

Joint Forecasting of Panoptic Segmentations with Difference Attention

Colin Gruber¹ Cyril Jazra¹ Wenjie Luo² Liangyan Gui¹ Alexander Schwing¹
¹ University of Illinois at Urbana-Champaign ² Waymo

Abstract

Forecasting of a representation is important for safe and effective autonomy. For this, panoptic segmentations have been studied as a compelling representation in recent work. However, recent state-of-the-art on panoptic segmentation forecasting suffers from two issues: first, individual object instances are treated independently of each other; second, individual object instance forecasts are merged in a heuristic manner. To address both issues, we study a new panoptic segmentation forecasting model that jointly forecasts all object instances in a scene using a transformer model based on ‘difference attention.’ It further refines the predictions by taking depth estimates into account. We evaluate the proposed model on the Cityscapes and AIDrive datasets. We find difference attention to be particularly suitable for forecasting because the difference of quantities like locations enables a model to explicitly reason about velocities and acceleration. Because of this, we attain state-of-the-art on panoptic segmentation forecasting metrics.

1. Introduction

Forecasting is needed for safe and effective autonomous systems [9, 26]. For this reason, forecasting has been studied in many different domains from computer vision and robotics to machine learning. In common across domains is the discussion about what representations are useful for forecasting. Representations which have been studied range from trajectories [10, 11, 13, 32, 53] and bounding boxes [29, 30, 39, 50, 51] to semantic segmentation [6, 25, 28, 34, 36], instance segmentation [8, 19, 27], images [12, 24, 52] and recently also panoptic segmentations [14, 38].

Each representation has applications which benefit from their use. We focus on panoptic segmentations as they naturally disentangle 1) objects which change position in an image due to observer motion; from 2) object instances which change position due to both observer and instance motion.

However, the state-of-the-art on panoptic segmentation forecasting [14] is challenged by two key issues. First, foreground predictions of individual instances are made independently of each other. This is suboptimal because the

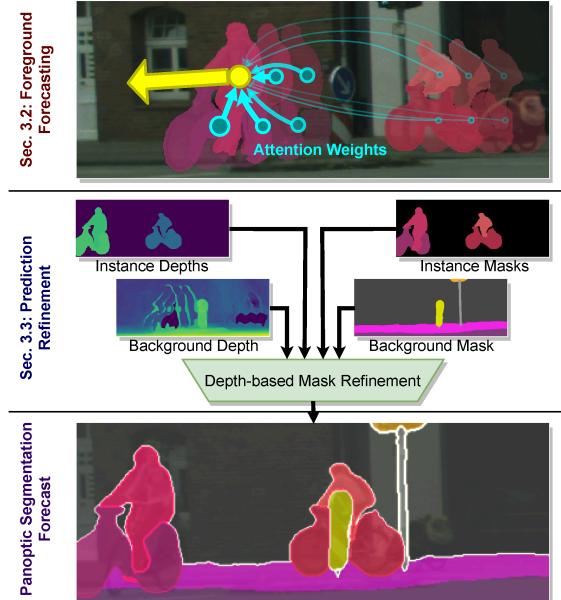


Figure 1. Our panoptic segmentation forecasting. We jointly reason about every instance in a scene to predict instance masks (top), and then reason about the relative depth of foreground and background components (middle) to produce an output (bottom).

movements of instances are clearly correlated, *e.g.*, when considering traffic patterns like the ones in the Cityscapes dataset [7]. Second, the method opted for a simple strategy to merge individual object instance segmentation forecasts with the background forecast. Specifically, in [14], object instance segmentation forecasts are always placed in front of the background segmentation forecast. This assumes that no background objects are located closer to the camera than any foreground entity, which is not true in practice.

In this work, we study a new method to address these two issues: 1) To jointly forecast object instance segmentations, we develop a modified attention module for transformer models. Specifically, instead of the inner-product attention in classical transformers, we propose ‘difference attention.’ This developed difference attention fits tasks like forecasting because it enables reasoning about velocities and acceleration, which is non-trivial with classical inner-product attention (see Fig. 1 top). 2) To properly reason about object and background placement, we develop a refinement head

which denoises background depth estimates and compares them against foreground predictions (see Fig. 1 middle).

We assess our method on the challenging Cityscapes [7] and AIDrive [46] datasets. We find difference attention and refinement to provide accurate results (see Fig. 1 bottom) which yield a new state-of-the-art of 37.6 PQ for mid-term forecasting on Cityscapes and 48.5 PQ on AIDrive. Code to reproduce results is available via <https://github.com/cgraber/psf-diffattn>.

2. Related work

Forecasting has been studied across communities [41].

Forecasting of non-semantic representations. Trajectories are arguably one of the representations for which forecasting has been studied most. Trajectories specify the future position of individual objects, either in 2D or 3D [10, 11, 32, 53]. For example, Hsieh *et al.* [18] disentangle position and pose of multiple moving objects – but only on synthetic data. Mittal *et al.* [33] forecast scene flow for point cloud data using self-supervision to reduce training data requirements. Kosiorek *et al.* [22] track instances to forecast their future. Several works have focused on anticipating future pose and location of specific object types, often people [13, 31]. However, arguably, a trajectory forecast provides little beyond position, velocity and acceleration.

To obtain more information, forecasting of future RGB frames has been studied [12, 24, 52]. Due to the high-dimensional space of the forecasts and because of the ambiguity in the forecasts, results often remain blurry, despite significant recent advances. For instance, recent work models uncertainty over future frames using, *e.g.*, latent variables [44, 52] or treats foreground and background separately [47]. Moreover, Ye *et al.* [52] forecast future RGB frames by modeling each foreground object separately. Note, all these methods differ from ours in architecture and output: we forecast a semantic representation.

Closer to our work is AgentFormer [54]. It also uses transformers to forecast and introduces an identity encoding via agent-aware attention. Our work differs in that we predict panoptic segmentations while they predict birds-eye-view locations. Additionally, we develop difference attention and auxiliary losses which we find to aid forecasting.

Forecasting semantic segmentations. Recently, methods have been studied to estimate semantic segmentations for future, unobserved frames. Luc *et al.* [28] use a deep-net to estimate a future semantic segmentation given the current RGB frame and its semantics as input. Nabavi *et al.* [34] use recurrent models with semantic maps as input. Chiu *et al.* [6] further use a teacher net to provide an additional supervision during training. Šarić *et al.* [36] use learnable deformations to help forecast future semantics given the observed frames. Lin *et al.* [25] design an autoencoder which 1) compresses input feature pyramids into a low-resolution

predictive feature map, 2) predicts this representation for a future frame, and 3) expands it back into a feature pyramid for decoding. However, importantly, these methods do not explicitly consider dynamics of the scene.

While Jin *et al.* [21] jointly predict flow and future semantic segmentations, recent work [37] explicitly warps deep features to obtain a future semantic segmentation. Similarly, Terwilliger *et al.* [40] use a long-short-term-memory (LSTM) module to estimate a flow field which is then used to warp the semantic output of a given input frame. However, by warping in output space, their model has a limited ability to cope with occlusions. While flow improves the modeling of the dynamic world, these methods only consider the dynamics at the pixel-level. Instead, we model dynamics at the object level.

Recent methods [17, 35, 43, 49] estimate future semantic segmentations by reasoning about shape, egomotion, and foreground motion separately. However, none of these methods reason explicitly about individual instances, while our method yields a full future panoptic segmentation forecast, *i.e.*, a prediction for every instance.

Forecasting future instance segmentations. Recent methods which forecast an instance segmentation use a conv net or an LSTM module to regress to the deep features which correspond to the future instance segmentation [19, 27]. For example, Couprie *et al.* [8] use a conv net to forecast future instance contours together with an instance-wise semantic segmentation to estimate future instance segmentation. However, their method only estimates foreground and not background semantics.

Unlike these works, we predict both instance segmentation masks for foreground objects and background semantics for future time steps.

Forecasting panoptic segmentations. In recent years, panoptic segmentation has become a popular scene understanding task [3–5, 23, 45]. Very recently [14, 38], it has been proposed as a useful representation for forecasting because it naturally disentangles 1) objects which move in an image just because of observer motion; from 2) object instances which move due to both observer and instance motion.

The state-of-the-art [14] forecasts the future position of individual object instances independently of each other via an encoder-decoder architecture which is executed separately for every object instance. Moreover, the obtained instance forecasts are combined in a heuristic manner by simply pasting objects in front of background without considering depth information of background objects.

In contrast, we propose a method for panoptic segmentation forecasting which jointly forecasts all detected object instances at once via a tailored transformer attention. This helps to benefit from correlations between instances. Moreover, we study how to combine the individual forecasts in a differentiable way. We discuss our method next.

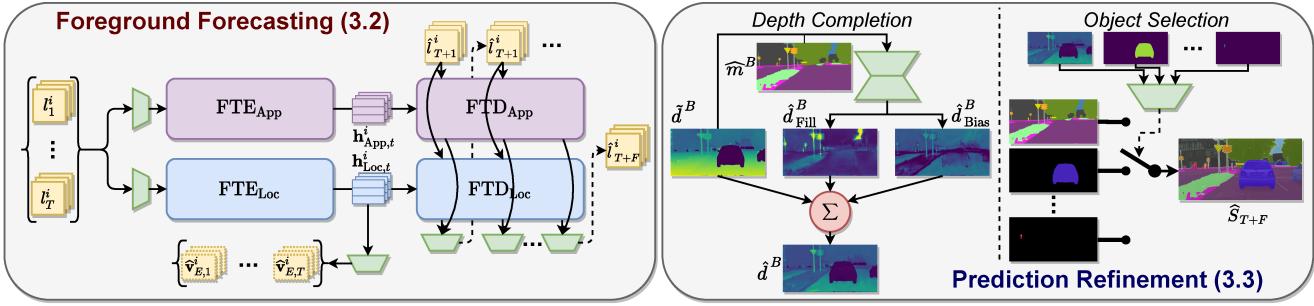


Figure 2. **Method overview.** The foreground forecasting component (Sec. 3.2) predicts future location and appearance jointly for all instances using our newly introduced difference attention transformer. This is followed by a prediction refinement stage (Sec. 3.3), which first completes/denoises input reprojected depths and then uses these along with predicted foreground instance depths to select which object is closest to the camera.

3. Method

In this section, we describe our method for joint forecasting of multiple object instances with the developed difference attention. We start by formalizing the forecasting task and by providing an overview of our approach. Next, we describe the developed difference attention (Sec. 3.1). We use this in our foreground forecasting module, which models interactions between individual instances (Sec. 3.2). Finally, we present the refinement head, which refines the initial foreground instance predictions by considering the background predictions and the depth (Sec. 3.3). An overview of our approach is presented in Fig. 2.

Forecasting task. Given T RGB images I_1, \dots, I_T of height H and width W , panoptic segmentation forecasting aims to predict the panoptic segmentation \hat{S}_{T+F} corresponding to an unobserved future frame I_{T+F} at a fixed number of timesteps F from the last observation recorded at time T . Each pixel in \hat{S}_{T+F} is assigned a class $c \in \{1, \dots, C\}$ and an instance ID. In addition to these inputs, we assume access to camera poses o_1, \dots, o_T and depth maps d_1, \dots, d_T for all input frames. We study the use of both camera poses from odometry sensors, and camera poses estimated using visual SLAM [2]. Following [14], we obtain depth maps from input stereo image pairs [15].

Overview. To address forecasting, we follow the paradigm introduced by Gruber *et al.* [14]. Specifically, they divide the task into two components: 1) the foreground component, which focuses on ‘things’ object instances annotated within the dataset; and 2) the background component, which focuses on all annotated ‘stuff’ object classes. These two components are modeled differently because the causes for the displacement of the corresponding objects in the image plane differ. Specifically, background objects such as buildings and poles shift due to camera motion, while foreground objects like cars and pedestrians move due to both camera motion as well as their own individual motion.

For a fair evaluation, we utilize the same background model as Gruber *et al.* [14], who lift background seman-

tics into a 3D point cloud using the estimated input depth, transform the depth based on the target frame camera information, project to the image plane, and refine the projected semantics using a semantic segmentation model. See Appendix A for more details.

However, the approach developed by Gruber *et al.* [14] has two primary drawbacks which we correct in this work:

First, their approach to forecast the foreground components of the scene uses an RNN-based encoder-decoder model which models the trajectory of each instance independently of all other instances. This is sub-optimal: in many cases, the movement of individual entities is correlated, *e.g.*, due to the flow of traffic. To enable modeling of this correlation, we develop a difference attention module which we detail in Sec. 3.1. It is particularly suitable for forecasting because of its innate ability to reason about the velocities of inputs. We use this difference attention transformer to jointly reason about the future trajectories of all entities in a scene, which we detail in Sec. 3.2.

Second, Gruber *et al.* [14] combine foreground and background predictions by “stacking” all predicted foreground instances on top of the predicted background. This approach assumes that no background objects are located closer to the camera than any foreground entity, which is not true in practice. Hence, in Sec. 3.3, we introduce our model to combine foreground and background predictions in a per-pixel fashion by reasoning about their depths.

3.1. Difference Attention for Transformers

To better address forecasting, we develop a difference attention module for transformers. We find this difference attention to be particularly suitable for forecasting because the difference of quantities like locations enables a model to easily reason about velocities and acceleration. In contrast, classical transformer attention is based on inner products which do not naturally encode these quantities.

Formally, the difference attention module operates on two d -dimensional inputs $\mathbf{X}_{\text{self}} \in \mathbb{R}^{M_1 \times d}$ and $\mathbf{X}_{\text{other}} \in \mathbb{R}^{M_2 \times d}$ of lengths M_1 and M_2 , respectively, reasons about

the differences between these inputs, and outputs representation $\mathbf{Y} \in \mathbb{R}^{M_1 \times d}$ which encodes these differences. For this, we first compute entity scores

$$\mathbf{Z} = \mathbf{Q} \mathbf{K}_R^T - \mathbf{1}_{M_1 \times 1} \text{diag}(\mathbf{K}_B \mathbf{K}_R^T)^T, \quad (1)$$

where $\mathbf{1}_{M_1 \times 1}$ is the $M_1 \times 1$ matrix filled with ones. \mathbf{Q} is computed from \mathbf{X}_{self} and \mathbf{K}_B and \mathbf{K}_R are computed from $\mathbf{X}_{\text{other}}$ with MLPs, *i.e.*, $\mathbf{Q} = f_Q(\mathbf{X}_{\text{self}})$, $\mathbf{K}_R = f_{K_R}(\mathbf{X}_{\text{other}})$, and $\mathbf{K}_B = f_{K_B}(\mathbf{X}_{\text{other}})$. Intuitively, this operation allows the entity score computation to be a function of the difference between the two inputs \mathbf{X}_{self} and $\mathbf{X}_{\text{other}}$. This is useful for forecasting, as the offset of input locations and their change over time is necessary to understand motion.

