$0.234\ m/s^2$
Model 2. Layered control with 2 stop motor primitives.	$w_{T,1} = 0.0509\ m^2s^{-6}$ $w_{T,2} = 0.989\ m^2s^{-6}$ $w_{A,2} = \infty$	$0.146\ m/s^2$
Model 3. Layered control with 3 stop motor primitives.	$w_{T,1} = 0.0509\ m^2s^{-6}$ $w_{T,2} = 0.283\ m^2s^{-6}$ $w_{A,2} = 0.0207\ s^{-1}$ $w_{T,3} = 4.669\ m^2s^{-6}$ $w_{A,3} = \infty$	$0.069\ m/s^2$
Intelligent Driver Model	$s_0 = 0.049\ m$ $s_1 = 0$ $\tau = 0.330\ s$ $b = 1.636\ ms^{-2}$	$0.092\ m/s^2$

C. Motor Primitive Fitting

The three models have been fitted onto the median acceleration curve. The best fitting parameters are listed in Table I.

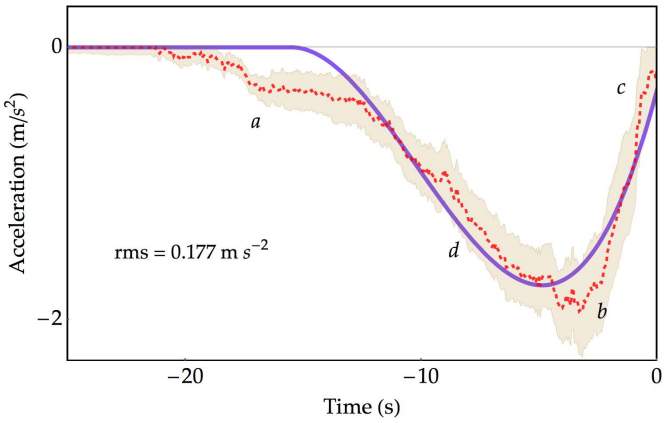

Fig. 10. Model 1 best fit.

For Model 1, the fit is shown in Fig.10. With only one brake primitive, Model 1 does not describe the initial gas release (Fig.10, label a). Also, the maximum deceleration (label b) and the final acceleration (label c) are not captured very well. The traffic light sub-set is fitted a little better than the roundabout sub-set (Table I) because there is less correction of the final speed. Recalling the interpretation given in section III.A.1, the parameter w_A is larger for roundabouts because the final acceleration a_f has to be closer to zero; w_T is smaller for traffic lights because deceleration begins earlier.

Fig.11, left, shows Model 2 fits, respectively for traffic lights (top), roundabouts (bottom) and all scenarios (center). For the latter, at time -23.15 s, the model switches from free-flow to deceleration; at time -16.05 s the gas pedal saturates. At time -13.45 s the brake primitive is engaged. Hence, compared to Model 1 (Fig.10), the deceleration phase produced with releasing the gas pedal (label a) is now modeled. Also, the fit in the brake phase and, in particular, with the final acceleration adjustment (d-c) is improved.

Modeling the gas release, which can be achieved with two stop primitives, may thus be an important modeling innovation for at least two reasons: the first is mirroring of the vehicle by a driver or passenger, who will understand that the automated vehicle is going to stop, even before brake activation; the second reason is energy saving (if the agent is used for automation). In fact, fuel consumption is cut when the gas is chocked. Hence maneuver in Fig.11 uses less fuel than that of Fig.10, which would instead continue to feed the engine until the brake phase begins.

Fig.11 (center column) shows the maneuver that is produced by Model 3, which obtains a further improvement thanks to modeling the final acceleration adjustments.

For the traffic light case (top row) Model 2 and Model 3 fit the same way. This is confirmed by similar residuals in Table I. In fact, the third stop primitive of Model 3 – $\{w_{T,3}, w_{A,3}\}$ in Table I – is not substantially different from the second primitive – $\{w_{T,2}, w_{A,2}\}$ – which is also very similar to $\{w_{T,2}, w_{A,2}\}$ of Model 2. Thus, both Model 3 and Model 2 have substantially only one brake phase, because there is no important correction of acceleration at the end for traffic light approaches.

