

Intelligent scheduling for in-car notifications

1st Author Name

Affiliation

Address

e-mail address

Optional phone number

2nd Author Name

Affiliation

Address

e-mail address

Optional phone number

3rd Author Name

Affiliation

Address

e-mail address

Optional phone number

ABSTRACT

The process of driving a car involves a cognitive load that varies over time. Additional load comes from secondary factors not directly associated with the driving process, including navigation devices, entertainment systems and the car's own warnings. In this paper, we present a framework for intelligent scheduling of in-car notifications based on the driver's estimated cognitive load. As the single channel for communication, it reschedules the notifications using a priority queue, and relays them to the driver based on the urgency of the notification and the overall estimated cognitive load being experienced by the driver at any given moment. We evaluate our system using a dataset collected from a car's CAN bus during multiple on-road trials and show that our proposed approach reduces the number of simultaneous calls on the driver's attention during the driving task. We also demonstrate that our intelligent scheduling significantly reduces the maximum cognitive load experienced by the driver and the frequency with which high loads occur.

INTRODUCTION

Cognitive load is commonly defined as the relationship between the cognitive demands placed on a user by a task and the user's cognitive resources [17]. For a driver, the process of controlling the vehicle safely along a road while avoiding collisions and obeying the necessary regulations is the primary load source, but increasingly there are other calls upon the driver's attention: entertainment systems, phone calls, social network alerts, satellite navigation instructions, and information from the car itself about fuel warnings, tyre pressures, and scheduled maintenance. Such 'secondary' events are generally, at present, simply presented to the driver at the time they occur. But the moment at which the driver is negotiating a difficult sliproad onto a fast-moving highway may not be the most appropriate time for a sudden audiovisual warning that they need to arrange a service in the next 1000 miles. Studies reported that when a secondary in-vehicle distractor is presented to the driver at a highly demanding driving instance, such as when approaching a light-controlled stop, the primary driving task is significantly affected [10, 16, 7], even

Paste the appropriate copyright statement here. ACM now supports three different copyright statements:

- ACM copyright: ACM holds the copyright on the work. This is the historical approach.
- License: The author(s) retain copyright, but ACM receives an exclusive publication license.
- Open Access: The author(s) wish to pay for the work to be open access. The additional fee must be paid to ACM.

This text field is large enough to hold the appropriate release statement assuming it is single spaced.

Every submission will be assigned their own unique DOI string to be included here.




Figure 1: System overview showing the main components of our proposed framework: message classification, priority queue and cognitive load estimation.

if the secondary task requires cognitive, rather than visual, attention.

As more functionality is added to in-vehicle systems, there will be an increasing load on the driver's attention resources [4, 3]. In this paper, we propose a framework for intelligently scheduling those secondary tasks, such that they do not coincide with the unavoidable primary tasks (or with each other), with the aim of improving the safety and comfort of the driver, passengers and other road users.

The main contributions of this paper can be summarised as follows:

1. Presenting a system for intelligent scheduling of in-car audio notifications based on the driver's estimated cognitive load. See Figure 1 for an overview diagram of our system, which has two main components: message classification (priority queue), and cognitive load estimation modeling, based on real-world recordings of controller area network-bus (CAN-bus) data.
2. Demonstrating that our scheduling approach significantly reduces the maximum cognitive load experienced by the driver and the frequency with which high loads occur.

RELATED WORK

In-vehicle information systems are believed to introduce increased cognitive loads on the driver [4], especially in the case of elderly drivers [8]. Previous work has emphasised the importance of detecting when a driver is distracted by those systems, typically with the aim of alerting the driver that they should pay more attention. However, there is no quantitative analysis of such distractions, especially for the auditory/cognitive distractions. (For an overview of the methods employed in past studies, see Bach *et al* [4].)

In-vehicle distractions can be divided into two types: visual and cognitive. Bach *et al* [4] describe them as *eyes-off-the-road* time and *mind-off-the-road* time (the latter being when the driver may be looking at the road, but their thoughts are elsewhere and hence reaction times are slower). Cognitive distraction may often be caused by auditory signals, and is particularly apparent during phone calls [9, 7, 16].