Given these entity scores \mathbf{Z} , we compute the final attended representation \mathbf{Y} which corresponds to \mathbf{X}_{self} via

$$\mathbf{Y} = \text{softmax}\left(\mathbf{Z}/\sqrt{d}\right) \mathbf{V}_O - \mathbf{V}_S, \quad (2)$$

where $\mathbf{V}_O = f_{V_O}(\mathbf{X}_{\text{other}})$ and $\mathbf{V}_S = f_{V_S}(\mathbf{X}_{\text{self}})$. Intuitively, this enables the final output \mathbf{Y} to encode the differences between the two inputs \mathbf{X}_{self} and $\mathbf{X}_{\text{other}}$. This is again suitable for forecasting, as it enables representations to encode the velocity of an instance, which is critical for reasoning about future motion. We now discuss how we use this difference attention for foreground forecasting.

3.2. Foreground Forecasting

Our forecasting model is tasked with predicting a panoptic segmentation \hat{S}_{T+F} for time $T+F$. This is done by forecasting representations for the N instances in the scene, followed by a final refinement. We represent each instance at all times during forecasting using three components $l_t^i := \{\mathbf{x}_t^i, \mathbf{r}_t^i, p_t^i\}$: a 5-dimensional vector $\mathbf{x}_t^i := [x_0, y_0, x_1, y_1, d]$ representing the upper-left and lower-right corners of the bounding box enclosing instance i as well as the estimated distance of the instance from the camera at time t , a feature tensor $\mathbf{r}_t^i \in \mathbb{R}^{256 \times 14 \times 14}$ representing the visual appearance of the instance at time t , and a binary value $p_t^i \in \{0, 1\}$ which indicates whether instance i is present in frame I_t . Additionally, given background prediction logits \hat{m}^B and background reprojected depths \tilde{d}^B , the final output of the forecasting model is

$$\hat{S}_{T+F} = \text{Ref}(\text{FD}(\text{FE}(\{l_t^i, c^i, o_t\}_{1:T}^{1:N}), \{o_t\}_{T+1:F}), \hat{m}^B, \tilde{d}^B). \quad (3)$$

Here, the forecasting encoder FE operates on input representations l_t^i , classes c^i , and odometry $o_t \forall i \in \{1, \dots, N\}, t \in \{1, \dots, T\}$ and computes embeddings $\mathbf{h}_{\text{Loc},t}^i$ and $\mathbf{h}_{\text{App},t}^i$ which encode locations and appearances, respectively. The forecasting decoder FD processes these embeddings to autoregressively compute embeddings $\tilde{\mathbf{h}}_{\text{Loc},t}^i$ and $\tilde{\mathbf{h}}_{\text{App},t}^i$, which are used to produce outputs \hat{l}_t^i . These outputs are subsequently combined with background semantics

\hat{m}^B and depths \tilde{d}^B using refinement model Ref to produce the final panoptic segmentation \hat{S}_{T+F} . We discuss the encoder and decoder which use difference attention next, and we detail refinement in Sec. 3.3.

Forecasting Transformer Encoder. The encoder FE produces two embeddings for every instance i at every time t : the first, $\mathbf{h}_{\text{Loc},t}^i \in \mathbb{R}^{d_e}$ where d_e is the size of the embedding, contains information about its location as well as its observed motion; the second, $\mathbf{h}_{\text{App},t}^i \in \mathbb{R}^{256 \times 14 \times 14}$, contains information about its appearance. These are obtained using two newly developed forecasting transformer encoders. The use of transformers for this task permits to jointly reason about every instance both as a function of time and as a function of the other instances present in the scene.

The first transformer encoder produces in parallel $\forall i, t$ the location encoding

$$\{\mathbf{h}_{\text{Loc},t}^i\}_{1:T}^{1:N} = \text{FE}_{\text{Loc}}(\{l_t^i, c^i, o_t\}_{1:T}^{1:N}). \quad (4)$$

For this, it uses all input instances at every point in time, *i.e.*, $\{l_t^i\}_{1:T}^{1:N}$, as well as classes c^i and odometry o_t . Different from classical transformer encoders, FE_{Loc} is trained via auxiliary losses to natively reason about both the velocity of each instance across time as well as the motion of each instance relative to each other. Hence, the embedding $\mathbf{h}_{\text{Loc},t}^i$ is trained to encode information about the velocity, which we show improves the ability of the decoder to anticipate the instances' future motion.

The second transformer encoder, which produces the appearance encoding

$$\mathbf{h}_{\text{App},t}^i = \text{FE}_{\text{App}}(\{l_t^i, c^i, o_t\}_{1:T}^{1:N}), \quad (5)$$

maintains the spatial structure of the input appearance features. This is beneficial for predicting a spatial output.

Both the location and the appearance components of the forecasting transformer encoder are comprised of the same general structure: first, a feature representation for every instance is produced as a function of its location, its appearance, its object class, the current camera motion, and the current time. Second, these feature representations are processed using our customized transformer encoders FTE_{Loc} and FTE_{App} . Letting $\beta \in \{\text{Loc}, \text{App}\}$ denote the modules for the location encoder and the appearance encoder, respectively, this is formally described as

$$\{\mathbf{h}_{\beta,t}^i\}_{1:T}^{1:N} = \text{FE}_\beta(\{l_t^i, c^i, o_t\}_{1:T}^{1:N}) \quad (6)$$

$$\Leftrightarrow \begin{cases} \bar{\mathbf{x}}_{\beta,t}^i = f_\beta(l_t^i, c^i, o_t, t) & \forall i, t \\ \{\mathbf{h}_{\beta,t}^i\}_{1:T}^{1:N} = \text{FTE}_\beta(\{\bar{\mathbf{x}}_{\beta,t}^i\}_{1:T}^{1:N}) \end{cases}, \quad (7)$$

where f_{Loc} uses multilayer perceptrons and f_{App} uses convolutional nets which are described fully in Appx. B. Note, depending on β , Eq. (6) refers to either Eq. (4) or Eq. (5). They perform the computations given in Eq. (7). All features $\{\bar{\mathbf{x}}_{\beta,t}^i\}$ are used as input into the transformer FTE_β .

For FTE_{Loc} , all self-attention modules use the difference attention formulation introduced in Sec. 3.1. This design facilitates the ability of the model to reason about the velocity of the entities, which can be represented by differences in input embeddings which correspond to the same instance at different points in time, as well as the relative offsets between different entities. We find that the use of this form of attention leads to improved forecasting results.

The appearance transformer encoder FTE_{App} is built using convolutional transformers. Specifically, it consists of a transformer whose linear projections have been replaced with convolutional layers. This enables a spatially meaningful representation at all stages during encoding.

For more about attention computation see Appendix C.

Forecasting Transformer Decoder. The decoder utilizes the representations produced by the encoder to predict the future location $\hat{\mathbf{x}}_i^t$, the future appearance $\hat{\mathbf{r}}_i^t$, and the future presence \hat{p}_i^t of each object i for future time steps $t \in \{T+1, \dots, T+F\}$. Predictions are computed autoregressively, starting with the most recent input locations $\hat{\mathbf{x}}_i^T := \mathbf{x}_i^T$ and appearance features $\hat{\mathbf{r}}_i^T := \mathbf{r}_i^T$.

For future time step $t \in \{T+1, \dots, T+F\}$, both the location decoder FD_{Loc} and the appearance decoder FD_{App} take the following structure, with $\beta \in \{\text{Loc}, \text{App}\}$:

$$\begin{aligned} \{\tilde{\mathbf{h}}_{\beta,t}^i\}_{1:N}^{1:N} &= \text{FD}_\beta(\{\hat{l}_t^i\}_{T:t-1}^{1:N}, \{c^i\}_{1:N}^{1:N}, \{o_t\}_{T+1:t}, \{\mathbf{h}_{\beta,t}^i\}_{1:T}^{1:N}) \\ &\Leftrightarrow \begin{cases} \hat{\mathbf{x}}_{\beta,t}^i = \tilde{f}_\beta(\hat{l}_t^i, c^i, o_t, t) & \forall i, t \\ \{\tilde{\mathbf{h}}_{\beta,t}^i\}_{T+1:T+F}^{1:N} = \text{FTD}_\beta(\{\hat{\mathbf{x}}_{\beta,t}^i\}_{T+1:T+F}^{1:N}) \end{cases} \end{aligned}$$

Similar to their corresponding encoder modules, the location transformer decoder FTD_{Loc} uses difference attention, the appearance transformer decoder FTD_{App} is a convolutional transformer, and both utilize agent-aware attention.

Final location, appearance, and presence predictions are obtained from the embeddings produced at each time via

$$\hat{\mathbf{x}}_t^i = f_{\text{LocOut}}(\tilde{\mathbf{h}}_{\text{Loc},t}^i) + \hat{\mathbf{x}}_{t-1}^i, \quad (8)$$

$$\hat{p}_t^i = f_{\text{POut}}(\tilde{\mathbf{h}}_{\text{Loc},t}^i), \quad (9)$$

$$\hat{\mathbf{r}}_t^i = f_{\text{AppOut}}(\tilde{\mathbf{h}}_{\text{App},t}^i), \quad (10)$$

where f_{LocOut} and f_{POut} are multilayer perceptrons and f_{AppOut} is a convolutional network.

Training. The foreground model is trained by providing it with input location and appearance features, predicting the future states of each of these, and regressing against pseudo-ground-truth future locations \mathbf{x}_t^{*i} , appearance features \mathbf{r}_t^{*i} , and presences p_t^{*i} which are obtained by running instance detection and tracking on future frames. We formally specify the losses in Appendix D.

In addition, we train the forecasting location encoder to estimate the velocity $\hat{\mathbf{v}}_{\text{E},t}^i$ of each instance via

$$\hat{\mathbf{v}}_{\text{E},t}^i = f_{\text{vel}}(\mathbf{h}_{\text{Loc},t}^i), \quad (11)$$

where f_{vel} is a multilayer perceptron. This auxiliary prediction task requires the encoder to include information about the motion of each instance within the representation it produces. We find this to lead to better forecasting results.

3.3. Prediction Refinement

To address the aforementioned second shortcoming of [14], we develop a refinement which combines foreground and background predictions as a function of their estimated depth. This allows foreground instances to be placed behind background objects, which yields more natural predictions.

While this would be easy if the depth signal was reliable, the only depth signal we have for the background is the depth of the reprojected points that are used as input for the background prediction model. These depths are both noisy and incomplete, *i.e.*, not every location will correspond to a reprojected point from an earlier frame. Hence, the refinement model has two primary jobs: first, it needs to complete as well as denoise the input depth; second, it needs to select which object is closest based on these depths as well as the depths of foreground instances.

Formally, the refinement head is provided with predicted foreground locations¹ $\hat{\mathbf{x}}^i$, appearances $\hat{\mathbf{r}}^i$, and presences \hat{p}^i for N instances. Given these components, if $p^i = 0$, then instance i is discarded, as the model anticipates that the object is not in frame I_{T+F} due to occlusions or leaving the scene; otherwise, the prediction mask \hat{m}^i is obtained via

$$\hat{m}^i = \text{MaskOut}(\hat{\mathbf{x}}^i, \hat{\mathbf{r}}^i), \quad (12)$$

where MaskOut predicts a fixed-size mask using MaskR-CNN’s mask head and then pastes it into the location specified by $\hat{\mathbf{x}}^i$. The prediction head additionally uses estimated instance depths $\{\hat{d}^i\}_{1:N}^{1:N}$, predicted background semantic logits $\hat{m}^B \in \mathbb{R}^{H \times W \times C_{\text{BG}}}$, where C_{BG} is the number of background classes, the reprojected background depths $\tilde{d}^B \in \mathbb{R}^{H \times W}$, and a binary mask $Q \in \{0, 1\}^{H \times W}$ which indicates for each pixel whether or not we have an input background depth. It outputs an object selection map $\hat{P} \in \{0, \dots, N\}^{H \times W}$ which specifies, for every pixel, whether the background is in front (represented by value 0) or one of the instances is in front (represented by values 1 through N). We get the final panoptic segmentation via

$$\hat{S}_{T+F} = \mathbf{1}[\hat{P} = 0] \arg \max(\hat{m}^B) + \mathbf{1}[\hat{P} > 0](\hat{P} + C_{\text{BG}}).$$

The refinement head is composed of two modules: the first produces completed/denoised background depth prediction \hat{d}^B , and the second uses this alongside the foreground instance information to compute the object selection map \hat{P} . We describe both components next.

Depth completion model. We formulate the depth completion model using two outputs. The first, $\hat{d}_{\text{Fill}}^B \in \mathbb{R}^{H \times W}$,

¹For readability, we drop subscript $T+F$ for predictions in this section.

represents an initial estimation of the depths for all input locations which are missing a depth, *i.e.*, where $Q = 0$. The second, $\hat{d}_{\text{Bias}}^B \in \mathbb{R}^{H \times W}$, represents an offset added to the input depths in order to refine and denoise them. Given these predictions, the output of this module is

$$\tilde{d}^B = Q\tilde{d}^B + (1 - Q)\hat{d}_{\text{fill}}^B + \hat{d}_{\text{Bias}}^B, \quad (13)$$

where \hat{d}_{fill}^B and \hat{d}_{Bias}^B are obtained using small convolutional networks specified in Appendix E.

Object selection model. Given the completed/denoised background depth prediction \tilde{d}^B , object selection determines for every output pixel which object is closest to the camera. We require that this module be fully differentiable such that gradients computed from its outputs can be propagated through to the depth completion model.