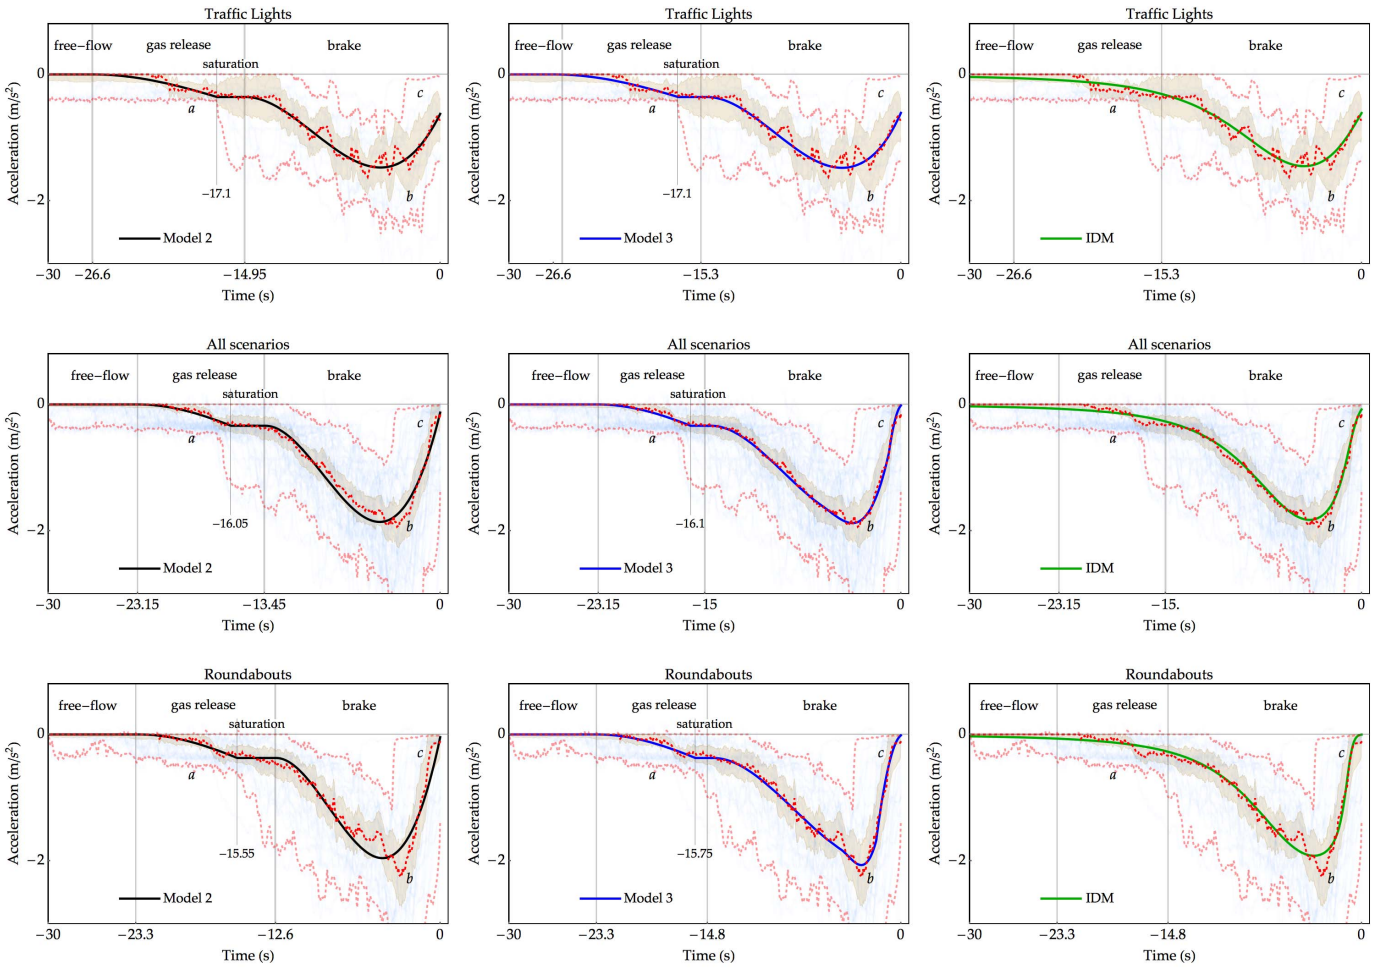

Fig. 11. Comparison of models with experimental data. Top row: traffic light subset. Center row: complete dataset. Bottom: roundabouts subset. Left: Model 2: layered control with two stop primitives (gas release and brake). Centre: Model 3: layered control with three stop primitives (gas release, brake, brake release) and simple rolling stop. Right: Intelligent Driver Model (IDM).

For roundabouts (bottom row) the final brake release (anticipating possible rolling stop) is instead important. Hence, Model 3, fits better: the triangular shape of the median acceleration is reproduced with greater accuracy, and in particular the sharp peak of deceleration (label b) and the final adjustment (label c). The root mean square fitting residuals in Table I confirm the better fit; with acceleration error of the order of $0.05 - 0.08\ m/s^2$ rms for Model 3 (part of which is data noise and only part remaining model approximations). The mid row of Fig.11 refers to the complete set. The fitting of Model 3 is remarkable, which is also because the complete set has more samples and thus less noise, allowing to better appreciate the modeling errors.

D. Residual Significance Analysis

We have shown that some models look to better fit the experimental data, whereas other models show systematic deviations in some zones of the curve (for example, Model 1 departs systematically from the fitting curve in zones labeled a, b, c, and d, Fig.10).

To test whether these deviations are *statistically significant* the time domain has been divided into bins of 0.1s. For each bin the distribution of all observed accelerations has been compared to model predictions. The null hypothesis was that the acceleration predicted by the model in the bin was the true median of the observed accelerations in the same bin (which is equivalent to testing whether the median of the residual in each bin is significantly different from zero).

For each bin the most appropriate statistical location test was carried out, depending on the population distribution (the Student t-test in bins where the distribution was normal; otherwise the Wilcoxson signed rank test or the Fisher sign test as a last resource).

Table II summarizes the results, listing the intervals where statistically significant deviations occurs (the p value is the probability that the acceleration of the model is the true median of the data). Each interval is located in one zone (Fig.10): zone a for model deficiency on the coasting phase, zone b for the peak acceleration, zone c for the final acceleration and zone d for the phase preceding the peak acceleration.

Model 1 departs from the experimental data in zone a for all the three datasets because it does not model the coasting phase at all. It also does not model well the peak and pre-peak acceleration (b, d) for the complete dataset and for the roundabout dataset. Lastly, it does not model well the final acceleration for the roundabout set, zone c.

