CAREER: B-Morpheus: Perceptual-Motor Active Inference Framework for Highly Automated Mobility Systems

Introduction

The long-term career goal of the Principal Investigator (PI) is to discover computational principles of general intelligence, replicate them to build Artificial Intelligence (AI) for highly automated mobility, and integrate them with lifelong learning through project-based approaches in K-12, undergraduate, and graduate students. The overarching objective of this project is to explore a novel Machine Learning (ML) approach and propose an alternative and more effective perspective in perception and control for highly automated mobility systems. This will be achieved by addressing existing problems in other methods, such as Imitation Learning (IL) and Reinforcement Learning (RL). The PI aims to establish a sturdy foundation in the field of intelligent machines that replicate the functionality of the brain for highly automated mobility through the proposed synergistic activities of interdisciplinary research and education.

Overview of Proposed Research and Education Plan

Research Overview: Recent advances in AI, particularly in ML and Computer Vision (CV), have demonstrated promising results in highly automated mobility systems. However, it is unlikely that the achievement of *Level 5 - Full Driving Automation* [1] will be realized in the near future, considering the recent Robotaxi accidents of Waymo and GM Cruise, as well as several Tesla Autopilot crashes [2–10]. This is primarily due to the inherent limitations of existing machine-learning methods, as evidenced by the causes of the accidents [4, 9, 11–15]. Traditional approaches rely on passive data collection and use them to train ML systems. These learned action policies may have limited adaptability to new driving scenarios and varying conditions, leading to performance degradation in real-world situations. In order to overcome these challenges, this project aims to design and build B-Morpheus, a novel Active Inference Framework (AIF) using perceptual-motor learning [16–18] for highly automated mobility systems. B-Morpheus has a generative internal model that learns the associative relationship between the actions of an agent and corresponding environmental changes by interacting with it through motor babbling (random movement of its bodies to obtain proprioception) [19–22]. The motor imagery capability of B-Morpheus uses the generative model to predict the sensory outcome of actions. By integrating the generative models and motor imagery into the AIF, B-Morpheus achieves goal-directed tasks more robustly in automated mobility. B-Morpheus actively explores and refines its perception and action strategies through perceptual-motor learning with active inference principles. This leads to more efficient learning, improved robustness in dynamic environments, reduced data dependency, and potentially hierarchical representation for structured learning. **Research Thrusts:** The following three research thrusts will be investigated to build B-Morpheus, an intelligent agent based on perceptual-motor AIF: (i) the development of generative models to realize perceptual-motor learning, (ii) the creation of a neural AIF to select optimal action policies through *covert* actions, and (iii) the validation of the perceptual-motor AIF with actual vehicles in a realistic environment.

Education Overview: The long-term educational goal is to attract a more diverse student population to Science, Technology, Engineering, and Math (STEM) disciplines, retain them by providing them with positive hands-on experiences, and encourage them to pursue their careers in STEM. The primary objectives are as follows. The first is to teach students, particularly those from underrepresented groups in the Detroit metropolitan area, and share the PI's passion for research on intelligent systems by organizing summer pre-college engineering programs for K-12 and developing online course modules for undergraduate and graduate students as well as working professionals. The PI will provide more opportunities for underrepresented groups by pursuing sponsorships (GM Foundation and The Donald Lee Smith Fund were sponsors for the PI's previous pre-college programs in multiple years). This initiative will facilitate diversity in STEM and prepare them for emerging industries that are yet to emerge in the United States. The second is to foster students' enjoyment of STEM through hands-on experience using the Project-Based Learning (PBL) methodology, which will introduce STEM to them with positive engagement. The third is to provide students with opportunities to engage in lifelong learning by cultivating self-directed learning through the PBL. Through positive engagement in learning, students will be prepared to study and comprehend new technologies even after their graduation. **Education Thrusts:** The PI's educational hypothesis is that teaching in engineering and science can be improved by means of PBL, hands-on experience, and lifelong learning. (i) PBL is a teaching method where students gain knowledge and skills by working to investigate a complex question, problem, or challenge. Faculty guides students on projects and helps them to apply classroom learning to current real-world challenges. Hands-on experience is based on the idea that learning is facilitated by body engagement. (ii) The preparation for lifelong learning necessitates the provision of a robust foundation for lifelong learning, whereby students can continue to refine and update their knowledge and skillsets.

Expected Significance

Intellectual Merit: The novelties of this project are (i) the first major study on neural AIF for highly automated mobility aiming to validate with actual vehicles in the real world, (ii) the first major investigation of driving behavior as motor skill development through perceptual-motor learning, and (iii) the elimination of the need to use RL frameworks required to train a controller. If successful, the B-Morpheus will provide deeper insight into how to build a robust and reliable intelligent control and perception system for mobility. The proposed perceptual-motor AIF will also provide a firm foundation for embodied cognition, which could be the next step towards Artificial General Intelligence (AGI). This integrated framework for perception and control provides alternative and superior solutions to existing approaches in IL and RL for automated mobility in how the machine learns a task. The novelty of B-Morpheus lies in learning perceptual-motor skills that refer to coordinating sensory information and motor actions to interact with dynamic environments instead of mapping sensory information to actions through sense-plan-act methods. An additional novelty is the way of training complex tasks. IL requires to collect task-specific demonstrations from experts. RL necessitates the definitions of reward, value, and cost functions and has difficulty in transferring policies to a new environment. In contrast, B-Morpheus learns perceptual-motor skills to construct causal representations from actions to sensory feedback and identify an optimal action policy through motor imagery to conduct goal-directed tasks. If successful, B-Morpheus will address the challenges of IL and RL with regard to high data dependency, long-tail distribution, and generalization challenges by learning perceptual-motor skills and integrating perception with action through the AIF.

Broader Impact Highlight: The proposed project, B-Morpheus, aims to discover the computational principles of general intelligence and replicate them to build AI for mobility systems. The results of this project will ensure the robustness and adaptability of highly automated mobility, so it will save more human lives and reduce greenhouse gases through better perception and more efficient control systems. B-Morpheus will also provide new knowledge and tools to educate future engineers and researchers. The education plan integrated with the research will lead to hands-on experience for K-12 and new research opportunities for undergraduate and graduate students. In addition to the proposed research, one hardware platform and full-stack software will be developed to be used in educating the multidisciplinary research topics involved in B-Morpheus. An open-source and modular robotic platform, B-ROVER (RObotic Vehicle for Education and Research), will be designed for the field of AI and autonomy education using a 1/4th scale Electric Vehicle (EV) with a Drive-By-Wire (DBW) system [23–25]. This programmable scaled vehicular platform will be used to introduce perceptual-motor AIF to K-12 and undergraduate students and test students' ideas on the platform without worrying about any safety issues. To provide B-ROVER with a full-stack software package, OSCAR (Open-Source Robotic Car Architecture for Research and Education) supporting Deep Neural Network (DNN)-based autonomous driving [26] will be further developed to support this platform. If successful, B-ROVER EV models with OSCAR will be deployed to provide an affordable and accessible platform to educate perceptual-motor AIF to a broader community, including underrepresented groups in Detroit, Michigan (77.9% of the residents are African American [27]). In addition, a Kia Soul EV with Sygnal DBW [28, 29] and a Chrysler Pacifica Hybrid

[30, 31] with Dataspeed DBW [32]) will be utilized as well to provide students with various opportunities to learn brain-inspired AI and practice their knowledge of autonomous driving.