More formally, we compute the aggregate depth tensor $\mathbf{D} \in \mathbb{R}^{H \times W \times (N+1)}$ whose 0-th channel is the completed background depth \tilde{d}^B and whose i -th channel for $i \in \{1, \dots, N\}$ is $\mathbf{1}[\hat{m}^i \geq 0.5]d^i + \mathbf{1}[\hat{m}^i < 0.5]d_{\text{fgmax}}$. Here, $\mathbf{1}[\cdot]$ is the indicator function applied to all spatial locations in \hat{m}^i . Further, d_{fgmax} is a large constant. We also construct a value tensor $\mathbf{V} \in \mathbb{R}^{H \times W \times (N+1)}$ whose n -th channel is computed by applying a convolutional net to the background logits \hat{m}^B and foreground probabilities \hat{m}^i . The final prediction is computed via

$$\tilde{P}_{i,j} = \text{softmax}(-\mathbf{D}_{i,j}) \circ \mathbf{V}_{i,j}, \quad (14)$$

where $\tilde{P} \in \mathbb{R}^{H \times W \times (N+1)}$ are object selection scores for each pixel, $\tilde{P}_{i,j} = \arg \max \tilde{P}_{i,j}$, and \circ is the Hadamard product. Specifically, for each pixel location (i, j) , we use the softmax function to determine the smallest depth and then multiply by the value vector to attain correct scaling of the output probabilities.

Training. For training, we compute the input instance masks \hat{m}_{T+F}^i using the pseudo-ground-truth locations \mathbf{x}_t^{*i} . We then obtain completed background depths and compute final object selection scores \tilde{P}_{T+F} . This is compared to the ground-truth object selection P_{T+F}^* using cross-entropy. We additionally apply a squared norm loss to the predicted depth bias \hat{d}_{Bias}^B such that the model is encouraged to trust the input depths where possible. Note that we do not supervise the depth completion model to predict globally accurate depths. Instead, we only require that the completed background depths have the correct relative value compared to the foreground depths, *i.e.*, the predicted depths lead to selecting either the foreground or the background correctly per-pixel.

4. Experiments

We demonstrate that the proposed difference attention and refinement lead to a new state-of-the-art for panoptic

segmentation forecasting. We additionally show the contribution each component makes to the final improvement via ablations. In addition, we demonstrate how these improvements carry over to related dense forecasting tasks. Following prior work [14], we test our forecasting model on the Cityscapes dataset [7]. We additionally run experiments on the recently-introduced AIDDrive dataset [46].

4.1. Cityscapes

Data. The Cityscapes dataset contains 5,000 sequences of 30 frames each, where ground-truth panoptic segmentations are provided for the 20th frame of each sequence. Here, we evaluate our forecasting model on panoptic segmentation forecasting. Additional results for instance segmentation and semantic segmentation forecasting can be found in Appendix H and Appendix I. We consider two types of forecasting: short-term and mid-term forecasting, each looking 3 and 9 frames into the future respectively. In both cases, we take every third frame as input to our model, hence matching the methods used in prior work [14, 27, 28, 37].

Metrics. Following prior work [14], we consider three metrics: *segmentation quality* (SQ), *recognition quality* (RQ), and *panoptic quality* (PQ). First, we match predicted and target segments, where true positive matches require the intersection over union (IoU) of the two segments to be at least 0.5. SQ corresponds to the average IoU of true matched positive segments. RQ corresponds to the F1 score computed over matches. Finally, PQ is the product of SQ and RQ. These metrics are computed for each individual class and then averaged over all classes.

Baselines. We compare against the baselines introduced in [14]. *Panoptic Deeplab (Oracle)* applies the Panoptic Deeplab model [3] on the target frame, and represents an upper bound on performance due to its access to oracle future information. *Panoptic Deeplab (Last Seen Frame)* applies this model to the most recently observed frame, which represents a model assuming no camera or instance motion. *Flow* computes optical flow [20] from the last two observed frames and then uses it to warp the panoptic segmentation obtained from the last observed frame. *Hybrid Semantic/Instance Forecasting* fuses a semantic segmentation forecast [40] with an instance segmentation forecast [27] to create a panoptic segmentation for the target frame. Finally, *IndRNN-Stack* is the model introduced by Graber *et al.* [14] which forecasts individual instances using an RNN encoder-decoder model and stacks all foreground components on top of all background components.

Results. The results for all models on the panoptic segmentation forecasting task are presented in Tab. 1. The proposed approach achieves state-of-the-art across both short- and mid-term settings on all metrics when compared to methods which don't access future information.

Fig. 3 presents a visual comparison. IndRNN-Stack is

	Short term: $\Delta t = 3$												Mid term: $\Delta t = 9$																							
	All			Things			Stuff			All			Things			Stuff			All			Things			Stuff											
	PQ	SQ	RQ	PQ	SQ	RQ	PQ	SQ	RQ	PQ	SQ	RQ	PQ	SQ	RQ	PQ	SQ	RQ	PQ	SQ	RQ	PQ	SQ	RQ												
	Panoptic Deeplab (Oracle)†	60.3	81.5	72.9	51.1	80.5	63.5	67.0	82.3	79.7	60.3	81.5	72.9	51.1	80.5	63.5	67.0	82.3	79.7	60.3	81.5	72.9	51.1	80.5	63.5	67.0	82.3	79.7								
Panoptic Deeplab (Last seen frame)	32.7	71.3	42.7	22.1	68.4	30.8	40.4	73.3	51.4	22.4	68.5	30.4	10.7	65.1	16.0	31.0	71.0	40.9	41.4	73.4	53.4	30.6	42.0	49.3	25.9	69.5	34.6	13.4	67.1	19.3	35.0	71.3	45.7			
Flow	43.2	74.1	55.1	35.9	72.4	48.3	48.5	75.3	60.1	29.7	69.1	39.4	19.7	66.8	28.0	37.0	70.8	47.7	49.0	74.9	63.3	40.1	72.5	54.6	55.5	76.7	69.5	36.3	71.3	43.9	72.9	56.2				
Hybrid [40] (bg) and [27] (fg)	50.2	75.7	64.3	42.4	74.2	56.5	55.9	76.8	70.0	37.6	71.4	49.5	28.6	69.0	40.1	44.1	73.2	56.4	43.2	74.1	55.1	35.9	72.4	48.3	48.5	75.3	60.1	29.7	69.1	39.4	19.7	66.8	28.0	37.0	70.8	47.7
Ours	50.2	75.7	64.3	42.4	74.2	56.5	55.9	76.8	70.0	37.6	71.4	49.5	28.6	69.0	40.1	44.1	73.2	56.4	43.2	74.1	55.1	35.9	72.4	48.3	48.5	75.3	60.1	29.7	69.1	39.4	19.7	66.8	28.0	37.0	70.8	47.7

Table 1. **Panoptic segmentation forecasting evaluated on the Cityscapes validation set.** † has access to the RGB frame at time $T + F$. Higher is better for all metrics.

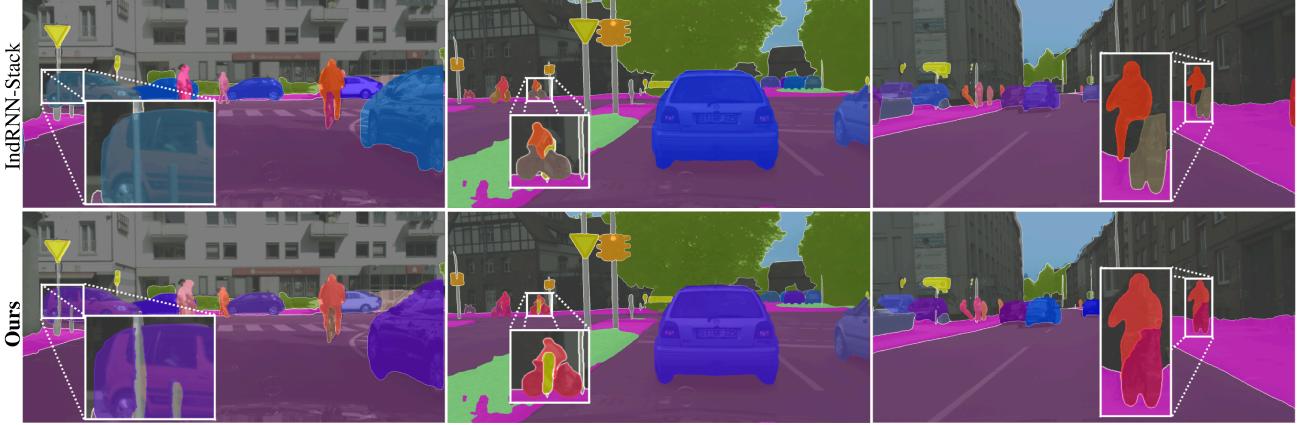


Figure 3. Mid-term panoptic segmentation forecasts on Cityscapes. Unlike IndRNN-Stack, our approach is able to properly place foreground instances behind background objects (left two columns). Additionally, our approach models interactions between objects, leading to additional improvements (right column).

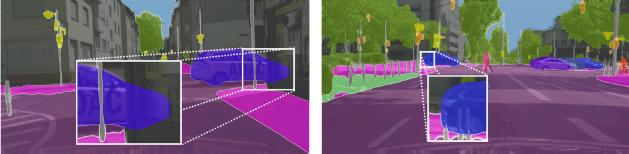


Figure 4. **Failure cases.** Left: the car is incorrectly predicted to be in front of the building on the right. Right: the car is incorrectly predicted to be in front of a few poles.

not capable of placing foreground instances behind background objects, which leads to missing segmentations such as poles in the left column and the street sign in the middle column. Our approach properly reasons about the depth of these objects and places the poles in front of the car and the street sign in front of the cyclist. Additionally, since IndRNN-Stack predicts instances independently, it can make trivial errors such as predicting a cyclist floating away from their bicycle (right column). Our approach, which models interactions among instances and can reason about the fact that cyclists should always move with their bicycles, does not make this error. Additional visualizations comparing these models are presented in Appendix J.

Limitations. Fig. 4 presents a few sequences where our model mispredicts the relative location of foreground and background components. The noisiness of the input point clouds can introduce error in depth reasoning, especially for far away objects which have similar depth. The fact that we only use one depth value for a foreground instance can in-

	$\Delta t = 3$			$\Delta t = 9$		
	PQ	SQ	RQ	PQ	SQ	RQ
Ours	50.2	75.7	64.3	37.6	71.4	49.5
1) w/o difference attention	49.9	75.6	64.0	36.8	71.6	48.3
2) w/o auxiliary encoder loss	49.1	75.3	63.0	36.5	71.5	47.9
3) w/o refinement	49.9	75.6	63.9	36.4	71.0	48.0
4) w/ ORB-SLAM odometry	49.6	75.7	63.5	37.2	71.5	49.0
w/ ground truth future odometry	50.5	75.8	64.7	39.7	72.0	52.1

Table 2. **Validating our design choices** using Cityscapes. Higher is better for all metrics. All approaches use predicted future odometry unless otherwise specified.

introduce errors for larger objects. Similar to IndRNN-Stack, the method struggles with instance detection and tracking errors as we assume these inputs to be correct.

Ablations. Tab. 2 summarizes results studying the impact of modeling decisions. 1) *w/o difference attention* uses standard dot product attention for all transformers in place of the difference attention module developed in Sec. 3.1. Our full model’s superior performance over 1) demonstrates that the difference attention model is able to better reason about instance motion. 2) *w/o auxiliary encoder loss* trains the forecasting model without applying a loss to the velocity output from Eq. (11). This leads to worse results, and shows that the auxiliary loss helps bias the encoder representations to encode motion information useful for forecasting. 3) *w/o refinement* does not use the refinement head, and instead stacks foreground predictions on top of background predictions, following Graber *et al.* [14]. This leads to missed

	PQID	All SQID	RQID	PQID	Things SQID	RQID	PQID	Stuff SQID	RQID	PQ	All SQ	RQ	PQ	Things SQ	RQ
Pan. Deeplab †	64.0	84.9	74.4	60.9	79.5	76.3	64.6	85.9	74.0	64.0	84.9	74.4	60.9	79.5	76.3
Pan. Deeplab*	37.5	75.8	47.5	16.1	70.8	21.5	41.4	76.7	52.2	37.7	75.5	47.8	17.5	68.8	23.8
Flow	40.8	75.7	51.6	19.1	69.9	25.7	44.7	76.7	56.3	40.9	75.4	51.7	19.8	68.5	26.9
IndRNN-Stack	45.1	73.9	57.4	24.1	68.7	32.8	49.0	74.9	61.9	45.3	73.8	57.7	25.1	68.1	34.5
Ours	45.1	73.9	57.2	25.8	69.9	34.3	48.6	74.7	61.4	45.2	73.8	57.5	26.6	69.0	35.8
IndRNN-Stack‡	48.3	75.1	61.1	24.7	68.9	33.3	52.6	76.2	66.1	48.5	75.0	61.3	25.8	68.3	35.1
Ours‡	48.5	75.1	61.3	26.2	69.3	34.9	52.6	76.1	66.1	48.7	75.0	61.5	26.9	68.8	36.2

Table 3. Metrics computed on the AIODrive dataset (0.5 second forecast). †: Oracle; *: Applied to last seen frame; ‡: use ground-truth depth inputs. IndRNN-Stack and our models use ground-truth odometry input.