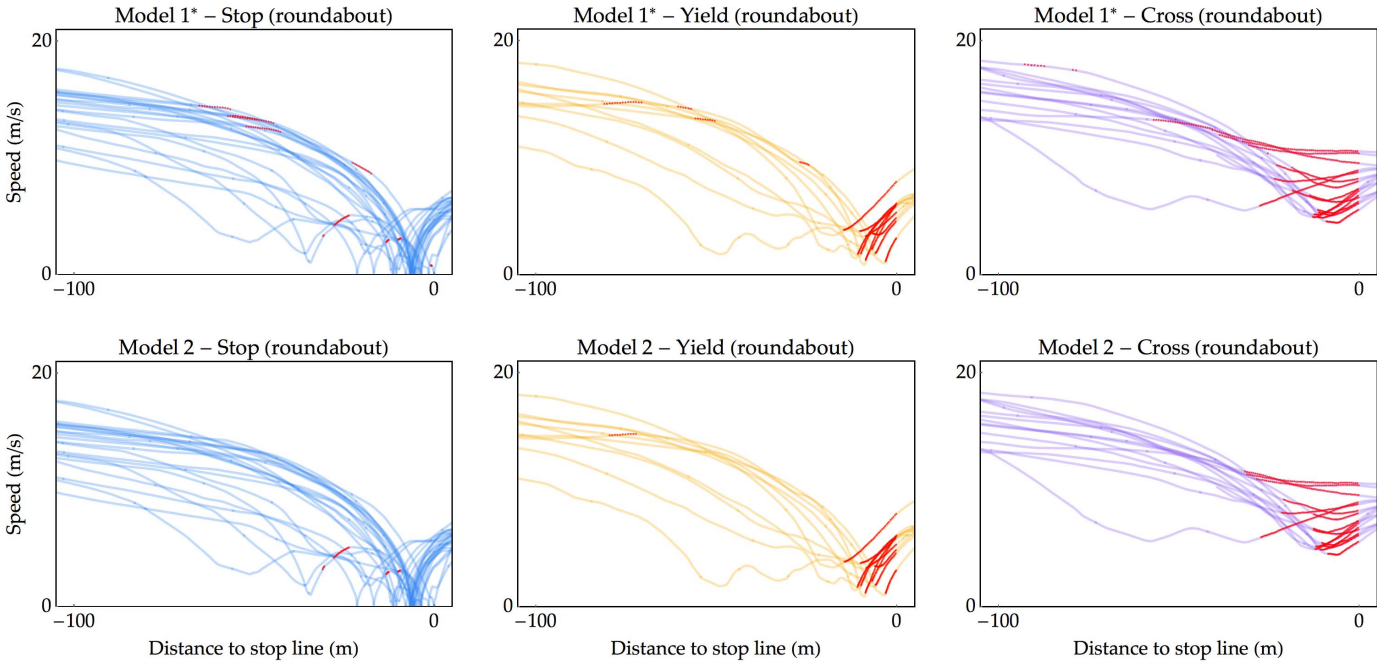

Fig. 12. Comparison of mirroring the stop actions by Model 1* and Model 2. The trajectories corresponding to the roundabout at location 1 (Fig.3) are shown. The three types of maneuvers, stop, yield and cross are shown from left to right columns. Red dots show where Model required jerk fall below the 0.05 quantile threshold. This is where a hypothetical advisory warning associated with that threshold would be issued.

Model 2 improves over Model 1. There is no statistically significant deviation for the traffic light dataset; there is a deviation in zone *d-b* for the complete dataset and for the roundabout dataset (related to the difficulty of modeling the peak acceleration and the preceding linear phase); and there is a deviation for the final acceleration in roundabouts (zone *c*).

Model 3 has no statistically significant deviations from data for all the three datasets.

E. Bootstrap Analysis

A bootstrap analysis [54] has been carried out to estimate the intervals of variation of the model parameters.

Since the complete dataset is a combination of heterogeneous situations (roundabouts, curves, traffic light, etc.) we have restricted the bootstrap analysis to the traffic light and roundabout cases separately. This way the estimated variations in the model parameters reflect variations due to different drivers, together with different executions for the same driver, in consistently similar situations (e.g., roundabouts vs. roundabouts). It is implicit that, with more data this same analysis could be carried out for individual drivers, thus finding personalized range of variations of the model parameters.

Table III gives an estimate of the range of variability of the model parameters. These has to be regarded as rough estimates, since the number of resamples that could be used was limited by the long computation time required by each re-fitting procedure for Models 2 and 3. For the same reason we report estimates for the 10 and 90 percentiles, rather than 5 and 95 percentiles, that turned out to be more stable within different resample numbers.

The range of variation of the model parameters has to be interpreted as the consequence of different driving styles (of both different drivers and of the same driver in different moments). Hence, the 10 percentiles tend to represent slow maneuvers (beginning early) and the 90 percentiles fast ones. Approximately the parameters vary by a factor of 2, except w_A. In evaluating the variations of the latter, one must however recall that it needs to vary between 0 and ∞ to switch between free and zero final acceleration (i.e., large variations in w_A are required to obtain small changes in the final acceleration).

VI. Comparison With the Intelligent Driver Model

The Intelligent Driver Model (IDM) is given in (9).

$$\dot{v} = a\left(1 - \left(\frac{v}{v_d}\right)^{\delta} - \left(\frac{s^*}{s}\right)^2\right) \tag{9}$$

where v is the velocity, $v_d$the desired velocity, s the gap with the preceding vehicle and s^* the desired gap:

$$s^* = s_0 + s_1\sqrt{\frac{v}{v_d}} + \tau v + \frac{v\,\Delta v}{\sqrt{4ab}} \tag{10}$$

where Δv is the velocity difference with the preceding vehicle (for modeling stop maneuvers $\Delta v = v$).

The model has parameters: $s_0, s_1, \tau, a, b, v_d, \delta$, which may be interpreted as modeling: a) the free flow (a, v_d, δ); in particular v_d has the same meaning of aimed velocity it has in the motor primitive (A1.2) and a, δ model the way v_d is approached, just like $w_{T,0}$ does in the motor primitive; b) the deceleration for car following or for stopping (s_0, s_1, τ, b); in particular s_0 means the clearance left at stop, τ the aimed time headway, b the maximum deceleration, s_1 models a velocity related clearance but is often set to zero.

Using the words of the authors [17], the model 'interpolates the tendency to accelerate with acceleration $a_f = a(1 - (v/v_d)^\delta)$ on a free road and the tendency to brake with deceleration $a_b = -a(s^*/s)^2$ when approaching a front vehicle'.

From the perspective of human behavior modeling, a_f, and a_b may be regarded as the response to two affordances (drive at desired speed and brake to keep relative position and velocity with respect to vehicle ahead). These are *added* in the resulting action $\dot{v}$ (9). However, in human behavior action selection almost always is not additive but exclusive (either one or another is realized).

With these observations, one can expect that (9) captures large scale phenomena resulting from 'combination' of free-flow and following behaviors – just like expected from the architecture of Fig.1 – but will miss features that arise from *competition* of behaviors, such as the *switching* between different phases/states, asymmetries of different motor channels and channel saturation.

The IDM has been fitted onto the same three data sets. Fitting parameters and residuals are given in Table I. For the actual fitting, parameters (s_0, s_1, τ, b) have been allowed to vary but parameters (a, v_d, δ), which are designed to model the free flow, have been set to the values indicated in [17] ($a = 0.73\ m/s^2, \delta = 4$), and v_d the initial velocity (in the IDM, there is coupling between parameters used to model the free flow acceleration a_f, and the car following acceleration a_b).

