

Does Computer Vision Matter for Action?

BRADY ZHOU^{1*}, PHILIPP KRÄHENBÜHL^{1,2}, AND VLADLEN KOLTUN¹

¹Intel Labs

²University of Texas at Austin

*Corresponding author: brady.zhou@utexas.edu

Controlled experiments indicate that explicit intermediate representations help action.

<https://doi.org/10.1126/scirobotics.aaw6661>

Biological vision systems evolved to support action in physical environments [1, 2]. Action is also a driving inspiration for computer vision research. Problems in computer vision are often motivated by their relevance to robotics and their prospective utility for systems that move and act in the physical world. In contrast, a recent stream of research at the intersection of machine learning and robotics demonstrates that models can be trained to map raw visual input directly to action [3–6]. These models bypass explicit computer vision entirely. They do not incorporate modules that perform recognition, depth estimation, optical flow, or other explicit vision tasks. The underlying assumption is that perceptual capabilities will arise in the model as needed, as a result of training for specific motor tasks. This is a compelling hypothesis that, if taken at face value, appears to obsolete much computer vision research. If any robotic system can be trained directly for the task at hand, with only raw images as input and no explicit vision modules, what is the utility of further perfecting models for semantic segmentation, depth estimation, optical flow, and other vision tasks?

We report controlled experiments that assess whether specific vision capabilities are useful in mobile sensorimotor systems that act in complex three-dimensional environments. To conduct these experiments, we use realistic three-dimensional simulations derived from immersive computer games. We instrument the game engines to support controlled execution of specific scenarios that simulate tasks such as driving a car, traversing a trail in rough terrain, and battling opponents in a labyrinth. We then train sensorimotor systems equipped with different vision modules and measure their performance on these tasks.

Our baselines are end-to-end pixels-to-actions models that are trained directly for the task at hand. These models do not rely on any explicit computer vision modules and embody the assumption that perceptual capabilities will arise as needed, in the course of learning to perform the requisite sensorimotor task. To these we compare models that receive as additional input the kinds of representations that are studied in computer vision research, such as semantic label maps, depth maps, and optical flow. We can therefore assess whether representations produced in computer vision are useful for sensorimotor challenges. In effect, we ask: What if a given vision task was solved? Would this matter for learning to act in complex three-dimensional environments?

Our first finding is that computer vision does matter. When agents are provided with representations studied in computer vision, they achieve higher performance in sensorimotor tasks. The effect is significant and is consistent across simulation platforms and tasks.

We then examine in finer granularity how useful specific computer vision capabilities are in this context. Our second finding is that some computer vision capabilities appear to be more impactful for mobile sensorimotor operation than others. Specifically, depth estimation

and semantic scene segmentation provide the highest boost in task performance among the individual capabilities we evaluate. Using all capabilities in concert is more effective still.

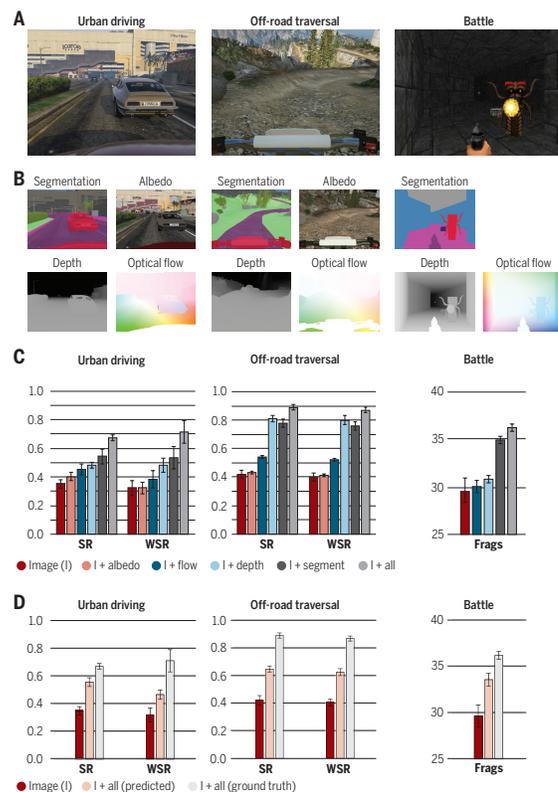


Fig. 1. Assessing the utility of intermediate representations for sensorimotor control. (A) Sensorimotor tasks. From left to right: urban driving, off-road trail traversal, and battle. (B) Intermediate representations. Clockwise from top left: semantic segmentation, intrinsic surface color (albedo), optical flow, and depth. (Albedo not used in battle.) (C) Main results. For each task, we compare an image-only agent with an agent that is also provided with ground-truth intermediate representations. The agent observes the intermediate representations during both training and testing. Success rate (‘SR’) is the fraction of scenarios in which the agent successfully reached the target location; weighted success rate (‘WSR’) is weighted by track length; ‘frags’ is the number of enemies killed in a battle episode. We show mean and standard deviation in each condition. (D) Supporting experiments. In the ‘Image + All (predicted)’ condition, the intermediate representations are predicted in situ by a convolutional network; the agent is not given ground-truth representations at test time. The results indicate that even predicted vision modalities confer a significant advantage.

We also conduct supporting experiments that aim to probe the role of explicit computer vision in detail. We find that explicit computer vision is particularly helpful in generalization, by providing abstraction that helps the trained system sustain its performance in previously unseen environments. Finally, we show that the findings hold even when agents predict the intermediate representations in situ, with no privileged information.

RESULTS

We perform experiments using two simulation environments: the open-world urban and suburban simulation Grand Theft Auto V [7–9] and the VizDoom platform for immersive three-dimensional battles [5, 10]. In these environments we set up three tasks: urban driving, off-road trail traversal, and battle. The tasks are illustrated in Figure S8A.

For each task, we train agents that either act based on the raw visual input alone or are also provided with one or more of the following intermediate representations: semantic and instance segmentation, monocular depth and normals, optical flow, and material properties (albedo). The intermediate representations are illustrated in Figure S8B. The environments, tasks, agent architectures, and further details are specified in the supplement.

Figure S8C summarizes the main results. Intermediate representations clearly help. The supplement reports additional experiments that examine these findings, possible causes, and alternative hypotheses.

ANALYSIS

Our main results indicate that sensorimotor agents can greatly benefit from predicting explicit intermediate representations of scene content, as posited in computer vision research. Across three challenging tasks, an agent that sees not just the image but also the kinds of intermediate representations that are pursued in computer vision research learns significantly better sensorimotor coordination. Even when the intermediate representations are imperfectly predicted in situ by a small, light-weight deep network, the improvements are significant (Figure S8D).