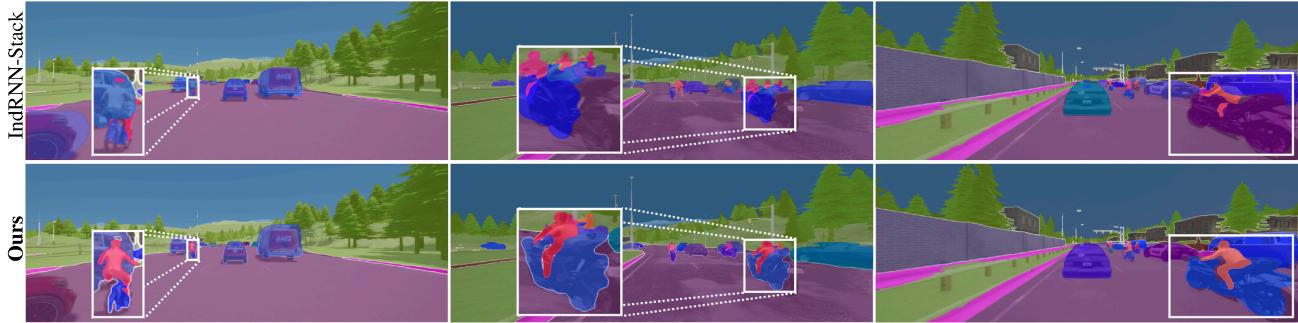


Figure 5. Panoptic segmentation forecasts on the AIODrive dataset. Our approach better captures the relationships between cyclists/their vehicles than IndRNN-Stack, which cannot model these relationships. Also, our model produces higher-fidelity pedestrian instance masks.

background objects which are occluded by foreground predictions, hence a drop in results. 4) w/ *ORB-SLAM odometry* uses input odometry obtained from [2], and shows that our method also works with odometry obtained from image data. The final ablation demonstrates that access to more accurate future camera motion leads to improvements.

4.2. AIODrive

Data. The AIODrive dataset [46] contains a large number of synthetically generated traffic scenarios and provides many inputs and annotations, including stereo images, LiDAR, ground-truth depth maps, panoptic segmentations, and more. The use of a simulator to obtain data and annotations results in AIODrive containing panoptic segmentation annotations, including instance tracks, for all frames. Here, we use the subset of the labels corresponding to Cityscapes classes, consisting of 2 “things” and 11 “stuff” classes. We use 5 frames of input and forecast the 5th frame into the future (corresponding to a 0.5s forecast). Additional details can be found in Appendix G.

Metrics. In addition to previously used metrics, we introduce metrics which account for object identity. Specifically, we evaluate using PQID, SQID, and RQID, which require matches computed between predicted and ground-truth objects to have the same instance ID. These metrics are more appropriate for the forecasting setting due to the fact that the previously used metrics can mark matches between different instances as true positives, meaning the motion of an instance was incorrectly predicted but the metric did not properly evaluate this. Note that we cannot compute these metrics on Cityscapes, as that data only contains annota-

tions for a single frame per sequence.

Results. The results for all models on the panoptic segmentation forecasting task on AIODrive are presented in Tab. 3. Because the All PQID is averaged over 2 “things” classes and 11 “stuff” classes, this metric is biased towards “stuff” performance. Hence, All PQID is comparable between IndRNN-Stack and our model. However, the differences are much clearer on the “things” metrics, as our approach is better able to reason about the motion of individual object instances. Furthermore, there is a small drop in performance between PQID and PQ, indicating that some of the true positive matches found when computing PQ are between incorrect ground-truth instances. Fig. 5 shows results for our method and IndRNN-Stack on AIODrive. Our approach produces better forecasts for cyclists and their bikes, due to the use of difference attention.

5. Conclusion

We introduce a new model for panoptic segmentation forecasting. It uses difference attention which we find to be more suitable to forecasting than standard attention as it can reason about velocities and acceleration. A new refinement head also merges predictions based on depth. This improves prior work on all panoptic forecasting metrics.

Acknowledgements: This work is supported in part by NSF #1718221, 2008387, 2045586, 2106825, MRI #1725729, NIFA 2020-67021-32799 and Cisco Systems Inc. (CG 1377144 - thanks for access to Arcetri).

References

- [1] Apratim Bhattacharyya, Mario Fritz, and Bernt Schiele. Bayesian prediction of future street scenes using synthetic likelihoods. *ICLR*, 2019. 7
- [2] Carlos Campos, Richard Elvira, Juan J. Gómez, José M. M. Montiel, and Juan D. Tardós. ORB-SLAM3: An accurate open-source library for visual, visual-inertial and multi-map SLAM. *arXiv preprint arXiv:2007.11898*, 2020. 3, 8, 5
- [3] Bowen Cheng, Maxwell D. Collins, Yukun Zhu, Ting Liu, Thomas S. Huang, Hartwig Adam, and Liang-Chieh Chen. Panoptic-DeepLab: A simple, strong, and fast baseline for bottom-up panoptic segmentation. In *CVPR*, 2020. 2, 6
- [4] Bowen Cheng, Ishan Misra, Alexander G. Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. *arXiv*, 2021. 2
- [5] Bowen Cheng, Alexander G. Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. 2021. 2
- [6] Hsu-Kuang Chiu, Ehsan Adeli, and Juan Carlos Niebles. Segmenting the future. *IEEE Robotics and Automation Letters*, 2020. 1, 2
- [7] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016. 1, 2, 6
- [8] Camille Couprie, Pauline Luc, and Jakob Verbeek. Joint future semantic and instance segmentation prediction. In *ECCV workshops*, 2018. 1, 2
- [9] Kenneth James Williams Craik. *The Nature of Explanation*. Cambridge University Press, 1943. 1
- [10] Qi Dai, Vaishakh Patil, Simon Hecker, Dengxin Dai, Luc Van Gool, and Konrad Schindler. Self-supervised object motion and depth estimation from video. In *CVPR Workshops*, 2020. 1, 2
- [11] Sébastien Ehrhardt, Oliver Groth, Aron Monszpart, Martin Engelcke, Ingmar Posner, Niloy Mitra, and Andrea Vedaldi. RELATE: Physically plausible multi-object scene synthesis using structured latent spaces. In *NeurIPS*, 2020. 1, 2
- [12] Hang Gao, Huazhe Xu, Qi-Zhi Cai, Ruth Wang, Fisher Yu, and Trevor Darrell. Disentangling propagation and generation for video prediction. In *ICCV*, 2019. 1, 2
- [13] Colin Gruber and Alexander Schwing. Dynamic neural relational inference. In *CVPR*, 2020. 1, 2
- [14] Colin Gruber, Grace Tsai, Michael Firman, Gabriel Brostow, and Alexander G. Schwing. Panoptic segmentation forecasting. In *CVPR*, 2021. 1, 2, 3, 5, 6, 7
- [15] Xiaodong Gu, Zhiwen Fan, Siyu Zhu, Zuozhuo Dai, Feitong Tan, and Ping Tan. Cascade cost volume for high-resolution multi-view stereo and stereo matching. In *CVPR*, 2020. 3, 6
- [16] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *ICCV*, 2017. 6
- [17] Lukas Hoyer, Patrick Kesper, Anna Khoreva, and Volker Fischer. Short-term prediction and multi-camera fusion on semantic grids. In *ICCV Workshops*, 2019. 2
- [18] Jun-Ting Hsieh, Bingbin Liu, De-An Huang, Li Fei-Fei, and Juan Carlos Niebles. Learning to decompose and disentangle representations for video prediction. In *NeurIPS*, 2018. 2
- [19] Jian-Fang Hu, Jiangxin Sun, Zihang Lin, Jian-Huang Lai, Wenjun Zeng, and Wei-Shi Zheng. Apanet: Auto-path aggregation for future instance segmentation prediction. *IEEE TPAMI*, 2021. 1, 2
- [20] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margaret Keuper, Alexey Dosovitskiy, and Thomas Brox. Flownet 2.0: Evolution of optical flow estimation with deep networks. In *CVPR*, 2017. 6
- [21] Xiaojie Jin, Huaxin Xiao, Xiaohui Shen, Jimei Yang, Zhe Lin, Yunpeng Chen, Zequn Jie, Jiashi Feng, and Shuicheng Yan. Predicting scene parsing and motion dynamics in the future. In *NeurIPS*, 2017. 2
- [22] Adam Kosiorek, Hyunjik Kim, Yee Whye Teh, and Ingmar Posner. Sequential attend, infer, repeat: Generative modelling of moving objects. In *NeurIPS*, 2018. 2
- [23] Yanwei Li, Hengshuang Zhao, Xiaojuan Qi, Liwei Wang, Zeming Li, Jian Sun, and Jiaya Jia. Fully convolutional networks for panoptic segmentation. In *CVPR*, 2021. 2
- [24] Xiaodan Liang, Lisa Lee, Wei Dai, and Eric P Xing. Dual motion GAN for future-flow embedded video prediction. In *ICCV*, 2017. 1, 2
- [25] Zihang Lin, Jiangxin Sun, Jian-Fang Hu, Qizhi Yu, Jian-Huang Lai, and Wei-Shi Zheng. Predictive feature learning for future segmentation prediction. In *ICCV*, 2021. 1, 2, 6, 7
- [26] Rodolfo R. Llinás. *I of the vortex: from neurons to self*. MIT Press, 2001. 1
- [27] Pauline Luc, Camille Couprie, Yann LeCun, and Jakob Verbeek. Predicting future instance segmentation by forecasting convolutional features. In *ECCV*, 2018. 1, 2, 6, 7, 8, 9, 12
- [28] Pauline Luc, Natalia Neverova, Camille Couprie, Jakob Verbeek, and Yann LeCun. Predicting deeper into the future of semantic segmentation. In *ICCV*, 2017. 1, 2, 6
- [29] Osama Makansi, Özgün Cicek, Yassine Marrakchi, and Thomas Brox. On Exposing the Challenging Long Tail in Future Prediction of Traffic Actors. In *ICCV*, 2021. 1
- [30] Srikanth Malla, Behzad Dariush, and Chiho Choi. Titan: Future forecast using action priors. In *CVPR*, 2020. 1
- [31] Karttikeya Mangalam, Ehsan Adeli, Kuan-Hui Lee, Adrien Gaidon, and Juan Carlos Niebles. Disentangling human dynamics for pedestrian locomotion forecasting with noisy supervision. In *WACV*, 2020. 2
- [32] Julieta Martinez, Michael J Black, and Javier Romero. On human motion prediction using recurrent neural networks. In *CVPR*, 2017. 1, 2
- [33] Himangi Mittal, Brian Okorn, and David Held. Just go with the flow: Self-supervised scene flow estimation. In *CVPR*, 2020. 2
- [34] Seyed Shahabeddin Nabavi, Mrigank Rochan, and Yang Wang. Future semantic segmentation with convolutional LSTM. In *BMVC*, 2018. 1, 2
- [35] Xiaojuan Qi, Zhengzhe Liu, Qifeng Chen, and Jiaya Jia. 3D motion decomposition for RGBD future dynamic scene synthesis. In *CVPR*, 2019. 2
- [36] Josip Šarić, Marin Oršić, Tonči Antunović, Sacha Vražić, and Siniša Šegvić. Single level feature-to-feature forecasting with deformable convolutions. In *German Conference on Pattern Recognition*, 2019. 1, 2, 7

- [37] Josip Saric, Marin Orsic, Tonci Antunovic, Sacha Vrazic, and Sinisa Segvic. Warp to the future: Joint forecasting of features and feature motion. In *CVPR*, 2020. [2](#), [6](#), [7](#), [8](#)
- [38] Josip Šarić, Sacha Vražić, and Siniša Šegvić. Dense semantic forecasting in video by joint regression of features and feature motion. *IEEE TNNLS*, 2021. [1](#), [2](#)
- [39] Oliver Styles, Victor Sanchez, and Tanaya Guha. Multiple object forecasting: Predicting future object locations in diverse environments. In *WACV*, 2020. [1](#)
- [40] Adam Terwilliger, Garrick Brazil, and Xiaoming Liu. Recurrent flow-guided semantic forecasting. In *WACV*, 2019. [2](#), [6](#), [7](#), [8](#), [9](#), [12](#)
- [41] Eugene Valassakis. Future object segmentation for complex correlated motions. Master's thesis, UCL, 2018. [2](#)
- [42] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan. N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, 2017. [2](#), [3](#)
- [43] Suhani Vora, Reza Mahjourian, Soeren Pirk, and Anelia Angelova. Future semantic segmentation using 3D structure. *arXiv:1811.11358*, 2018. [2](#)
- [44] Jacob Walker, Carl Doersch, Abhinav Gupta, and Martial Hebert. An uncertain future: Forecasting from static images using variational autoencoders. In *ECCV*, 2016. [2](#)
- [45] Huiyu Wang, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Max-deeplab: End-to-end panoptic segmentation with mask transformers. In *CVPR*, 2021. [2](#)
- [46] Xinshuo Weng, Yunze Man, Jinhyung Park, Ye Yuan, Dazhi Cheng, Matthew O'Toole, and Kris Kitani. All-In-One Drive: A Large-Scale Comprehensive Perception Dataset with High-Density Long-Range Point Clouds. *arXiv*, 2021. [2](#), [6](#), [8](#)
- [47] Yue Wu, Rongrong Gao, Jaesik Park, and Qifeng Chen. Future video synthesis with object motion prediction. In *CVPR*, 2020. [2](#)
- [48] Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang, Yanyan Lan, Liwei Wang, and Tieyan Liu. On layer normalization in the transformer architecture. In *ICML*, 2020. [2](#)
- [49] Jingwei Xu, Bingbing Ni, Zefan Li, Shuo Cheng, and Xiaokang Yang. Structure preserving video prediction. In *CVPR*, 2018. [2](#)
- [50] Takuma Yagi, Karttikeya Mangalam, Ryo Yonetani, and Yoichi Sato. Future person localization in first-person videos. In *CVPR*, 2018. [1](#)
- [51] Yu Yao, Mingze Xu, Chiho Choi, David J Crandall, Ella M Atkins, and Behzad Dariush. Egocentric vision-based future vehicle localization for intelligent driving assistance systems. In *ICRA*, 2019. [1](#)
- [52] Yufei Ye, Maneesh Singh, Abhinav Gupta, and Shubham Tulsiani. Compositional video prediction. In *ICCV*, 2019. [1](#), [2](#)
- [53] Raymond A. Yeh, Alexander G. Schwing, Jonathan Huang, and Kevin Murphy. Diverse Generation for Multi-agent Sports Games. In *CVPR*, 2019. [1](#), [2](#)
- [54] Ye Yuan, Xinshuo Weng, Yanglan Ou, and Kris Kitani. Agentformer: Agent-aware transformers for socio-temporal multi-agent forecasting. In *ICCV*, 2021. [2](#), [3](#)

Supplementary Material: Joint Forecasting of Panoptic Segmentations with Difference Attention

This appendix is structured as follows: Appendix A details the background prediction approach which we use to obtain preliminary background class predictions. Appendix B provides specific model architectural details for the forecasting transformer encoder and decoder. Appendix C explains in detail the agent-aware attention approach we use which allows for identity information to be encoded in the model. Appendix D describes the specific losses computed during training of the foreground forecasting model. Appendix E presents the model architecture used by the depth completion model introduced in Sec. 3.3. Appendix F describes additional details of implementation and model training. Appendix G contains additional information about the AIDrive dataset and experiments. Appendix H contains instance segmentation forecasting experimental results for Cityscapes. Appendix I contains semantic segmentation forecasting experimental results for Cityscapes. Appendix J presents additional model visualizations on Cityscapes for both the short- and mid-term settings. Appendix K contains the full per-class breakdown of the panoptic segmentation metrics presented in Tab. 1. Appendix L describes the major code libraries used to implement our model. Finally, Appendix M discusses potential negative societal impacts that could arise from the implementation of this work in practice.

