

Neuro-Cognitive Reward Modeling for Human-Centered Autonomous Vehicle Control

Zhuoli Zhuang, Yu-Cheng Chang, Yu-Kai Wang, Thomas Do, Chin-Teng Lin
University of Technology Sydney

{zhuoli.zhuang, fred.chang, yukai.wang, thomas.do, chin-teng.lin}@uts.edu.au

Abstract

Recent advancements in computer vision have accelerated the development of autonomous driving. Despite these advancements, training machines to drive in a way that aligns with human expectations remains a significant challenge. Human factors are still essential, as humans possess a sophisticated cognitive system capable of rapidly interpreting scene information and making accurate decisions. Aligning machine with human intent has been explored with Reinforcement Learning with Human Feedback (RLHF). Conventional RLHF methods rely on collecting human preference data by manually ranking generated outputs, which is time-consuming and indirect. In this work, we propose an electroencephalography (EEG)-guided decision-making framework to incorporate human cognitive insights without behaviour response interruption into reinforcement learning (RL) for autonomous driving. We collected EEG signals from 20 participants in a realistic driving simulator and analyzed event-related potentials (ERP) in response to sudden environmental changes. Our proposed framework employs a neural network to predict the strength of ERP based on the cognitive information from visual scene information. Moreover, we explore the integration of such cognitive information into the reward signal of the RL algorithm. Experimental results show that our framework can improve the collision avoidance ability of the RL algorithm, highlighting the potential of neuro-cognitive feedback in enhancing autonomous driving systems. Our project page is: <https://alex95gogo.github.io/Cognitive-Reward/>.

1. Introduction

Autonomous intelligent vehicles (AIVs) have greatly benefited from deep learning in recent years. Thanks to deep neural networks' strong feature extraction capabilities, end-to-end autonomous driving (E2E-AD) has emerged as a

promising approach. This approach directly maps raw sensor inputs, such as camera images, to control signals. In the CARLA driving simulator [15], several E2E-AD models [16, 44, 45, 61] have achieved driving performance comparable to that of rule-based expert systems, as demonstrated in the CARLA leaderboard 1.0 [52]. Despite these advancements, recent studies have shown that E2E-AD models primarily replicate expert trajectories rather than exhibiting human-like reasoning abilities. This limitation becomes evident in out-of-distribution scenarios, where models may fail due to their inability to generalize beyond their training data. For instance, [27] highlights how E2E-AD models often cut turns incorrectly to reach a target point, revealing a lack of true decision-making skills. In addition, a recent study [29] found that top-performing E2E-AD models from the CARLA leaderboard struggle in interactive driving scenarios, such as emergency braking, likely due to the lack of explicit guidance on decision-making and interaction dynamics. A possible cause of these failures may be the limitation of imitation learning (IL), a predominant training paradigm in E2E-AD. IL-based methods rely on direct supervision of control signals or trajectory replication. However, IL faces a notable limitation of the distribution shift problem [7], which means IL models do not learn from sufficient unforeseen failure cases as they learn a specific distribution from an expert [19, 54, 65, 66].

Unlike IL, reinforcement learning (RL) has a learning paradigm from trial and error, dynamically adapting to new environments based on a reward signal [39], and therefore has the potential to mitigate the distribution shift problem. However, without explicit alignment mechanisms, neither RL nor IL guarantees alignment with human expectations. To address this gap, reinforcement learning with human feedback (RLHF) has emerged as a powerful approach to align AI behavior with human preferences [8]. Traditional RLHF typically involves human annotators ranking AI-generated outputs or comparing trajectory segments based on replayed demonstrations. While effective, this process is time-consuming and relies on explicit feedback, which may not fully capture users' cognitive responses. In this

work, we introduce a novel method for collecting human preference signals by leveraging natural neural responses recorded via electroencephalography (EEG), allowing for a more intuitive and efficient alignment of AI decision-making with human cognition.

EEG is a non-invasive method for monitoring neural activity, such as event-related potentials (ERPs), which are time-locked responses to specific cognitive or sensory events. One of the most prominent ERP components is the P3, a positive peak occurring 300-500 ms after stimulus onset. Research has shown that an unexpected, unusual, or surprising stimulus, even if task-irrelevant, can elicit a frontal P3-like response within an attended sequence [10, 42, 48]. The human brain contains billions of neurons, and the firing of these neurons produces electrical responses. ERPs originate as postsynaptic potentials, which occur when similarly oriented neurons fire synchronously in response to visual or auditory stimuli [36]. Since the discovery that brain wave activity peaks approximately 300 milliseconds after exposure to an unpredictable stimulus [49], ERPs have been extensively studied in neuroscience as reliable biomarkers due to their millisecond-level temporal resolution [25, 26, 37]. A classic experimental design for eliciting ERPs is the “oddball paradigm”, wherein the ERP amplitude is typically larger in response to relatively infrequent and unexpected events [37]. Furthermore, ERP amplitudes have been shown to increase with task difficulty [26]. ERPs also provide complementary information to eye-tracking data. For instance, covert attention, the ability to focus on objects outside the fovea, cannot be detected with eye-tracking but is identifiable through ERP analysis [25].

This study leverages the advantages of ERPs, specifically their reliable millisecond-level representation of natural brain responses. It extends classical ERP analysis by exploring the potential of using ERP data to train RL models for AIV tasks. In this study, we collect EEG data from 20 participants as they actively drive in a realistic driving simulator. By analyzing their ERP responses, we find a correlation between ERP amplitude and driver reaction time measurements. To leverage this neuro-cognitive insight, we design and train a human cognitive reward model to predict ERP strength from scene images, allowing us to integrate this prediction into the reward signal for reinforcement learning. Rather than traditional RL with EEG feedback, which requires the collection of EEG data during inference [56–58, 63], this approach does not require human data during inference, and is therefore more scalable. This approach enables RL agents to incorporate human cognitive feedback, improving their decision-making in autonomous driving.

Our main contributions are summarized as follows:

- We present a novel dataset collected using a realistic

VR driving simulator, which captures drivers’ behavioral data, eye-tracking data, EEG signals, and corresponding scene images.

- We develop a human cognitive reward model that predicts EEG features from scene images.
- We propose an RLHF framework that integrates reward signals from the cognitive reward model and demonstrate that it can enhance the performance in two complex driving scenarios.