Research-Education Integration: The proposed research will contribute to a deeper scientific understanding of brain-inspired perception and control methods that closely resemble how an organism sees and acts to accomplish goal-directed tasks. Using the outputs of the proposed multidisciplinary research, the PI will create educational modules in multidisciplinary areas (brain science, cognitive science, CV, and ML). The synergistic integration of research and education will enable the project to attract more K-12, undergraduate, and graduate students and the general public through in-person summer camps and online and in-person workshops.

PI Qualifications and Long-Term Career Goal

PI Qualifications: The PI has a wealth of multidisciplinary experience in industry and academia, spanning product research and development, computational neuroscience, robotic microscopes for brains, neuro-anatomical data analyses and their visualization, robotics, and autonomous vehicles: (i) computational neuroscience [33–37]; (ii) robotic microscope, brain scanner, and neuronal data analysis tools [38–46]; (iii) robotics and mobility [47–54]; (iv) autonomous vehicles [55–61]; (v) cybersecurity and energy systems [62–65]. The PI also has actively participated with students in various types of collegiate challenges: (i) Indy Autonomous Challenge in 2020 - 2021 [66] (advanced to the second round), which is a collaborative effort to challenge university students to imagine, invent and prove a new generation of automated vehicle software (ii) SAE Mobility Forward: AI Mini Challenge in 2021 (3rd Places in Showcase Booth, Solutions Presentation, and Solutions Report). (iii) Battery Workforce Challenge in 2023 - 2026, sponsored by the American Society for Engineering Education (ASEE).

Long-Term Career Goal: The long-term research goal of the PI is to investigate true intelligence through embodied AI and build a robust intelligent machine that replicates the functionality of the brain. The PI will discover computational principles of general intelligence and apply them to highly automated mobility. The PI posits that embodied AI encompasses a diverse array of systems that interact with the physical world by employing sensors and actuators. The PI also believes that embodied AI can be a new frontier that goes beyond traditional AI systems that interact with information-centric environments. The long-term education goals are to attract more students from underrepresented groups to the field of intelligent mobility, which is the future growth engine, and increase enrollment and retention of engineering students who are underrepresented in STEM areas. This will be achieved through various methods, including the NSF Research Experience for Undergraduates (REU), the Summer Undergraduate Research Experience (SURE) for undergraduate students, and summer prep-college engineering camps for K-12. This will prepare them for lifelong learning through open-source hardware and software packages, thereby enabling students to continuously update their knowledge base and skills.

Research Background and Related Work

Research into how intelligent systems can optimally choose actions to control highly automated mobility is making rapid and considerable progress, in large part thanks to the use of ML. Algorithms based on RL using DNN are popular choices and have demonstrated success. Yet, many DNN-based RL algorithms are still suffering from the following problems: (i) sample inefficiency, (ii) difficulty in designing reward functions, (iii) higher potential in local optima even with a good reward function, and (iv) instability and hard-to-reproducible results [67]. To address the aforementioned problems, the PI proposes B-Morpheus, a perceptual-motor AIF that learns associative relationships from actions (cause) to sensory changes (effect). B-Morpheus employs a forward internal model to generate corresponding sensory changes caused by actions and leverages the model to efficiently learn behaviors through model predictions. The perceptual-motor learning thus provides the means for B-Morpheus to achieve goal-directed tasks.

AIF is a normative model that explains sentient behavior. This neuroscience theory characterizes perception, planning, and action as Bayesian inference [68, 69]. Perception in AIF is understood as a dynamic process of minimizing prediction errors in the brain. The prediction error is defined as a difference between the information obtained through sensory input and the output of the brain's internal model of the external world. According to AIF, the brain attempts to reduce these prediction errors by two mechanisms. First, it improves a generative model, which is the probabilistic model of the world. Second, it acts in the world to make sensory inputs match the predictions of the internal model. In AIF, perception is considered an active and inferential process that aims to minimize prediction errors. The internal generative and probabilistic models undergo updates based on the discrepancies between observations from the sensory input and expectations from the models, and further information is actively sought to refine the internal generative models. This active engagement in perception ensures that organisms can effectively understand the environment and respond to changes. In AIF, the outside world, which is called the *veridical* world and more often called a *hidden* state, in the active inference context, cannot be acquired as it is. The veridical world must be mapped to a *recognized* state through sensors. This mapped information is then represented in a compressed and abstracted form. Fig. 1 illustrates the overarching concept of the perception and action cycle in AIF. The observation o, then, should be encoded (or abstracted) into a latent vector through a neural network (commonly, a Convolutional Neural Network (CNN) is used as an encoder). The agent should be able to generate a slightly different state caused by an action. Thus, the extracted features in the form of a latent vector must be able to reconstruct many aspects of the input. The *recognition* density, denoted by $q_\theta(s)$, is a neural mechanism that estimates the hidden state of the world. To accurately approximate $p(s^*)$, this recognition density must be adequately trained.

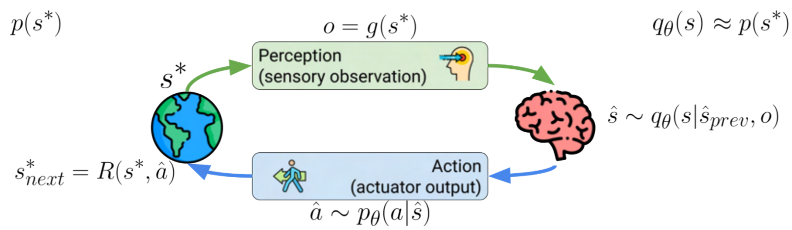

Fig. 1: Perception and action cycle in the AIF. R is the physical world dynamics. s^* is a hidden state. $\hat{s}$ is a recognized state. $g(\cdot)$ is a mapping function from a state to an observation. $p(\cdot)$ is a probabilistic density. $q(\cdot)$ is an internal model that is expected to be an approximation of $p(\cdot)$.

The concept of adjusting a model in response to environmental stimuli in dynamic systems can be traced back to traditional control theories. The Proportional Integral Derivative (PID) control is a linear and closed-loop feedback control system [70]. An error value is calculated by the difference between the model's output and the system's actual output. Subsequently, the actions are modulated with respect to the magnitude of the immediate error, the magnitude of the accumulated errors, and the ratio of the error changes. A more sophisticated model was introduced by the Markov Decision Process (MDP), which is a stochastic decision-making process [71]. The MDP relies on the environment, the agent's action, and the reward of the action from the environment. The Kalman Filter and its extensions, including the Extended Kalman Filter (EKF) and Unscented Kalman Filter (UKF), [72, 73] estimate a posterior by updating a predicted prior to be proportional to the estimation error in the previous step with a Kalman gain. Model Predictive Control (MPC) uses a model, objectives, and constraints to predict future states and control actions to achieve desired future outcomes [74]. More recently, RL has demonstrated more promising results in complex problems. RL is founded upon the Bellman equation [75], which is a process of identifying the optimal policy through an expected value of the action-value function of each state. Deep RL employs a DNN to approximate the Q-function [76] for extensive state spaces.

