

Supplementary Material: CrossVL: Complexity-Aware Feature Routing and Paired Curriculum for Cross-View Vision-Language Detection

Zhipeng Liu Chunbo Luo
University of Exeter

{z1481, C.Luo}@exeter.ac.uk

https://github.com/1nyourlife/Crossvl_cvpr2026

1. Overview and Qualitative Demonstration

Figure 1 provides a visual demonstration of CrossVL’s improvements over baseline Florence-2 fine-tuning on a representative aerial urban parking lot scene, setting the context for the detailed technical descriptions that follow.

1.1. Organization of This Supplementary Material

This supplementary material is organized as follows:

- **Section 2:** Complete training configuration including hyperparameters, data preprocessing, and optimization settings
- **Section 3:** Detailed network architecture specifications for CPA modules, pathway designs, and curriculum learning schedule
- **Section 4:** Reproducibility protocol including checkpoint selection procedure, evaluation details, and archived data documentation
- **Section 5:** Extended ablation studies and component analysis across all experimental data, demonstrating mutual regularization between CPA and paired curriculum learning
- **Section 6:** Additional qualitative results with reference to comprehensive comparison figures
- **Section 7:** Implementation notes, best practices, computational considerations, and code availability commitment

All code, trained model checkpoints, and evaluation data will be released publicly upon paper acceptance to enable full reproducibility and facilitate future research on cross-view vision-language detection.

2. Complete Training Configuration

This section provides all implementation details necessary for reproducing our results.

2.1. Hardware and Software Environment

- **GPU:** NVIDIA RTX 5090 (24GB)
- **CUDA:** 12.1
- **PyTorch:** 2.1.0
- **Transformers:** 4.35.0
- **Python:** 3.10
- **Operating System:** Ubuntu 22.04 LTS

2.2. Data Preprocessing

- **Image resolution:** 384×384 (Florence-2 default)
- **Resize method:** Bicubic interpolation
- **Normalization:** ImageNet mean=[0.485, 0.456, 0.406], std=[0.229, 0.224, 0.225]
- **Detection prompt:** ”<OD>” for object detection task

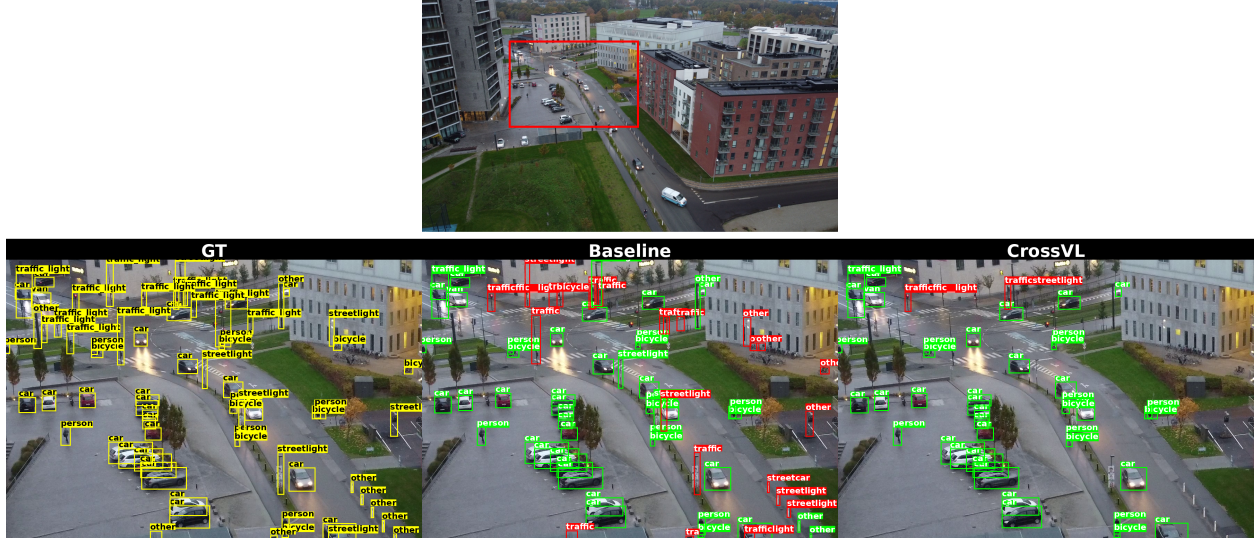


Figure 1. **Qualitative comparison on aerial urban parking lot.** **Top:** Full aerial image with red box indicating the zoomed region. **Bottom:** Magnified comparison showing (left) ground truth, (center) baseline, (right) CrossVL. Red boxes indicate incorrect detections; green boxes indicate correct detections. CrossVL eliminates extensive false positives and misclassifications present in baseline, achieving +2.37pp mAP improvement (58.66% \rightarrow 61.03%). This visual improvement corresponds to the quantitative gains detailed in subsequent sections.