A. Background Model

In this work, we utilize the background semantic prediction model introduced by Gruber *et al.* [14]. This approach lifts background semantics into a 3D point cloud using the estimated input depth, transforms the point cloud based on camera movement, projects to the image plane, and refines the projected semantics using a semantic segmentation model. Formally, this model estimates the semantics of background object classes for unseen future frame $T + F$ as

$$\hat{m}_{T+F}^B = \text{BGRef}(\{\text{proj}(m_t, d_t, K, H_t, u_t)\}_{1:T}), \quad (15)$$

where K represents camera intrinsic parameters, H_t is the 6-dof camera transform from input frame t to target frame $T + F$, m_t is the semantic segmentation for frame t which is obtained from a pre-trained model, d_t is the input depth map at time t , and u_t denotes the coordinates of all of the pixels in m_t which correspond to background semantic classes. Proj refers to the step which creates the sparse reprojected semantic map for frame $T + F$ given inputs for frame t , and BGRef refers to the background refinement model which produces a complete background prediction from the output of Proj .

The first step of the background model is to produce reprojected semantic point clouds $(\tilde{m}_t^B, \tilde{d}_t^B)$ which are processed by BGRef . These are obtained for each time $t \in \{1, \dots, T\}$ by applying Proj to the corresponding input frame I_t . Given per-pixel semantic prediction m_t and depth map d_t , Proj back-projects, transforms, and reprojects the pixels from input frame t to target frame $T + F$. This process is summarized as

$$\begin{bmatrix} x_t \\ y_t \\ z_t \end{bmatrix} = H_t \begin{bmatrix} K^{-1} \begin{bmatrix} u_t \\ 1 \\ 1 \end{bmatrix} \text{diag}(d_t) \end{bmatrix}, \quad (16)$$

$$\begin{bmatrix} u_{T+F} \\ 1 \end{bmatrix} = K \begin{bmatrix} x_t/z_t \\ y_t/z_t \\ 1 \end{bmatrix}, \quad (17)$$

$$\tilde{m}_t^B(u_{T+F}) = m_t^B(u_t), \quad (18)$$

$$\tilde{d}_t^B(u_{T+F}) = z_t, \quad (19)$$

where u_t is a vector whose entries dictate the pixel locations in m_t which correspond to background object classes and u_{T+F} is the vector which contains the location of these pixels in the target frame at time $T + F$. During this, we maintain the semantic class obtained from m_t and the projected depth of each pixel location. Whenever multiple pixels u_t from an input frame are projected to the same pixel u_{T+F} in the target frame, the depth and the semantic label of the pixel with the smallest depth is kept, as it is closest to the camera.

Given reprojected semantics \tilde{m}_t^B and depths \tilde{d}_t^B from the previous step, the background refinement model is tasked with predicting a final semantic output. This is done by concatenating the input from all frames and feeding them into a semantic

segmentation model, which can be described as

$$\begin{aligned}\hat{m}_{T+F}^{\text{Prob}} &= \text{BGRef}([\{\tilde{m}_t^B(u_{T+F})\tilde{d}_t^B(u_{T+F})\}_{1:T}]) \\ \hat{m}_{T+F}^B &= \arg \max_c(\hat{m}_{T+F}^{\text{Prob}}),\end{aligned}$$

where $\hat{m}_{T+F}^{\text{Prob}} \in \Delta_{C_{\text{BG}}}^{H \times W}$ represents the C_{BG} -dimensional output probability map per pixel, one for each background class, and the final output \hat{m}_{T+F}^B is obtained per-pixel by choosing the class with the largest probability.

The refinement network is trained using the cross-entropy loss

$$\mathcal{L}_{\text{bf}} := \frac{1}{\sum_{x,y} \mathbf{1}_{T+F}^{\text{bg}}[x,y]} \sum_{x,y} \mathbf{1}_{T+F}^{\text{bg}}[x,y] \sum_c m_{T+F}^{B*}(x,y,c) \log(\hat{m}_{T+F}^{\text{Prob}}(x,y)). \quad (20)$$

Here, $\mathbf{1}_{T+F}^{\text{bg}}[x,y]$ is an indicator function specifying whether pixel coordinates (x,y) correspond to background semantic classes for frame $T+F$, and $m_{T+F}^{i*}(x,y,c) = 1$ if the correct class for pixel (x,y) is c and 0 otherwise. For all experiments presented in this work, we use the specific background prediction model trained by Graber *et al.* [14]. Further implementation details related to model architecture and training can be found in the Appendix of [14].

B. Architecture details for Forecasting Transformer Encoder and Decoder

The feature model f_{Loc} processes input locations \mathbf{x}_t^i , appearances \mathbf{r}_t^i , instance classes c^i , odometry o_t , and time t to produce an embedding $\bar{\mathbf{x}}_{\text{Loc},t}^i$ which is processed by the transformer FTE. f_{Loc} can be fully specified by the following model components:

$$\mathbf{x}'^i_t = f_b([\mathbf{x}_t^i, \text{onehot}(c^i)]), \quad (21)$$

$$\mathbf{r}'^i_t = \text{AvgPool}(f_f(\mathbf{r}_t^i)), \quad (22)$$

$$\bar{\mathbf{x}}_{\text{Loc},t}^i = f_{e2}([f_{e1}([\mathbf{x}'^i_t, \mathbf{r}'^i_t, o_t]), \tau_t]). \quad (23)$$

First, an initial location embedding \mathbf{x}'^i_t is produced, where f_b is a linear layer and *onehot* represents a vector whose c^i -th element is set to one and whose other entries are set to zero. Similarly, initial appearance embedding \mathbf{r}'^i_t is produced, where f_f is a small convolutional network and *AvgPool* averages the result over the spatial dimensions. These two embeddings are concatenated with odometry o_t , passed through linear layer f_{e1} , concatenated with temporal encoding τ_t , and passed through the final linear layer f_{e2} . Specifically, the temporal encoding $\tau^t \in \mathbb{R}^{d_\tau}$ provides information to the model about the temporal location of the given instance in the sequence and whose k -th element is defined as

$$\tau^t(k) = \begin{cases} \sin(t/1000^{k/d_\tau}), & k \text{ is even} \\ \cos(t/1000^{(k-1)/d_\tau}), & k \text{ is odd} \end{cases}. \quad (24)$$

Here d_τ is the size of the temporal encoding and is set to 256 everywhere in this work. All linear layers in f_{Loc} have an output embedding size of 256, and f_f contains two 2D convolutional layers with a kernel size of 3, output channel size of 256, and ReLU activations after each.

The feature model f_{App} produces appearance embedding $\bar{\mathbf{x}}_{\text{App},t}^i$ as a function of the input appearances \mathbf{r}_t^i as well as input time t , and can be fully specified by the following model components:

$$\bar{\mathbf{x}}_{\text{App},t}^i = f_{ae2}([f_{ae1}(\mathbf{r}_t^i), \tilde{\tau}_t]), \quad (25)$$

where f_{ae1} is a 3×3 convolutional layer with output dimension 256, f_{ae2} is a 1×1 convolutional layer with output dimension 256, and $\tilde{\tau}_t \in \mathbb{R}^{d_\tau \times 14 \times 14}$ is equivalent to τ_t copied across spatial dimensions to match the size of \mathbf{r}_t^i .

The location transformer encoder FTE_{Loc} consists of two stacks of transformer encoder modules as originally defined in [42] consisting of layer norm, multi-head self-attention, feed-forward networks, and residual connections. Specifically, all transformers in this work use the Pre-LN construction [48], where the Layer Norm module is placed before the multi-head attention and feed-forward network, as we observed improved convergence. As specified in Sec. 3.2, the multi-head attention modules use both difference attention (Sec. 3.1) and agent-aware attention (Appendix C). The embedding dimension of all

keys, queries, and values as well as the output $\mathbf{h}_{\text{Loc},t}^i$ is 256, the hidden dimension of feedforward modules is 512, the dropout rate used is 0.1, and the number of heads used for multi-head attention is 8.

The appearance transformer encoder FTE_{App} additionally consists of two stacks of transformer encoder modules. However, unlike FTE_{Loc} , the standard dot-product attention formulation is used, and all linear projections in both the multihead attention modules as well as the feedforward network are replaced with 2D convolutional layers with a filter size of 3×3 . All embeddings maintain the same spatial dimensions of 14×14 during computation, the channel dimension used is 256, the hidden channel dimension of the feedforward modules is 512, the dropout rate used is 0.1, and the number of heads used for multi-headed attention is 8.

Note, for readability we formulate all models assuming every instance i is present at every input time step $t \in \{1, \dots, T\}$. However, in practice, some instances will not be present in some input time steps due to occlusions or instances entering/leaving the frame, *i.e.*, there are instances i and input frames t for which $p_t^i = 0$. For all i, t such that $p_t^i = 0$, we do not compute $\tilde{\mathbf{x}}_{\text{Loc},t}^i$ or $\tilde{\mathbf{x}}_{\text{App},t}^i$ since there are no inputs from which we can compute these. Consequentially, neither FTE_{Loc} nor FTE_{App} receive input representing instance i for time t and thus do not produce encoder representations $\mathbf{h}_{\text{Loc},t}^i$ and $\mathbf{h}_{\text{App},t}^i$ for them.

The decoder location feature model \tilde{f}_{Loc} produces the feature representation $\tilde{\mathbf{x}}_{\text{Loc},t}^i$ containing information about the most recently predicted location, odometry, and the corresponding instance class. \tilde{f}_{Loc} can be fully specified by the following model components:

$$\mathbf{x}''_t^i = f_{\text{d1}}([\tilde{\mathbf{x}}_{t-1}^i, \text{onehot}(c^i), o^t]), \quad (26)$$

$$\tilde{\mathbf{x}}_{\text{Loc},t}^i = f_{\text{d2}}([\mathbf{x}''_t^i, \tau_t]). \quad (27)$$

First, an initial representation \mathbf{x}''_t^i is computed from the previous location prediction $\tilde{\mathbf{x}}_{t-1}^i$ using linear layer f_{d1} , corresponding instance class c^i , and odometry o_t . This is concatenated with temporal encoding τ_t and passed through a second linear layer f_{d2} . Both f_{d1} and f_{d2} use output dimension equal to 256.

The decoder appearance feature model \tilde{f}_{App} produces the feature representation $\tilde{\mathbf{x}}_{\text{App},t}^i$ containing information about the most recently predicted appearance. \tilde{f}_{App} can be fully specified by the following model components:

$$\tilde{\mathbf{x}}_{\text{App},t}^i = f_{\text{ad2}}([f_{\text{ad1}}(\tilde{\mathbf{r}}_{t-1}^i), \tilde{\tau}_t]), \quad (28)$$

where f_{ad1} and f_{ad2} are convolutional layers with the same structure as f_{ae1} and f_{ae2} , respectively.

Both the location and appearance transformer decoders FDE_{Loc} and FDE_{App} use the same construction and hyperparameters as their encoder counterparts. The primary difference is that they are transformer decoders as defined in [42] and hence additionally introduce cross attention layers which operate on the encoder representations $\{\mathbf{h}_{\text{Loc},t}^i\}_{1:T}^{1:N}$ and $\{\mathbf{h}_{\text{App},t}^i\}_{1:T}^{1:N}$, respectively. Output decoder representations $\tilde{\mathbf{h}}_{\text{Loc},t}^i$ and $\tilde{\mathbf{h}}_{\text{Loc},t}^i$ are computed autoregressively; *e.g.*, previous predictions $\{\tilde{\mathbf{x}}_t^i\}_{T:t'-1}^{1:N}$ for times T through $t' - 1$ are used to compute the outputs $\{\tilde{\mathbf{h}}_{\text{Loc},t'}^i\}$. These embeddings are then used to produce $\{\tilde{\mathbf{x}}_t^i\}_{1:N}^{1:N}$ for time t' , and these new predictions are fed back into the model to produce output for the next time step $t' + 1$, and so on. Decoder attention is masked to maintain causality, *i.e.*, embeddings representing a given time t are prevented from attending to representations for future time steps $t' > t$.

f_{LocOut} , f_{POut} , and f_{vel} are all 3-layer multilayer perceptrons with hidden sizes [512, 256] and ReLU activations. f_{AppOut} is a 3×3 convolutional layer.

C. Agent-aware Attention

Due to their permutation-invariance with respect to their inputs, transformers do not have the inherent capacity to reason about the identity of the entities corresponding to input trajectories. To address this problem, Yuan *et al.* [54] introduced agent-aware attention. This approach allows transformers to encode the identity of its inputs within the model, which makes it easier for these models to reason about the trajectories of individual entities and leads to better forecasting performance.