Efforts have been made in the field of RL to mitigate the inherent limitations. The *World Model* was introduced to reduce direct interaction with the environment by generating models of the environment where the agent can simulate actions and future states [77]. *PlaNet* is a model-based RL learning approach that constructs a model of the environment similar to the *World Model*. It employs the Recurrent State Space Model (RSSM) to combine both deterministic and stochastic components of the environment [78]. This allows *PlaNet* to deal with more intricate dynamic environments. The descendants of *PlaNet* are *Dreamer, V2, and V3,* enhanced scalability and stability in more complex environments [79–81]. This lineage of efforts is more refined by Action Perception Divergence (APD) by divergence minimization [82] and *Director* via hierarchical planning and multi-level decision-making [83]. As RL methods are more sophisticated, they start sharing components with AIF [84]. For example, APD embraced the concept of a unified objective for action and perception by extending representation learning and control. An agent operating under the auspices of APD tries to minimize the joint divergence between the world and a target distribution, a methodology analogous to that employed by AIF. Despite these achievements in RL, the majority of work remains concentrated on relatively simple simulated tasks, including the inverted pendulum [69], mountain car [84–86], frozen lake [87] problems, and video games (Atari, Minecraft, or Super Mario) [79–81, 88]. Very recently, *DreamerV3* was utilized for quasi-realistic autonomous driving, but only simple Bird Eye View diagrams were employed [89].

An effort was also made to assess the capacity of AIF in perceptual-motor learning [91]. In this study, however, they designed a perceptual-motor problem solely for intercepting a moving target to investigate the role of anticipation in the visual guidance of action. World models of active inference agents were explored to model human driving behavior via recorded data [90], but the results showed that this approach is sensitive to input features. Recently, the majority of efforts utilizing AIFs have embraced RL to train its controller [69, 84, 86, 87, 91–93], thereby perpetuating the inherent issues associated with RL.

B-Morpheus is a deep neural AIF that does not utilize RL frameworks for policy optimization. Conversely, B-Morpheus employs perceptual-motor learning through motor imagery to conduct goal-directed tasks. This novel approach mitigates the inherent limitations of RL. Specifically, it eliminates the need for reward or value functions to be defined. Instead, B-Morpheus will introduce a desired state or *preference* [91] to guide action policies. No RL frameworks are required for searching policy. B-Morpheus is capable of internally simulating its future state resulting from action through an internal forward model [94–99], which is implemented by generative models. Furthermore, a more transparent decision-making process is possible with B-Morpheus because actions are selected through a deterministic process by a gradient descent of the difference between the *preference* and an expected state. It is noteworthy that several approaches from disparate disciplines appear to converge towards a similar idealization with embodied AI, which is a long-term career goal of the PI. This provides compelling evidence that this CAREER proposal is on the right track.

Research Plan

The proposed research will be driven by three key thrusts. Thrust 1 is to develop a forward internal model to realize perceptual motor learning. Thrust 2 is to construct an AIF to select the optimal action policy through the use of covert actions and their causal sensory representations. Thrust 3 validates the integrated work of Thrusts 1 and 2 with EVs in realistic environments. Any issues identified during the validation are fed back to refine B-Morpheus, if necessary. See Fig. 2 for the overall research plan.

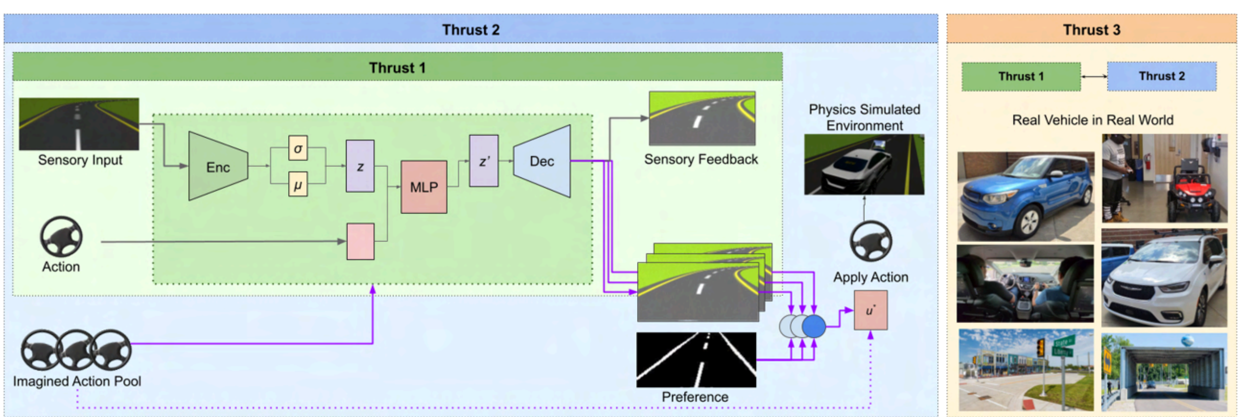

Fig. 2: Research plan in three thrusts. Thrust 1 is the forward internal model to learn perceptual motor skills. Thrust 2 is action selection through covert actions and their sensory feedback. Thrust 3 is to validate the proposed perceptual-motor AIF with EVs in realistic environments.

Research Thrust 1: Perceptual-Motor Learning

Introduction: AIF is a unified framework for perception and action. It elucidates the self-organization of living systems [68, 69] through the Free Energy Principle (FEP) [85, 100–102]. According to FEP, an agent needs to minimize a *surprise* to maintain its internal stability. The *surprise*, however, cannot be directly accessed, so variational free energy was proposed as a tractable proxy of *surprise*, which is the Evidence Lower Bound (ELBO) [103]. Fig. 3 (a) illustrates AIF, which posits that action generates predictions that align with prior expectations and perceptions, while perception ensures that the internal model of the world is consistent with observations to maintain internal equilibrium. For a more comprehensive understanding of the perception and action cycle in AIF, please refer to Fig. 3 (b).

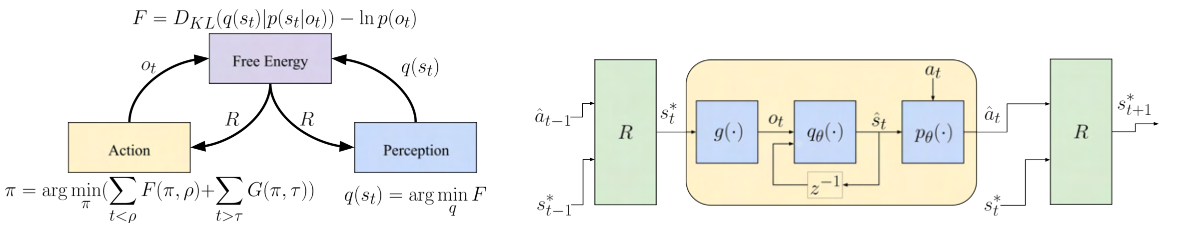

Fig. 3: Active Inference Framework. (a) AIF extends the variational free energy with self-evidencing. F is the free energy. D_{KL} is the Kullback Leibler divergence. π represents action policies through which sensations change to maximize evidence. $q(s_t)$ represents changing beliefs to minimize divergence. $G(\cdot)$ is the expected free energy. Adapted from [104] and modified. (b) Perception and action cycle in the active inference framework. R is the physical world dynamics. s^* is a veridical or hidden state. $s^{\hat{}}$ is a recognized state. $g(\cdot)$ is a mapping function from a state to an observation. $p(\cdot)$ is a probabilistic density. $q(\cdot)$ is an internal model that is expected to be an approximation of $p(\cdot)$.

This perception and action cycle can be realized through the use of a forward internal model [97–99, 105, 106] that predicts future states based on the current state and actions that cause the change. The forward model in motor control was initially proposed by theoretical and computational studies on cognitive science. These studies indicate that the central nervous system internally simulates the behaviors of the motor system in control. This forward model is considered a causal representation of the motor apparatus, meaning that the next state can be predicted by using the current state and the current motor *command*. In this project, the PI will extend the original motor-centric forward model into a world model. A state in the original forward model means an arm's position and its joint angles. However, a state in B-Morpheus refers to the environment, which is a high-dimensional dataset. In a highly automated mobility context, a state can be defined as a piece of visual information sensed by cameras, LIDARs (Light Detecting And Ranging), and radars. In contrast to RL frameworks, the forward internal model in B-Morpheus does not learn a specific task directly. Instead, the proposed perceptual-motor AIF learns how the world changes when an action is performed by an agent. This generative forward model, which is independent of a specific task, ensures the versatility of the proposed approach in performing goal-direct tasks.