The benefits of explicit vision are particularly salient when it comes to generalization. Equipping a sensorimotor agent with explicit intermediate representations of the scene leads to more general sensorimotor policies. As reported in the supplement, in urban driving, the performance of image-only and image+vision agents is nearly tied on the training set. However, when we test generalization to new areas, the image+vision agent outperforms the image-only agent on the test set even with an order of magnitude less experience with the task during training.

This generalization is exhibited not only by agents equipped with ground-truth representations, but also by agents that predict the intermediate representations in situ, with no privileged information at test time. An agent that explicitly predicts intermediate representations of the scene and uses these explicit representations for control generalizes better to previously unseen test scenarios than an end-to-end pixels-to-actions agent.

CONCLUSION

Computer vision produces representations of scene content. Much computer vision research is predicated on the assumption that these intermediate representations are useful for action. Recent work at the intersection of machine learning and robotics calls this assumption into question by training sensorimotor systems directly for the task at hand, from pixels to actions, with no explicit intermediate representations. Thus the central question of our work: Does computer vision matter for action? Our results indicate that it does. Models equipped with

explicit intermediate representations train faster, achieve higher task performance, and generalize better to previously unseen environments.

REFERENCES

1. J. J. Gibson, *The Ecological Approach to Visual Perception* (Houghton Mifflin, 1979).
2. P. S. Churchland, V. S. Ramachandran, T. J. Sejnowski. A critique of pure vision. *Large-scale neuronal theories of the brain*, C. Koch, J. L. David, eds. (MIT Press, 1994).
3. V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, others, Human-level control through deep reinforcement learning, *Nature* **518** (2015).
4. S. Levine, C. Finn, T. Darrell, P. Abbeel, End-to-end training of deep visuomotor policies, *Journal of Machine Learning Research* **17** (2016).
5. A. Dosovitskiy, V. Koltun, Learning to act by predicting the future, *ICLR* (2017).
6. F. Codevilla, M. Müller, A. López, V. Koltun, A. Dosovitskiy, End-to-end driving via conditional imitation learning, *ICRA* (2018).
7. S. R. Richter, V. Vineet, S. Roth, V. Koltun, Playing for data: Ground truth from computer games, *ECCV* (2016).
8. S. R. Richter, Z. Hayder, V. Koltun, Playing for benchmarks, *ICCV* (2017).
9. P. Krähenbühl, Free supervision from video games, *CVPR* (2018).
10. M. Kempka, M. Wydmuch, G. Runc, J. Toczek, W. Jaśkowski, VizDoom: A Doom-based AI research platform for visual reinforcement learning, *IEEE Conference on Computational Intelligence and Games* (2016).
11. P. Anderson, A. X. Chang, D. S. Chaplot, A. Dosovitskiy, S. Gupta, V. Koltun, J. Kosecka, J. Malik, R. Mottaghi, M. Savva, A. R. Zamir, On evaluation of embodied navigation agents, *arXiv:1807.06757* (2018).
12. M. Lin, Q. Chen, S. Yan, Network in network, *ICLR* (2014).
13. S. Ioffe, C. Szegedy, Batch normalization: Accelerating deep network training by reducing internal covariate shift, *ICML* (2015).
14. D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, *ICLR* (2015).
15. O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, *MICCAI* (2015).

ACKNOWLEDGMENTS

Author contributions. B.Z., P.K., and V.K. formulated the study, developed the methodology, and designed the experiments. B.Z. and P.K. developed the experimental infrastructure and performed the experiments. B.Z., P.K., and V.K. analyzed data and wrote the paper. **Competing interests.** The authors declare that they have no competing interests. **Material availability.** For data not presented in this paper and/or in the supplementary materials, please contact the corresponding author. Code and data for reproducing the results will be released upon publication.

SUPPLEMENTARY MATERIALS

Supplement S1. Synopsis of findings.
 Supplement S2. Materials and methods.
 Supplement S3. Experiments and analysis.

SUPPLEMENTARY MATERIALS

S1. SYNOPSIS OF FINDINGS

We report controlled experiments that assess whether specific vision capabilities, as developed in computer vision research, are useful in mobile sensorimotor systems that act in complex three-dimensional environments. To conduct these experiments, we use realistic three-dimensional simulations derived from immersive computer games. We instrument the game engines to support controlled execution of specific scenarios that simulate tasks such as driving a car, traversing a trail in rough terrain, and battling opponents in a labyrinth. We then train sensorimotor systems equipped with different vision modules and measure their performance on these tasks.

Our baselines are end-to-end pixels-to-actions models that are trained directly for the task at hand. These models do not rely on any explicit computer vision modules and embody the assumption that perceptual capabilities will arise as needed, in the course of learning to perform the requisite sensorimotor task. To these we compare models that receive as additional input the kinds of representations that are produced by modules developed in computer vision research, such as semantic label maps, depth maps, and optical flow. We can therefore assess whether representations produced in computer vision are useful for the sensorimotor challenges we study. In effect, we ask: What if a given vision task was solved? Would this matter for learning to act in complex three-dimensional environments?

Our first finding is that computer vision does matter. When agents are provided with representations studied in computer vision research, such as semantic label maps and depth maps, they achieve higher performance in sensorimotor tasks. The effect is significant and is consistent across simulation platforms and tasks.

We then examine in finer granularity how useful specific computer vision capabilities are in this context. To the extent that the computer vision research community aims to support mobile systems that act in three-dimensional environments, should it invest in improving semantic segmentation accuracy? Optical flow estimation? Intrinsic image decomposition? Do all these tasks matter to the same extent? Our second finding is that some computer vision capabilities appear to be more impactful for mobile sensorimotor operation than others. Specifically, semantic scene segmentation and depth estimation provide the highest boost in task performance among the individual capabilities we evaluate. (Using all capabilities in concert is more effective still.) This can further motivate research on image- and video-based depth estimation, which has gained momentum in recent years, as well as increased investment in solving multiple vision tasks in concert.

We also conduct supporting experiments that aim to probe the role of explicit computer vision in detail. We find that explicit computer vision is particularly helpful in generalization, by providing abstraction that helps the trained system sustain its performance in previously unseen environments. We also find that explicit vision boosts training efficiency, but the performance advantages are maintained even if the baselines are allowed to train ad infinitum. That is, vision does not merely accelerate training, it improves task performance even if training costs are not a factor. We trace this phenomenon to generalization: without the abstraction provided by explicit vision, end-to-end pixels-to-actions models are liable to overfit the training environments. Finally, we show that the findings hold even when agents predict the intermediate representations in situ, with no privileged information.

S2. MATERIALS AND METHODS

Environments

We perform experiments using two simulation environments: The open-world urban and suburban simulation Grand Theft Auto V (GTAV) and

the VizDoom platform for immersive three-dimensional battles.