Fig.11, right, plots the IDM model in the three scenarios, for comparison with Model 2 and Model 3. As expected, the IDM captures the large-scale phenomena that are described by the two considered affordances, a_f and a_b, and it does this remarkably well, but misses the details that are due to the exclusive selection of one or another motor primitive. In the traffic light case, the modeling of gas release and the saturation of acceleration is not captured (label a). In the roundabout case, the peak deceleration (b) and the final correction (c), which are the result of switching between two brake modalities cannot be captured likewise. In addition, while the parameters for Models 2 and 3 in Table I have been easily interpreted for the three stop scenarios, the variation of the parameters of the IDM, for the same three scenarios, is more difficult to interpret (e.g., what does the time headway parameter τ mean for a stop condition? and how can the three different values for the three scenarios be explained?).

Table II shows that statistically significant deviations exist for the IDM in two parts of zone a (before and after the gas pedal release), which again depends on the fact that "interpolating" between behavior is not biologically plausible. The IDM also does not model well the peak and pre-peak acceleration for the complete dataset and the roundabouts dataset (zone d-b). For the latter dataset, there also is a mismatch at the final (rolling) phase, zone c.

Concerning Table III, the variation of the model parameters is typically much larger than the variations of the other models.

TABLE II
RESIDUAL SIGNIFICANCE ANALYSIS

Model	Statistically significant deviations		Zones (as in Fig.10)
	Complete data set (all location types)		
Model 1.	$t \in [-4.1s, -2.0s]$,	$p = 0.007$	b
	$t \in [-9.2s, -6.7s]$,	$p = 0.011$	d
	$t \in [-21.5s, -11.85s]$,	$p < 0.001$	a
Model 2.	$t \in [-8.75s, -4.7s]$,	$p = 0.016$	d
Model 3.			
Intelligent Driver Model	$t \in [-11.15s, -9.3s]$,	$p = 0.015$	d
	$t \in [-16.95s, -15.95s]$,	$p = 0.024$	a
	$t \in [-25s, -20.95s]$,	$p = 0.009$	a
	Traffic light data subset (location 12 and 16)		
Model 1.	$t \in [-21.5s, -14.7s]$,	$p = 0.027$	a
Model 2.			
Model 3.			
Intelligent Driver Model	$t \in [-25, -22.8s]$,	$p = 0.029$	a
	Roundabout data subset (locations 1, 2, 3, 4, 9, 11, 13, 14, 15).		
Model 1.	$t \in [-0.8s, 0s]$,	$p = 0.014$	c
	$t \in [-1.8s, -3.4s]$,	$p = 0.004$	b
	$t \in [-6.55s, -4.85s]$,	$p = 0.007$	d
	$t \in [-21.5s, -11.5s]$,	$p < 0.001$	a
Model 2.	$t \in [-0.75s, 0s]$,	$p = 0.013$	c
	$t \in [-6.55s, -4.45s]$,	$p = 0.003$	$d - b$
Model 3.			
Intelligent Driver Model	$t \in [-0.6s, 0s]$,	$p < 0.001$	c
	$t \in [-5.35s, -4.45s]$,	$p = 0.030$	$d - b$
	$t \in [-17s, -15.6s]$,	$p = 0.015$	a
	$t \in [-25s, -21.95s]$,	$p = 0.025$	a

VII. IMPACT

One first impact of this work is definitely methodological: a framework based on bio-inspired principles has been introduced for both the modelling of drivers and for developing agents for driver support/interaction. Indeed, the models presented here are the logical application –to the specific case of vehicle stopping– of the affordance competition hypothesis and of known human motor optimality principles (these principles constrain the models before fitting to experimental data is even attempted). One can expect that the adoption of such principles may allow to design better agents for human driver support beyond the application domain specifically studied in this paper.

Driver models derived in this way have perhaps the main strength in human vehicle interactions. The similarity of sensorimotor control sought in this paper approach is indeed finalized to mirroring (both robot mirroring the human and vice-versa). We provide here one example; however, on this aspect the reader is invited to consider our previous work which present more extensive explanations [46]–[48].

1) Using the Models for Mirroring: The basic idea of mirroring is using one agent as a model for another agent.

TABLE III
INTERVALS OF VARIATION OF MODEL PARAMETERS

Model		10 percentile	90 percentile
Traffic light data subset (location 12 and 16)			
Model 1.	$w_T\,(m^2 s^{-6})$	0.053	0.26
	$w_A\ (s^{-1})$	0.048	0.56
Model 2.	$w_{T,1}(m^2 s^{-6})$	0.016	0.039
	$w_{T,2}(m^2 s^{-6})$	0.21	0.43
	$w_{A,2}\ (s^{-1})$	0.38	1.3
Model 3.	$w_{T,1}(m^2 s^{-6})$	0.016	0.035
	$w_{T,2}(m^2 s^{-6})$	0.16	0.34
	$w_{A,2}\ (s^{-1})$	0.14	1.3
	$w_{T,3}(m^2 s^{-6})$	0.30	0.98
	$w_{A,3}\ (s^{-1})$	0.33	1.7
Intelligent Driver Model.	$b\ (ms^{-2})$	1.03	6.8
	$\tau\ (s)$	0.50	4.0
	$s_0(m)$	0.85	20.
Roundabout data subset (locations 1, 2, 3, 4, 9, 11, 13, 14, 15).			
Model 1.	$w_T\,(m^2 s^{-6})$	0.49	0.94
	$w_A\ (s^{-1})$	3.1	∞
Model 2.	$w_{T,1}(m^2 s^{-6})$	0.034	0.063
	$w_{T,2}(m^2 s^{-6})$	0.74	1.3
	$w_{A,2}\ (s^{-1})$	∞	∞
Model 3.	$w_{T,1}(m^2 s^{-6})$	0.028	0.065
	$w_{T,2}(m^2 s^{-6})$	0.25	0.49
	$w_{A,2}\ (s^{-1})$	0.	0.19
	$w_{T,3}(m^2 s^{-6})$	4.2	6.1
	$w_{A,3}\ (s^{-1})$	∞	∞
Intelligent Driver Model.	$b\ (ms^{-2})$	1.5	1.8
	$\tau\ (s)$	0.17	0.44
	$s_0(m)$	0.012	0.091

Facing the same situation, the model is expected to produce affordances that can be compared to the actual human driver behavior.