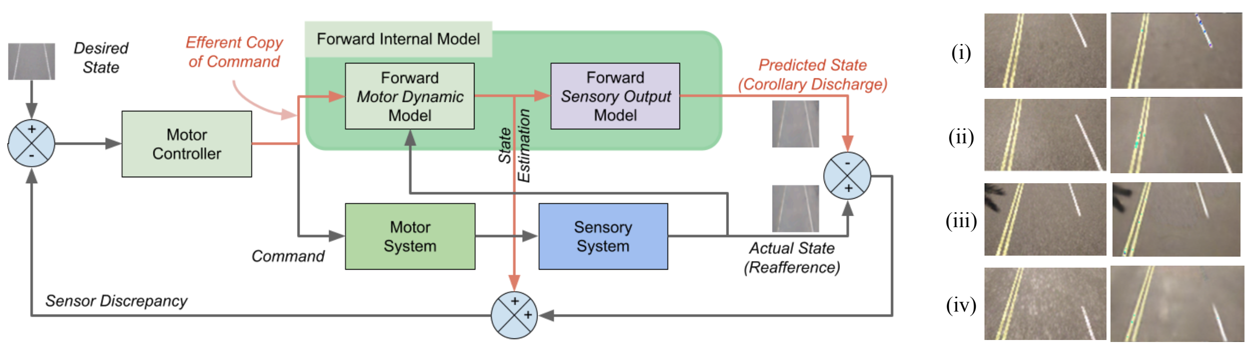

Fig. 4: Image prediction for LFA using CNN-based U-Net. (a) System overview (b) Ground truth (left) vs. predicted (right) images from a model at the training epoch 330. From the top, predictions after (i) 0.03 sec, (ii) 0.15 sec, (iii) 0.3 sec, and (iv) 0.45 sec later. Adapted from [55].

Preliminary Work: The PI initiated research in the field of image prediction context to mitigate control latency [55, 107]. A delayed control signal resulting from inference time can be mitigated by modulating a prediction of a future state to keep stable vehicle control. The latency issue was investigated with the

Adaptive Neural Ensemble Controller (ANEC) [107]. Although ANEC does not directly generate a future state, the fundamental objective of these efforts is to estimate the future state, which constitutes an essential component of the forward internal model, a generative model for the AIF. To further investigate the state estimation caused by an action, the PI proposed to use a Smith predictor [108, 109] and a forward internal model for vision-based Lane-Following Assist (LFA) [55] (Fig. 4). In this study, the PI used a U-Net structure [110] based on CNN for a forward internal model implementation. The system receives the desired state for its task. For LFA, the desired state is the ideal front camera image when the vehicle is driving at the center of the lane. When a driver applies throttle and steering controls, the system feeds these control commands to the forward internal model and motor system. The motor system then modifies the vehicle's speed and orientation based on the driver's motor commands. The sensory system is able to perceive the changes in the state by motor commands and subsequently generates the front camera image output of the changed state. The forward internal model uses an efference copy of motor commands and produces the prediction of the next state. A discrepancy between the sensor readings and the next lane image prediction will be added with state estimation results from the forward internal model in order to adjust the estimation. The image prediction results from the epoch 330 in the training are shown in Fig. 4 (b). The prediction results show some blurred artifacts, but the qualitative validation of the predicted images is promising. Through these preliminary studies, the PI validated the overall research direction and established a firm foundation for this research thrust.

Approaches: The promising preliminary studies have indicated that the forward internal model can be implemented as a generative model for perceptual-motor AIF. To ensure the diversity of a future state, the U-Net architecture will be replaced by a type of conditional Variational Auto Encoder (cVAE) [111–113]. In a VAE, the feature vectors are pre-defined probability distributions (Gaussian distribution is commonly used), represented by means and variances of the distributions. Therefore, by slightly adjusting the values of the feature vectors of a cVAE, the decoder of the cVAE can generate an output that is different but similar to the input conditioned by the agent action and state. The DNN architecture of the cVAE for the forward model is depicted in Fig. 5 (a), wherein a steering angle value, the current speed, and a time are provided as inputs to predict with a front camera image [114, 115]. The prediction results are illustrated in Fig. 5 (b) (c), which depict the generated images at 0.1 sec and 0.3 sec later when the current input and control signal are given. The predicted images exhibit some degree of blurring when compared to the ground truth images, yet they retain a considerable number of salient features pertaining to road images. Indirect validation of the generated images' quality can be conducted by using them for End-to-End (E2E) IL to infer control signals. A video demonstration of this approach can be found at [116].

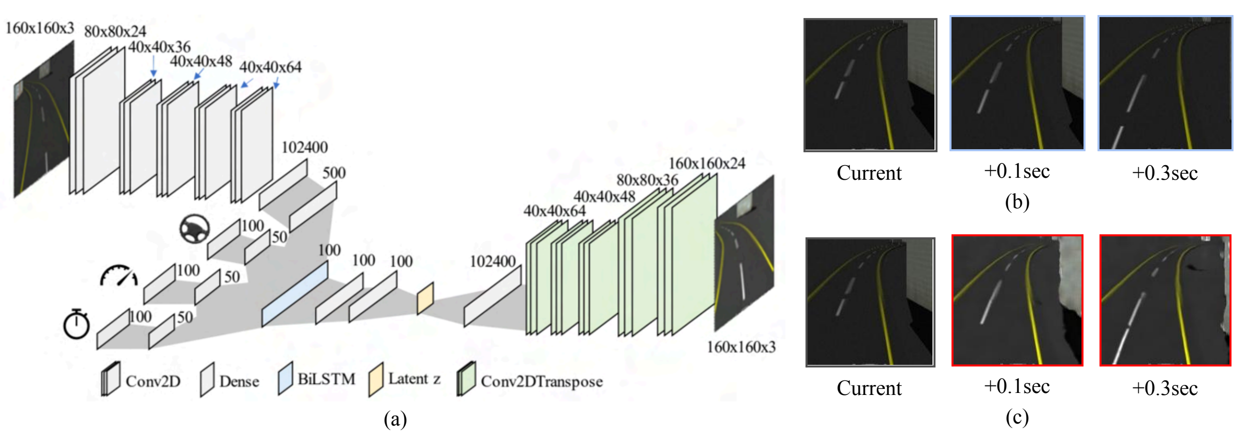

Fig. 5: cVAE model architecture for generating a future road image caused by an action and the current state and prediction results. (a) DNN architecture for a cVAE model. (b) Ground truth images after 0.1 sec and 0.3 sec. (c) Corresponding predicted images.

The forward internal model based on VAE for perceptual-motor learning should be trained through motor babbling [19–22]. In contrast to behavior cloning or RL frameworks, the training dataset for perceptual-motor learning is neither specific to a task nor to expert demonstrations. Conversely, perceptual-motor skill is to learn the causal representations of the agent's action rather than learning a specific task. In order to learn the causality, the training data must have an *action-babbling* type of data, which means a vehicle must be driven to collect more diverse cases by moving wobbly in a mobility application. The training dataset must have sufficient cases from which causality can be inferred. In a preliminary study, the PI collected 744,180 images to train a cVAE model. The performance of steering angle predictions using the generated image was measured in MSE, which was 3.02°, indicating that the generated image is suitable for use in steering angle prediction for an IL [115].