2. Related Work

2.1. RL with human feedback

RL has demonstrated remarkable success in complex tasks, such as playing Atari games [39] and predicting protein structures [30]. However, RL typically requires carefully crafted reward functions, and the reward signals are often sparse, which hinders RL’s development and application. In addition, RL may not align with human expectations because it optimizes behavior purely based on predefined reward functions, which often fail to capture the complexity of human values and expectations. To address this challenge, RLHF [8] was introduced. RLHF trains a reward model by learning from human preferences, typically by having humans rank or compare model outputs, and then uses this model to guide the RL agent. RLHF has since become a paradigm for fine-tuning large language models (LLMs), aligning them more closely with human preferences [40]. It has been successfully applied in state-of-the-art LLMs such as ChatGPT [1], LLAMA [17], and Gemini [53]. This work builds on these advances by exploring the application of brain modeling methods in reward modeling for AIV tasks, which models a more intuitive or natural process, and extends the human-feedback-driven reinforcement learning.

2.2. RL with EEG feedback

Human physiological signals offer a more natural source for guiding RL models. For instance, leveraging neural networks to approximate eye-tracking features has been explored to enhance RL policy models [35, 67]. Yet, brain signals, a more direct source of human feedback, remain largely unexplored. One such signal is the error-related potential (ErrP), an ERP involuntarily elicited when a human perceives unexpected errors in an environment [21, 55]. ErrP has been explored in various human-machine shared autonomy tasks. For example, ErrP has been utilized to provide implicit feedback to RL agents when a human observes an error, demonstrating effectiveness in grid-world navigation tasks [56, 57]. Additionally, ErrP has been employed to train human policy models, which offer feedback to RL agents when the robot’s uncertainty is high [58]. Similarly, ErrP has been applied for reward shaping in human-in-the-

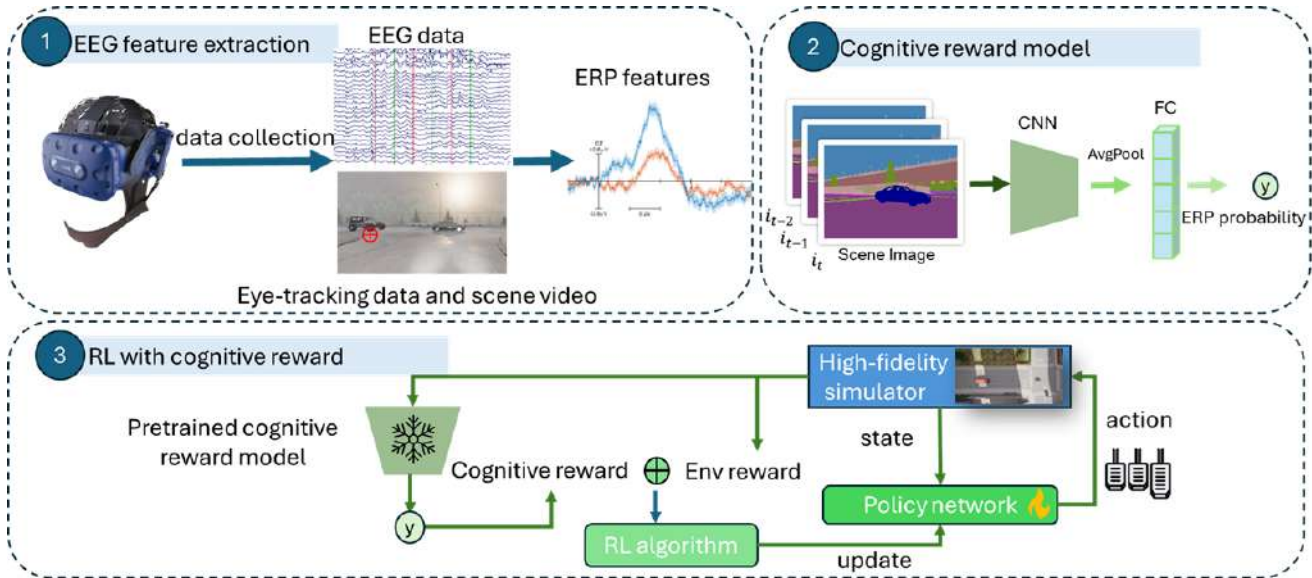


Figure 1. The framework of the human cognitive reward model. First, raw EEG data are preprocessed, and the most prominent feature, which is the ERP, is extracted. Second, an EEG feature prediction model predicts the ERP from the scene images. Lastly, the prediction probability is used as part of the reward of reinforcement learning for autonomous driving tasks.

loop RL to accelerate training, proving beneficial in navigation games [63]. Building on these theoretical groundings, this work leverages ErrP to guide RL agents. Prior approaches face two main limitations: they require real-time EEG data collection during inference, and ErrP signals occur only after an error event, which may introduce latency in corrective actions. To address these challenges, we propose a novel method that predicts EEG features directly from image data. This approach enables scalable and efficient training by eliminating the need for continuous EEG recordings, making RLHF more practical and accessible.

2.3. Predicting brain response from images

Advances in deep learning have made it possible to model brain responses to visual stimuli. For example, magnetoencephalography (MEG), electroencephalography (EEG), and functional magnetic resonance imaging (fMRI) activities have been modeled using latent representations of visual stimuli [9, 22, 64]. For instance, linear encoding models could predict EEG responses to arbitrary images, highlighting the possibility of developing models that bridge visual inputs and brain activity [22]. In addition, a tool called FactorTopy is used to map features from pre-trained deep neural networks onto brain activity patterns measured via fMRI [64]. In line with this, the researchers found that deep neural networks could effectively model the temporal progression of scene processing in the brain in an EEG study [9]. This work aims to build on these advances by introducing a framework that predicts brain responses to images for reward modeling.