Grand Theft Auto V. GTAV is a highly realistic simulation of a functioning city and its surroundings, based on Los Angeles [7, 8]. We instrumented GTAV to extract multiple vision modalities by intercepting and modifying DirectX rendering commands as the game is played [9]. The game has an active modding community with tools to access and manipulate game state. The game’s internal scripting interface provides access to control inputs, such as steering, acceleration, and braking of the player’s vehicle. The GTAV environment features a rich set of weather conditions, a full day-night cycle, and a large and diverse game world. We make use of this during both training and evaluation. See Figures S6a and S6b for examples.

VizDoom. The VizDoom environment is a lightweight wrapper around the 1990s video game Doom [10]. The environment uses a 2.5-dimensional software renderer that stores intermediate representations in memory buffers. It allows multiple simulations to run in parallel on the same processor, each running at up to $200\times$ real time. This makes the VizDoom environment particularly suitable for data-hungry sensorimotor learning algorithms. VizDoom allows for detailed control over scenarios via a custom scripting interface [10] and level designer. See Figure S6c for a visual reference.

Sensorimotor tasks

We set up three major tasks across the two simulators: urban driving (GTAV), off-road trail traversal (GTAV), and battle (VizDoom).

Urban driving. In the urban driving task, an agent controls the continuous acceleration, braking, and steering of a car. The agent needs to navigate a functioning urban road network and avoid obstacles by processing a diverse visual input stream. We use the same car model for all experiments. An agent starts at one of a set of predefined locations on a street in an urban area. The goal of the task is to reach a corresponding predefined target location without accidents (e.g., collisions with other vehicles, pedestrians, or the environment). The start and target locations are directly connected by a major road. If the route contains intersections, the agent can reach its target by always turning right at intersections. This avoids higher-level navigation and path finding issues [6]. The task ends if the agent comes within 5 meters of the target, or damages the vehicle in any way. We measure performance on this task via two metrics: success rate (SR) and success rate weighted by track length (WSR). An agent succeeds in scenario i if it comes within 5 meters of the target; this is indicated by S_i . We report both success rate $\frac{1}{N} \sum_i S_i$ and success rate weighted by track length $\sum_i S_i l_i / \sum_i l_i$, where l_i is the length of track i . Our success rate is analogous to the SPL metric [11], since we expect the agent to follow a single path from the start location to the target. Both metrics are normalized, where 0 means failure on all tracks and 1 means that the agent reached the target in all tested scenarios. Figure S6a shows some input images for the urban driving task.

Off-road trail traversal. In this task, the agent operates a dirt bike on a trail in rough terrain. The agent controls acceleration, brake, and steering. In comparison to urban driving, this task has fewer dynamic obstacles, but is more challenging to navigate. The overall evaluation protocol is analogous to the urban driving task: the agent is tested on a set of scenarios (start/goal locations) and is evaluated in terms of SR and WSR. The goal in each scenario is to reach a goal from a start location. The agent perceives the world via a head-mounted camera. Figure S6b shows some example input images.

Battle. This task uses the VizDoom environment. We follow the task setup of Dosovitskiy and Koltun [5] (specifically, D3 in their nomenclature). In each scenario, the goal is to eliminate as many

enemies as possible while navigating a labyrinth. Enemies directly engage and potentially kill the agent. Secondary goals include the collection of ammunition and health packs. The agent controls steering and forward/backward/left/right motion, and can fire a weapon. We measure success by the number of enemies killed (frags) per episode, aggregating statistics over 200,000 time-steps in total. Figure S6c shows some example views seen by the agent.

Further details. For GTAV, we use geographically non-overlapping regions for training, validation, and testing. An agent never drives or rides through a test area during training. All parameters are tuned on a separate validation set. The training set contains 100 scenarios (start/goal pairs), validation 20, test 40. For VizDoom, we follow the standard experimental setup, using the default parameters of Dosovitskiy and Koltun [5], with a train/test split across labyrinth layouts.

Modalities

We study five input modalities: raw RGB image, monocular depth estimate, semantic and instance segmentation, optical flow, and material estimate (albedo).

RGB image. In GTAV, we use the raw RGB image before the head-up display (HUD) is drawn; i.e., without a map, cursor, and other auxiliary information. We resize the image to 128×128 pixels and scale the color range to $[-1, 1]$. Figure S7a shows an example. In VizDoom, we use the simulator output in grayscale normalized to $[-1, 1]$, following Dosovitskiy and Koltun [5].

Depth. We use both the absolute depth estimate obtained directly from the rendering engine and a finite-difference surface normal estimate. In GTAV, the depth estimate is derived from the depth buffer of the DirectX rendering pipeline. In VizDoom, the software renderer produces an 8-bit depth estimate. While the depth image captures the global structure of the scene, the surface normals bring out local detail. Figure S7b shows an example.

Segmentation. We compute both semantic and instance segmentation. In GTAV, we produce instance-level segmentation for pedestrians, vehicles, and animals, and class-level semantic segmentation for road, sidewalk, terrain, and objects such as signs and trees. We use the game’s scripting engine to track all moving objects, and intercept the rendering commands to segment them [9]. For semantic segmentation, we manually labelled all texture IDs in the environment [7, 8]. For VizDoom, we use ‘item’, ‘enemy’, and other object classes, as provided by the simulator. The label map contains a class-indicator vector for each pixel. In addition, we use an instance boundary map to represent instance segmentation. Figure S7c shows an example.

Flow. We use three different representations of optical flow: Plain optical flow, as the motion of each pixel between consecutive captured frames; a static motion estimate, caused by parallax or ego-motion; a dynamic motion estimate, caused by object motion or deformation. Static and dynamic motion sum up to optical flow. Each motion field contains a vertical and horizontal component per pixel. In GTAV, we track all moving objects using the game’s scripting engine to compute the dynamic flow component. Static flow is derived directly from the 3D location of a pixel and the relative camera motion. For VizDoom, we track all moving objects as 2D sprites, and compute the static flow component from the depth map and relative camera motion. We follow the video understanding literature and clip motion estimates into the range $[-32, 32]$ and divide by 64 to normalize. We use backward flow from the current frame to the previous frame, as it is more complete for agents moving forward. Figure S7e shows an example.

Albedo. We compute the intrinsic surface color of each pixel, known as albedo. The albedo factors out the effect of lighting and represents the intrinsic material properties of objects. Albedo is represented as an RGB image normalized to $[-1, 1]$. This modality is only available in GTAV, where it is explicitly represented in the rendering pipeline. Figure S7d shows an example.

Agents

We parameterize all agents as reactive deep networks. They observe the current input modalities and predict the next action.