Anticipated Results: A forward internal model employing a generative model in the form of a cVAE-like neural network. If successful, perceptual-motor learning will be prepared for Research Thrust 2.

Success Metrics: The output of the forward internal model should be sufficiently accurate to be used as input to an E2E IL-based controller as an indirect measure. The driving performance with the generated images from the forward model will be compared with the ground truth dataset. Driving performance tests must demonstrate at least 95% of the baseline to be considered a success.

Potential Challenges: VAEs are known for blurry output due to the randomness of the source information [117, 118]. If the output is not sufficiently clear, the option of using a Generative Adversarial Network (GAN) [119] or a diffusion models [120, 121] to enhance the image quality will be explored.

Research Thrust 2: Action Selection via Motor Imagery

Introduction: This research thrust aims to develop a method for finding an action policy for AIF. With this thrust, the highly automated mobility agent can perceive the environment and generate control signals by integrating generative models with active inference. To investigate how to determine that one action can be more desirable than others, the core idea of active inference must be explored. The brain is a predictive machine that predicts future sensory information and infers its causes. To minimize prediction errors, the brain builds prior knowledge in the form of an internal model representing the world, which will be investigated through Research Thrust 1. This theory is called the Bayesian brain hypothesis [122], which formulates perception as a constructive process based on internal *generative* models [85, 100]. According to the FEP [101], any adaptive system must minimize its free energy by resisting a tendency of disorder to be stable in its environment. The Free Energy (FE) can be defined as follows. $F(o,\zeta) = D_{KL}(q(s;\zeta)||p(s|o)) - \ln p(o)$, where o is an observation, ζ represents a brain state, D_{KL} is the Kullback Leibler divergence, *q(.)* is called recognition density, meaning an internal probabilistic density in the brain about the world, *p(s|o)* is a posterior probability, s is a state, *p(o)* is an observation density. The first term of the equation, $D_{KL}(q(s;\zeta)||p(s|o))$, represents the perceptual inference, which can be improved by a more accurate internal model: $\arg\min_\zeta F(o,\zeta)$. The second term, $-\ln p(o)$, which means a *surprise* in the sensory observation, can be minimized by acting in the environment. In other words, an active inference agent can minimize this term by acting on the environment to avoid this surprise: $\arg\min_u F(o,\zeta)$. Note that action u is not a parameter of the FE, but it can indirectly influence the observation o. The fundamental principle underlying this approach is that the agent initiates a prediction, which subsequently guides its actions to facilitate the realization of the prediction. Assuming that the brain makes an accurate estimation $q(s;\zeta) \simeq p(s|o)$ (this can be achieved through Research Thrust), the active inference component of the FE can be minimized by lowering the surprise, $-\ln(p(o))$. To achieve this, actions must be carried out with an expected observation state, which is referred to as *preference* [91], in the environment. The *preference* can be defined as a desired state ("*sensorial states that an agent would like to obtain* [91]").

Preliminary Work: To validate this approach, the PI has applied the perceptual-motor AIF to lane-keeping and lane-changing tasks. In this preliminary work, the PI assumed that a vehicle is programmable through a DBW system and equipped with a front-facing camera. In the perceptual-motor AIF, the lane-keeping task is understood as follows: (i) A desired state is set to keep a vehicle in the center of a lane. (ii) A stream of images is transmitted to the AIF. (iii) The forward internal model in the AIF generates sensory feedback resulting from the covert action pool. (iv) The similarity distances between a generated state and the desired state are calculated. (v) A covert action that minimizes the minimum distance between sensory feedback and the desired state is selected as the optimal control signal to be applied. Subsequently, the aforementioned steps are repeated throughout the task. Fig. 8 shows the steering angle changes the sensory feedback after 0.3 seconds. This figure clearly demonstrates the appropriate changes based on different steering angles. The next pictures in Fig. 9 depict scenarios of steering angles over time spans between 0.1 and 0.6 seconds. Fig. 9 (a) illustrates that left-turn steering causes the vehicle to deviate slightly to the right, and Fig. 9 (b) shows when right-turn steering angles are applied. In the case of lane-changing, the task is compartmentalized into sub-tasks: (i) Approaching the lane where the vehicle is going to be located, (ii) crossing a line of a lane, and (iii) keeping the new lane. The work is still at an early stage, but the results are promising. For a video demonstration of this method, refer to [116].

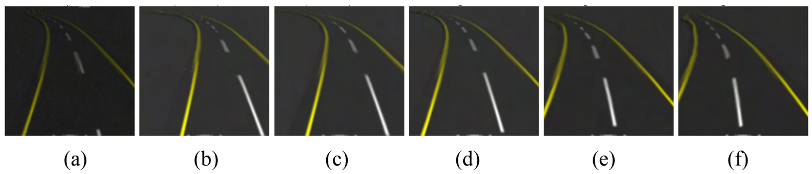

Fig. 8: Sensory feedback by a steering angle for 0.3 sec. (a) Input. (b) 100° left. (c) 70° left. (d) Straight. (e) 70° right. (f) 100° right turn.

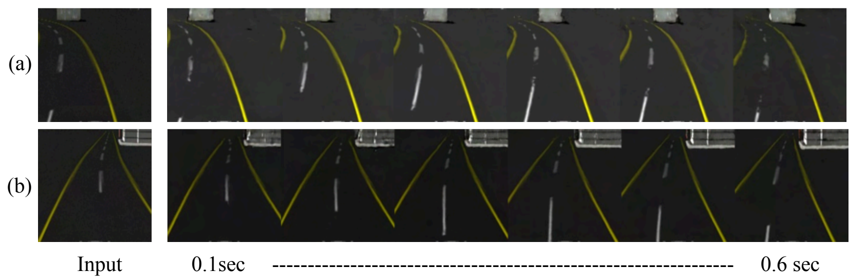

Fig. 9: Various cases of steering angles over time spans between 0.1 and 0.6 seconds. (a) Left-turn steering. (b) Right-turn steering.

To further validate the action selection in driving situations, the PI also employed CARLA [123], an open-source autonomous driving simulator, to apply the proposed method to LFA. For a video demonstration, refer to [124]. An IL agent was used as a baseline to compare driving performance. In these preliminary experiments, agents based on perceptual-motor AIF demonstrated superior performance in terms of adaptability, generalization, and data efficiency compared to the IL baseline and a subset of CoRL2017 benchmarks [125, 126] with IL and RL methods [127].

Approaches: In conventional AIFs, an internal model referred to as a recognition density is expected to be capable of capturing how the environment changes as a result of an action. The construction of a precise recognition density is challenging, given that the environment is comprised of high-dimensional signals. In this project, however, the PI proposes to use a forward internal model trained by perceptual-motor learning instead of capturing the recognition density of the world. This approach will make it much easier for a generative model to be trained even with the same high-dimensional signals, such as images from cameras and 3D point clouds from LIDARs. The generative model will be trained to learn how to make small but meaningful changes caused by immediate action rather than capturing the environment. The identification of an appropriate action to achieve a desired state from the current state can be achieved through the use of an inverse model in control theory or robotics [105, 106, 128]. The training of an inverse model is challenging due to the virtually infinite number of state transitions that must be learned. The PI proposes an alternative approach to the identification of an inverse model. The proposed method employs the simulation theory of cognitive function [121–123], which postulates three components: (i) simulation of action, (ii) simulation of perception, and (iii) anticipation. According to this theory, the pre-motor cortex in the frontal lobes of the brain can be activated to perform without any overt movement. The brain is also capable of generating perceptual activity in the absence of external stimuli.