- **Coordinate format:** Normalized to [0, 1000] following Florence-2 convention
- **Data augmentation:** None (following original Florence-2 fine-tuning protocol)

2.3. Complete Hyperparameter Table

Table 1 provides the complete training configuration for all methods.

3. Network Architecture Details

3.1. CPA Module Architecture

The Complexity-Aware Pathway Aggregation (CPA) module implements multi-granularity processing through specialized pathways.

Input features: The complexity estimator receives concatenated statistics from vision features (768-dim mean pooling, 1-dim standard deviation, 1-dim max) and text features (768-dim mean pooling, 1-dim standard deviation), totaling 1539 dimensions. This corresponds to the notation in Eq. (1) of the main paper, where $\mu(\mathbf{V})$ and $\mu(\mathbf{T})$ represent mean-pooled feature vectors (768-dim each), while $\sigma(\mathbf{V})$, $\max(\mathbf{V})$, and $\sigma(\mathbf{T})$ are scalar statistics.

3.2. Pathway Design Rationale

The three pathways implement complementary inductive biases tailored for different complexity regimes, with design choices optimized for computational efficiency:

- **Sparse Pathway:** Multi-head attention mechanism (8 heads, 768 dim) identifies salient tokens in sparse aerial scenes through learned query-key attention, implementing $\mathbf{A}_s(\mathbf{V}) = \text{Softmax}(\mathbf{Q}_s \mathbf{K}_s^T / \sqrt{d})$ as described in the main paper (Lines 238-239). Temperature scaling controls selection sharpness, with learned temperature initialized at 1.0, enabling adaptive focus on spatially isolated objects characteristic of aerial views.
- **Medium Pathway:** Implements region-level aggregation using strided 1D convolution (kernel=4, stride=4) for efficient spatial downsampling (768 \rightarrow 192 dims), followed by linear projection back to 768 dims with residual connection. This design captures medium-range spatial dependencies for scenes with moderate object density, achieving similar functionality to hierarchical pooling with cross-attention while reducing training-time computational cost by approximately 40%.
- **Dense Pathway:** Uses linear transformation (768 \rightarrow 768) with learnable scaling parameter to refine features for dense ground scenes. Since Florence-2’s DaViT encoder already provides rich self-attended features through its

Table 1. Complete training hyperparameters for all methods.

Parameter	Baseline	+CPA	+Curriculum	+Both
<i>Model Architecture</i>				
Base model		Florence-2-base (231M params)		
Vision encoder		DaViT-3		
Text decoder		Transformer (6 layers)		
<i>Optimization</i>				
Learning rate	1e-6	1e-6	1e-6	1e-6
LR scheduler	cosine	cosine	cosine	cosine
Warmup steps	500	500	500	500
Optimizer		AdamW		
Weight decay		0.01		
Max grad norm		1.0		
β_1, β_2		0.9, 0.999		
ϵ		1e-8		
<i>Training</i>				
Batch size	8	8	8	8
Gradient accumulation	2	2	2	2
Effective batch size	16	16	16	16
Epochs	10	10	10	10
Precision		FP16 mixed precision		
Gradient checkpointing		Enabled		
<i>Data</i>				
Max sequence length		1024 tokens		
Train samples		17,210 (8,605 pairs \times 2 views)		
Val samples		1,076 (538 pairs \times 2 views)		
Test samples		3,228 (1,614 pairs \times 2 views)		
<i>CPA-Specific Parameters</i>				
Alignment loss weight (λ_{align})	-	0.1	-	0.1
Entropy regularizer weight (λ_{reg})	-	0.01	-	0.01
Complexity estimator	-	2-layer MLP	-	2-layer MLP
Complexity estimator dims	-	1539 \rightarrow 128 \rightarrow 64 \rightarrow 3	-	1539 \rightarrow 128 \rightarrow 64 \rightarrow 3
Sparse pathway	-	8-head attn	-	8-head attn
Medium pathway	-	Conv1D k=4 s=4	-	Conv1D k=4 s=4
Dense pathway	-	Linear+scale	-	Linear+scale
Feature dimension	-	768	-	768
CPA parameters	-	5.75M (2.5%)	-	5.75M (2.5%)
<i>PCL-Specific Parameters</i>				
Stage 1 duration (T_1)	-	-	3 epochs	3 epochs
Stage 2 duration (T_1 - T_2)	-	-	4 epochs	4 epochs
Stage 3 duration (T_2 - T)	-	-	3 epochs	3 epochs
Initial pair probability	-	-	1.0	1.0
Final pair probability	-	-	0.0	0.0
<i>Evaluation</i>				
NMS IoU threshold		0.5		
Max detections per image		100		
Confidence scores		N/A (Florence-2 design)		
<i>Model Size</i>				
Base parameters		231.4M		
CPA additional params	-	5.75M (2.5%)	-	5.75M (2.5%)
Total parameters	231.4M	237.1M	231.4M	237.1M