Let $\mathbf{X}_{\text{self}} \in \mathbb{R}^{M_1 \times d}$ and $\mathbf{X}_{\text{other}} \in \mathbb{R}^{M_2 \times d}$ of lengths M_1 and M_2 , respectively, be the input sequences with embedding dimension d . For self-attention, both input sequences are the same and represent the input trajectories of a number of agents, while for cross-attention, the first input sequence corresponds to a future trajectory forecast and the second corresponds to

input trajectories. The agent-aware attention output $\mathbf{Y} \in \mathbb{R}^{M_1 \times d}$ is then computed as

$$\mathbf{Z} = \mathbf{M} \odot (\mathbf{Q}_{\text{agent}} \mathbf{K}_{\text{agent}}^T) + (1 - \mathbf{M}) \odot (\mathbf{Q}_{\text{context}} \mathbf{K}_{\text{context}}^T), \quad (29)$$

$$\mathbf{Y} = \text{softmax} \left(\mathbf{Z} / \sqrt{d} \right) \mathbf{V}, \quad (30)$$

where \odot represents element-wise multiplication. Specifically, agent-aware attention first computes two sets of keys $\mathbf{K}_{\text{agent}} = f_{K,\text{agent}}(\mathbf{X}_{\text{other}})$, $\mathbf{K}_{\text{context}} = f_{K,\text{context}}(\mathbf{X}_{\text{other}})$ and queries $\mathbf{Q}_{\text{agent}} = f_{Q,\text{agent}}(\mathbf{X}_{\text{self}})$, $\mathbf{Q}_{\text{context}} = f_{Q,\text{context}}(\mathbf{X}_{\text{self}})$ from the original inputs. It then computes two sets of attention scores from the *agent* keys/queries and from the *context* keys/queries and selects between them using mask $\mathbf{M} \in \{0, 1\}^{M_1 \times M_2}$. This mask encodes identity information: $\mathbf{M}_{ij} = 1$ if entity i in the first input sequence and entity j in the second input sequence correspond to the same agent, and $\mathbf{M}_{ij} = 0$ otherwise. In other words, two sets of attention parameters are computed, and one set is used for input pairs corresponding to the same agent while the other is used for all pairs corresponding to different agents, *i.e.*, the context for this agent. Value aggregation proceeds as in standard attention from this step.

We additionally use agent-aware attention within the difference attention module defined in Sec. 3.1. This is implemented in a similar fashion, where separate attention parameters are computed for input pairs corresponding to the same agent and for input pairs corresponding to different agents. We formally specify this as

$$\mathbf{Z} = \mathbf{M} \odot \left(\mathbf{Q}_{\text{agent}} \mathbf{K}_{R,\text{agent}}^T - \mathbf{1}_{M_1 \times 1} \text{diag} \left(\mathbf{K}_{B,\text{agent}} \mathbf{K}_{R,\text{agent}}^T \right)^T \right) + \quad (31)$$

$$(1 - \mathbf{M}) \odot \left(\mathbf{Q}_{\text{context}} \mathbf{K}_{R,\text{context}}^T - \mathbf{1}_{M \times 1} \text{diag} \left(\mathbf{K}_{B,\text{context}} \mathbf{K}_{R,\text{context}}^T \right)^T \right), \quad (32)$$

$$\mathbf{Y} = \text{softmax} \left(\mathbf{Z} / \sqrt{d} \right) \mathbf{V}_O - \mathbf{V}_S, \quad (33)$$

with $\mathbf{K}_{R,\text{agent}} = f_{K,R,\text{agent}}(\mathbf{X}_{\text{other}})$, $\mathbf{K}_{R,\text{context}} = f_{K,R,\text{context}}(\mathbf{X}_{\text{other}})$, $\mathbf{K}_{B,\text{agent}} = f_{K,B,\text{agent}}(\mathbf{X}_{\text{other}})$, $\mathbf{K}_{B,\text{context}} = f_{K,B,\text{context}}(\mathbf{X}_{\text{other}})$, $\mathbf{V}_O = f_{V_O}(\mathbf{X}_{\text{other}})$, and $\mathbf{V}_S = f_{V_S}(\mathbf{X}_{\text{self}})$.

D. Losses for Foreground Forecasting

The loss used by the foreground forecasting model are

$$\mathcal{L}_{\text{FG}} = \mathcal{L}_{\text{Loc}} + \mathcal{L}_{\text{P}} + \mathcal{L}_{\text{App}} + \mathcal{L}_{\text{Vel}}. \quad (34)$$

The location loss \mathcal{L}_{Loc} trains the bounding box predictions $\hat{\mathbf{x}}_{\text{Box},t}^i := [\hat{x}_{0,t}, \hat{y}_{0,t}, \hat{x}_{1,t}, \hat{y}_{1,t}]$ and depth predictions \hat{d}_t^i to match the target boxes $\mathbf{x}_{\text{Box},t}^{*i}$ and depths d_t^{*i} . This is specified as

$$\mathcal{L}_{\text{Loc}} := \frac{1}{\sum_{i=1}^N \sum_{t=T+1}^{T+F} p_t^{*i}} \sum_{i=1}^N \sum_{t=T+1}^{T+F} p_t^{*i} \left(\lambda_1 \text{SmoothL1}(\hat{\mathbf{x}}_{\text{Box},t}^i, \mathbf{x}_{\text{Box},t}^{*i}) + \lambda_2 \text{SmoothL1}(\hat{d}_t^i, d_t^{*i}) + \lambda_3 \text{IoU}(\hat{\mathbf{x}}_{\text{Box},t}^i, \mathbf{x}_{\text{Box},t}^{*i}) \right), \quad (35)$$

where p_t^{*i} is ground-truth presence, *i.e.*, equals 1 if instance i is present in frame t and 0 otherwise, IoU is bounding box intersection-over-union, SmoothL1 is the function

$$\text{SmoothL1}(\mathbf{a}, \mathbf{b}) := \sum_j \text{SmoothL1Fn}(\mathbf{a}_j, \mathbf{b}_j), \quad (36)$$

$$\text{SmoothL1Fn}(a, b) := \begin{cases} \frac{1}{2}(a - b)^2, & \text{if } |a - b| < 1, \\ |a - b| - \frac{1}{2} & \text{otherwise} \end{cases}, \quad (37)$$

and coefficients $\lambda_1 = 1$, $\lambda_2 = 10$, $\lambda_3 = 100$ are used to balance the magnitudes of the losses.

The presence loss \mathcal{L}_{P} trains the presence predictions $\hat{p}_t^i \in \mathbb{R}$ to correctly indicate whether a given instance i is present in frame t , and is computed as

$$\mathcal{L}_{\text{P}} := \frac{\lambda_4}{NF} \sum_{i=1}^N \sum_{t=T+1}^{T+F} p_t^{*i} \log \sigma(\hat{p}_t^i) + (1 - p_t^{*i}) \log(1 - \sigma(\hat{p}_t^i)), \quad (38)$$

where σ is the sigmoid function and $\lambda_4 = 10$.

The appearance loss \mathcal{L}_{App} trains the appearance predictions $\hat{\mathbf{r}}_t^i$ for instance i at frame t to match the target features \mathbf{r}_t^{*i} extracted for this instance at frame t , and consists of the mean-squared error of the features for all valid instance/time pairs, *i.e.*,

$$\mathcal{L}_{\text{App}} := \frac{\lambda_5}{\sum_{i=1}^N \sum_{t=T+1}^{T+F} J p_t^{*i}} \sum_{i=1}^N \sum_{t=T+1}^{T+F} \sum_{j=1}^J p_t^{*i} (\hat{\mathbf{r}}_{j,t}^i - \mathbf{r}_{j,t}^{*i})^2, \quad (39)$$

where j indexes over all spatial dimensions of the feature tensors, $J = 256 \times 14 \times 14$ is the total number of elements of the feature tensors, and $\lambda_5 = 10$.

The encoder velocity loss \mathcal{L}_{Vel} trains the velocity predictions $\hat{\mathbf{v}}_{E,t}^i \in \mathbb{R}^4$ to match the ground-truth velocities $\mathbf{v}_t^{*i} := \mathbf{x}_{t+1}^{*i} - \mathbf{x}_t^{*i}$, and is computed as

$$\mathcal{L}_{\text{Vel}} := \frac{\lambda_6}{\sum_{i=1}^N \sum_{t=1}^T p_t^{*i} p_{t+1}^{*i}} \sum_{i=1}^N \sum_{t=1}^T p_t^{*i} p_{t+1}^{*i} \text{SmoothL1}(\hat{\mathbf{v}}_{E,t}^i, \mathbf{v}_t^{*i}), \quad (40)$$

where $\lambda_6 = 1$.

E. Depth Completion Model

The depth completion model operates on noisy and incomplete reprojected background depth \tilde{d}^B along with depth mask Q and background class probabilities \hat{m}^{Prob} and produces depth maps \hat{d}_{Fill}^B and \hat{d}_{Bias}^B . This model can be formally represented using the following components:

$$d_1 = f_{dc1}([\tilde{d}^B, Q, \hat{m}^{\text{Prob}}]), \quad (41)$$

$$d_2 = d_1 + \text{Upsample}(f_{dc2}(\text{Downsample}(d_1))), \quad (42)$$

$$\hat{d}_{\text{Fill}}^B = f_{\text{fill}}(d_2), \quad (43)$$

$$\hat{d}_{\text{Bias}}^B = f_{\text{Bias}}(d_2). \quad (44)$$

First, reprojected background depth \tilde{d}^B , depth mask Q , and predicted background class probabilities \hat{m}^{Prob} are concatenated together and processed with convolutional layer f_{dc1} which uses a kernel size of 3 and has output channel dimension 32. The output of this, d_1 , is downsampled by a factor of 2 using bilinear interpolation, fed into convolutional network f_{dc2} , upsampled to the original resolution using bilinear interpolation, and added with d_1 to produce the second intermediate output d_2 . f_{dc2} contains 2 convolutional layers with a kernel size of 3, output channel dimension of 32, and a ReLU activation between them. The outputs \hat{d}_{Fill}^B and \hat{d}_{Bias}^B are then obtained from d_2 using convolutional networks f_{fill} and f_{bias} , respectively. Both of these networks contain two convolutional layers with a ReLU activation between them, where the first layer uses a kernel size of 3 and an output channel size of 32 and the second layer uses a kernel size of 1. The final background depth estimate \hat{d}^B is obtained from outputs \hat{d}_{Fill}^B and \hat{d}_{Bias}^B as specified in Eq. (13).

F. Additional Implementation Details

The overall approach is trained in two stages: first, the foreground prediction model is trained; afterwards, the corresponding parameters are frozen, and the refinement model is trained.

The foreground model is trained for 48000 steps using the ADAM optimizer; the initial learning rate is set to 10^{-4} , and it is lowered to 10^{-5} after 36000 optimization steps. All odometries o_t are normalized by subtracting the training data set mean and then dividing by training data set standard deviation before being used as input. All location inputs \mathbf{x}_t^i are normalized to lie within $[-1, 1]$; furthermore, location outputs $\hat{\mathbf{x}}_t^i$ are made at this normalized scale and unnormalized before being used at later stages. During training and inference, forecasts are only predicted for instances present in the most recent input frame, *i.e.*, for instances i such that $p_T^i = 1$. During training of the foreground model, ground-truth future odometry is used. During evaluation, unless otherwise noted, the egomotion estimation module described by Graber *et al.* [14] was used to obtain future odometry which was used as input. We use the same odometry representation as Graber *et al.* [14] consisting of a five-dimensional vector containing speed and yaw rate of the ego-vehicle at time t as well its top-down displacement and angular displacement between steps t and $t - 1$.

Ablation 4 in this work uses odometry during inference that was obtained using ORB-SLAM3 [2]. This was run using stereo images, where each sequence of 30 frames was treated as its own SLAM session providing 6-dof poses for all frames in the sequence.

The refinement model is trained for 24000 steps using the ADAM optimizer; the initial learning rate is set to 10^{-4} , and it is lowered to 10^{-5} after 18000 optimization steps. During training, the inputs are scaled to a spatial resolution of $\frac{H}{4} \times \frac{W}{4}$, and the loss is additionally computed at this scale. During inference, inputs are scaled to the final spatial resolution, *i.e.*, $H \times W$.

To process a Cityscapes sequence, the model needs 560 ms on average using an NVIDIA A6000, which is on par with the 700 ms required by Gruber *et al.* [14]. This can be significantly reduced by further engineering effort.

G. Additional AIODrive Details

The AIODrive sequences are annotated using 23 object classes. To facilitate comparison against results on the Cityscapes dataset, we only train and evaluate using background classes which are also present in Cityscapes. This leaves 11 background “stuff” classes and 2 foreground “things” classes (the only annotated “things” instances in AIODrive are “vehicles” and “pedestrians”). As annotations are only provided for the trainval dataset, we split this into a training dataset containing all annotated sequences for towns 1 through 5 and a validation dataset containing all annotated sequences for town 6. We use 5 frames as input and forecast the 5th frame into the future, corresponding to 0.5 seconds of input and a 0.5 second forecast (which is comparable to the Cityscapes mid-term setting).

During both training and evaluation, we only consider instances whose masks have an area of at least 400 pixels in an attempt to filter out distant, imperceptible instances. For evaluation, we use non-overlapping sequences of 10 frames from each validation sequence. Additionally, the data contains some periods of time with little to no motion, which skews the evaluation metrics artificially high. To ensure that the metrics properly capture the ability of the models to anticipate motion, we filter out validation sequences where the recording vehicle is moving less than 1 m/s at all points in the input sequence and where at least half of the instance mask centers move less than 10 pixels. This leaves 814 sequences with motion for evaluation purposes. To ensure that the tracking-based metrics can be computed, we use ground-truth instance bounding boxes and ids as input to the forecasting models.

The base semantic and instance segmentation models are the same as that used for Cityscapes, *i.e.*, MaskRCNN [16] for instance segmentation and Panoptic Deeplab [3] for semantic segmentation. For both, we initialize from the Cityscapes pre-trained model and finetune on AIODrive. For the models that use predicted depth, we use Cascade-Stereo [15] on the stereo input images. We do not finetune the depth model on this dataset.

H. Cityscapes Instance Segmentation

We also evaluate our Cityscapes-trained model on instance segmentation. Here, we consider only ‘things’ instances during evaluation, and hence we disregard the pixels corresponding to the ‘stuff’ classes.

Metrics. We evaluate instance segmentation using the standard metrics [7]: 1) *Average Precision* (AP) computes true positives using a number of overlap thresholds, averages over these thresholds, and then averages over classes; 2) AP50 computes average precision with an overlap threshold of 0.5 and then averages across classes.