Imagining to perceive something is essentially the same as actual perception. Fig. 6 illustrates the process of internalization of action and perception.

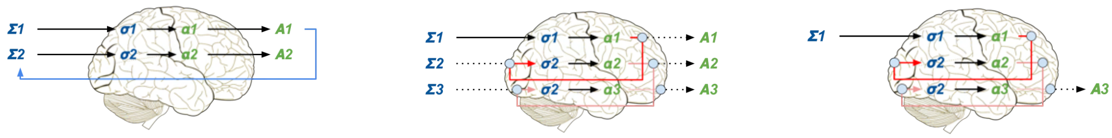

(a) (b) (c)

Fig. 6: Internal simulation theory. (a) Sensorimotor loop. *Σ* is sensory input from an environment. σ is an internal representation of the sensory input. *a* is an internal representation of an action. *A* is an actual motor output. The loop of *Σ*→σ→*a*→*A* is repeated as new sensor input comes. (b) As the sensorimotor loop repeats, action feedback starts being internalized, and the sensory input caused by the action starts being internalized. (c) Simulated perception after internalization is established [129].

The proposed approach involves the use of simulated perception caused by a covert action in optimal action policy to achieve goal-directed tasks. This idea is similar to the *Dreamer* series [79–81, 83], but they employed a model-based RL framework, which necessitates the definition of a model and reward/value functions. In contrast, B-Morpheus utilizes cover action and simulated sensorial feedback. Since the actions and their sensorial consequences are covert, they can be internally simulated as much as possible without concerns for physical consequences. As long as a preferred sensorial consequence can be defined, goal-directed tasks can be accomplished through nested covert actions and their simulated sensorial consequences. In practice, the selection of an action is based on the similarity distance between the sensorial consequences and a desired state. Fig. 7 illustrates the action selection process in B-Morpheus that can be understood as a minimization of the Expected Free Energy (EFE), $\mathcal{G}_t(\pi) = \mathbb{E}_{q_\phi(o_t,x_t|\pi)}[\ln q_\phi(x_t|\pi) - \ln p_\theta(o_t,x_t)]$, where π is an action policy [68, 102] in AIF. Structural SIMilarity (SSIM) [130] is used to measure the distance between a desired state and a state: $\mathrm{SSIM}(\mathbf{x},\mathbf{y}) = [l(\mathbf{x},\mathbf{y})]^\alpha \cdot [c(\mathbf{x},\mathbf{y})]^\beta \cdot [s(\mathbf{x},\mathbf{y})]^\gamma$, where **x** and **y** are signals to be compared, $l(.)$ is luminance, $c(.)$ is contrast, and $s(.)$ is structure. $\alpha > 0$, $\beta > 0$, and $\gamma > 0$. The PI also defines STRC (x, y), which is $s(\mathbf{x}, \mathbf{y})$ to consider only structure without using luminance and contrast. At present, the focus is on immediate action in a deterministic policy. However, the PI will also significantly expand the action selection method through (i) the gradient descent of the similarities to achieve further goals and (ii) hierarchical planning.

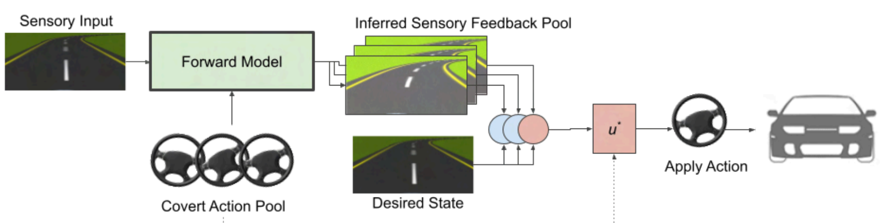

Fig. 7: Action selection through sensorial consequences from simulated actions.

Anticipated Results: Optimal action policy for conducting given goal-directed tasks.

Success Metrics: If an action policy can generate action sequences up to 600 ms in the future, that will be considered successful. Action sequences will be tested in driving behaviors such as lane keeping and lane changing. The results will be compared with multiple behavior cloning approaches in terms of driving performance. ±5% in the performance comparison will be considered successful.

Potential Challenges: SSIM and STRC may not be sufficient to measure similarity distances. Also, generating a longer action sequence for goal-directed behavior is challenging. Peak Signal-to-Noise Ratio (PSNR) [131, 132] can be an alternative metric for similarity comparison. The challenges in generating a longer but stable simulated action and perception loop can be explored through hierarchical planning [81].

Research Thrust 3: Validation in the RealWorld and Refinement B-Morpheus

Introduction: This thrust is to validate B-Morpheus with real vehicles in the real world and provide feedback on any issues to refine the proposed perceptual-motor AIF. The algorithms developed and tested in a simulated environment must be tested in a real vehicle to ensure validity. It should be noted that the PI does not aim to address the Simulation-to-Reality (Sim2Real) [133] problem itself. Instead, to ensure its validity, the PI proposes to test the perceptual-motor AIF with real vehicles in the real world. B-Morpheus will be developed through Thrusts 1 and 2. Yet, its validity in real-world performance must be explored and investigated. Using a real vehicle with an immature algorithm can imply safety concerns for researchers and students. So, one popular choice is to use a small model car. There are numerous efforts to design a model car that is small but can be a viable option for research. Following are some of them: MIT RACECAR [134], F1TENTH [135], MuSHR [136], Go-CHART [137], Duckiebots of Duckietown [138], Donkey Car [139], Amazon DeepRacer [140], NVIDIA Jetson [141] Nano-based two-wheel differential driving robots, and JetBots [142]. However, two significant limitations of these vehicular platforms have been identified: (i) The lack of reproducibility owing to heavy craftsmanship requirements due to extensive modifications of the vehicular platform (removal and replacement of motors, installation of a new Electronic Speed Controller (ESC), custom Printed Circuit Boards (PCBs), etc. (ii) The restricted onboard processing capabilities due to the platform size.

Preliminary Results: For the first and second phases, the PI has been developing the 1/4th scaled EV with a software suite to support research and education for highly automated mobility. This model EV vehicle is named the B-ROVER (Fig. 10 (a)), utilizing Pixhawk PX4 [143] and PX4-Autopilot [144], an open standard for aerial, underwater, and ground vehicles. The new design ensures reproducibility and overcomes the restriction of onboard computing. The PI has also been developing an open-source project, OSCAR [26], a full-stack software package with two simulated vehicles to support this platform. To validate the proposed design's capability, E2E IL utilizing a DNN was tested (Fig. 10 (b) and (c)). For the third phase, the PI has been working to integrate the Sygnal DBW [28] that is designed to be used in the Automotive Safety Integrity Level (ASIL)-D [145] with the Kia Soul EV. The initial setup is shown in [29].