3. Dataset

3.1. Data collection setup

This study recruited 32 healthy participants (mean age: 25.6 ± 4.6 years), all of whom were rigorously screened to ensure optimal health status. Data from 20 participants were included in the final analysis, while the remaining 12 were excluded due to VR-induced motion sickness or insufficient engagement in the driving tasks. Participants exhibited normal or corrected-to-normal auditory and visual acuity, with pre-participation evaluations confirming the absence of neurological or psychiatric conditions, as well as current alcohol or substance misuse. These criteria were established to maintain the integrity and reliability of the data collected for the study. Participants were compensated for their involvement at a rate exceeding the national minimum wage. The participants had an average duration of driver’s license possession of 4.9 ± 3.9 years. Before the experiment, they completed the Driver Skill Inventory (DSI) questionnaire [32], achieving a mean perceptual-motor skills score of 4.1 ± 0.3 and a mean safety skills score of 3.7 ± 0.3 on a scale ranging from 1 to 5. Informed written consent was obtained from all participants before their involvement in the study. The research protocol was reviewed and approved by the University’s Human Research Ethics Committee, ensuring compliance with ethical standards for studies involving human subjects. We used an HTC Vive Pro Eye VR headset ($2 \times 1440 \times 1600$ resolution, 90 Hz refresh rate, 3D spatial sound) and the CARLA simulator to develop the scenarios, and Logitech G923 Racing Wheel and Pedal to

collect active driving data. The EEG signals were recorded by the Synamps2 system using 64 channels with Ag/AgCl electrodes according to the international 10–20 system, following the setting in [67]. The sampling rate of the raw EEG data is 1000 Hz, and the impedances of the electrodes were kept below $15k\Omega$ during the experiment. Each scenario was conducted in a separate session lasting approximately 6 minutes to minimize participant fatigue. Participants were allowed to take sufficient rest between sessions. Participants were instructed to perform the driving tasks to the best of their ability, with the goal of reaching the destination as fast as possible while avoiding collisions. To the best of our knowledge, our dataset is the first dataset to include EEG, gaze, and active control data from drivers.

3.2. Scenario Design

Emergency Braking Scenario: Inspired by [29], we developed a dynamic and interactive emergency braking scenario that presents a significant challenge to current autonomous driving models. This scenario is particularly designed to induce failure cases, thereby testing the robustness of both human and AI driving systems. Additionally, our emergency-braking task is modeled after the adaptive cruise control (ACC) paradigm, as described by [59]. In this task, the ego vehicle is instructed to closely follow a lead vehicle in the same lane, maintaining a position that prevents trailing vehicles from overtaking. Participants were informed that a trailing vehicle was closely following them, creating a sense of urgency to maintain speed and avoid being overtaken. The lead vehicle travels at speeds of up to 8 m/s and may brake unexpectedly, requiring the ego vehicle to respond promptly to avoid collisions while maintaining effective following distance. This scenario presents substantial cognitive demands, as it requires continuous attention and real-time adjustments over extended periods to track the behavior of the lead vehicle accurately. To prevent participants from anticipating the event, the lead vehicle braked at a random interval of 4–7 seconds after the previous event. Participants completed two sessions of active driving under different lighting conditions, one during the day and one at night. In addition, they participated in two control sessions where an AI system managed acceleration and braking, and they only needed to control the steering wheel, providing a baseline condition for ERP analysis.

Left Turn Scenario: The left turn scenario followed the design in [60]. Here, the ego vehicle is required to execute a left turn at an intersection while avoiding collisions with oncoming vehicles from the right. These oncoming vehicles approach at variable speeds ranging from 3m/s to 5m/s and do not yield to the ego vehicle. The ego vehicle is afforded the autonomy to merge into the traffic flow at any opportune moment to reach the designated goal point as swiftly as possible. To test the generalization of our algorithm, we trained

the models in Town01 and tested the models in Town05.



Figure 2. Upper: The dataset collection environment. The HTC VIVE Pro Eye VR headset and Logitech G923 Racing Wheel and Pedal give the subject a more realistic driving experience. Lower: the example of scene image of driver’s view, camera’s view, and bird’s eye view.

4. Methods

4.1. EEG preprocessing

The EEG preprocessing was done via the following pipeline in EEGLAB v2020.0 [11] using MATLAB version 2022a. First, the EEG raw data were filtered using a bandpass filter with cutoff frequencies of 0.5 and 100 Hz, employing a zero-phase digital filter. Second, the EEG data was down-sampled from 1000 Hz to 250 Hz. Next, the 50-Hz power line noise was removed via the CleanLine EEGLAB plugin using multi-tapering. Next, automated artifact rejection was applied via EEGLAB CleanArtifacts plugin. Bad channels were removed based on their standard deviations and

Dataset	Active ctr.	Vehicle data	Camera	Hazards	View	#subjects	# Frames	EEG	Gaze
Ours	+	+	$S^{rgb,depth,sem}$	+	VR	20	720K	+	+
Collision Threat [38]	+	-	-	-	screen	25	-	+	-
Sustained-attention [6]	+	+	-	+	screen	27	-	+	-
MPDB [51]	+	-	-	+	screen	35	-	+	+
DrFixD(night) [14]	-	-	S^{rgb}	-	screen	31	67K*	-	+
LBW [31]	-	-	$S^{rgb,depth}$	-	screen	28	123K*	-	+
CoCAtt [46]	+	+	S^{rgb}	-	screen	11	17K*	-	+
MAAD [23]	-	-	S^{rgb}	-	screen	23	60K	-	+
TrafficGaze [13]	-	-	S^{rgb}	-	screen	28	77K*	-	+
DADA-2000 [18]	-	-	S^{rgb}	+	screen	20	658K	-	+
DR(eye)VE [2]	+	-	S^{rgb}	-	on-road	8	555K	-	+
BDD-A [62]	-	+	S^{rgb}	+	screen	45	378K*	-	+
C42CN [50]	+	-	S^{rgb}	+	screen	68	-	-	+
TETD [12]	-	+	S^{rgb}	-	screen	20	100	-	+
3DDS [5]	+	-	S^{rgb}	-	screen	10	192K	-	+

Table 1. Comparison between our dataset and other driving datasets. Ours contains the largest number of frames and is collected using a realistic VR driving simulator. Abbreviations: Camera – S: scene-facing camera, RGB: 3-channel image, sem: semantic segmentation mask. Frame counts marked with * are estimates based on video duration and frame rate.

correlation with other channels. Bad portions of data were rejected with the Artifact Subspace Reconstruction (ASR) algorithm [4]. Next, we replace the removed channels by interpolating them from the neighborhood channels. Independent Component Analysis (ICA) was then performed using the AMICA function, and dipole fitting was applied to identify neural components. Lastly, ICA components classified as non-brain sources with a probability of brain component less than 50% (e.g., muscle, eye, or electrical noise) were removed using ICLabel [41] and dipole fitting, with components exhibiting a residual variance above 15% being discarded. To analyze the ERPs, we defined the onset of the front vehicle’s braking as the event marker. We discard trials if the subject did not focus on the car as revealed by the eye-tracking.