Visual distractions can be measured through the use of head- and gaze-tracking [6]. Liang *et al* [13] compared different approaches of measuring off-road glancing behaviour in order to estimate a driver’s visual distraction. Cognitive distractions, on the other hand, are harder to detect or estimate, as measuring cognitive load in real-world situations is a difficult task. A common technique for measuring cognitive load is to challenge the participant to perform a secondary task at the same time as a primary one, and measure the degree to which the primary, or secondary, performance is impaired. This is not safe to apply in our real-world on-road driving trials. Other studies have involved manual annotation of the estimated load based on video recordings watched after the drive. These perception methods are useful in measuring visual, rather than cognitive distractions [6]. Finally, biometric sensors have also been tried [5, 15], but their inconvenience means that any system depending on them could not readily be deployed for general use.

Many studies explored the relationship between different driving parameters, which can be obtained from vehicle data and the driver’s cognitive load. The driving speed has been found to be inversely correlated with the cognitive load [11]. This is because highway driving imposes fewer distractions than urban driving, or because drivers only speed up in the absence of such challenges. Another study suggests that the driver, under identical road conditions, would slow down when the cognitive load was introduced, perhaps because they are attempting to free-up resources for the secondary task by simplifying the primary (driving) task [12]. Moreover, drivers’ cognitive load has been found to be highest while approaching and driving through junctions [11], which indicates that the steering wheel angle, acceleration and deceleration are useful metrics for estimating cognitive load.

Although several studies reported the effect of in-car notifications, only few investigated redistributing, rescheduling or controlling the flow of the messages. Pompei *et al* [14] developed a platform for investigating the use of driver activity monitoring with the aim of controlling information flow in the vehicle and reducing the driver’s cognitive load. However, no qualitative experiments were reported.

DATA COLLECTION

The dataset we use is collected in real-driving scenarios. Eight participants took an instrumented vehicle around the same 90-minute route, chosen to include urban, rural and highway sections to get diverse driving conditions. Activity inside and outside the car was captured by several cameras, and data from the car’s CAN-bus was recorded, including pedal pressures, the use of controls such as indicators, and steering wheel angle. Since we aim to inform real systems




Figure 2: Partitioning of the message queue into five smaller independent queues according to importance classification.

deployed in production vehicles, we have avoided more intrusive measures such as the collection of biometric data through ECGs, EEGs, or skin conductivity, focussing on data sources which can readily be captured in a production vehicle.

SYSTEM DESIGN

In the Apollo space missions, all communications to the astronauts from mission control went through a single individual, to avoid the possible confusion of multiple simultaneous communications. This individual, designated the ‘Capsule Communicator’ or CAPCOM, was typically a trained astronaut who had some understanding of what the crew was going through.

The system we are proposing for in-car communications closely resembles that of the CAPCOM, in that messages from all driver aids are passed to it, analysed and placed at the appropriate position in a priority queue. Then, the message at the head of the queue is relayed at a time when the driver’s estimated cognitive load is below a certain threshold. This threshold is set according to the importance of the message.

Message Classification and Priority Queue

To classify the different messages and define the mechanics of the queue, we associate two different values with each message: *Message Importance* and *Message Priority*.

Message Importance is the primary classification; a five-tier system with the following categories: Critical, Very high, High, Medium and Low. Once assigned, Message Importance remains static. The message queue is also partitioned into five small, independent queues, one for each category, illustrated in Figure 2. Partition independence guarantees that a message with critical importance cannot be held up by one with low importance even if the latter has been in the queue for some considerable time.

We define the criteria for a message to be placed in one of the above categories as follows:

- *Critical* messages are those requiring a prompt response from the driver. Relayed immediately, they generally convey information about safety risks such as seatbelts not being in use, or urgent navigation instructions.
- *Very High* importance messages alert the driver to situations where action is required in the short term. These messages include fuel and vehicle fault warnings.
- *High* importance messages convey information which does not require immediate action, such as traffic warnings.




Figure 3: Modeling of Cognitive Load (CL) as a function of different car parameters.

- *Medium* importance messages contain notifications of future events, such as a service due date.
- *Low* importance messages are purely for the driver’s entertainment. They include social media and interesting landmark notifications.

Message Priority is a secondary ranking mechanism which defines the order of messages within a given partition of the queue. This is necessary because some messages are time- or distance-dependent and may ‘time-out’ (become redundant) if held up for too long. An example would be an incoming phone call.

Each type of message has a specified initial priority value, which gives a natural ordering should multiple messages enter the queue in quick succession. If the message is time- or distance-dependent it will also have a specified time-out value. The priority of each message increases linearly over time so that, all other things being equal, messages which have been waiting longest will be relayed first. Messages with a time-out value increase priority at a faster rate, depending on its expiry, and get discarded if they are not delivered before becoming redundant.