To implement the complete mirroring process, the agent should produce a complete set of affordances. This is not the case for the models of this paper, which produce affordances for only two possible goals (free flow and stop). Hence, Models 1-3 can be used in a restricted way here: only to assess whether a stop behavior falls within the actions that can be human-directed.

We give an example comparing Model 2 to Model 1 in a roundabout. Choice of Model 2 is motivated because it the simplest model involving both the gas pedal and brake pedal as distinct motor channels. Choice of Model 1 is because it is the simples model, with only one primitive for stopping.

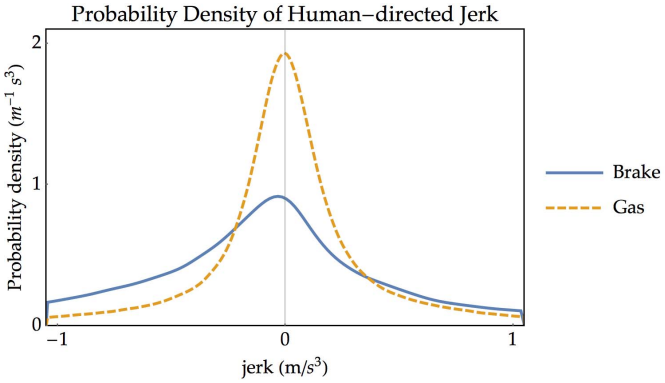

Fig. 13. Distribution of human-directed jerk produced with the use of the gas or brake pedal. Driving data from the InteractIVe dataset.

TABLE IV
HUMAN-DIRECTED JERK

Quantile	0.005	0.05	0.95	0.995
Brake pedal	$-3.1\ m/s^3$	$-1.6\ m/s^3$	$1.5\ m/s^3$	$3.7\ m/s^3$
Gas pedal	$-2.5\ m/s^3$	$-0.89\ m/s^3$	$0.9\ m/s^3$	$2.5\ m/s^3$
Total	$-2.6\ m/s^3$	$-1.1\ m/s^3$	$1.0\ m/s^3$	$2.7\ m/s^3$

Moreover, for Model 1 we use the parametrization of [48] ($w_T = 0$, $w_A = \infty$, instead of the figures given in Table I[2]) for comparison with that work, which further develops the notion of mirroring in intersection scenarios. We call this Model 1*.

Mirroring in practice is implemented using the driver models in parallel with the human driving. With Algorithms 1 and 2, at any time the actions that the models would implement for stopping can be computed (we ignore the free flow primitive m_0 because we aim to assess only whether a stop behavior is 'feasible' for the human driver). To decide whether Model 1* and Model 2 actions can be human directed we compare the required longitudinal control of the models $j_i(0)$ to the distribution of longitudinal jerk observed in human driving. Fig.13 shows the longitudinal jerk distributions for human drivers, making distinction between gas and brake pedal.

Table IV reports the jerk of different quantiles. A rational approach is thus setting a probability threshold (e.g., 0.05 quantiles) and deem that, with that probability, the human driver will not carry out the action required by the model (for advisory and cautionary warnings two distinct thresholds can be selected).

Fig.12 compares mirroring with model 1* and 2 at 0.05 quantiles. Red dots show where the model required jerk falls below the chosen quantile threshold.

Model 1*, which uses only one stop primitive tend to confuse fast approaches with violations of the stop line (as shown by dots on the highest curves). This does not happen for Model 2, which 'is aware' that actions occurring on the brake can have larger jerk.

[2]The difference in actual trajectories is not large: the main difference is that the final acceleration is exactly zero.

In case of yield violations (Fig.12, center, right) both models detect the critical condition almost at the same time.

An interesting comment is due for the alarms that are produced when drivers 'adjust' their position (Fig.12, left, bottom curves), i.e., when the driver resumes speed to move closer to the stop line. This situation, in both models, and in a few cases, uses barely[3] enough acceleration to be confused with the intention to trespass the stop line. What happens is that both models interpret the driver behavior in the library of actions that they have, which do not include any behavior like 'move closer to the stop line'. If such a motor behavior were added, the models would find that both 'trespass the stop line' and 'move closer to the stop line' actions would be a potential fit, and if the fitness were comparable, the latter would win because being the most logical in the context (of course, if the driver insists with the acceleration at some point the intention of trespassing the line would be clearer).

VIII. Conclusions

The main contribution of this paper is the demonstration that drivers can be modelled with layered control architectures, including bio-inspired selection mechanisms, conforming to recent developments in cognitive science.

The engineering of an agent having motor behaviors similar to humans is relatively simple: motor primitives may be produced (with optimal control) to represents simple motor units; and, at higher level, many motor units may be simultaneously primed, and the most suited one selected, thus implementing adaptive behavior.

In order to correctly produce the architecture (both primitives and selection mechanisms) some basic understanding of human motor primitives and selection criteria is necessary.

This paper casts light on stop maneuvers. However, with similar knowledge of how humans respond to other types of affordances, layered architecture could scale to more complex scenarios.

The paper has shown that stopping is indeed made of 3-4 different motor units: deceleration by releasing the gas pedal, braking, releasing the brake to 'smooth' the stop or releasing the brake to produce rolling stop are the main units of this vocabulary. These motor units are not rigidly programmed in sequences because their activation is the result of a competition process (nor any finite state machine is needed).

Model 3 fits the median behaviors in three scenarios with remarkable accuracy. Model 2, missing the motor primitive designed for the final correction at stop, is somewhat less accurate. Both however model the transition from gas pedal release to brake activation which may be important for enabling mirroring processes.

Compared with the state of the art Intelligent Driver Model, both Model 2 and Model 3 describe features (transitions and states) that IDM misses. Model 3 also achieves the best fits ever (Table I).

[3]If the quantile thresholds id lowered only a little these situations disappear because they are borderline.

The main application domain of such kind of modelling is designing agents for human-vehicle interactions, as shown with one example. Models 1 and 2 have been parametrized (including jerk distributions) using data from a naturalistic driving set that involved 25 different drivers and one vehicle. To parametrize the models in other situations (e.g., different countries, different types of vehicles, etc.) or for driver personalization, similar (small size) experimental campaigns may be necessary.