4.2. EEG feature prediction model

Since EEG signals are usually absent during both training and inference, we develop an EEG feature prediction network that estimates EEG activations from scene images. We focus on event-related potentials (ERPs), identified as the most salient EEG feature. To classify ERP-inducing trials, we apply a moving average filter with a window size of 20 samples and a peak-to-peak threshold of 1.7 μ V, which was selected to balance the training set, resulting in a 50/50 class split, and also corresponds to minimum ERP peak amplitudes reported in [47]. Based on this threshold, we partition the dataset into high-ERP and low-ERP trials and perform five-fold cross-validation. A lightweight CNN is designed to predict whether a trial induces an ERP. The architecture has three layers of convolutional layers and an

average pooling layer to reduce the number of parameters and avoid overfitting, as summarized in Fig. 1. To train the model, we use the Binary Cross-Entropy (BCE) loss, given by:

$$\mathcal{L}_{BCE} = -\frac{1}{N} \sum_{i=1}^N [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)] \quad (1)$$

where y_i is the ground truth label (1 for high ERP, 0 for low ERP), and \hat{y}_i is the predicted probability. The model processes sequences of segmentation images and outputs a binary classification indicating ERP presence. Its lightweight design enables real-time inference, supporting efficient RL training.

4.3. Reinforcement Learning Problem Formulation

We frame our task as a goal-oriented collision-avoidance problem, where the ego vehicle must control its acceleration and braking to safely navigate a dynamic environment. This control task is modeled as a Markov Decision Process (MDP), formally defined by the tuple $\{S, A, P, R\}$. At each timestep t , the agent observes a state $s_t \in S$, chooses an action $a_t \in A$, and transitions to a new state $s_{t+1} \in S$ based on the transition dynamics P . The agent receives a scalar reward r_t determined by the reward function $R(s, a)$. The objective of the reinforcement learning (RL) agent is to learn an optimal policy π that maximizes the expected discounted cumulative reward $\sum_{k=0}^{\infty} \gamma^k r_{t+k}$, where $\gamma \in (0, 1)$ denotes the discount factor used to prioritize immediate rewards over distant ones.

State Space: The state space S comprises temporal sequences of visual observations. Specifically, each state is

defined as $S = \{s_t \mid s_t = \{I_{t-2}, I_{t-1}, I_t\}, t \in \mathbb{N}\}$, where each $I \in \mathbb{R}^{h \times w \times 1}$ represents a single-channel semantic segmentation frame captured by a front-facing camera with a 90° field of view, simulated within the CARLA environment [15]. These image sequences provide the temporal context necessary for decision-making.

Action Space: The action space A consists of continuous scalar values representing longitudinal control inputs: $A = \{a_t \mid a_t \in [-1, 1], t \in \mathbb{N}\}$ where $a_t = -1$ corresponds to maximum braking, and $a_t = 1$ indicates full throttle.

Reward Function: The agent is incentivized to avoid collisions and maintain safe following distances. The total reward at each timestep is composed of cognitive reward and environment reward:

$$r_t = R(\cdot | s_t, a_t) = \beta r_{cog}(s_t) + r_{collide} \cdot (s_t \in C_{collide}) + \omega r_{idle} \cdot (s_t \in C_{idle}) + \delta r_{gap}(s_t) \quad (2)$$

where the term $r_{cog}(s_t)$ represents the predicted cognitive reward associated with state s_t , as estimated by the EEG feature predictor (\hat{y}_i , defined in Equation 1). A negative weight $\beta = -1$ is applied to penalize states associated with increased cognitive load or perceived difficulty. $C_{collide}$ and C_{idle} denote the collision and idle states, respectively. The environment reward components are defined as follows: A severe penalty $r_{collide} = -100$ is applied upon collision. A mild penalty $r_{idle} = -1$ is incurred when the ego vehicle’s speed falls below 0.2 m/s. r_{gap} provides a time-gap-based reward that encourages maintaining an optimal following distance.

4.4. Policy network




Figure 3. The policy network determines the vehicle’s policy from a sequence of three segmentation images, starting with a CNN to extract features. These features are then processed through a self-attention layer. The policy network includes two MLP prediction heads: one for estimating the throttle and brake strength and another for predicting TTC, which aids in training regularization.

We design a policy network that incorporates a self-attention mechanism (see Fig. 3). The model takes as input a sequence of three consecutive semantic segmentation images, forming the state representation $s_t \in \mathbb{R}^{h \times w \times 3}$, where h and w denote the height and width of the images, respectively. The policy network jointly predicts two outputs

within a unified framework: an action $a_t \in [-1, 1]$ representing throttle and brake control, and the time-to-collision (TTC) in the range $[0, 5]$. A shallow CNN first encodes the input images into a feature map $F \in \mathbb{R}^{h/16 \times w/16 \times f}$. A self-attention layer is applied to enlarge the receptive field as follows. The feature map is then flattened into $N \in \mathbb{R}^{n \times f}$, where $n = \frac{h}{16} \times \frac{w}{16}$, before being projected into query Q , key K , and value V matrices via fully connected layers f_Q, f_K , and f_V by $Q = f_Q(N), K = f_K(N), V = f_V(N)$, where $Q, K, V \in \mathbb{R}^{n \times d}$, and d represents the dimensionality of the latent space for each token. These matrices are used to compute self-attention as follows: $\text{SelfAttention} = \text{softmax}\left(\frac{QK^\top}{\sqrt{d}}\right)V$. In this formulation, $QK^\top \in \mathbb{R}^{n \times n}$ represents the pairwise attention scores between elements, which are scaled by \sqrt{d} to enhance numerical stability. The softmax function normalizes these scores, converting them into a probability distribution where each row sums to 1. These attention weights are applied to V , producing the final contextualized feature representation. The output of the self-attention layer is then flattened and passed through multilayer perceptrons (MLPs) to predict the action $a_t \in [-1, 1]$ and TTC, where -1 corresponds to maximum braking and 1 represents maximum throttle. The TTC prediction is computed as:

$$TTC = \text{clip}(Dis / (V_{ego} - V_{front}), 0, 5) \quad (3)$$

where Dis is the distance to the closest vehicle, V_{ego} is the ego vehicle’s speed, and V_{front} is the speed of the nearest vehicle ahead. We clip TTC to the range $[0, 5]$ seconds to ensure the model focuses on critical situations. The mean squared error (MSE) loss \mathcal{L}_{mse} aligns the predicted TTC with the ground truth. The overall loss function of the policy network is: $\mathcal{L}_{total} = \mathcal{L}_\pi + \alpha \mathcal{L}_{mse}$, where \mathcal{L}_π is the RL policy loss, \mathcal{L}_{mse} is an auxiliary loss that regularizes TTC prediction. $\alpha = 0.1$ is a weighting factor that balances the RL loss and the MSE loss.

5. Results

5.1. EEG analysis

We further investigated whether the ERP peak correlated with reaction time, which was defined as the interval between the front vehicle’s braking and the subject’s initiation of braking. The Pearson correlation between ERP peak latency and reaction time yielded a statistically significant result with $p = 0.0438$, indicating a significant positive association. Additionally, we compared the ERP waveform under active reaction conditions to a baseline condition in which the subjects were not required to respond. The ERP amplitude was substantially higher when subjects responded to the emergency braking of the front vehicle, as shown in Fig. 4. To assess statistical significance, we conducted a

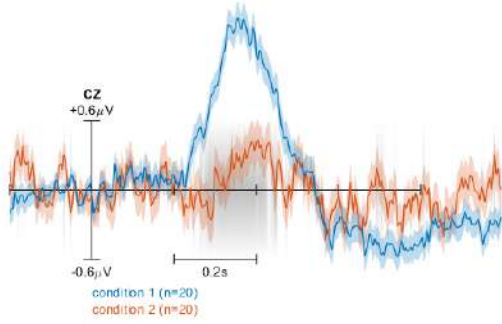


Figure 4. The average ERP wave from 20 participants. Condition 1: participants actively react to the emergency braking. Condition 2: participants are not required to react actively. The gray region indicates a significant difference between conditions.




Figure 5. Machine attention visualization of the policy network of RL across three time steps in the emergency braking scenario. Our model consistently focuses on the lead vehicle.

non-parametric permutation test with 10,000 iterations. The shaded region between 300 and 500 ms revealed a significant difference. This finding aligns with our expectations, as the P3 amplitude is known to increase with greater cognitive effort, supporting the interpretation that P3 amplitude may serve as an index of resource allocation [26].

5.2. EEG feature prediction model

To eliminate the need for EEG data collection during inference, we propose a human cognitive reward model that predicts whether the trial will induce ERP directly from vi-

sual input. To ensure robustness across subjects, we conducted five-fold cross-validation. We first compared our EEG feature prediction model with several commonly used image classification baseline methods, including ResNet-18 [24], Swin-ViT [33], and ConNeXt [34]. Our model achieved a slightly higher average accuracy of 82%. Notably, it reached a processing speed of 204 frames per second (fps), making it more suitable for use as a pretrained model within the RLHF framework. When training the RL with 1M steps, it can save 2.1 hours of training time compared with Resnet-18.

	F ₁	F ₂	F ₃	F ₄	F ₅	Mean	FPS
Resnet-18	82	89	79	79	76	81	73
Swin-ViT	80	85	75	80	77	79	62
ConvNeXt	82	89	75	81	76	80	74
Ours	80	85	77	86	81	82	204

Table 2. Five-fold cross-validation result and the FPS of the EEG feature prediction model.

5.3. Driving performance analysis

5.3.1. Baseline methods

Vanilla RL: For reinforcement learning, we adopt the TD3 algorithm [20], an off-policy approach known for its robustness and improved sample efficiency in continuous control tasks. **Behavior Cloning (BC):** In addition to using the standard TD3 algorithm, we include a baseline that leverages BC [3] to initialize the policy. Specifically, we construct a supervised learning model trained on human driving demonstrations, using the architecture shown in Fig. 3. The model is optimized using an MSE loss. **Prioritized Human In-the-Loop (PHIL) RL:** PHIL utilized the Prioritized Experience Replay (PER) [43] mechanism, which organizes the replay buffer based on the magnitude of Temporal Difference (TD) errors. PHIL replaces the exploratory experience of RL with human guidance into the experience replay buffer, and prioritized experience with a high TD error. **TD3-lag:** TD3-lag is a safe RL method based on TD3, augmented with a Lagrangian-based approach to model collision as a cost function to enforce safety constraints [28]. **RLHF:** A custom interface was developed to present side-by-side pairs of clips sampled from our RL agents exhibiting diverse driving behaviors, as shown in supplementary materials. Each video takes 2 seconds following the RLHF framework [8]. Three participants completed 2000 pairwise queries, each selecting one of three options: “prefer left,” “prefer right,” or “cannot tell.” The whole data collection takes about 10 hours of manual labelling. The responses were stored in csv format and used to train a human preference reward model with the Bradley-Terry loss [8]. This

	Emergency Braking			Left Turn		
	Route Completion \uparrow	Driving Score \uparrow	Infraction \uparrow	Route Completion \uparrow	Driving Score \uparrow	Infraction \uparrow
Vanilla	23 \pm 27	16 \pm 19	0.66 \pm 0.07	60 \pm 16	45 \pm 19	0.68 \pm 0.09
BC	65 \pm 31	55 \pm 29	0.72 \pm 0.08	48 \pm 5	29 \pm 3	0.62 \pm 0.04
PHIL	59 \pm 23	44 \pm 29	0.67 \pm 0.14	38 \pm 32	31 \pm 30	0.68 \pm 0.11
RLHF	73 \pm 32	66 \pm 39	0.80 \pm 0.19	63 \pm 21	49 \pm 28	0.71 \pm 0.15
TD3-lag	44 \pm 33	35 \pm 34	0.72 \pm 0.16	40 \pm 28	32 \pm 22	0.68 \pm 0.06
Ours	85 \pm 43	79 \pm 31	0.84 \pm 0.14	67 \pm 8	57 \pm 10	0.77 \pm 0.05

Table 3. Summary statistics for Route Completion, Driving Scores, Infraction Scores for emergency braking and left-turn scenario.

reward model was then integrated into RL training under identical conditions to our ERP-based predictor.