When the time comes to feed another message to the driver, which in our system is done by a text-to-speech module, the queues are consulted in turn: from the ‘Critical’ down to the ‘Low’ importance. Messages from a later queue will only be processed when the earlier queues are empty. Then, within a queue, the Message Priorities will be used to select the message to be played first.

Cognitive Load Estimation

The cognitive load (CL) resulting from the driving task was estimated using the CAN data from the vehicle. Although correlations have been presented, very little work has been done previously to quantify the factors which affect the driver’s CL. We have defined and modelled five factors believed to affect the driver’s CL (summerised in Figure 3). Each of these factors was defined on a scale between 0 and 1 and then scaled by a constant factor based on estimates of how large an effect it would have on the driver’s CL. Finally, all of these scaled values are summed to give a total CL value. The factors we considered are:

Speed – It has been shown that CL is inversely correlated with driving speed, v [11, 12]. Building on this, we modelled CL as a Gaussian function of car speed as illustrated in

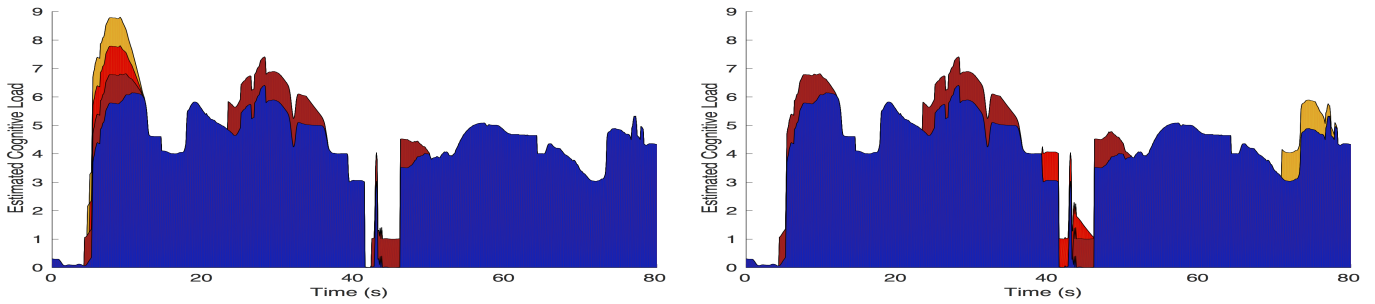
Figure 3a. When driving, if $v \leq 20\text{mph}$, one of two scenarios is taking place: 1. There is a speed limit of 20mph, or 2. Traffic conditions prohibit driving any faster. Both of these imply a very high hazard density [2] and hence a heightened CL for the driver, so when $v \leq 20\text{mph}$ we assume CL due to speed is maximum. When $v = 30\text{mph}$, CL starts decreasing but stays relatively high because 30mph speed limits are imposed in ‘built-up’ areas, meaning the hazard density remains high [2]. As speed increases above $v = 30\text{mph}$, the CL decreases rapidly because of a sharp decrease in hazards posed by pedestrians and cyclists. This trend continues to the point where CL is minimum when $v = 70\text{mph}$, which is the maximum speed limit in UK [1].

Steering Angle – Previous work showed that CL is increased when driving around corners or performing other turning manoeuvres [5]. These driving events can be captured by measuring the steering wheel angle, θ . We observed that $|\theta| < 15^\circ$ is consistent with driving on straight roads, or roads which have a very large radius of curvature. Steering angles in the range $15^\circ \leq |\theta| < 25^\circ$ indicated stretches of road with a much smaller radius of curvature, but not corners or roundabouts. $|\theta| \geq 25^\circ$ was found to be broadly consistent with turning corners and negotiating roundabouts. Based on this, we modelled $CL(\theta)$ as an inverted Gaussian function with mean 0° and standard deviation 20° , illustrated in Figure 3b.

Steering Instability – Besides steering angle, we defined steering instability in terms of a sliding window of previous 3 seconds of driving. That is because a steering angle of 0 that happens while exiting a roundabout, for example, might incorrectly indicate that the driver is having a low CL if only the instantaneous angle is considered, although the opposite is true [11]. Therefore, we modelled CL due to steering instability as a Gaussian function with mean 22.5 and standard deviation 2.0, depicted in Figure 3c. This function changes significantly over a small critical region, chosen such that the CL values obtained would show whether the driver had turned a corner (or not) in the last three seconds based on the steering angle threshold $|\theta| \geq 25^\circ$, discussed in the previous section.