GTAV. We train the GTAV agent by imitation learning. For each task, we construct 100/20/40 scenarios for our train/validation/test sets respectively. The three sets are located in different geographic regions of the map. Each scenario ranges from 50 to 400 meters in length, with an average of 250 meters. To collect training data for imitation learning, we use the game’s internal autopilot through the built-in scripting library. The autopilot has access to the internal game state and completes each scenario perfectly. The autopilot exposes all the control commands through the scripting interface, but does not allow a player (or agent) to take control once engaged. We record the autopilot’s actions on each training scenario and capture the input modalities, actions (speed, steering angle), and other game state information at 8 frames per second (FPS). Since the autopilot drives close to perfectly, this basic training scheme never learns to recover from mistakes. To counteract this, we inject noisy actions following the setup of Codevilla et al. [6]. We inject 0.25 seconds of noise for every 5 seconds of driving. Frames with injected noise are not used for training, the recovery is.

The agent uses a deep fully-convolutional control network F to output the steering angle θ , given the input modalities of the current time step. We use a network-in-network model [12], with the addition of batch normalization [13] before every major convolution block, and leaky ReLU with slope 0.2. These two modifications stabilize and accelerate training. The architecture is visualized in Figure S2.

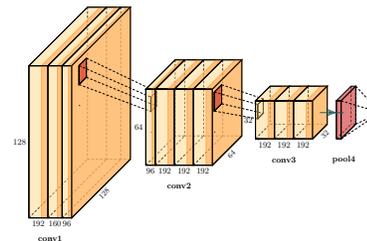


Fig. S2. Imitation learning agent architecture. Figure S10 provides more architecture details.

Each agent sees 100K training frames, the equivalent of 3.5 hours of game time, unless otherwise stated. Training tracks are visited multiple times under different time-of-day and weather conditions.

The control network is trained using the Adam optimizer [14], with batch size 64 and learning rate 10^{-4} , using an \mathcal{L}_1 loss between the network output $F(x)$ and the autopilot’s output θ . We experimented with several different loss functions, such as mean squared error or binned angular classification, but the simple \mathcal{L}_1 performed best. We train each network for 100K iterations, decaying the learning rate by $\frac{1}{2}$ every 25k iterations. Full convergence takes ~ 1 day.

We use a PID controller for acceleration and braking. The PID controller targets a constant velocity for each task. The controller accelerates (positive control) or brakes the vehicle (negative control) to maintain a constant speed of 15 miles per hour for the mountain biking task, and 25 miles per hour for the urban driving task. The agent

updates the control input 8 times per second. In between updates the control signal is repeated.

VizDoom. The VizDoom agent is trained with Direct Future Prediction (DFP) [5], a sensorimotor learning algorithm that uses a stream of measurements as supervisory signal. A DFP agent takes in three inputs: an image, state measurements (health, ammo, etc.), and a goal vector of the desired state measurements (maximize health, ammo, frags). A separate neural-network submodule processes each input into an intermediate representation. A prediction network produces two outputs from this intermediate representation: the expected value of the future measurements over all actions, and the differences in future measurements for all actions. We use a 6-layer convolutional network modeled after the Atari architecture of Mnih et al. [3]. DFP uses an experience buffer starting from a random policy and actions sampled using an ϵ -greedy policy. ϵ decays over training epochs. Given this experience buffer, DFP trains the prediction network by minimizing the \mathcal{L}_1 difference between the predictions and the future measurements. See Dosovitskiy and Koltun [5] for more details and hyperparameters.

The DFP agent predicts its own future state conditioned on the current action, and chooses the optimal action every four frames. Each agent sees a total of 50 million training frames (the equivalent of 2.2 months of subjective time) unless otherwise stated. We normalize all input modalities and stack them as channels in the input image.

Vision networks

To predict the vision modalities, we use a U-net [15], an encoder-decoder architecture with skip connections. To save computation, we use a single encoder, with a separate decoder for each predicted modality. The architecture is specified in detail in Figure S3. The prediction network takes two images as input, representing the current and previous time steps. (The previous time step is only needed for optical flow.) We use a separate loss function for each modality. Depth, normal, factored flow, and albedo use the \mathcal{L}_1 distance between the prediction and the ground truth. We normalize the ground-truth data for training. Segmentation and boundary prediction are trained using a cross-entropy loss. All loss functions are summed without weighting them individually. The network jointly infers all modalities in roughly 100ms.

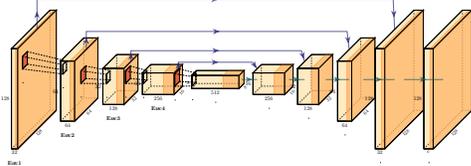


Fig. S3. Network architecture used to infer the vision modalities. Vision modalities are inferred by a U-net. Figure S11 provides further architecture details.

S3. EXPERIMENTS AND ANALYSIS

Figure S8 summarizes the main results on the three tasks. For each task, we compare an image-only agent with an agent for which various aspects of computer vision are solved. The agent observes the ground-truth vision modalities during both training and testing. For GTAV, we report the success rate (SR) and weighed success rate (WSR) across all test scenarios. We ran the entire evaluation 3 times for different daytime and weather conditions. We report the mean and standard deviation for each metric. For VizDoom, we report the average number of enemies killed (frags) over 3 x 50,000 steps of testing. Here as well we report mean and standard deviation.

The results indicate that ground-truth vision modalities help significantly. On urban driving and off-road traversal, the success rate roughly doubles when all modalities are used. In VizDoom, the number of frags increases by more than 20%. All results are statistically significant.

This increased performance could come from various sources. One possibility is that vision-equipped agents generalize better. Another possibility is that the vision modalities provide information that is not otherwise present in the image. We investigate each of these hypotheses independently.

First, we evaluate how much of the vision-equipped agents' performance comes from the ground-truth modalities, and how much can be inferred by a computer vision model that runs in situ with no access to ground-truth data at test time. This is the 'predicted' condition. Figure S9 summarizes the results. Even with a simple U-net model, the inferred vision modalities significantly improve upon the performance of the image-only agent across all tasks.

Next, we investigate if a sufficiently powerful image-only network with the right structure could learn a good generalizable representation given enough data. We train a sensorimotor agent using the same network structure as the model used in the 'predicted' condition, but without ground-truth supervision during training. This is the 'unsupervised' condition. This agent is trained end-to-end using the same imitation learning objective as the image-only agent, and does not receive any additional supervision. Figure S5 shows the performance of a ground truth agent, compared to a 'predicted' and 'unsupervised' agent as the training set size increases on Urban driving and Off-road traversal. The performance of the unsupervised agent increases with the size of the training set, but requires an order of magnitude more data to reach the performance of the predicted vision agent.

Next, we compare the performance of unsupervised, predicted, and ground-truth agents on the training set. The results are reported in Table S1. In order to report compatible statistics, we included a subset of the test tracks in the training set for these experiments. Table S1 reports performance on these tracks. All agents attain nearly perfect performance in this condition. This indicates that the differences between these agents in other experiments can be attributed to different generalization abilities.