Appendix I
Stop Motor Primitives

The solution of the Optimal Control Problem (1), (2), (3), (4) yields the following three outcomes:

1) The optimal movement time T, which is *the smallest real positive root* of the following 8th order polynomial equations.

$$\begin{aligned} &w_A w_T^2 T_8 + 18 w_A w_T T^7 + 9(9 w_T - a_0^2 w_A^2) T^6 + \\ &- 144 a_0 w_A (a_0 + v_0 w_A) T^5 - 72(9 a_0^2 + 8 v_0^2 w_A^2 + \\ &- 5 a_0 w_A (s_f w_A - 6 v_0)) T^4 + 144(-7 a_0 (9 v_0 - 5 s_f w_A) + \\ &+ 4 v_0 w_A (-14 v_0 + 5 s_f w_A)) T^3 + 144(135 a_0 s_f - 216 v_0^2 + \\ &+ 260 s_f v_0 w_A - 25 s_f^2 w_A^2) T^2 - 43200 s_f (-3 v_0 + s_f w_A) T + \\ &- 129600 s_f^2 = 0 \end{aligned} \tag{11}$$

We know (11) has at least one real positive root (the polynomial is negative for $T = 0$ and tend to $+\infty$ for $T \to +\infty$). The meaning of the first root, and of possible further roots was already discussed in [47]. It may be sufficient to say here that further real roots of (11), if they exist, represent additional strategies to get to the stop point, which however may imply inversion of movement or slow-down and re-acceleration (in both cases taking longer time). Our main interest here is to find the shortest time to stop; hence the first root. For the computation of the polynomial roots, the Jenkins-Traub algorithm has been reliably used [55].

2) The optimal final acceleration:

$$a_f = \frac{3(8 T v_0 + a_0 T^2 - 20 s_f)}{T^2 (9 + T w_A)} \tag{12}$$

3) The motor primitive (and derivatives):

$$\begin{aligned} s(t) &= c_1 t + \frac{1}{2} c_2 t^2 + \frac{1}{6} c_3 t^3 + \frac{1}{24} c_4 t^4 + \frac{1}{120} c_5 t^5 \\ v(t) &= c_1 + c_2 t + \frac{1}{2} c_3 t^2 + \frac{1}{6} c_4 t^3 + \frac{1}{24} c_5 t^4 \\ a(t) &= c_2 + c_3 t + \frac{1}{2} c_4 t^2 + \frac{1}{6} c_5 t^3 \\ j(t) &= c_3 + c_4 t + \frac{1}{2} c_5 t^2 \end{aligned} \tag{13}$$

with:

$$\begin{aligned} c_1 &= v_0, \quad c_2 = a_0 \\ c_3 &= 60 \frac{s_f}{T^3} - 36 \frac{v_0}{T^2} + 3 \frac{a_f - 3 a_0}{T} \\ c_4 &= -360 \frac{s_f}{T^4} + 192 \frac{v_0}{T^3} + 6 \frac{6 a_0 - 4 a_f}{T^2} \\ c_5 &= 720 \frac{s_f}{T^5} - 360 \frac{v_0}{T^4} + 60 \frac{a_f - a_0}{T^3} \end{aligned} \tag{14}$$

APPENDIX II
FREE FLOW MOTOR PRIMITIVES

The solution of the Optimal Control Problem (1), (2'), (3), (4) yields the following three outcomes[4]:

The primitives have the same structure given in (12), but with coefficients given by [47]:

$$c_1 = v_0, \quad c_2 = a_0, \; c_5 = 0$$
$$c_3 = 6\frac{(v_d - v_0)}{T^2} - 4\frac{a_0}{T}, \quad c_4 = 6\frac{a_0}{T^2} + 12\frac{(v_d - v_0)}{T^3} \tag{15}$$

where the movement time is given by the first real positive root of the following polynomial:

$$w_T T^4 - 4a_0^2 T^2 + 24a_0 T(v_d - v_0) - 36(v_d - v_0)^2 = 0 \tag{16}$$

REFERENCES

[1] C. C. Macadam, "Understanding and modeling the human driver," *Vehicle Syst. Dyn.*, vol. 40, nos. 1–3, pp. 101–134, Jan. 2003.

[2] M. Plöchl and J. Edelmann, "Driver models in automobile dynamics application," *Veh. Syst. Dyn.*, vol. 45, no. 7/8, pp. 699–741, 2007.

[3] F. Biral, D. Bortoluzzi, V. Cossalter, and M. Da Lio, "Experimental study of motorcycle transfer functions for evaluating handling," *Vehicle Syst. Dyn.*, vol. 39, no. 1, pp. 1–25, 2003.

[4] P. Cisek, "Cortical mechanisms of action selection: The affordance competition hypothesis," *Philos. Trans. Roy. Soc. London A, Math. Phys. Sci.*, vol. 362, no. 1485, pp. 1585–1599, 2007.

[5] L. Cattaneo and G. Rizzolatti, "The mirror neuron system.," *Arch. Neurol.*, vol. 66, no. 5, pp. 557–560, 2009.

[6] S. Hurley, "The shared circuits model (SCM): How control, mirroring, and simulation can enable imitation, deliberation, and mindreading.," *Behav. Brain Sci.*, vol. 31, no. 1, pp. 1–22, Feb. 2008.

[7] D. M. Wolpert, K. Doya, and M. Kawato, "A unifying computational framework for motor control and social interaction.," *Philos. Trans. Roy. Soc. London A, Math. Phys. Sci.*, vol. 358, no. 1431, pp. 593–602, Mar. 2003.

[8] G. Hesslow, "The current status of the simulation theory of cognition," *Brain Res.*, vol. 1428, pp. 9–71, Jan. 2012.

[9] G. Hesslow, "Conscious thought as simulation of behaviour and perception," *Trends Cogn. Sci.*, vol. 6, no. 6, pp. 242–247, Jun. 2002.

[10] J. Decety and J. Grézes, "The power of simulation: Imagining one's own and other's behavior," *Brain Res.*, vol. 1079, no. 1, pp. 4–14, Mar. 2006.

[11] A. Meltzoff, "Understanding the intentions of others: Re-enactment of intended acts by 18-month-old children," *Develop. Psychol.*, vol. 31, no. 5, pp. 838–850, 1995.

[12] Y. Demiris, "Prediction of intent in robotics and multi-agent systems.," *Cognit. Process.*, vol. 8, no. 3, pp. 151–158, Sep. 2007.