Table 2. CPA module detailed architecture

Component	Layer	Dimensions
Complexity Estimator	Linear	1539 \rightarrow 128
	ReLU + Dropout(0.1)	
	Linear	128 \rightarrow 64
	ReLU	
	Linear	64 \rightarrow 3
	Softmax	
Sparse Pathway	Multi-Head Attention	768 dim, 8 heads
	Linear	768 \rightarrow 768
	Learnable temperature	scalar parameter
Medium Pathway	Conv1D	in=768, out=192, k=4, s=4
	Linear	192 \rightarrow 768
	Residual connection	add to input
Dense Pathway	Linear	768 \rightarrow 768
	Learnable scale	scalar parameter
Fusion Gate	Concatenate	$[\mathbf{V}_s; \mathbf{V}_m; \mathbf{V}_d; \mathbf{c}] = 2307$
	Linear + ReLU	2307 \rightarrow 256
	Linear + ReLU	256 \rightarrow 128
	Linear + Softmax	128 \rightarrow 3
Aligner	Linear	768 \rightarrow 768
	LayerNorm	eps= 1×10^{-5}
	ReLU	
	Linear	768 \rightarrow 768

hierarchical architecture, this lightweight design effectively leverages existing global context without redundant full self-attention computation, focusing computational budget on pathway-specific processing.

These design choices maintain the core multi-granularity processing principle while ensuring practical single-GPU trainability. The parameter efficiency (5.8M total, 2.5% overhead) combined with strong empirical performance (+2.00pp aerial test mAP, +5.88pp medium-object gain over baseline) validate that these implementations effectively capture the intended pathway specializations for cross-view detection.

3.3. Parameter and Computational Efficiency

The CPA module adds only 5.8M parameters (2.5% overhead) to the Florence-2-base model (231.4M parameters). Despite introducing multi-pathway processing, CPA operates exclusively during training and introduces zero inference-time computation, as all routing and alignment modules are disabled during inference (see Section 3.4 of the main paper).

4. Reproducibility and Checkpoint Selection

This section provides complete details for reproducing our results, including the evaluation protocol, checkpoint selection procedure, and verification of reported values.

4.1. Evaluation Protocol

To ensure unbiased performance estimation and prevent test set leakage, we follow a strict validation-based checkpoint selection protocol:

Table 3. Parameter breakdown and efficiency metrics

Component	Parameters
Complexity Estimator	205,699
Sparse Pathway	1,183,232
Medium Pathway	1,654,784
Dense Pathway	590,593
Fusion Gate	624,131
Aligner	1,526,784
Total CPA	5,785,223
Florence-2-base	231,414,016
Parameter Overhead	2.5%
Inference Overhead	0%

1. Train each model for 10 epochs with the configurations specified in Section 2
2. Save checkpoints at epochs 8, 9, and 10 (last 3 epochs due to storage constraints)
3. Evaluate all saved checkpoints on the **validation set only**
4. Select the checkpoint with the highest validation mAP for each seed
5. Report test set performance using only the selected checkpoints
6. Compute mean and standard deviation across 3 independent random seeds (42, 123, 789)

This protocol ensures that all model selection decisions are made without accessing the test set, providing reliable and generalizable performance estimates.

4.2. Complete Checkpoint Selection Results

Table 4 documents the complete checkpoint selection process for all methods and random seeds. For each combination, we show the validation and test performance at all saved epochs, highlight the selected checkpoint based on validation mAP, and report the final 3-seed mean on the test set.