Finally, Figure S4 shows the performance for the Battle task as the training set grows. We compare image-only and image+vision agents. As the training set grows, the image+vision agent quickly begins to outperform the image-only agent, reaching higher performance with less data.

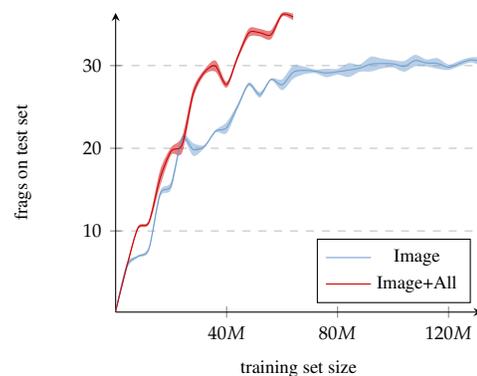


Fig. S4. Performance of image-only and vision-equipped agents as a function of training set size on battle. The shaded area shows the standard deviation of the evaluation metric across different evaluation runs.

	Urban driving		Off-road traversal	
	SR	WSR	SR	WSR
Unsupervised	0.97 ± 0.05	0.98 ± 0.04	1.0 ± 0.0	1.0 ± 0.0
Predicted	1.0 ± 0.0	1.0 ± 0.0	0.94 ± 0.01	0.95 ± 0.01
Ground Truth	0.97 ± 0.05	0.95 ± 0.08	0.98 ± 0.03	0.99 ± 0.02

Table S1. Performance of image-only and vision-equipped agents on training tracks. All models fit the training data almost perfectly.

Discussion

Our main results indicate that sensorimotor agents can greatly benefit from predicting explicit intermediate representations of scene content, as posited in computer vision research. Across three challenging tasks, an agent that sees not just the image but also the kinds of intermediate representations that are pursued in computer vision performs significantly better at sensorimotor control.

Which modalities help most? Across all tasks, semantic and instance segmentation help significantly. Semantic segmentation clearly highlights the drivable or walkable area, and obstacles in the agent’s path show up in the semantic label map and the instance boundaries. The drivable area in particular can be challenging to infer from the raw image in an urban environment, while being clearly delineated in the semantic segmentation. The size and shape of instances in the instance segmentation provides useful cues about the distance and heading of obstacles and other agents.

Depth and normal maps also boost sensorimotor performance. Nearby obstacles are clearly visible in the depth map and the normal map. Depth also provides a direct distance estimate to obstacles, allowing the agent to prioritize which scene elements to focus on.

Optical flow and intrinsic material properties appear less helpful.

Agents with explicit vision generalize better. The benefits of explicit vision are particularly salient when it comes to generalization. Equipping a sensorimotor agent with explicit intermediate representations of the scene leads to more general sensorimotor policies. In urban driving and off-road traversal, the training set performance of unsupervised, predicted, and ground-truth agents is nearly tied in success rate. They all fit the training set well, as shown in Table S1. However, when we test generalization to new areas, the ground-truth agent outperforms the unsupervised agent on the test set even with an order of magnitude less experience with the task during training (Figure S5).

Does more data help? Could it be that an unsupervised agent simply needs to undergo more training to learn to act well? Perhaps with sufficient training, the performance of the image-only agent will catch up with its vision-equipped counterparts? Experiments reported in Figure S5 indicate that this is not straightforward. Even when provided with an order of magnitude more training data, the unsupervised agent does not learn to perform as well as the vision-equipped agent.

REFERENCES

1. J. J. Gibson, *The Ecological Approach to Visual Perception* (Houghton Mifflin, 1979).
2. P. S. Churchland, V. S. Ramachandran, T. J. Sejnowski. A critique of pure vision. *Large-scale neuronal theories of the brain*, C. Koch, J. L. David, eds. (MIT Press, 1994).
3. V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, others, Human-level control through deep reinforcement learning, *Nature* **518** (2015).
4. S. Levine, C. Finn, T. Darrell, P. Abbeel, End-to-end training of deep visuomotor policies, *Journal of Machine Learning Research* **17** (2016).
5. A. Dosovitskiy, V. Koltun, Learning to act by predicting the future, *ICLR* (2017).
6. F. Codevilla, M. Müller, A. López, V. Koltun, A. Dosovitskiy, End-to-end driving via conditional imitation learning, *ICRA* (2018).
7. S. R. Richter, V. Vineet, S. Roth, V. Koltun, Playing for data: Ground truth from computer games, *ECCV* (2016).
8. S. R. Richter, Z. Hayder, V. Koltun, Playing for benchmarks, *ICCV* (2017).
9. P. Krähenbühl, Free supervision from video games, *CVPR* (2018).
10. M. Kempka, M. Wydmuch, G. Runc, J. Toczek, W. Jaśkowski, ViZDoom: A Doom-based AI research platform for visual reinforcement learning, *IEEE Conference on Computational Intelligence and Games* (2016).
11. P. Anderson, A. X. Chang, D. S. Chaplot, A. Dosovitskiy, S. Gupta, V. Koltun, J. Kosecka, J. Malik, R. Mottaghi, M. Savva, A. R. Zamir, On evaluation of embodied navigation agents, *arXiv:1807.06757* (2018).
12. M. Lin, Q. Chen, S. Yan, Network in network, *ICLR* (2014).
13. S. Ioffe, C. Szegedy, Batch normalization: Accelerating deep network training by reducing internal covariate shift, *ICML* (2015).
14. D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, *ICLR* (2015).
15. O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, *MICCAI* (2015).

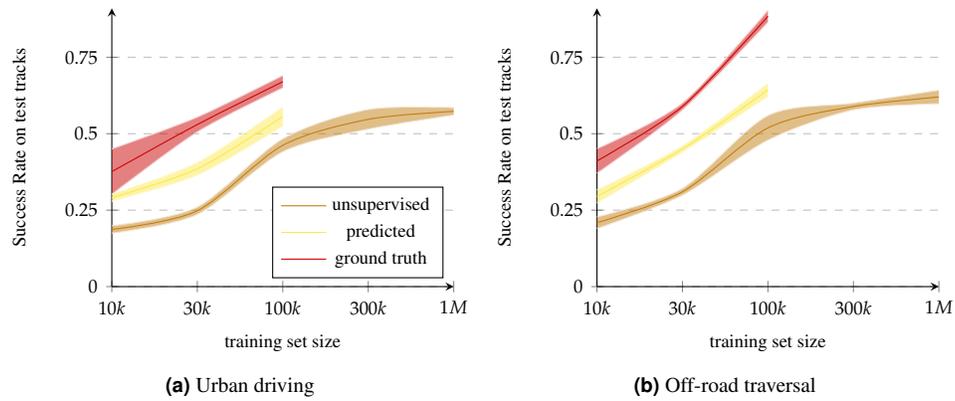


Fig. S5. Performance of unsupervised, predicted, and ground-truth vision agents as a function of training set size. The shaded area shows the standard deviation of the evaluation metric across different evaluation runs. All models share the same control network architecture. Unsupervised and predicted conditions share the same vision network architecture.