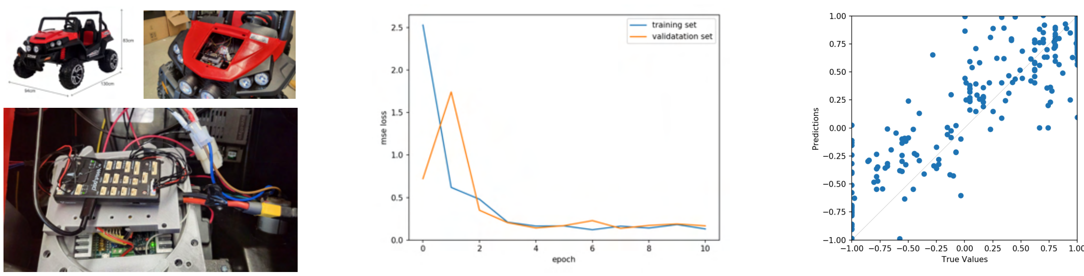

(a) (b) (c)

Fig. 10: B-ROVER design and preliminary DNN-based controller test. (a) Driving controller with a custom-designed housing (b) MSE loss in training (c) Steering predictions. The plot indicates the DNN is well-trained, showing the predictions are close to the ground truth values

Approaches: The PI proposes using three phases of the vehicles to provide the complete test and validation platform: (i) simulated vehicles, (ii) 1/4th scaled EVs, and (iii) full-size real EVs. For the first phase, the PI has been developing OSCAR [26], a full-stack software package based on Robot Operating System (ROS) [146] and Gazebo [147] to support DNN-based E2E IL, that has two simulated vehicles: the Ford Fusion and the Polaris Ranger. The second phase is an affordable and accessible mesoscale EV with a DBW. To address the issues of reproducibility of the design and onboard processing capability restriction in the existing platforms, the PI aims to build a 1/4th scaled EV without extensive modification and provide full-stack software for AI-based perception, planning, and control. The third phase is to use a full-size real vehicle in a realistic environment. The PI has two full-size EVs with DBWs and a sensor suite. These EVs will be utilized to assess the efficacy of B-Morpheus at MCity [148], a world-class facility for testing the performance and safety of connected and automated vehicles at the UM Ann Arbor campus. The PI has access to Mcity as a faculty member in the UM system. Through the Hardware-In-the-Loop (HIL) tests in realistic environments, B-Morpheus will be more validated as a theory and further refined as an applicable framework for highly automated mobility systems in the real world. In addition, a Chrysler Pacifica Hybrid with Dataspeed DBW [32] is ready to be used (Fig. 11).

(a) (b)

Fig. 11: (a) Kia Soul EV with a Sygnal DBW system and (b) Chrysler Pacifica Hybrid with a sensor suite. The DBW kit includes brake, throttle, steer, and shift-by-wire controller modules. It also has the following features for research and development for autonomous driving: driver override by pressing the brake, throttle, shifting, or turning the steering wheel. The sensor suite includes the following components: one high-resolution 64-channel LIDAR, two 32-channel LIDAR, four RGB cameras, and one GPS antenna

Anticipated Results: (i) Two simulated vehicles in ROS. (ii) programmable 1/4th scaled EVs (iii) A full-stack software package supporting DNN-based controllers (iv) Datasets of testing and validating B-Morpheus with real vehicles in realistic dynamic environments. These will provide valuable insights into potential issues in the proposed approaches. These results will refine perceptual-motor AIF for real-world applications.

Success Metrics: A performance metric must achieve a level of at least 90% in a real-world environment with real vehicles to be considered a success. This is due to the challenges inherent in transferring data from a simulated environment to the real world.

Potential Challenges: Latency in perception and control can be a potential challenge. Performance degradation due to perception latency can be addressed by delay compensation through future state estimation and modulation of actions [33, 34, 107, 149].

Integrated Education Plan

The PI is to integrate his research thrusts with education and outreach activities. There are two educational thrusts to propel the integration of research and education, where underrepresented groups, including African American high school students in Detroit Metropolitan areas and undergraduate and graduate women in STEM, will be prioritized. The PI will also apply for Dearborn Discovery Core [150, 151] with new undergraduate courses developed through this proposal. The two key thrusts in education and outreach are as follows: (i) Project-Based Learning (PBL): Faculty guides students on projects and helps them to apply classroom learning to current real-world challenges. UM-Dearborn has been a strong supporter of PBL for its colleges. Every year, the College of Engineering and Computer Science (CECS) exhibits exemplary work on PBL through faculty members. Hands-on experience is a critical component of PBL. When studying engineering, it is crucial to begin practicing from the outset and continue doing so throughout one's academic career. Learning is facilitated by body engagement, according to research on embodied cognition [152, 153]. The PI hypothesizes that STEM education should be more akin to musical instrument training. Charles Kettering once said, *"If we taught music the way we try to teach engineering, in an unbroken four-year course, we could end up with all theory and no music. When we study music, we start to practice from the beginning, and we practice for the entire time"* [154]. (ii) Preparation for Lifelong Learning: STEM learning must be a self-initiated education, which is referred to as lifelong learning. The world is changing at near-lightning speed. In STEM areas, it is even faster. Scientific facts may remain consistent for longer periods, yet the ways of learning and utilizing them have undergone significant changes. In the near future, it is likely that there will no longer be a need to write entry-level code due to the advent of generative AI, such as ChatGPT [155], Gemini [156], and LLaMA [157].

Education Thrust 1: Project-Based Learning through Hands-on Experience

Project-Based Learning (PBL): PBL promotes critical thinking, teamwork, and practical skills to solve real-world problems by providing students with the opportunity to work on a project over several weeks during a semester. Faculty guides students on these projects, helping them to apply classroom learning to current real-world challenges. However, due to the additional workload and commitment to time and effort required for a successful project experience, it is challenging to implement PBL for existing courses, even if there is a strong institutional initiative. Yet, the PI believes the topics of STEM education must be connected to real-world applications and has a long history of working on PBL, aiming at introducing the core concept of PBL, which is learning by doing, to undergraduate and graduate students. Almost all courses taught by the PI have final project components, including ECE-450/650 Mobile Robotics, CE-491 App Development for Mobile Devices, CE-426/626 Real-Time Embedded Systems at Kettering University and ECE-3641 Robotic Manipulation, ECE-5831 Pattern Recognition and Neural Networks at the UM-Dearborn. The PI's home institute has been a strong supporter of PBL. CECS has held an annual PBL showcase to publicize the work. The PI participated in the 2023 PBL showcase with course projects [158]. In 2023, CECS sent a group of faculty to the PBL workshop hosted by the Center for PBL at Worcester Polytechnic Institute (WPI) [159]. The PI volunteered to attend the 2.5-day intensive in-person workshop at WPI. The proposed research thrusts will be integrated with the PBL efforts. The machine intelligence of causality and adaptability will be incorporated into potential student project topics in ECE-5831 Pattern Recognition and Neural Networks and ECE-505 Introduction to Embedded Systems for graduate students, and ECE-3641 Robotic Manipulation, ECE-434 Introduction to Machine Learning, and ECE-4641 Mobile Robotics for undergraduate students. If successful, the PI will seek to significantly expand the PBL efforts to other undergraduate and graduate courses and will also engage more actively with the institutional efforts on PBL.

(a) (b) (c)

Fig. 12: Computer engineering summer camp. (a) Robot programming and smartphone app for controlling the robot through Bluetooth. (b) Programming example code. (c) A group picture after the camp.