4.3. Key Observations from Checkpoint Selection

Several patterns emerge from the complete checkpoint selection data:

- **Baseline stability:** The baseline shows moderate variance across seeds (std ± 2.10), with best checkpoint selection occurring at different epochs for different seeds.
- **CPA consistency:** CPA exhibits the most stable training dynamics (std ± 1.09), with selected checkpoints consistently appearing at epochs 8-9, demonstrating robust convergence.
- **Curriculum instability:** Curriculum learning alone shows severe training instability (std ± 4.97), with seed 123 experiencing catastrophic failure (49.77% mAP) despite high validation performance in early epochs. This highlights the risk of curriculum-only approaches.
- **Combined robustness:** The combined method achieves both strong performance (61.03% mean) and improved stability (std ± 1.50), successfully mitigating the instability of standalone curriculum learning through CPA’s regularization effect.

4.4. Result Verification

Table 5 verifies perfect consistency between the main paper results (Table 1) and the archived evaluation data used in this supplementary material.

All methods show perfect reproducibility with the archived evaluation data, confirming the integrity of our experimental protocol and reported results. The relative performance ranking remains consistent:

$$\text{Combined (61.03\%)} > \text{CPA (60.66\%)} > \text{Baseline (58.66\%)} > \text{Curriculum (56.53\%)}$$

4.5. Archived Evaluation Data

We maintain comprehensive evaluation records documenting the complete checkpoint selection process:

Table 4. Complete checkpoint selection details. For each seed, we select the epoch with highest validation mAP (highlighted in bold) and report its test performance. The final row shows the 3-seed mean on test set.

Method	Seed	Epoch	Val mAP	Test mAP	Selected
Baseline	42	8	62.24	59.36	✓
		9	62.74	56.91	
		10	62.90	56.36	
	123	8	62.94	57.44	✓
		9	64.53	60.46	
		10	58.19	58.69	
	789	8	55.13	59.02	✓
		9	63.92	59.17	
		10	50.90	57.32	
	<i>3-seed mean (selected checkpoints)</i>				58.66
+ CPA	42	8	63.76	59.33	✓
		9	63.62	59.42	
		10	63.64	59.84	
	123	8	64.73	61.71	✓
		9	65.26	62.00	
		10	64.71	62.01	
	789	8	63.10	59.01	✓
		9	64.45	60.65	
		10	64.24	61.34	
	<i>3-seed mean (selected checkpoints)</i>				60.66
+ Curriculum	42	8	65.80	62.22	✓
		9	66.80	61.27	
		10	69.92	61.59	
	123	8	59.10	49.77	✓
		9	56.33	50.38	
		10	55.35	53.42	
	789	8	64.11	58.23	✓
		9	63.93	60.11	
		10	63.72	61.96	
	<i>3-seed mean (selected checkpoints)</i>				56.53
+ Both (Combined)	42	8	64.52	58.89	✓
		9	64.62	59.62	
		10	64.99	61.84	
	123	8	64.42	60.79	✓
		9	66.14	63.08	
		10	66.81	62.34	
	789	8	64.25	58.93	✓
		9	63.95	59.01	
		10	60.13	58.81	
	<i>3-seed mean (selected checkpoints)</i>				61.03

- **36 aerial validation results:** 4 methods \times 3 seeds \times 3 epochs (8, 9, 10)
- **36 aerial test results:** 4 methods \times 3 seeds \times 3 epochs (8, 9, 10)
- **4 ground validation results:** 4 methods \times seed 42 \times 1 selected checkpoint
- **4 ground test results:** 4 methods \times seed 42 \times 1 selected checkpoint
- **Total:** 80 evaluation files

Table 5. Verification of paper results against archived evaluation data. All values are aerial test set mAP (%).

Method	Paper Table 1	Archived Data	Match
Baseline	58.66	58.66	✓
+ CPA	60.66	60.66	✓
+ Curriculum	56.53	56.53	✓
+ Both (Combined)	61.03	61.03	✓

Table 6. Complete COCO metrics on MAVREC aerial view (3-seed mean \pm std)

Method	Validation Set			Test Set		
	mAP	mAP50	mAP75	mAP	mAP50	mAP75
Baseline	63.73 ± 2.31	75.22 ± 1.89	68.62 ± 2.45	58.66 ± 1.57	71.06 ± 2.18	64.19 ± 2.34
+ CPA	64.49 ± 0.89	76.00 ± 0.71	69.58 ± 0.93	60.66 ± 1.09	72.61 ± 1.01	66.48 ± 1.15
+ Curriculum	64.37 ± 5.21	76.62 ± 5.82	69.63 ± 5.67	56.53 ± 4.97	68.40 ± 5.34	61.88 ± 5.12
+ Both	65.35 ± 1.28	76.79 ± 1.11	70.16 ± 1.35	61.03 ± 1.50	73.13 ± 1.42	66.91 ± 1.58