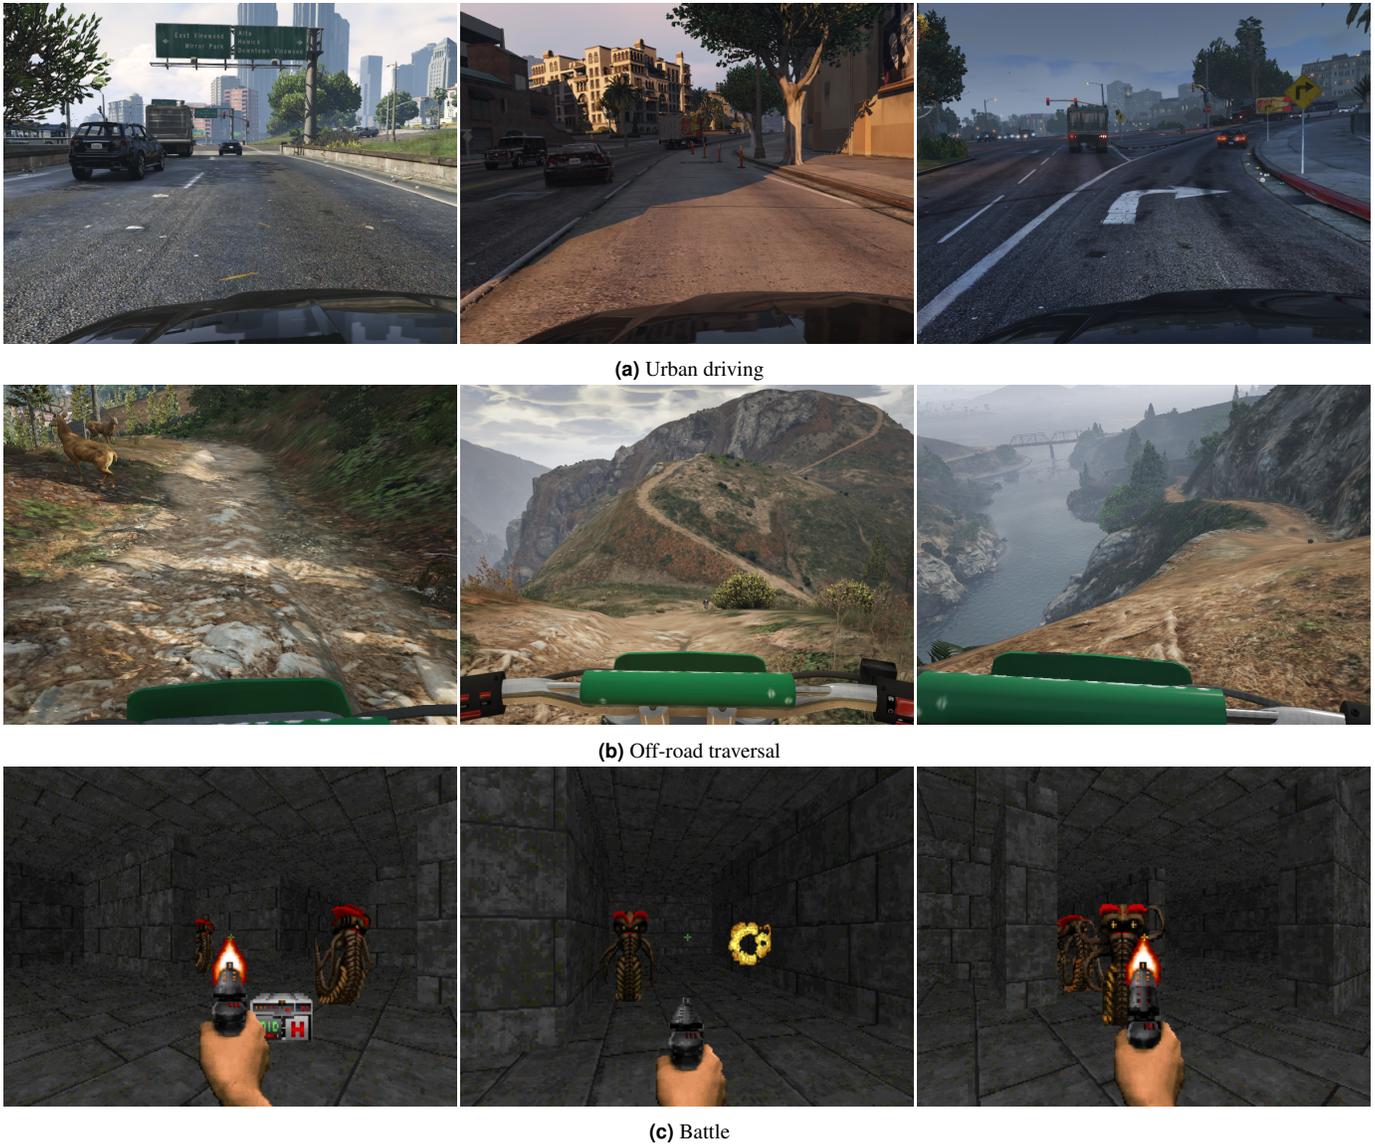


Fig. S6. Three sensorimotor tasks used in our experiments. For each task, we show three views from the agent's viewpoint.