5.3.2. Evaluation metrics

We used three commonly used evaluation metrics in CARLA, including *driving score*, *route completion*, and *infraction score* as our evaluation metrics.

Driving Score: The driving score is defined as a weighted sum of the product of route completion and infraction penalty: $\text{Driving Score} = \sum_i w_i \cdot (R_i \cdot P_i)$. Let R_i be the percentage of completion of the i -th route, and P_i the infraction penalty of the i -th route. The driving score has a maximum value of 100. **Route Completion (R_i):** Percentage of the route distance completed by an agent before a collision or reaching the destination. It has a maximum value of 100. **Infraction Penalty (P_i):** Each collision will introduce a discount factor of 0.6. Agents start with an ideal base score of 1, which is multiplicatively reduced each time an infraction is committed, down to a minimum value of 0.

5.3.3. Evaluation results

To ensure the robustness and reproducibility of our findings, we trained five models using different random seeds. To evaluate the generalization capability of the RL agents, we trained and evaluated them in different towns. Specifically, we trained the agents in CARLA town 7 and evaluated them in town 4 for the emergency braking scenarios, and trained in town 1 and evaluated in town 5 for the left-turn scenario. Our method achieved the highest route completion score of 85 and 67, infraction penalty score of 0.84 and 0.77, and the driving score of 79 and 57, showing the effectiveness of our framework. We investigate whether the cognitive reward model enhances the internal representation of the policy network by visualizing its machine attention. Specifically, we examine what the policy network has learned through the self-attention mechanism, computed as follows: $\text{MachineAttention} = \text{softmax}\left(\frac{QK^T}{\sqrt{d}}\right)$. The resulting machine attention matrix is averaged across all rows to produce a mean attention vector, which is subsequently reshaped into a $h/16 \times w/16$ feature map, reflecting the spa-

tial distribution of aggregated attention over the 2D layout. As illustrated in Fig. 5, our model consistently attends to the lead vehicle across all time steps, whereas the machine attention in other baseline models appears more dispersed.

6. Limitations

This study includes only two scenarios involving 20 participants, as some individuals experienced motion sickness, limiting broader participation. In future work, the integration of foveated rendering technology may help mitigate motion sickness, enabling more diverse and extensive data collection. Additionally, our current EEG feature prediction model is scenario-specific. Nevertheless, it serves as a foundational step toward leveraging EEG-based feature prediction as implicit feedback for RL agents. With the low-cost data collection system proposed in this work, there is potential to train a more generalizable model in the future as larger and more varied datasets become available.

7. Conclusion

In this work, we present an RL framework that integrates human cognitive reward feedback derived from EEG signals to enhance intelligent vehicle performance. To the best of our knowledge, this is the first study to introduce EEG-based reward modeling in the context of autonomous driving. First, we collected a multimodal dataset from 20 participants engaged in a simulated driving task with eye-tracking, EEG, and control data. ERP analysis revealed a prominent P3 component. To eliminate the need for collecting EEG data during inference, we developed a lightweight EEG feature prediction model capable of estimating ERP occurrence from scene images. Finally, we integrated the predicted cognitive signals into an RL framework and evaluated our approach in the CARLA driving simulator. The results demonstrate that agents trained with cognitive reward feedback exhibit safer driving behavior. Overall, our findings suggest that human cognitive signals, even when indirectly inferred, can provide valuable guidance for training safer and more human-aligned autonomous systems.

8. Acknowledgement

This work was supported in part by the Australian Research Council (ARC) under discovery grant DP250103612 and DP260101395, ARC Research Hub for Human-Robot Teaming for Sustainable and Resilient Construction (ITRH) grant IH240100016, and Australian National Health and Medical Research Council (NHMRC) Ideas Grant APP2021183. Research was also sponsored in part by the Australia Advanced Strategic Capabilities Accelerator (ASCA) under Contract No. P18-650825 and ASCA EDT DA ID12994, and the Australian Defence Science Technology Group (DSTG) under Agreement No: 12549. Last but not least, many thanks to Mr. Cheng-you Lu for the insightful discussion and Mrs. Haiting Lan for proofreading.