Brake Pressure – Since using the brakes is an indication of increased CL, we modelled CL as an inverse function of brake pressure as shown in Figure 3d with a $p > 8$ indicating that the driver intends to slow down significantly or bring the vehicle to a halt ,i.e. the CL increases significantly.

Cruise – We define the term *cruising* to describe a situation where a low CL is imposed upon the driver. While hard to



(a) Estimated cognitive load experienced by the driver from the primary driving task (in blue) + additional cognitive load imposed by simulated notifications (in red, orange and yellow) without the smart scheduling intervention. Messages of all categories, from the critical (in red) to those of lower importance (orange and yellow) are delivered at the time they occur, even if they overlap.

(b) Cognitive load redistribution with the smart scheduling intervention. Critical messages are still presented immediately, while less urgent ones are delayed until times of lower CL. All messages are queued up to avoid audio overlap, though our model of the cognitive load induced does extend a little beyond the duration of the message, so some overlap of CL may be seen in our analysis.

Figure 4: A case study showing the operation of our scheduling system on a segment of the route containing urban driving.

define precisely, phrases like ‘cruising along’ or ‘motorway cruising’ imply that the driver is on a straight road, travelling at a reasonably high, steady speed – not in stop-start traffic – for an extended period of time. To capture this idea, the likelihood of being in a *cruising* state is calculated as a function of instantaneous high speed and lack of braking pressure and recent stability of both steering and speed over a window of 4 seconds. CL is then modelled as an inverse function of the likelihood of being in a *cruising* state, as shown in Figure 3e.

The total CL value at a given point is calculated as a weighted sum of all of the above mentioned factors and used to determine the categories of the messages which can be relayed to the driver at this point in time.

EVALUATION

To evaluate our framework performance when applied to data captured in real driving scenarios, we created simulations where a range of events/notifications were generated and injected into the recorded driving data as if they had been generated by in-car systems. Our intelligent scheduling framework takes these, queues them up, prioritises them and delivers them at the appropriate time according to the approach described above. To assess the effectiveness of our system, we modelled the *additional* cognitive load resulting from the messages being fed to the driver.

In our implementation, the total CL value varied from 0 to 7. The CL thresholds chosen for different message categories were as follows: *Critical* messages can be relayed in all CL values, CL value must be below 4.0 in order to relay *Very High* importance messages, for *High* importance messages the CL must be less than 3.2, *Medium* importance messages can only be relayed when CL is below 3.1, and *Low* messages are relayed when the CL is below 3.0. These values were chosen empirically based on our data. In the future, the deciding of these thresholds could be done by safety authorities or vehicle manufacturers. Moreover, since our notifications consist of spoken audio of different lengths, we assume that they impose a single unit of load for the duration of the audio plus two further seconds, then that the load decays to zero over the following three seconds. This was chosen as representing the

hearing, parsing, understanding and processing of the audio messages.

In this section we first present a case study illustrating the impact of using our system on message delivery during a short section of the route during urban driving. We show how our system redistributes the cognitive load and reduces the number of simultaneous calls on driver’s attention. Then we apply the intelligent scheduling to the full dataset recorded by multiple participants, showing the effect of our system on the overall estimated cognitive load on longer journeys.

Case study analysis

Given our models of the baseline load caused by driving and the additional load resulting from notifications, we can add these together to get the approximate overall load experienced by the driver at any given moment. It may be argued that the load resulting from context-switching in the case of multiple simultaneous tasks is greater than simply the sum of those tasks when experienced individually, but we have not assumed any multi-tasking overhead in our model.

Our case study presented in Figure 4 shows in detail how our system affects the scheduling of secondary event notifications on top of the driving baseline in a short segment of urban driving route representing 80 seconds where the driver navigates a roundabout then enters a highway. For the purposes of demonstration, we inject a few sample notifications, including three overlapping at the beginning of the route segment.

We measure the estimated cognitive load with and without our scheduling system. Figure 4a shows the notifications at the time they occur, without any scheduling intervention, while Figure 4b shows how the notifications are redistributed with the scheduler. Some notifications stay in their place because they are in ‘high importance’ categories, such as satellite navigation instructions. Other lower-importance messages are delayed, and so end up in the troughs, when the baseline load is lower. High peaks of overall cognitive load are therefore redistributed or ‘smoothed out’.