[13] M. Jeannerod, "The representing brain: Neural correlates of motor intention and imagery," *Behav. Brain Sci.*, vol. 17, no. 2, pp. 187–202, 1994.

[14] M. Jeannerod, "Neural simulation of action: A unifying mechanism for motor cognition." *NeuroImage*, vol. 14, no. 1, p. S103–S109, Jul. 2001.

[15] A. N. Meltzoff, "The 'like me' framework for recognizing and becoming an intentional agent," *Acta Psychol.*, vol. 124, no. 1, pp. 26–43, 2007.

[16] M. Saifuzzaman and Z. Zheng, "Incorporating human-factors in car-following models: A review of recent developments and research needs," *Transp. Res. C, Emerg. Technol.*, vol. 48, pp. 379–403, Nov. 2014.

[17] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations," *Phys. Rev. E, Stat. Phys. Plasmas Fluids Relat. Interdiscip. Top.*, vol. 62, no. 2, pp. 1805–1824, Aug. 2000.

[18] M. Liebner, F. Klanner, M. Baumann, C. Ruhhammer, and C. Stiller, "Velocity-based driver intent inference at urban intersections in the presence of preceding vehicles," *IEEE Intell. Transp. Syst. Mag.*, vol. 5, no. 2, pp. 10–21, May 2013.

[19] G. S. Aoude, V. R. Desaraju, L. H. Stephens, and J. P. How, "Driver behavior classification at intersections and validation on large naturalistic data set," *IEEE Trans. Intell. Transp. Syst.*, vol. 13, no. 2, pp. 724–736, Jun. 2012.

[4]The optimality criterion is like (3), but without the terminal cost $w_A a_f^2$, which has no meaning given hat the final acceleration is zero in (2').

[20] E. Ward and J. Folkesson, "Multi-classification of driver intentions in yielding scenarios," in *Proc. IEEE Conf. Intell. Transp. Syst. (ITSC)*, Oct. 2015, pp. 678–685.

[21] T. J. Prescott, P. Redgrave, and K. N. Gurney, "Layered control architectures in robots and vertebrates," *Adapt. Behav.*, vol. 7, pp. 99–127, Jan. 1999.

[22] J. H. Jackson, "Evolution and dissolution of the nervous system," in *Selected Writings of John Hughlings Jackson*, vol. 2, J. Taylor, Ed. London, U.K.: Staples Press, 1958, pp. 45–75.

[23] R. Brooks, "A robust layered control system for a mobile robot," *IEEE J. Robot. Autom.*, vol. RA-2, no. 1, pp. 14–23, Mar. 1986.

[24] P. Redgrave, T. J. Prescott, and K. Gurney, "The basal ganglia: A vertebrate solution to the selection problem?" *Neuroscience*, vol. 89, no. 4, pp. 1009–1023, 1999.

[25] J. A. Michon, "A critical view of driver behavior models: What do we know, what should we do?" *Human Behavior and Traffic Safety*. New York, NY, USA: Plenum Press, 1985, pp. 485–520.

[26] D. Windridge, A. Shaukat, and E. Hollnagel, "Characterizing driver intention via hierarchical perception–action modeling," *IEEE Trans. Human–Mach. Syst.*, vol. 43, no. 1, pp. 17–31, Jan. 2013.

[27] E. Hollnagel and D. D. Woods, *Joint Cognitive Systems*. New York, NY, USA: Taylor & Francis, 2005.

[28] G. Marti, A. H. P. Morice, and G. Montagne, "Drivers' decision-making when attempting to cross an intersection results from choice between affordances," *Front. Hum. Neurosci.*, vol. 8, p. 1026, Jan. 2015.

[29] D. Liu and E. Todorov, "Evidence for the flexible sensorimotor strategies predicted by optimal feedback control," *J. Neurosci.*, vol. 27, no. 35, pp. 9354–9368, Aug. 2007.

[30] E. Todorov and M. I. Jordan, "Optimal feedback control as a theory of motor coordination," *Nature Neurosci.*, vol. 5, no. 11, pp. 1226–1235, 2002.

[31] E. Todorov, "Optimality principles in sensorimotor control.," *Nature Neurosci.*, vol. 7, no. 9, pp. 907–915, Sep. 2004.

[32] P. M. Bays and D. M. Wolpert, "Computational principles of sensorimotor control that minimize uncertainty and variability," *J. Physiol.*, vol. 578, pp. 387–396, Jan. 2007.

[33] C. M. Harris and D. M. Wolpert, "The main sequence of saccades optimizes speed-accuracy trade-off," *Biol. Cybern.*, vol. 95, no. 1, pp. 9–21, Jul. 2006.

[34] P. Viviani and T. Flash, "Minimum-jerk, two-thirds power law, and isochrony: Converging approaches to movement planning," *J. Exp. Psychol. Hum. Perception Perform.*, vol. 21, no. 1, pp. 32–53, Mar. 1995.

[35] C. M. Harris, "Signal-dependent noise determines motor planning," *Nature*, vol. 394, pp. 780–784, Aug. 1998.

[36] C. M. Harris, "Biomimetics of human movement: Functional or aesthetic?" *Bioinspiration Biomimetics*, vol. 4, no. 3, p. 033001, Sep. 2009.

[37] T. Flash, Y. Meirovitch, and A. Barliya, "Models of human movement: Trajectory planning and inverse kinematics studies," *Robot. Autonomous Syst.*, vol. 61, no. 4, pp. 330–339, Apr. 2013.

[38] D. Kim, C. Jang, and F. C. Park, "Kinematic feedback control laws for generating natural arm movements," *Bioinspiration Biomimetics*, vol. 9, no. 1, p. 016002, 2014.

[39] A. J. Nagengast, D. A. Braun, and D. M. Wolpert, "Optimal control predicts human performance on objects with internal degrees of freedom.," *PLoS Comput. Biol.*, vol. 5, no. 6, p. e1000419, Jun. 2009.

[40] A. Simpkins and E. Todorov, "Complex object manipulation with hierarchical optimal control," in *Proc. IEEE SSCI Symp. Ser. Comput. Intell. (ADPRL)*, Apr. 2011, pp. 338–345.