References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 2
- [2] Stefano Alletto, Andrea Palazzi, Francesco Solera, Simone Calderara, and Rita Cucchiara. Dr (eye) ve: a dataset for attention-based tasks with applications to autonomous and assisted driving. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 54–60, 2016. 5
- [3] Michael Bain and Claude Sammut. A framework for behavioural cloning. In *Machine Intelligence 15*, pages 103–129, 1995. 7
- [4] Sarah Blum, Nadine SJ Jacobsen, Martin G Bleichner, and Stefan Debener. A riemannian modification of artifact subspace reconstruction for eeg artifact handling. *Frontiers in human neuroscience*, 13:141, 2019. 5
- [5] Ali Borji, Dicky N Sihite, and Laurent Itti. Computational modeling of top-down visual attention in interactive environments. In *BMVC*, pages 1–12, 2011. 5
- [6] Zehong Cao, Chun-Hsiang Chuang, Jung-Kai King, and Chin-Teng Lin. Multi-channel eeg recordings during a sustained-attention driving task. *Scientific data*, 6(1):19, 2019. 5
- [7] Jonathan Chang, Masatoshi Uehara, Dhruv Sreenivas, Rahul Kidambi, and Wen Sun. Mitigating covariate shift in imitation learning via offline data with partial coverage. *Advances in Neural Information Processing Systems*, 34:965–979, 2021. 1
- [8] Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30, 2017. 1, 2, 7
- [9] Radoslaw Martin Cichy, Aditya Khosla, Dimitrios Pantazis, and Aude Oliva. Dynamics of scene representations in the human brain revealed by magnetoencephalography and deep neural networks. *NeuroImage*, 153:346–358, 2017. 3
- [10] Eric Courchesne, Steven A Hillyard, and Robert Galambos. Stimulus novelty, task relevance and the visual evoked potential in man. *Electroencephalography and clinical neurophysiology*, 39(2):131–143, 1975. 2
- [11] Arnaud Delorme and Scott Makeig. Eeglab: an open source toolbox for analysis of single-trial eeg dynamics including independent component analysis. *Journal of neuroscience methods*, 134(1):9–21, 2004. 4
- [12] Tao Deng, Kaifu Yang, Yongjie Li, and Hongmei Yan. Where does the driver look? top-down-based saliency detection in a traffic driving environment. *IEEE Transactions on Intelligent Transportation Systems*, 17(7):2051–2062, 2016. 5
- [13] Tao Deng, Hongmei Yan, Long Qin, Thuyen Ngo, and BS Manjunath. How do drivers allocate their potential attention? driving fixation prediction via convolutional neural networks. *IEEE Transactions on Intelligent Transportation Systems*, 21(5):2146–2154, 2019. 5
- [14] Tao Deng, Lianfang Jiang, Yi Shi, Jiang Wu, Zhangbi Wu, Shun Yan, Xianshi Zhang, and Hongmei Yan. Driving visual saliency prediction of dynamic night scenes via a spatio-temporal dual-encoder network. *IEEE Transactions on Intelligent Transportation Systems*, 2023. 5
- [15] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator. In *Conference on robot learning*, pages 1–16. PMLR, 2017. 1, 6
- [16] Yiqun Duan, Zhuoli Zhuang, Jinzhao Zhou, Yu-Cheng Chang, Yu-Kai Wang, and Chin-Teng Lin. Enhancing end-to-end autonomous driving systems through synchronized human behavior data. In *Proceedings of the 1st International Workshop on Brain-Computer Interfaces (BCI) for Multimedia Understanding*, pages 1–8, 2024. 1
- [17] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024. 2
- [18] Jianwu Fang, Dingxin Yan, Jiahuan Qiao, Jianru Xue, He Wang, and Sen Li. Dada-2000: Can driving accident be predicted by driver attentionf analyzed by a benchmark. In *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pages 4303–4309. IEEE, 2019. 5
- [19] Angelos Filos, Panagiotis Tigkas, Rowan McAllister, Nicholas Rhinehart, Sergey Levine, and Yarin Gal. Can autonomous vehicles identify, recover from, and adapt to distribution shifts? In *International Conference on Machine Learning*, pages 3145–3153. PMLR, 2020. 1
- [20] Scott Fujimoto, Herke Hoof, and David Meger. Addressing function approximation error in actor-critic methods. In *International conference on machine learning*, pages 1587–1596. PMLR, 2018. 7
- [21] William J Gehring, Brian Goss, Michael GH Coles, David E Meyer, and Emanuel Donchin. A neural system for error detection and compensation. *Psychological science*, 4(6): 385–390, 1993. 2
- [22] Alessandro T Gifford, Kshitij Dwivedi, Gemma Roig, and Radoslaw M Cichy. A large and rich eeg dataset for modeling

- human visual object recognition. *NeuroImage*, 264:119754, 2022. 3
- [23] Deepak Gopinath, Guy Rosman, Simon Stent, Katsuya Terahata, Luke Fletcher, Brenna Argall, and John Leonard. Maad: A model and dataset for “attended awareness” in driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3426–3436, 2021. 5
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 7
- [25] HJ Heinze, George R Mangun, W Burchert, H Hinrichs, M Scholz, TF Münte, A Gös, M Scherg, S Johannes, H Hundseshagen, et al. Combined spatial and temporal imaging of brain activity during visual selective attention in humans. *Nature*, 372(6506):543–546, 1994. 2
- [26] Jack B Isreal, Gregory L Chesney, Christopher D Wickens, and Emanuel Donchin. P300 and tracking difficulty: Evidence for multiple resources in dual-task performance. *Psychophysiology*, 17(3):259–273, 1980. 2, 7
- [27] Bernhard Jaeger, Kashyap Chitta, and Andreas Geiger. Hidden biases of end-to-end driving models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8240–8249, 2023. 1
- [28] Jiaming Ji, Borong Zhang, Jiayi Zhou, Xuehai Pan, Weidong Huang, Ruiyang Sun, Yiran Geng, Yifan Zhong, Josef Dai, and Yaodong Yang. Safety gymnasium: A unified safe reinforcement learning benchmark. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. 7
- [29] Xiaosong Jia, Zhenjie Yang, Qifeng Li, Zhiyuan Zhang, and Junchi Yan. Bench2drive: Towards multi-ability benchmarking of closed-loop end-to-end autonomous driving. *Advances in Neural Information Processing Systems*, 37:819–844, 2025. 1, 4
- [30] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *nature*, 596(7873):583–589, 2021. 2
- [31] Isaac Kasahara, Simon Stent, and Hyun Soo Park. Look both ways: Self-supervising driver gaze estimation and road scene saliency. In *European Conference on Computer Vision*, pages 126–142. Springer, 2022. 5
- [32] Timo Lajunen and Heikki Summala. Driving experience, personality, and skill and safety-motive dimensions in drivers’ self-assessments. *Personality and individual differences*, 19(3):307–318, 1995. 3
- [33] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 7
- [34] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 7
- [35] Angela Lopez-Cardona, Carlos Segura, Alexandros Karatzoglou, Sergi Abadal, and Ioannis Arapakis. Seeing eye to ai: Human alignment via gaze-based response rewards for large language models. *arXiv preprint arXiv:2410.01532*, 2024. 2
- [36] Steven J Luck. *An introduction to the event-related potential technique*. MIT press, 2014. 2
- [37] Steven J Luck and Steven A Hillyard. Electrophysiological correlates of feature analysis during visual search. *Psychophysiology*, 31(3):291–308, 1994. 2
- [38] Gustav Markkula, Zeynep Uludağ, Richard McGilchrist Wilkie, and Jac Billington. Accumulation of continuously time-varying sensory evidence constrains neural and behavioral responses in human collision threat detection. *PLoS Computational Biology*, 17(7):e1009096, 2021. 5
- [39] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013. 1, 2
- [40] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022. 2
- [41] Luca Pion-Tonachini, Ken Kreutz-Delgado, and Scott Makeig. Iclabel: An automated electroencephalographic independent component classifier, dataset, and website. *NeuroImage*, 198:181–197, 2019. 5
- [42] John Polich and Marco D Comerchero. P3a from visual stimuli: typicality, task, and topography. *Brain topography*, 15: 141–152, 2003. 2
- [43] Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized experience replay. *arXiv preprint arXiv:1511.05952*, 2015. 7
- [44] Hao Shao, Letian Wang, Ruobing Chen, Hongsheng Li, and Yu Liu. Safety-enhanced autonomous driving using interpretable sensor fusion transformer. In *Conference on Robot Learning*, pages 726–737. PMLR, 2023. 1
- [45] Hao Shao, Letian Wang, Ruobing Chen, Steven L Waslander, Hongsheng Li, and Yu Liu. Reasonnet: End-to-end driving with temporal and global reasoning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13723–13733, 2023. 1
- [46] Yuan Shen, Niviru Wijayaratne, Pranav Sriram, Aamir Hasan, Peter Du, and Katherine Driggs-Campbell. Cocatt: A cognitive-conditioned driver attention dataset. In *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, pages 32–39. IEEE, 2022. 5
- [47] Archisman Shubhadarshan and Uneza Gaiwale. Effect of advancing age on event-related potentials (p300) measures. *The Egyptian Journal of Otolaryngology*, 40(1):107, 2024. 5
- [48] Maryam Soltani and Robert T Knight. Neural origins of the p300. *Critical Reviews™ in Neurobiology*, 14(3-4), 2000. 2
- [49] Samuel Sutton, Margery Braren, Joseph Zubin, and ER John. Evoked-potential correlates of stimulus uncertainty. *Science*, 150(3700):1187–1188, 1965. 2