Overall cognitive load analysis

For a broader evaluation of the effect of the scheduler on the overall cognitive load of the driver, we analysed eight CAN data files from the dataset described in ‘Data collection’ above, one for each of 8 different participants driving our full route (10 hrs of driving in total).

An appropriate number of messages were generated for the length of each journey and were randomly distributed through the duration of each driver’s route. To simulate real-world scenarios, messages were injected at random locations on the route, at rates depending on the type of message as follows: Sat-Nav messages on average once per minute, Personal messages (emails/social notifications) on average once every two minutes, and High-importance messages on average every half-hour. In real life, these frequencies can differ according to each driver’s situation and notification preferences, but overall this does not have a significant effect on our estimated cognitive load results, as the same frequencies of messages were used for all participants’ data.

The first evaluation metric was to check if our scheduling approach affects the maximum cognitive load on the driver. The peak cognitive load can be the point of driving where an accident is highly likely to occur, especially due to unforeseen circumstances. At this time the driver’s attention is already significantly occupied.

Participant	Max CL (baseline)	Max CL w/ scheduling
P1	12.0	9.0
P2	12.0	8.8
P3	13.0	8.8
P4	9.3	8.0
P5	10.2	8.9
P6	15.2	8.4
P7	11.9	8.9
P8	9.8	8.4

Table 1: Comparing the estimated maximum cognitive load for each participant with and without our system. Our scheduling approach significantly decreases the maximum value of the estimated cognitive load ($p < 0.005$).

In our analysis, we calculated the total estimated cognitive load (baseline estimated cognitive load from CAN-bus data plus additional cognitive load due to simulated notifications) at each point in the route. We then compared the peak values of estimated cognitive load with and without our scheduling approach for each of the eight participants. As shown in Table 1, our scheduling approach significantly decreases the peak value for all of the eight participants’ data. Performing a paired sample t-test demonstrated that the estimated cognitive load peaks were significantly lower when our scheduling approach was used ($p < 0.005$). This shows that using our scheduling approach, the worst case events of high peaks of CL can have less negative effect, as they will be associated with lower values of CL.

As a second evaluation metric, we wanted to consider our system’s effect in reducing other high values of the driver’s cognitive load: how *much* of the journey involved CL peaks

Participant	Percentage of the drive above different thresholds					
	1xSD w/out	1xSD with	2xSD w/out	2xSD with	3xSD w/out	3xSD with
P1	16.32	16.79	0.89	0.40	0.21	0.0
P2	11.78	12.88	1.04	0.50	0.26	0.0
P3	12.78	13.90	1.00	0.68	0.45	0.0
P4	14.25	14.54	0.54	0.29	0.0	0.0
P5	15.41	16.53	0.80	0.53	0.05	0.0
P6	15.83	16.36	0.89	0.56	0.37	0.0
P7	10.88	11.42	0.59	0.54	0.28	0.0
P8	12.76	13.92	0.80	0.47	0.01	0.0

Table 2: Comparing the percentage of the total route where the driver’s estimated cognitive load exceeds different thresholds (as multiples of Standard Deviations (SD) above the mean). Note that our scheduling framework reduces the amount of time spent with a load greater than 2SD and 3SD above the mean.

exceeding certain thresholds? We defined the thresholds to be multiples of standard deviations above the mean. Table 2 shows the percentage of the journey spent above these thresholds for each of the eight participants’ data, comparing between the two scenarios: with and without our scheduling system.

The comparison shows that the amount of time spent with a load greater than 2 or 3 standard deviations above the mean is reduced by our scheduler. The time spent above 1 S.D. is slightly increased, but this is expected since the load removed from the higher peaks is redistributed, in part, to these less-stressful times.

CONCLUSIONS AND FUTURE WORK

In this paper, we have presented an intelligent scheduling system for in-car notifications which aims to reduce the overall cognitive load experienced by drivers from in-vehicle systems. We described the two components of our system, namely: message classification using priority queues, and cognitive load estimation. We evaluated our approach using CAN-bus data collected from a dataset of real-world recordings. We demonstrated that our scheduling approach significantly reduces the peak cognitive load experienced by the driver and the frequency with which high loads occur. Thus, our system smooths out the cognitive load on the driver by presenting less urgent messages at a time when the driver is not occupied with difficult driving situations.