[41] T. Flash and B. Hochner, "Motor primitives in vertebrates and invertebrates.," *Current Opinion Neurobiol.*, vol. 15, no. 6, pp. 660–666, Dec. 2005.

[42] F. A. Mussa-Ivaldi and S. A. Solla, "Neural primitives for motion control," *IEEE J. Ocean. Eng.*, vol. 29, no. 3, pp. 640–650, Jul. 2004.

[43] V. Cossalter, M. Da Lio, R. Lot, and L. Fabbri, "A general method for the evaluation of vehicle manoeuvrability with special emphasis on motorcycles," *Vehicle Syst. Dyn., Int. J. Vehicle Mech. Mobility*, vol. 31, no. 2, pp. 113–135, 1999.

[44] E. Bertolazzi, F. Biral, M. D. Lio, A. Saroldi, and F. Tango, "Supporting drivers in keeping safe speed and safe distance: The SASPENCE subproject within the European framework programme 6 integrating project PReVENT," *IEEE Trans. Intell. Transp. Syst.*, vol. 11, no. 3, pp. 525–538, Sep. 2010.

[45] R. S. Sharp and H. Peng, "Vehicle dynamics applications of optimal control theory," *Vehicle Syst. Dyn.*, vol. 49, no. 7, pp. 1073–1111, 2011.

[46] M. Da Lio *et al.*, "Artificial Co-drivers as a universal enabling technology for future intelligent vehicles and transportation systems," *IEEE Trans. Intell. Transp. Syst.*, vol. 16, no. 1, pp. 244–263, Feb. 2015.

[47] P. Bosetti, M. D. Lio, and A. Saroldi, "On curve negotiation: From driver support to automation," *IEEE Trans. Intell. Transp. Syst.*, vol. 16, no. 4, pp. 2082–2093, Aug. 2015.
[48] M. Da Lio, A. Mazzalai, and M. Darin, "Cooperative intersection support system based on mirroring mechanisms enacted by bio-inspired layered control architecture," *IEEE Trans. Intell. Transp. Syst.*, to be published. [Online]. Available: http://ieeexplore.ieee.org/document/8011474/
[49] R. Bogacz and K. Gurney, "The basal ganglia and cortex implement optimal decision making between alternative actions," *Neural Comput.*, vol. 19, no. 2, pp. 442–477, 2007.
[50] E. Todorov and M. I. Jordan, "A minimal intervention principle for coordinated movement," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 15. 2003, pp. 27–34.
[51] F. J. Valero-Cuevas, M. Venkadesan, and E. Todorov, "Structured variability of muscle activations supports the minimal intervention principle of motor control," *J. Neurophysiol.*, vol. 102, no. 1, pp. 59–68, 2009.
[52] *interactIVe—Accident Avoidance by Intervention for Intelligent Vehicles*. Accessed: Feb. 28, 2017. [Online]. Available: http://www.interactive-ip.eu/
[53] P. Bosetti, M. Da Lio, and A. Saroldi, "On the human control of vehicles: An experimental study of acceleration," *Eur. Transp. Res. Rev.*, vol. 6, no. 2, pp. 157–170, 2013.
[54] B. Efron and R. Tibshirani, *An Introduction to the Bootstrap*. New York, NY, USA: Chapman & Hall, 1994.
[55] M. A. Jenkins and J. F. Traub, "A three-stage algorithm for real polynomials using quadratic iteration," *SIAM J. Numer. Anal.*, vol. 7, no. 4, pp. 545–566, 1970.

Mauro Da Lio received the Laurea degree in mechanical engineering from the University of Padova, Italy, in 1986. He is currently a Full Professor of mechanical systems with the University of Trento, Italy. His earlier research activity was on modeling, simulation, and optimal control of mechanical multibody systems, in particular vehicle and spacecraft dynamics. More recently, his focus shifted to the modeling of human sensory-motor control, in particular drivers and motor impaired people. Prior to his academic career, he worked for an offshore oil research company in underwater robotics (a EUREKA Project). He was involved in several EU framework programme 6 and 7 projects (PReVENT, SAFERIDER, interactIVe, VERITAS, AdaptIVe, and NoTremor). He is currently the coordinator of the EU Horizon 2020 Dreams4Cars Research and Innovation Action: a collaborative project in the Robotics domain which aims at increasing the cognition abilities of artificial driving agents by means of offline simulation mechanisms broadly inspired by the human dream state.

Alessandro Mazzalai received the B.S. degree in industrial engineering and the M.Sc. degree in mechatronics engineering from the University of Trento, Trento, Italy, in 2011 and 2013, respectively, where he is currently pursuing the Ph.D. degree in mechatronic engineering.

Since 2014, he has been a Teaching Assistant with the Department of Industrial Engineering, University of Trento. Since 2014, he has also been involved in the European FP7 project AdaptIVe. His research interests include automated driving in urban scenarios, the development of an artificial driver, and the study of vehicle dynamics.

Kevin Gurney was born in West London, U.K., in 1956. He received the B.Sc. degree in mathematical physics from the University of Sussex in 1977, and the M.Sc. degree in digital systems and the Ph.D. degree in neural networks from Brunel University in 1986 and 1989, respectively.

After his Ph.D., he held a post-doctoral position at the Department of Human Sciences, Brunel University, before taking a lectureship at the Department of Psychology, The University of Sheffield, in 1995. Since then, he has held the posts of a Senior Lecturer and a Reader with The University of Sheffield, where he is currently a Professor of computational neuroscience. He has published in a wide variety of journals, including *Nature Reviews Neuroscience*, the IEEE TRANSACTIONS ON NEURAL NETWORKS, and *PLoS Biology*. He has previously conducted research on neural networks and vision. He is currently researching the neuroscience of decision making and action selection.

Dr. Gurney is a fellow of the Royal Society of Biology.

Andrea Saroldi received the Laurea degree in physics from the University of Torino, Italy, in 1985. He is currently a Senior Researcher with the Centro Ricerche Fiat, where he leads the group on sensors and technologies inside the unit dealing with advanced driver assistance systems. He has been a Coordinator for development of CRF demonstrator vehicles inside PReVENT, interactIVe, and currently AdaptIVe Integrated Project. Main field of activity refers to processing of sensor data and development of driver assistance functions for preventive safety and autonomous driving.