- [50] Salah Taamneh, Panagiotis Tsiamyrtzis, Malcolm Dcosta, Pradeep Buddharaju, Ashik Khatri, Michael Manser, Thomas Ferris, Robert Wunderlich, and Ioannis Pavlidis. A multimodal dataset for various forms of distracted driving. *Scientific data*, 4(1):1–21, 2017. 5
- [51] Xiaoming Tao, Dingcheng Gao, Wenqi Zhang, Tianqi Liu, Bing Du, Shanghang Zhang, and Yanjun Qin. A multimodal physiological dataset for driving behaviour analysis. *Scientific data*, 11(1):378, 2024. 5
- [52] CARLA Team. Carla autonomous driving leaderboard. <https://leaderboard.carla.org/>, 2025. 1
- [53] Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, et al. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*, 2024. 2
- [54] Marin Toromanoff, Emilie Wirbel, and Fabien Moutarde. End-to-end model-free reinforcement learning for urban driving using implicit affordances. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7153–7162, 2020. 1
- [55] Hein T van Schie, Rogier B Mars, Michael GH Coles, and Harold Bekkering. Modulation of activity in medial frontal and motor cortices during error observation. *Nature neuroscience*, 7(5):549–554, 2004. 2
- [56] Xiaofei Wang, Hsiang-Ting Chen, and Chin-Teng Lin. Error-related potential-based shared autonomy via deep recurrent reinforcement learning. *Journal of Neural Engineering*, 19(6):066023, 2022. 2
- [57] Xiaofei Wang, Hsiang-Ting Chen, Yu-Kai Wang, and Chin-Teng Lin. Implicit robot control using error-related potential-based brain–computer interface. *IEEE Transactions on Cognitive and Developmental Systems*, 15(1):198–209, 2022. 2
- [58] Zizhao Wang, Junyao Shi, Iretiayo Akinola, and Peter Allen. Maximizing bci human feedback using active learning. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 10945–10951. IEEE, 2020. 2
- [59] Zhe Wang, Helai Huang, Jinjun Tang, Xianwei Meng, and Lipeng Hu. Velocity control in car-following behavior with autonomous vehicles using reinforcement learning. *Accident Analysis & Prevention*, 174:106729, 2022. 4
- [60] Jingda Wu, Zhiyu Huang, Wenhui Huang, and Chen Lv. Prioritized experience-based reinforcement learning with human guidance for autonomous driving. *IEEE Transactions on Neural Networks and Learning Systems*, 2022. 4
- [61] Penghao Wu, Xiaosong Jia, Li Chen, Junchi Yan, Hongyang Li, and Yu Qiao. Trajectory-guided control prediction for end-to-end autonomous driving: A simple yet strong baseline. *Advances in Neural Information Processing Systems*, 35:6119–6132, 2022. 1
- [62] Ye Xia, Danqing Zhang, Jinkyu Kim, Ken Nakayama, Karl Zipsper, and David Whitney. Predicting driver attention in critical situations. In *Computer Vision–ACCV 2018: 14th Asian Conference on Computer Vision, Perth, Australia, December 2–6, 2018, Revised Selected Papers, Part V 14*, pages 658–674. Springer, 2019. 5
- [63] Duo Xu, Mohit Agarwal, Ekansh Gupta, Faramarz Fekri, and Raghupathy Sivakumar. Accelerating reinforcement learning using eeg-based implicit human feedback. *Neurocomputing*, 460:139–153, 2021. 2, 3
- [64] Huzheng Yang, James Gee, and Jianbo Shi. Brain decodes deep nets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23030–23040, 2024. 3
- [65] Lantao Yu, Tianhe Yu, Jiaming Song, Willie Neiswanger, and Stefano Ermon. Offline imitation learning with sub-optimal demonstrations via relaxed distribution matching. In *Proceedings of the AAAI conference on artificial intelligence*, pages 11016–11024, 2023. 1
- [66] Boyuan Zheng, Sunny Verma, Jianlong Zhou, Ivor W Tsang, and Fang Chen. Imitation learning: Progress, taxonomies and challenges. *IEEE Transactions on Neural Networks and Learning Systems*, (99):1–16, 2022. 1
- [67] Zhuoli Zhuang, Cheng-You Lu, Yu-Cheng Fred Chang, Yu-Kai Wang, Thomas Do, and Chin-Teng Lin. Aegis: Human attention-based explainable guidance for intelligent vehicle systems. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pages 1–19, 2025. 2, 4