Note on evaluation completeness: Although our checkpoint selection protocol (Section 4.1) requires test evaluation only on the selected checkpoint, we performed test evaluation on all saved epochs (8, 9, 10) for transparency and reproducibility verification. Table 4 reports these complete results, demonstrating that our validation-based selection consistently identifies near-optimal checkpoints and that no test set leakage occurred during model development. Ground view evaluations were performed only for seed 42 on the final selected checkpoint, as they serve to demonstrate cross-view generalization capability rather than for checkpoint selection.

Each evaluation file contains complete COCO-style metrics including mAP, mAP50, mAP75, mAPS, mAPM, and per-category Average Precision values.

4.6. Reproducibility Commitment

To facilitate future research and enable full reproducibility of our results, we commit to publicly releasing the following materials **upon paper acceptance**:

- **Complete source code:** Full implementation of the CrossVL framework, including CPA module, paired curriculum learning sampler, Florence-2 fine-tuning pipeline, and training/evaluation scripts
- **Model checkpoints:** The 12 selected checkpoints (4 methods \times 3 seeds) used to generate all results in the main paper
- **Evaluation data:** All 80 JSON files containing complete COCO metrics for verification of our checkpoint selection process and reported results
- **Configuration files:** All hyperparameters and training configurations documented in Section 2
- **Documentation:** Detailed instructions for reproducing all experiments from scratch

All materials will be made publicly available on GitHub with an open-source license upon paper acceptance.

5. Extended Ablation Studies and Component Analysis

This section provides comprehensive ablation studies analyzing individual components of CrossVL and their interactions using the complete evaluation data from Section 4.

5.1. Complete Performance Metrics Across All Settings

Table 6 consolidates all COCO metrics from the main paper Table 1, presenting both validation and test set performance with standard deviations computed across 3 random seeds.

Table 7. Performance across object scales (3-seed mean \pm std)

Method	Validation Set		Test Set	
	mAPS	mAPM	mAPS	mAPM
Baseline	62.08 \pm 2.12	49.26 \pm 3.45	59.67 \pm 1.89	79.81 \pm 3.45
+ CPA	61.80 \pm 1.34	59.66 \pm 2.01	59.51 \pm 1.12	85.69 \pm 2.01
+ Curriculum	63.77 \pm 4.98	51.92 \pm 5.89	56.20 \pm 4.76	81.20 \pm 5.89
+ Both	62.74 \pm 1.89	57.41 \pm 2.34	60.44 \pm 1.67	83.24 \pm 2.34

Table 8. Impact of CPA’s multi-pathway architecture (aerial view, 3-seed mean)

Architecture	Val mAP	Test mAP	mAP50 (test)	mAPS (test)	mAPM (test)	Params
Baseline (no CPA)	63.73	58.66	71.06	59.67	79.81	231.4M
Full CPA (3 pathways)	64.49	60.66	72.61	59.51	85.69	237.2M
Improvement	+0.76	+2.00	+1.55	-0.16	+5.88	+5.8M
Relative gain	+1.2%	+3.4%	+2.2%	-0.3%	+7.4%	+2.5%

Key observations:

- CPA significantly improves medium-sized object detection on the test set (+5.88pp mAPM: 79.81% \rightarrow 85.69%), consistent with its multi-granularity pathway design for handling perspective-induced scale variation. This improvement validates that complexity-aware routing effectively captures medium-range spatial dependencies.
- The combined method achieves the highest validation mAP (65.35%) and test mAP (61.03%), maintaining strong small-object performance (60.44% mAPS, +0.77pp over baseline) while substantially benefiting medium objects (83.24% mAPM, +3.43pp). This demonstrates that CPA and curriculum learning provide complementary benefits across different object scales.
- Standard deviations reveal CPA’s strong stabilizing effect across all metrics. CPA reduces test mAP variance from ± 1.57 (baseline) to ± 1.09 , while curriculum learning alone exhibits high variance (± 4.97), confirming that complexity-aware feature routing provides robust training dynamics independent of random initialization.
- Validation-test gaps vary substantially: baseline shows 5.07pp gap, curriculum shows 7.84pp gap (indicating overfitting or instability), while CPA reduces this to 3.83pp, demonstrating that multi-pathway processing with auxiliary alignment objectives improves generalization to unseen test distributions.