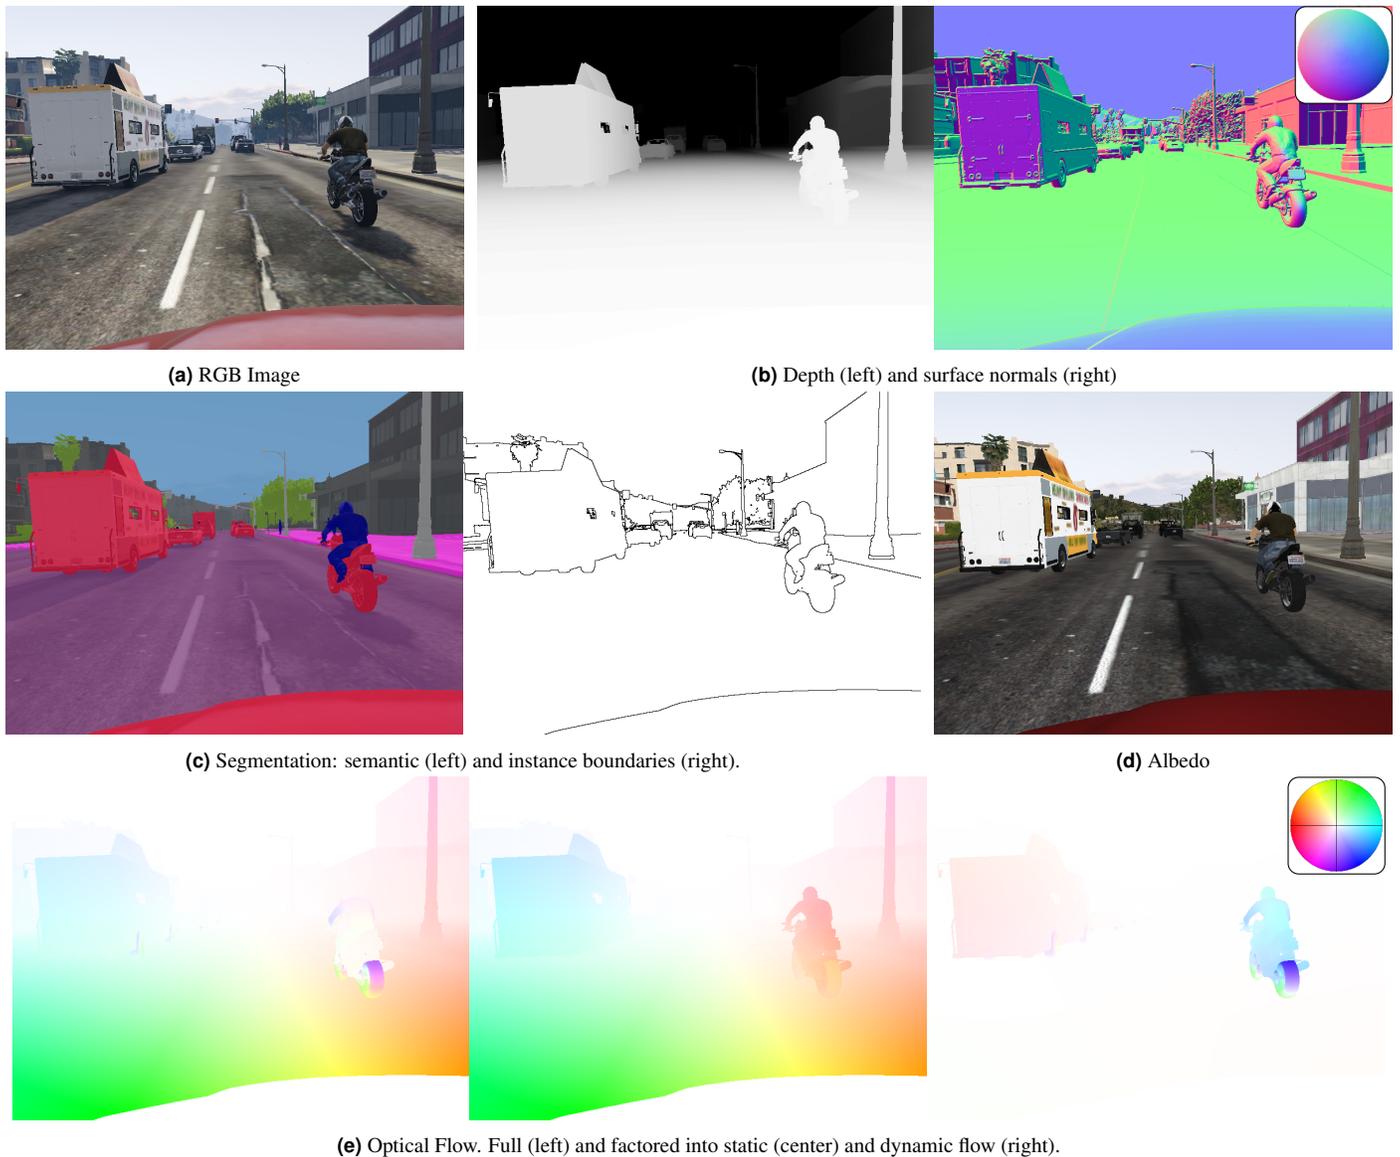


Fig. S7. Different computer vision modalities used in our experiments, illustrated on the urban driving task. For normal maps, the inset shows the different normal directions projected onto a virtual sphere. For optical flow, the inset shows the flow direction as an offset to the center pixel.

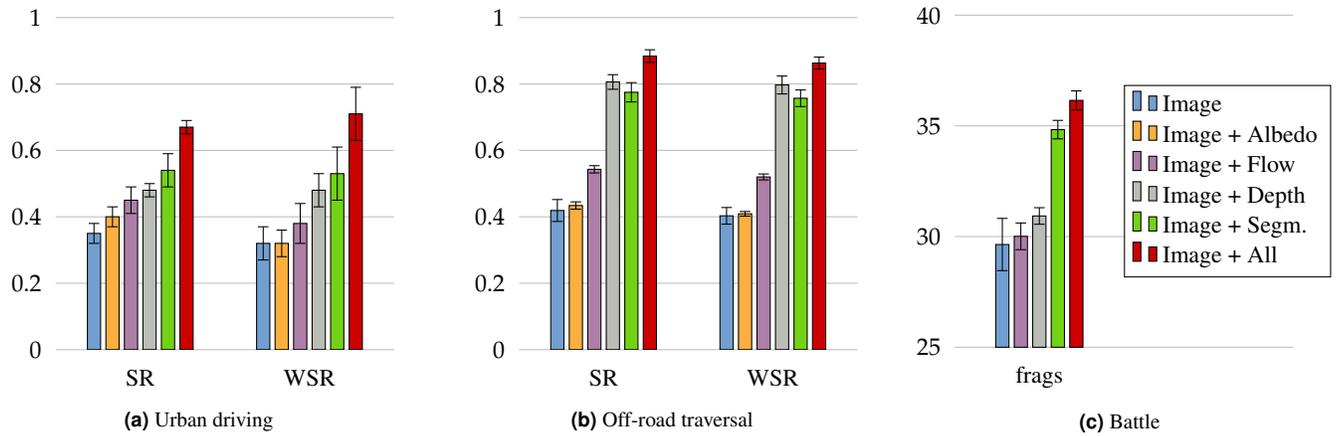


		Image	I. + Albedo	I. + Flow	I. + Depth	I. + Label	I. + All
Urban driving	SR	0.35 ± 0.03	0.40 ± 0.03	0.45 ± 0.04	0.48 ± 0.02	0.54 ± 0.05	0.67 ± 0.02
	WSR	0.32 ± 0.05	0.32 ± 0.04	0.38 ± 0.06	0.48 ± 0.05	0.53 ± 0.08	0.71 ± 0.08
Off-road traversal	SR	0.42 ± 0.03	0.43 ± 0.01	0.54 ± 0.01	0.81 ± 0.02	0.78 ± 0.03	0.89 ± 0.02
	WSR	0.40 ± 0.03	0.41 ± 0.01	0.52 ± 0.01	0.80 ± 0.03	0.76 ± 0.03	0.87 ± 0.02
Battle	Frag	29.6 ± 1.2	-	30.0 ± 0.6	30.9 ± 0.4	34.8 ± 0.4	36.2 ± 0.4

(d) Corresponding numeric results

Fig. S8. Performance of agents equipped with different input representations. The black whiskers show standard deviations across different evaluation runs.