Hands-on Experience: STEM education necessitates hands-on experiential training as most STEM fields assume that theoretical knowledge will be applied to solve real-world problems. In this context, hands-on educational modules will be developed as part of this proposal. STEM education should be more akin to musical instrument training. While theoretical knowledge is essential for musicians to play instruments, it is not sufficient. Musicians must engage in hands-on practice with the instrument. Similarly, hands-on experiences are vital through PBL. Likewise, hands-on experiences are essential through PBL. It is often the case that established courses in higher education lack systematic coverage of robotic hardware and software. It is, therefore, evident that hands-on experience constitutes an essential component of PBL, inclusive teaching [160], and active learning in the flipped classroom [161]. The laboratory materials currently utilized in existing courses will be augmented and further refined to enhance the hands-on learning experience. The courses include but are not limited to ECE-3641 Robotic Manipulation, ECE-4641 Mobile Robotics, ECE-505 Introduction to Embedded Systems, and ECE-5831 Pattern Recognition and Neural Networks. In addition, the PI will organize and host a week-long summer engineering camp for 20 high school students. From 2011 to 2019, the PI hosted two week-long pre-college summer camps at Kettering University, Flint, Michigan. Curriculum materials developed by the PI include robotics, computer engineering, and smartphone programming [162]. A small-scale mobile robot was designed and built with DC motors, and students wrote a mobile app to control the robot through Bluetooth connections. According to a post-program survey in the study [162], more than 90% of participants showed satisfied or positive responses. Fig. 12 shows pictures from the computer engineering summer camp prior to the onset of the COVID-19 pandemic. The PI will reinstate these efforts by expanding outreach opportunities to Detroit Metropolitan areas and showcase the work at the Michigan Pre-College & Youth Outreach Conference [163]. After validating algorithms from the Research Thrusts in a simulated environment, 1/4th scale EVs and two full-size EVs that the PI owned will be used to provide students with hands-on experience. Additionally, they will be utilized in an NSF REU Site program, which will be integrated into the initiative to develop hands-on educational modules.

Education Thrust 2: Preparation of Lifelong Learning

To facilitate lifelong learning, it is essential that educational modules be made available beyond the conclusion of this project. Furthermore, the modules must be straightforward to use, requiring minimal cost and effort. The PI will provide the design of a mesoscale EV hardware platform along with open-source software that can be built at a low cost and utilized for robust intelligence studies on a smaller scale. The PI has a strong record in designing and implementing a ride-on EV [23–25], which can be controlled by high-level signals on the ROS, an open-source robotics middleware suite. These educational modules will include documents, videos, and codes. The educational materials will be created and managed using online documentation tools (ReadTheDocs [164] and GitBook [165]), ensuring that they remain up-to-date. Instruction videos will be posted and shared via YouTube, and code will be shared via GitHub [166], an online code repository management tool for broader dissemination. The PI will focus on the reproducibility of the design to ensure the educational materials can be used for lifelong learning. A 1/4th scaled model EV with a DBW based on Pixhawk [143] and PX4-Autopilot [144], B-ROVER [25], will be used as a standard hardware platform for lifelong learning. A full-stack software package to support this hardware, OSCAR [167], will be further developed to support perceptual-motor learning in addition to DNN-based E2E controllers for autonomous driving.

Assessments and Evaluation Plan

To improve the integration of the proposed research and education, an assessment of educational thrusts will be conducted. The PI will collect course evaluation results from the courses in which PBL materials were introduced and assess their effectiveness. Additional survey-type evaluation metrics will be developed. The PI will also work with the Diversity, Equity & Inclusion (DEI) office to improve engagement levels of underrepresented students. The results will be disseminated through ASEE annual conferences so that the PI can share its efforts with broader communities.

Broader Impacts

Scientific/Technological Impact: The proposed project, B-Morpheus, will contribute to a more in-depth scientific understanding of computational principles of general intelligence and ensure robustness in intelligent systems for highly automated mobility. This project will ensure the robustness and adaptability of highly automated mobility, thereby saving more human lives and reducing greenhouse gases through better perception and more efficient control systems. The principles of intelligence and the novel AI method will also provide new knowledge and tools to educate future engineers and researchers.

Education: The proposed education plan will provide students with hands-on experience at K-12 and undergraduate levels and will facilitate new research opportunities for undergraduate and graduate students. In addition to the proposed novel research, this proposal will result in the development of one hardware platform, B-ROVER, which is a 1/4th scale EV base with a custom-designed DBW system [23–25], and one full-stack software package, OSCAR supporting DNN-based E2E IL [26] will be developed. They will be utilized to introduce brain-inspired machine learning for automated mobility to K-12 and undergraduate students and test students' ideas on the platform without concern for safety issues. Additionally, the PI will be a facilitator of campus-wide PBL efforts and will deploy the PBL components to undergraduate and graduate courses. The PI also will develop and host a week-long summer engineering camp for twenty high school students in the Detroit area. Detroit, Michigan, is known as the automotive headquarters, and 77.9% of the residents are African American [27]. In Michigan, 26 automotive-related manufacturers, including the Big Three (General Motors, Ford Motor Company, and Stellantis), have headquarters or technology centers, and they employ more than 100,000 engineers [168]. For lifelong learning for undergraduate and graduate students, several media types (text, pictures, videos) of educational materials, including hardware and software design and source code for the proposed research and education, will be open-sourced.

Integration of Research and Education: The PI will transform novel findings, insights, and datasets from the proposed research into course modules in various disciplines, including brain, cognitive, computer science/engineering, CV, ML, and AI. The students taught by the modules will be inspired to find new research opportunities. The interdisciplinary nature of the proposed research and education will spark interest in STEM [169, 170], particularly among students not initially drawn to STEM, helping them perceive connections between disciplines and identify potential areas of interest in the studies where they can derive enjoyment from learning.

Effectiveness of Evaluation and Assessment Plan: The proposed projects will be evaluated and assessed in both qualitative and quantitative terms, with a clear metric table to ensure their effectiveness. This will entail the following: (i) surveys of at least 90% of participants to measure overall performance; (ii) structured interviews with at least five participants to review; and (iii) dissemination of results through at least one publication in an educational conference or journal per year.

Project Timeline

	Tasks	CAREER Grant Years					Beyond CAREER
		1	2	3	4	5	6 ~ 10 Years
	Research						
1	Generative models to realize perceptual-motor learning						
	Design and development of forward and Inverse model						
	Test and validate in a physics-simulated environment						
2	Perceptual-motor AIF to select optimal action policies through covert actions						
	Develop computational models of simulation theory						
	Integrate perceptual-motor learning to AIF						
3	Validation of perceptual-motor AIF in realistic environments						
	Test and validate in robots and vehicles						
	Education						
1	Project-Based Learning through Hands-on Experience						
	Design and Development of PBL modules for courses						
	PBL workshop for UM-Dearborn						
	REU Site proposal and implementation						
2	Lifelong Learning						
	Open-source hardware and software packages						
	Online documentation and video instructions						

Prior NSF Support

NSF Major Research Instrument Award #2214830: (PI, $244,610.00, 9/01/2022 - 8/31/2025), "*MRI: Acquisition of Autonomous Plug-In Hybrid Vehicle Platform for Multidisciplinary Research and Education at the University of Michigan-Dearborn*." This project aims to acquire a high-performance autonomous EV platform with a sensor suite to advance fundamental science and engineering research and education. **Broader Impacts:** The platform is crucial instrumentation to significantly enhance interdisciplinary research and education at UM-Dearborn in the activities, including embodied cognitive vehicle, in-vehicular network security, energy consumption, environmental perception, cybersecurity, and driver behavior analyses in electric and advanced mobilities. **Intellectual Merits:** The proposed platform will enable collaborative research in a realistic environment with a full-scale programmable vehicle in the aforementioned emerging research areas. **Results:** The procurement process was completed on 7/20/2023, and the equipment was delivered on 1/19/2024. Publications: Kim, D., Khalil, A., Nam, H., and Kwon, J., "OPEMI: Online Performance Evaluation Metrics Index for Deep Learning-Based Autonomous Vehicles" *IEEE Access*, v.11, 2023 https://doi.org/10.1109/ACCESS.2023.3246104.