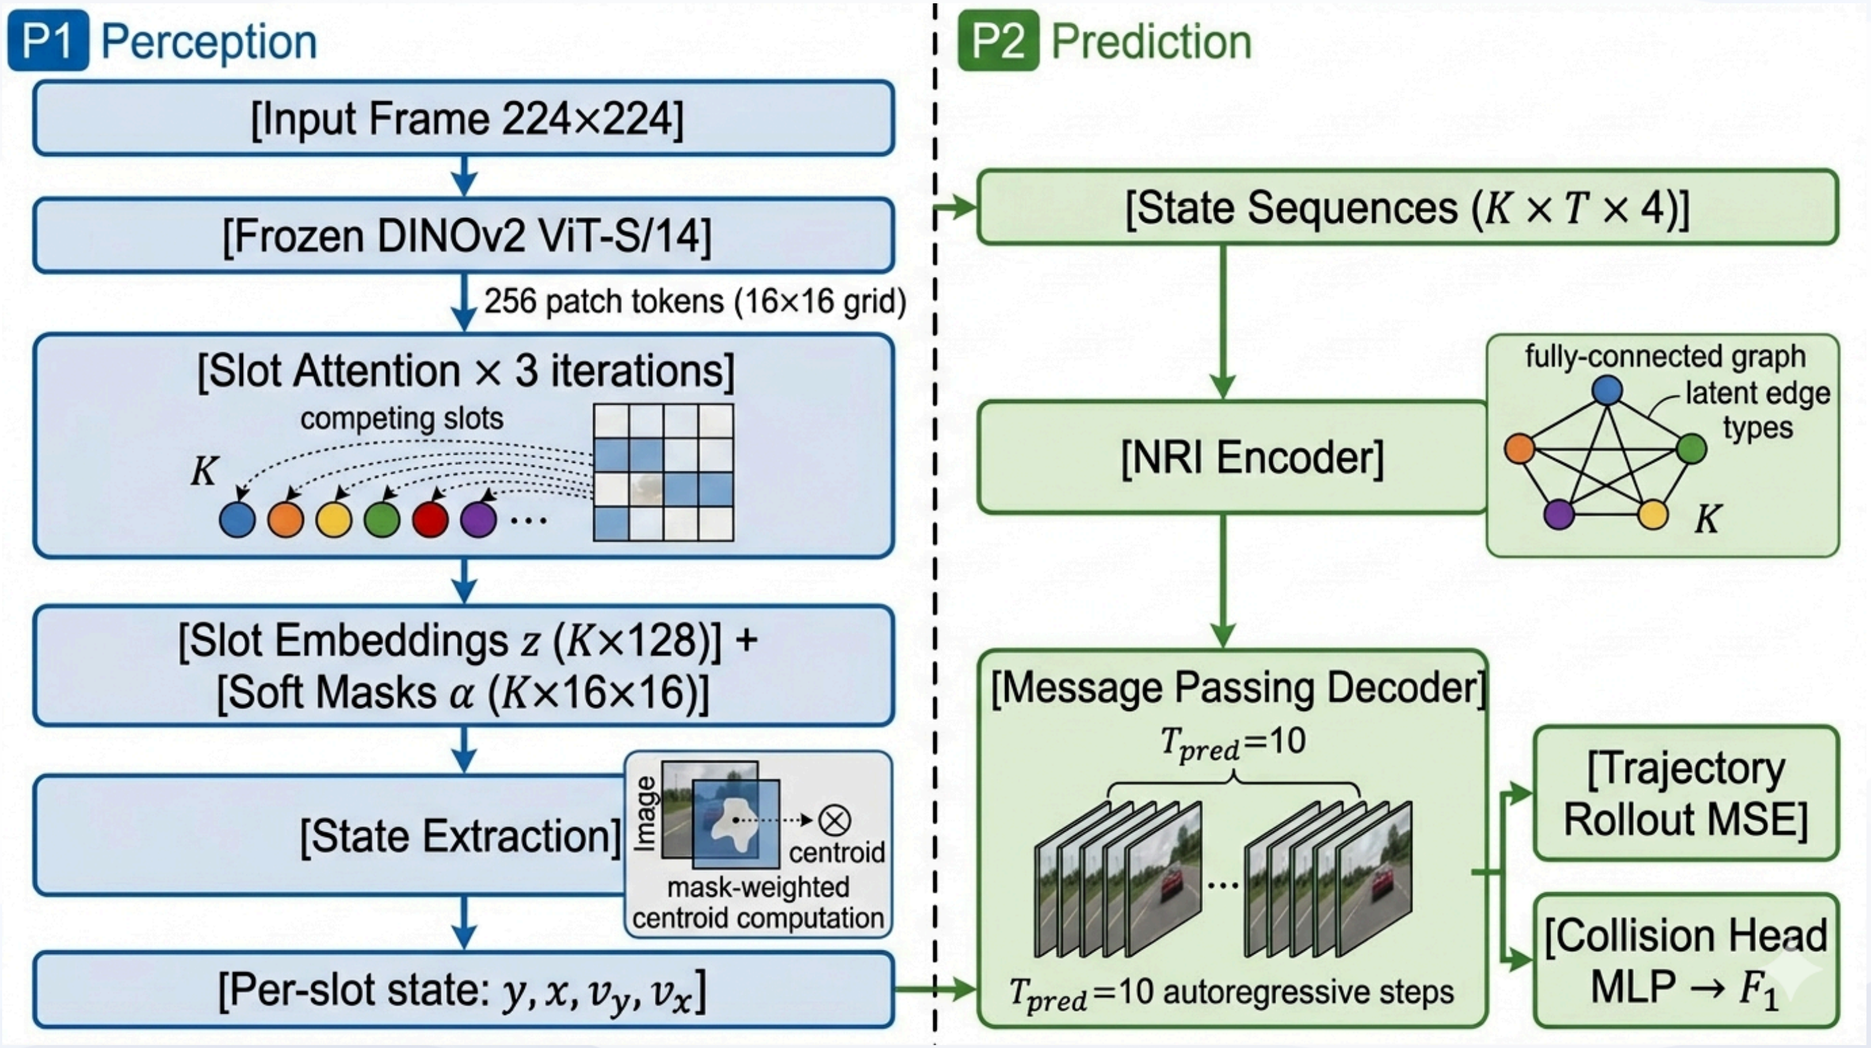
P2P Pipeline



Training Objective

1. Feature reconsrtuction

$$\mathcal{L}_{\text{P1-B}} = \mathcal{L}_{\text{feat}} + \mathcal{L}_{\text{usage}} + \mathcal{L}_{\text{entropy}} + \mathcal{L}_{\text{ground}}$$

2. JEPA + feature reconsrtuction

$$\mathcal{L}_{\text{P1-D}} = \mathcal{L}_{\text{JEPA}} + \alpha \mathcal{L}_{\text{feat}} + \text{regularizers}$$

3. Motion-weighted reconstruction

$$\mathcal{L} = \sum_i (1 + \lambda \hat{\Delta}_i) \, \| \hat{f}_i - f_i \|^2$$

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{P2}} + \beta \mathcal{L}_{\text{coll}}$$

- Input: slot-derived states [y,x,v_y,v_x]
- NRI graph network predicts future rollouts
- Main loss:

$$\mathcal{L}_{\text{P2}} = \sum_{t=T+1}^{T+H} \| \hat{\mathbf{s}}_t - \mathbf{s}_t \|^2$$

- Pairwise future trajectories go into an MLP collision head
- Binary cross-entropy loss:

$$\mathcal{L}_{\text{coll}} = \text{BCE}(\hat{\mathbf{e}}, \mathbf{e})$$