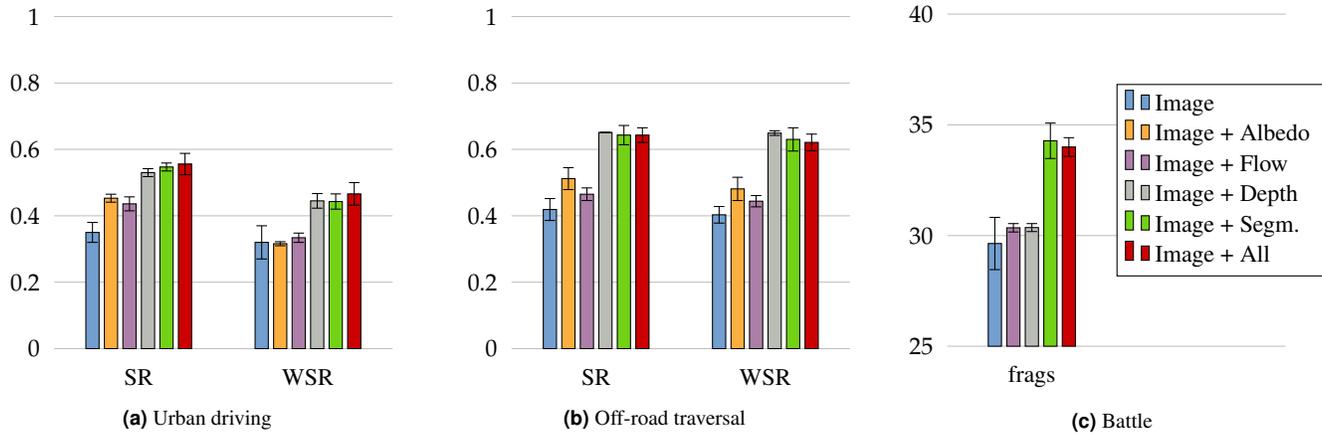


		Image	I. + Albedo	I. + Flow	I. + Depth	I. + Label	I. + All
Urban driving	SR	0.35 ± 0.03	0.45 ± 0.01	0.44 ± 0.02	0.53 ± 0.01	0.55 ± 0.01	0.56 ± 0.03
	WSR	0.32 ± 0.05	0.32 ± 0.01	0.33 ± 0.01	0.45 ± 0.02	0.44 ± 0.02	0.47 ± 0.03
Off-road traversal	SR	0.42 ± 0.03	0.51 ± 0.03	0.47 ± 0.02	0.65 ± 0.00	0.64 ± 0.03	0.64 ± 0.02
	WSR	0.40 ± 0.03	0.48 ± 0.04	0.44 ± 0.02	0.65 ± 0.01	0.63 ± 0.04	0.62 ± 0.03
Battle	Frag	29.6 ± 1.2	-	30.35 ± 0.19	30.36 ± 0.18	34.28 ± 0.80	34.00 ± 0.42

(d) Corresponding numeric results

Fig. S9. Performance of agents equipped with different predicted input representations. The black whiskers show standard deviations across different evaluation runs.

Layer	Output shape
Image	$128 \times 128 \times c$
Conv 5x5 + LReLU	$128 \times 128 \times 192$
Conv 1x1 + LReLU	$128 \times 128 \times 160$
Conv 1x1 + LReLU	$128 \times 128 \times 96$
Max-Pool 3x3	$64 \times 64 \times 96$
Dropout	$64 \times 64 \times 96$
Batch Norm	$64 \times 64 \times 96$
Conv 5x5 + LReLU	$64 \times 64 \times 192$
Conv 1x1 + LReLU	$64 \times 64 \times 192$
Conv 1x1 + LReLU	$64 \times 64 \times 192$
Avg-Pool 3x3	$32 \times 32 \times 192$
Dropout	$32 \times 32 \times 192$
Batch Norm	$32 \times 32 \times 192$
Conv 3x3 + LReLU	$32 \times 32 \times 192$
Conv 1x1 + LReLU	$32 \times 32 \times 192$
Conv 1x1	$32 \times 32 \times 1$
Global Average Pooling	1

Fig. S10. Imitation learning agent architecture. The number of input channels is dependent on which vision modalities the agent has access to. The last layer of the network averages the values over all pixel locations to produce the final action.

Layer	Output shape
Current + Prev RGB	$128 \times 128 \times 6$
Conv 3x3 (Encoder 1)	$128 \times 128 \times 32$
Max-Pool 3x3	$64 \times 64 \times 32$
ConvBlock (Encoder 2)	$64 \times 64 \times 64$
Max-Pool 3x3	$32 \times 32 \times 64$
ConvBlock (Encoder 3)	$32 \times 32 \times 128$
Max-Pool 3x3	$16 \times 16 \times 128$
ConvBlock (Encoder 4)	$16 \times 16 \times 256$
Max-Pool 3x3	$8 \times 8 \times 256$
ConvBlock	$8 \times 8 \times 512$

(a) Encoder architecture.

Layer	Output shape
Encoder Features	$8 \times 8 \times 512$
Conv 3x3	$8 \times 8 \times 256$
Upsample	$16 \times 16 \times 256$
Concat Encoder 4	$16 \times 16 \times 512$
ConvBlock	$16 \times 16 \times 256$
Conv 3x3	$16 \times 16 \times 128$
Upsample	$32 \times 32 \times 128$
Concat Encoder 3	$32 \times 32 \times 256$
ConvBlock	$32 \times 32 \times 128$
Conv 3x3	$32 \times 32 \times 64$
Upsample	$64 \times 64 \times 64$
Concat Encoder 2	$64 \times 64 \times 128$
ConvBlock	$64 \times 64 \times 64$
Conv 3x3	$64 \times 64 \times 32$
Upsample	$128 \times 128 \times 32$
Concat Encoder 1	$128 \times 128 \times 64$
ConvBlock	$128 \times 128 \times 32$
Conv 3x3	$128 \times 128 \times c$

(b) Decoder architecture.

Fig. S11. Network architecture used to infer the vision modalities. A ConvBlock consists of BN-Conv3x3-LReLU-Conv3x3-LRelu-Conv1x1. In the encoder, the ConvBlock output has twice as many channels as the input. In the decoder, the ConvBlock outputs have half as many channels as the input. One decoder network is trained for each predicted modality. All modalities share the encoder. The depicted structure is used in urban driving and off-road traversal. Battle agents use an analogous structure with one less layer in the encoder and the decoder.