Baselines. We compare against the baselines presented by Gruber *et al.* [14]. F2F is introduced by Luc *et al.* [27] and predicts the features of a future scene using a convolutional model. It then obtains instances by passing these features through MaskRCNN heads. IndRNN-Stack is the independent RNN and stacking model by Gruber *et al.* [14]. PFA, introduced by Lin *et al.* [25], compresses input feature pyramids into a low-resolution feature map for forecasting.

Results. The results for this task are presented in Tab. 4. We outperform F2F and IndRNN-Stack in the mid-term setting but PFA performs better. This is to be expected because PFA was directly trained on instance segmentation while we directly apply the model trained on panoptic segmentation, *i.e.*, we don’t retrain our model specifically for instance segmentation.

I. Cityscapes Semantic Segmentation

Following prior work [14], we also evaluate our model on semantic segmentation forecasting. In this context, we do not care about specific instances. Hence, for each pixel, we discard all predicted identity information.

Metrics. Semantic segmentation forecasting is evaluated using the standard *intersection over union* (IoU) metric computed between predictions and ground truth per class and averaged over classes. IoU (MO), meanwhile, computes an average IoU over ‘things’ classes only.

	Short term: $\Delta t = 3$		Mid term: $\Delta t = 9$	
	AP	AP50	AP	AP50
Oracle	34.6	57.4	34.6	57.4
Last seen frame	8.9	21.3	1.7	6.6
F2F [27]	19.4	39.9	7.7	19.4
IndRNN-Stack [14]	17.8	38.4	10.0	22.3
PFA [25]	24.9	48.7	14.8	30.5
Ours	19.9	39.9	11.2	25.2

Table 4. **Instance segmentation forecasting on the Cityscapes Validation dataset.** Higher is better for all metrics.

Accuracy (mIoU)	Short term: $\Delta t = 3$		Mid term: $\Delta t = 9$	
	All	MO	All	MO
Oracle	80.6	81.7	80.6	81.7
Copy last	59.1	55.0	42.4	33.4
Bayesian S2S [1]	65.1	/	51.2	/
DeformF2F [36]	65.5	63.8	53.6	49.9
LSTM M2M [40]	67.1	65.1	51.5	46.3
F2MF [37]	69.6	67.7	57.9	54.6
IndRNN-Stack [14]	67.6	60.8	58.1	52.1
PFA [25]	71.1	69.2	60.3	56.7
Ours	67.9	61.2	58.1	51.7

Table 5. **Semantic forecasting results on the Cityscapes validation dataset.** Baseline numbers, besides oracle and copy last, are from [37]. Higher is better for all metrics. Our model exploits stereo and odometry, which are provided by typical autonomous vehicle setups and are included in Cityscapes.

Baselines. Many of the baselines operate by predicting the features of a future scene [1, 25, 36, 37]. LSTM M2M [40] warps input semantics using a predicted optical flow between the most recent frame and the target frame. Note that these approaches do not use depth inputs, and all except Bayesian S2S [1] do not use egomotion as input.

Results. The results for this task are given in Tab. 5. We outperform IndRNN-Stack by a small margin in the short-term setting, and have comparable results in the mid-term setting. We additionally outperform most other baselines. Note that this metric does not care about the boundaries between individual instances and hence weights some types of errors differently than the other metrics we use. These other metrics more properly evaluate whether specific instances are localized in the correct places, which we argue better captures the goals of forecasting. Note that PFA is directly trained on semantic segmentation forecasting while our approach was trained on forecasting of panoptic segmentations.

J. Additional Cityscapes Visualizations

Fig. 6 presents a visual comparison between our approach and IndRNN-Stack for the short-term setting for the sequences which were shown for the mid-term setting in Fig. 3. We present additional visualizations for the mid-term setting in Fig. 7 and for the short term setting in Fig. 8.

K. Additional Cityscapes Metrics

Tabs. 6 to 8 present metrics computed for the Cityscapes test dataset using the mid-term setting for panoptic, instance, and semantic segmentation forecasting, respectively. We outperform all other approaches for panoptic and instance segmentation forecasting on the test data. On semantic segmentation, we outperform IndRNN-Stack on the test data, whereas F2MF [37] outperforms our approach. However, note that the F2MF model used for test evaluation was trained on both the training and validation datasets, while the other models were trained only on the training data.

Tabs. 9 to 14 contain the per-class breakdown of all panoptic segmentation metrics shown in Tab. 1. The results shown in Tab. 1 consist of the average of these metrics taken over the values obtained for each individual class. Our model is better

	PQ	All SQ	RQ	Things			Stuff		
				PQ	SQ	RQ	PQ	SQ	RQ
Flow	25.6	70.1	34.0	12.4	66.3	18.1	35.3	72.9	45.5
Hybrid [40] (bg) and [27] (fg)	29.4	69.8	38.5	18.0	67.2	25.7	37.6	71.6	47.8
IndRNN-Stack	35.7	72.0	46.5	24.0	69.0	33.7	44.2	74.2	55.8
Ours	36.9	72.7	48.0	26.7	70.3	37.0	44.4	74.4	55.9

Table 6. **Panoptic segmentation forecasting evaluated on the Cityscapes test set, mid-term.** Higher is better for all metrics.

	AP	AP50	Accuracy (mIoU)	All	MO
F2F [27]	6.7	17.5	F2MF [37]*	59.1	56.3
IndRNN-Stack	8.4	19.8	IndRNN-Stack	57.7	48.8
Ours	9.9	20.7	Ours	58.3	50.0

Table 7. **Instance segmentation forecasting on the Cityscapes Test dataset, mid-term.** Higher is better for all metrics.

Table 8. **Semantic segmentation forecasting results on the Cityscapes test dataset.** Baseline numbers, are from [37]; the * indicates training on both train and validation data. Higher is better for all metrics.

	road	sidewalk	building	wall	fence	pole	traffic light	traffic sign	vegetation	terrain	sky	person	rider	car	truck	bus	train	motorcycle	bicycle	mean
Deeplab (Oracle)†	97.9	78.2	88.5	29.4	38.9	60.0	55.6	74.5	89.5	36.1	87.9	50.8	46.4	67.3	51.5	66.6	37.8	44.2	44.1	60.3
Deeplab (Last seen frame)	94.3	52.4	71.1	11.3	19.4	6.1	12.9	15.0	72.1	16.9	72.7	10.3	8.0	29.6	35.1	51.7	24.2	9.8	7.9	32.7
Flow	95.6	61.5	79.8	17.3	28.6	8.7	26.2	36.8	80.7	26.9	79.7	21.0	14.0	43.4	40.6	56.8	26.7	23.2	18.7	41.4
Hybrid [40] (bg) and [27] (fg)	96.2	63.4	81.4	23.1	23.7	7.1	19.1	36.9	82.3	20.3	79.8	26.8	21.8	46.4	42.2	60.0	41.4	25.6	22.5	43.2
IndRNN-Stack	96.2	66.1	83.5	26.1	27.4	31.7	37.0	49.9	84.8	26.1	82.0	31.8	31.5	48.8	42.2	61.2	47.0	31.4	27.3	49.0
Ours	96.2	66.3	83.8	25.9	27.4	34.4	37.0	50.3	84.8	26.5	82.1	34.9	36.7	51.2	41.2	63.1	47.6	32.4	32.0	50.2

Table 9. **Per-class results for Panoptic Quality on Cityscapes validation dataset (short-term).**

on average for every metric than all other approaches, and it is additionally better than prior approaches for every metric for most classes.

L. Code Details

All models are implemented using PyTorch v. 1.10.0², which is made available for use with a custom BSD-style license.³ We additionally use the Detectron2 framework (version 0.4.1)⁴, which is released under the Apache 2.0 license.⁵ Code implementing our models and experiments can be found at <https://github.com/cgraber/psf-diffattn>.

M. Potential Negative Societal Impact

One of the primary applications of this work is to better enhance the ability of autonomous agents to anticipate the future and respond appropriately to a dynamic environment. In this context, problems can arise if an agent makes a decision based on a faulty prediction – for example, if a self-driving car does not anticipate a pedestrian stepping into the street, it could unintentionally hurt the pedestrian if they step out in front of the car. For such a system, the consequence of prediction errors can be injury or death. It is thus critical that appropriate care be taken before deployment of such a system to ensure that not only are prediction errors sufficiently low across a variety of environments but also that proper failsafes are put in place to minimize the negative consequences of acting upon a misprediction.

²<https://pytorch.org/>

³<https://github.com/pytorch/pytorch/blob/v1.10.0/LICENSE>

⁴<https://github.com/facebookresearch/detectron2>

⁵<https://github.com/facebookresearch/detectron2/blob/v0.4.1/LICENSE>



Figure 6. Short-term panoptic segmentation forecasts on Cityscapes.

	road	sidewalk	building	wall	fence	pole	traffic light	traffic sign	vegetation	terrain	sky	person	rider	car	truck	bus	train	motorcycle	bicycle	mean
Deeplab (Oracle)†	97.9	78.2	88.5	29.4	38.9	60.0	55.6	74.5	89.5	36.1	87.9	50.8	46.4	67.3	51.5	66.6	37.8	44.2	44.1	60.3
Deeplab (Last seen frame)	90.4	32.5	57.6	7.6	10.6	4.6	8.9	7.4	55.1	8.8	57.3	5.3	2.5	13.2	19.2	27.3	10.1	4.7	3.0	22.4
Flow	90.5	35.8	66.2	7.7	15.0	4.6	11.9	11.1	65.6	11.6	64.4	5.9	2.5	19.0	21.5	27.7	13.5	11.8	5.3	25.9
Hybrid [40] (bg) and [27] (fg)	93.2	44.9	70.5	12.4	14.8	1.2	8.0	10.8	69.7	13.9	67.2	8.0	4.5	27.3	33.5	41.7	27.9	8.3	6.1	29.7
IndRNN-Stack	93.9	50.8	76.4	18.2	19.9	8.7	18.7	28.5	77.0	18.6	72.7	16.2	12.0	33.3	36.1	53.0	29.8	14.1	12.6	36.3
Ours	93.9	50.9	76.5	18.1	20.8	9.3	18.8	28.6	77.0	18.6	72.7	19.9	14.6	39.5	38.8	56.9	26.2	18.6	14.5	37.6

Table 10. Per-class results for Panoptic Quality on Cityscapes validation dataset (mid-term).

	road	sidewalk	building	wall	fence	pole	traffic light	traffic sign	vegetation	terrain	sky	person	rider	car	truck	bus	train	motorcycle	bicycle	mean
Deeplab (Oracle)†	98.0	85.6	90.5	74.3	74.8	69.7	73.5	80.1	90.9	75.7	92.6	76.0	70.8	84.2	88.4	90.8	87.6	73.8	72.1	81.5
Deeplab (Last seen frame)	94.4	71.5	78.8	65.4	65.6	67.0	68.3	67.4	77.8	67.7	83.0	64.4	60.1	69.2	74.7	76.7	75.7	62.7	63.4	71.3
Flow	95.6	76.0	83.2	68.5	68.3	65.0	65.9	67.3	83.4	69.1	86.6	65.6	61.4	75.8	77.5	80.0	74.1	66.1	64.4	73.4
Hybrid [40] (bg) and [27] (fg)	96.3	77.2	84.9	70.0	69.0	59.5	63.6	65.9	84.6	70.8	86.5	66.8	61.9	77.2	80.3	83.1	80.5	65.6	63.8	74.1
IndRNN-Stack	96.3	77.0	86.3	71.1	69.4	61.4	65.4	70.6	86.6	71.3	88.3	67.7	63.8	77.7	81.4	81.4	74.8	67.6	65.8	74.9
Ours	96.3	77.1	86.4	71.4	69.8	61.7	65.4	70.8	86.6	70.9	88.3	69.4	66.4	78.9	81.7	84.1	77.9	68.1	67.2	75.7

Table 11. Per-class results for Segmentation Quality on Cityscapes validation dataset (short-term).

	road	sidewalk	building	wall	fence	pole	traffic light	traffic sign	vegetation	terrain	sky	person	rider	car	truck	bus	train	motorcycle	bicycle	mean
Deeplab (Oracle)†	98.0	85.6	90.5	74.3	74.8	69.7	73.5	80.1	90.9	75.7	92.6	76.0	70.8	84.2	88.4	90.8	87.6	73.8	72.1	81.5
Deeplab (Last seen frame)	90.7	68.2	72.6	63.4	62.4	66.1	72.7	73.0	71.2	64.0	77.3	63.7	61.3	66.8	62.9	70.8	74.3	56.4	64.4	68.5
Flow	90.8	68.6	76.0	66.1	64.1	64.1	69.0	67.2	75.0	64.5	78.5	63.5	60.4	69.1	70.2	74.3	75.8	60.2	63.0	69.5
Hybrid [40] (bg) and [27] (fg)	93.3	69.7	77.9	66.6	65.3	59.9	62.9	61.9	76.9	65.1	79.6	63.7	58.4	71.5	72.6	72.2	73.7	62.1	60.6	69.1
IndRNN-Stack	94.1	71.3	81.5	68.4	66.8	59.0	64.1	65.1	80.9	68.1	83.0	64.3	61.4	73.4	76.9	76.1	74.4	62.5	62.9	71.3
Ours	94.2	71.5	81.7	69.1	66.1	60.1	64.2	65.7	81.0	68.5	83.0	64.7	61.1	74.7	76.5	79.4	67.7	63.6	64.7	71.4

Table 12. Per-class results for Segmentation Quality on Cityscapes validation dataset (mid-term).