Future work includes comparing our system’s rescheduling of messages to the timing that human passengers would use based on their judgment of the driver’s cognitive load. We also plan to incorporate additional modalities for cognitive load estimation, such as gaze and head pose. Finally, the scheduling could also be made dependent on external sources such as the proximity of other vehicles, and events likely to be encountered in the future, such as approaching traffic congestion or upcoming navigation directions.

REFERENCES

1. UK government accidents statistics, GOV.UK. <https://www.gov.uk/government/statistical-data-sets/road-accidents-and-safety-statistical-tables-index>.
2. UK government speed limits, GOV.UK. <https://www.gov.uk/speed-limits>.
3. Angkititrakul, P., Kwak, D., Choi, S., Kim, J., PhucPhan, A., Sathyanarayana, A., and Hansen, J. H. Getting start with utdrive: driver-behavior modeling and assessment of distraction for in-vehicle speech systems. In *INTERSPEECH* (2007), 1334–1337.
4. Bach, K. M., Jæger, M. G., Skov, M. B., and Thomassen, N. G. Interacting with in-vehicle systems: understanding, measuring, and evaluating attention. In *Proceedings of the 23rd British HCI Group Annual Conference on People and Computers: Celebrating People and Technology*, British Computer Society (2009), 453–462.
5. Berka, C., Levendowski, D. J., Cvetinovic, M. M., Petrovic, M. M., Davis, G., Lumicao, M. N., Zivkovic, V. T., Popovic, M. V., and Olmstead, R. Real-time analysis of eeg indexes of alertness, cognition, and memory acquired with a wireless eeg headset. *International Journal of Human-Computer Interaction* 17, 2 (2004), 151–170.
6. Busso, C., and Jain, J. Advances in Multimodal Tracking of Driver Distraction. In *Digital Signal Processing for In-Vehicle Systems and Safety*. 2012, 253–270.
7. Caird, J. K., Willness, C. R., Steel, P., and Scialfa, C. A meta-analysis of the effects of cell phones on driver performance. *Accident Analysis & Prevention* 40, 4 (2008), 1282–1293.
8. Green, P. Variations in task performance between younger and older drivers: Umtri research on telematics. In *Association for the Advancement of Automotive Medicine Conference on Aging and Driving*, Citeseer (2001).
9. Hancock, P., Lesch, M., and Simmons, L. The distraction effects of phone use during a crucial driving maneuver. *Accident Analysis & Prevention* 35, 4 (2003), 501–514.
10. Hancock, P., Simmons, L., Hashemi, L., Howarth, H., and Ranney, T. The effects of in-vehicle distraction on driver response during a crucial driving maneuver. *Transportation Human Factors* 1, 4 (1999), 295–309.
11. Harms, L. Variation in drivers' cognitive load. effects of driving through village areas and rural junctions. *Ergonomics* 34, 2 (1991), 151–160.
12. Jamson, A. H., and Merat, N. Surrogate in-vehicle information systems and driver behaviour: Effects of visual and cognitive load in simulated rural driving. *Transportation Research Part F: Traffic Psychology and Behaviour* 8, 2 (2005), 79–96.
13. Liang, Y., Lee, J. D., and Yekhshatyan, L. How dangerous is looking away from the road? algorithms predict crash risk from glance patterns in naturalistic driving. *Human Factors: The Journal of the Human Factors and Ergonomics Society* 54, 6 (2012), 1104–1116.
14. Pompei, F. J., Sharon, T., Buckley, S. J., and Kemp, J. An automobile-integrated system for assessing and reacting to driver cognitive load. *Proceedings of Convergence* (2002), 411–416.
15. Reimer, B., Mehler, B., Coughlin, J. F., Godfrey, K. M., and Tan, C. An on-road assessment of the impact of cognitive workload on physiological arousal in young adult drivers. In *Proceedings of the 1st international conference on automotive user interfaces and interactive vehicular applications*, ACM (2009), 115–118.
16. Strayer, D. L., Cooper, J. M., and Drews, F. A. What do drivers fail to see when conversing on a cell phone? In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 48, SAGE Publications (2004), 2213–2217.
17. Wickens, C. D. Multiple resources and performance prediction. *Theoretical issues in ergonomics science* 3, 2 (2002), 159–177.