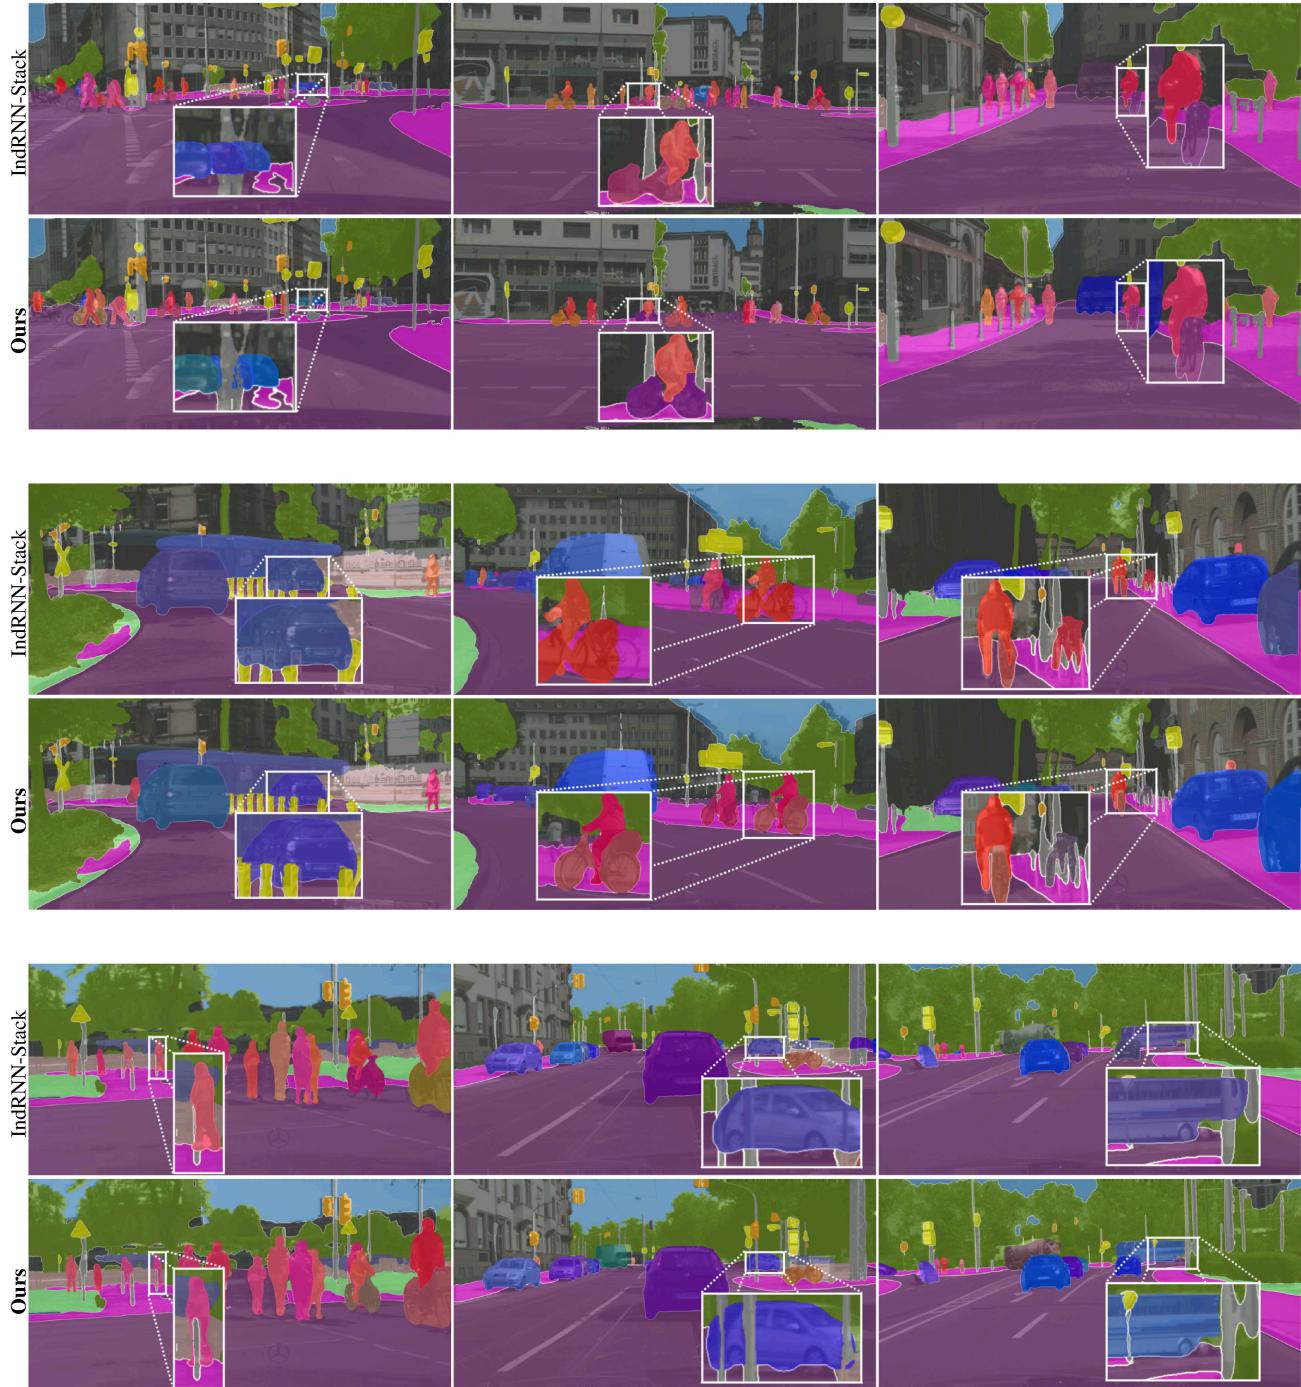


Figure 7. Additional mid-term visualizations.

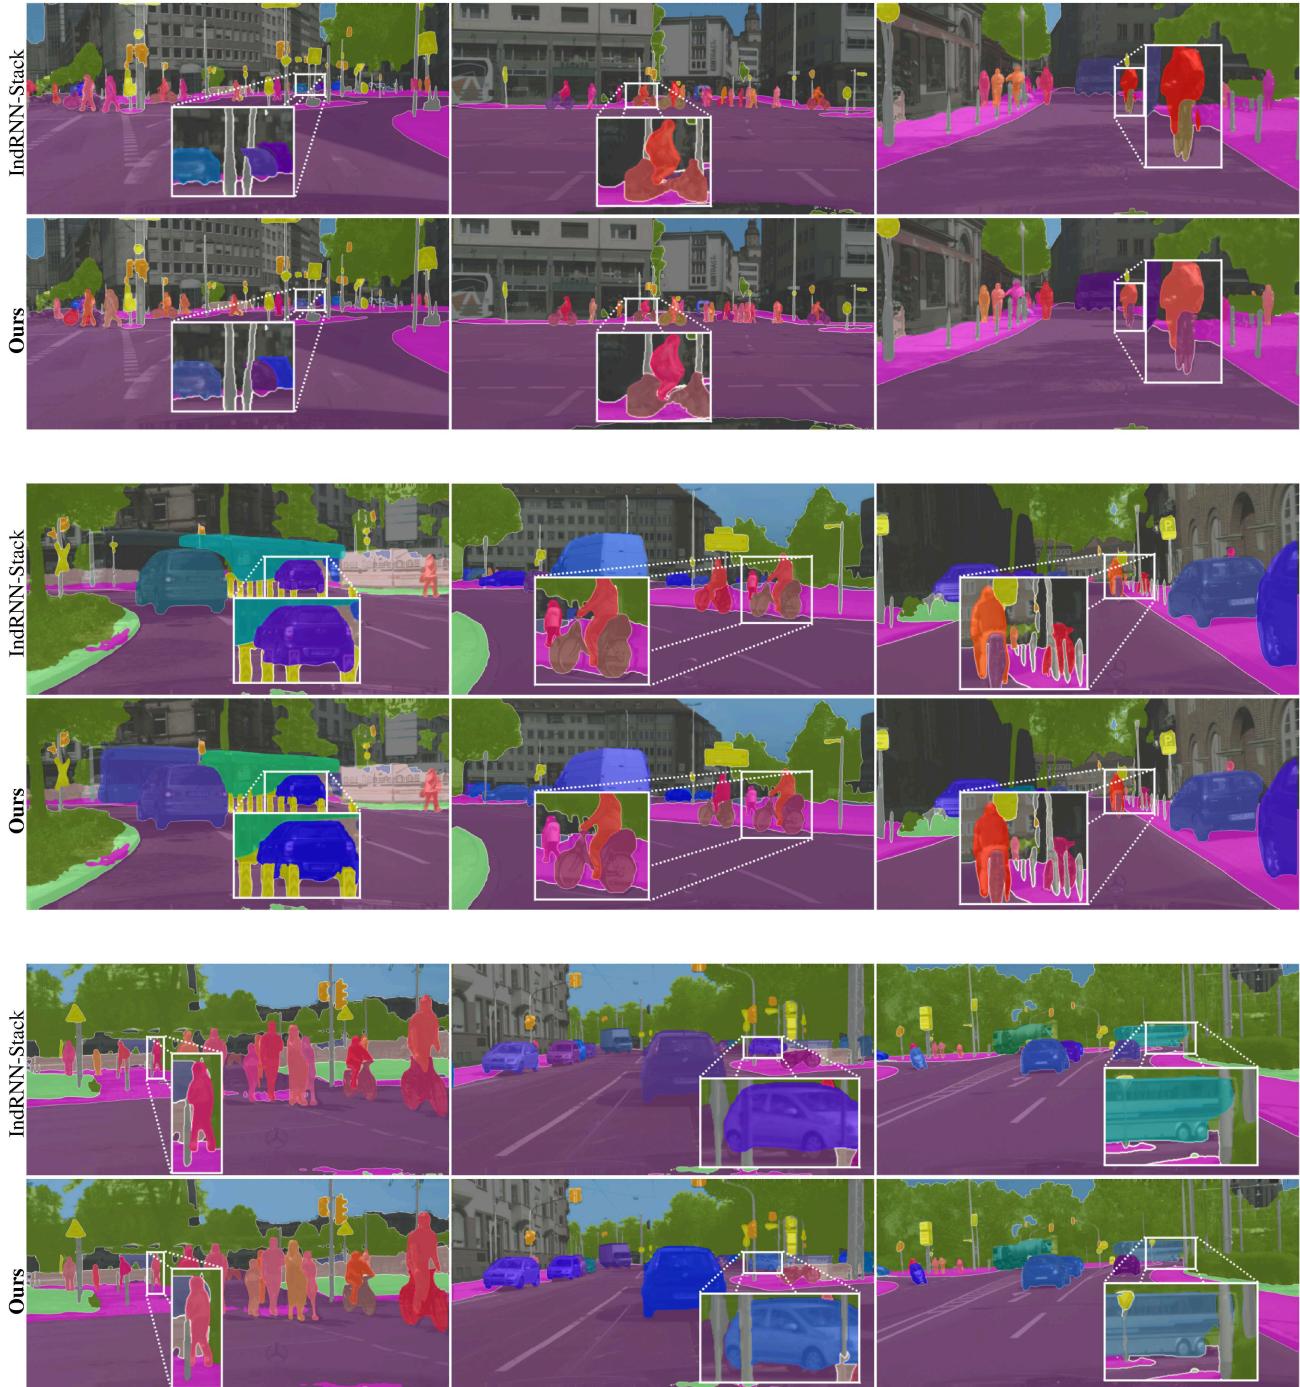


Figure 8. Additional short-term visualizations.

	road	sidewalk	building	wall	fence	pole	traffic light	traffic sign	vegetation	terrain	sky	person	rider	car	truck	bus	train	motorcycle	bicycle	mean
Deeplab (Oracle)†	99.9	91.3	97.8	39.5	52.1	86.1	75.6	93.0	98.5	47.7	94.9	66.9	65.5	80.0	58.2	73.4	43.1	59.9	61.2	72.9
Deeplab (Last seen frame)	99.9	73.4	90.2	17.3	29.7	9.1	18.9	22.2	92.7	25.0	87.6	16.0	13.2	42.8	47.0	67.4	32.0	15.6	12.4	42.7
Flow	99.9	81.0	95.9	25.2	41.9	13.4	39.7	54.7	96.8	39.0	92.0	32.1	22.8	57.3	52.4	71.0	36.0	35.1	29.0	53.4
Hybrid [40] (bg) and [27] (fg)	99.9	82.1	95.8	33.0	34.4	11.9	30.0	56.0	97.3	28.8	92.2	40.1	35.3	60.2	52.6	72.2	51.4	39.1	35.3	55.1
IndRNN-Stack	99.9	85.8	96.7	36.7	39.4	51.6	56.6	70.7	97.9	36.6	92.9	47.0	49.3	62.9	51.9	75.2	62.9	46.3	41.5	63.3
Ours	99.9	86.0	97.0	36.2	39.3	55.7	56.6	71.1	97.9	37.4	92.9	50.3	55.3	64.9	50.4	75.0	61.1	47.5	47.6	64.3

Table 13. Per-class results for Recognition Quality on Cityscapes validation dataset (short-term).

	road	sidewalk	building	wall	fence	pole	traffic light	traffic sign	vegetation	terrain	sky	person	rider	car	truck	bus	train	motorcycle	bicycle	mean
Deeplab (Oracle)†	99.9	91.3	97.8	39.5	52.1	86.1	75.6	93.0	98.5	47.7	94.9	66.9	65.5	80.0	58.2	73.4	43.1	59.9	61.2	72.9
Deeplab (Last seen frame)	99.7	47.6	79.3	12.1	17.0	7.0	12.2	10.1	77.4	13.7	74.1	8.3	4.2	19.8	30.6	38.5	13.6	8.3	4.7	30.4
Flow	99.7	52.2	87.1	11.7	23.4	7.2	17.3	16.5	87.4	18.0	82.1	9.2	4.2	27.5	30.6	37.3	17.8	19.6	8.4	34.6
Hybrid [40] (bg) and [27] (fg)	99.9	64.5	90.4	18.7	22.7	2.1	12.7	17.4	90.5	21.4	84.4	12.5	7.7	38.2	46.2	57.9	37.8	13.4	10.1	39.4
IndRNN-Stack	99.7	71.2	93.7	26.6	29.8	14.7	29.2	43.8	95.1	27.3	87.6	25.2	19.5	45.4	47.0	69.6	40.0	22.6	20.0	47.8
Ours	99.7	71.1	93.7	26.2	31.4	15.5	29.3	43.5	95.1	27.1	87.6	30.8	23.9	52.9	50.7	71.7	38.7	29.2	22.4	49.5

Table 14. Per-class results for Recognition Quality on Cityscapes validation dataset (mid-term).