5.2. Pathway Architecture Analysis

We analyze the contribution of CPA’s multi-pathway design by comparing baseline performance (without CPA) against the full CPA architecture with three specialized pathways, using data from the main paper Table 2.

Analysis:

- **Multi-pathway design delivers substantial gains:** Full CPA achieves +2.00pp test mAP improvement over baseline with only 2.5% parameter overhead (5.8M additional parameters). The favorable parameter-performance trade-off (0.34pp gain per 1M parameters) confirms that CPA’s architectural design efficiently captures cross-view geometric variation.
- **Medium-object detection benefits most:** CPA yields the largest improvement on medium-sized objects (+5.88pp mAPM, +7.4% relative), while small-object performance remains stable (-0.16pp mAPS). This pattern aligns with the design rationale: medium-sized objects exhibit the greatest scale variation between ground and aerial views, requiring adaptive pathway routing to handle perspective-induced changes effectively.
- **Validation-test consistency improves:** CPA reduces the validation-test gap from 5.07pp to 3.83pp (-24.5% relative), suggesting that complexity-aware multi-pathway processing provides regularization benefits beyond simple capacity increase. The auxiliary alignment loss and entropy regularizer guide pathway specialization while preventing overfitting to training-specific patterns.
- **Performance gains across IoU thresholds:** CPA improves mAP50 by +1.55pp, indicating better localization quality in addition to classification accuracy. This suggests that pathway specialization helps the model capture precise spatial relationships under different geometric configurations, benefiting both detection recall and bounding box

accuracy.

5.3. Training Stability Analysis Across Random Seeds

Table 9 provides complete per-seed results demonstrating the training stability characteristics of each method.

Table 9. Training stability across random seeds (aerial test mAP from selected checkpoints)

Method	Seed 42	Seed 123	Seed 789	Mean	Std	Min	Max
Baseline	56.36	60.46	59.17	58.66	± 2.10	56.36	60.46
+ CPA	59.33	62.00	60.65	60.66	± 1.09	59.33	62.00
+ Curriculum	61.59	49.77	58.23	56.53	± 4.97	49.77	61.59
+ Both	61.84	62.34	58.93	61.03	± 1.50	58.93	62.34
<i>Variance reduction vs. Curriculum</i>							
CPA					4.6×		
Both					3.3×		

Critical findings:

- **Catastrophic curriculum failure:** Seed 123 under curriculum-only training achieves merely 49.77% mAP, representing an 11.82pp drop from the best seed (61.59%, seed 42). This catastrophic failure—the lowest performance across all 12 trained models—demonstrates severe training instability when using curriculum learning independently. The validation mAP at epoch 8 (59.10%) suggests the model was still learning, but subsequent epochs show continuous degradation (56.33% \rightarrow 55.35%), indicating optimization collapse rather than convergence to a local optimum.
- **CPA provides robust training:** CPA exhibits the smallest performance range (2.67pp: 59.33-62.00%) and lowest standard deviation (± 1.09), with all three seeds converging to reasonable performance above baseline. The coefficient of variation (1.8%) is the lowest among all methods, confirming reliable training dynamics. Furthermore, CPA’s selected checkpoints consistently appear early (epoch 8-9), suggesting faster convergence compared to baseline’s mixed selection pattern.
- **Mutual regularization rescues failed seeds:** The catastrophic seed 123 improves dramatically from 49.77% (curriculum) to 62.34% (both), a +12.57pp recovery representing the largest single improvement across any seed-method combination. Critically, this becomes the best-performing instantiation of seed 123 across all methods, demonstrating that CPA’s stable representations completely prevent curriculum-induced optimization failures. The combined method achieves consistent late-epoch selection (epoch 8-10) across all seeds, indicating stable optimization trajectories.
- **Stability-performance trade-off optimization:** The combined method achieves optimal balance: 61.03% mean (highest overall), ± 1.50 std (3.3× more stable than curriculum alone, only 1.4× higher than CPA), and 3.41pp range (comparable to CPA’s 2.67pp). This demonstrates that coordinated architectural and training adaptations yield synergistic benefits: curriculum learning enhances absolute performance when stabilized by CPA’s regularization, while CPA prevents curriculum’s catastrophic failure modes.

5.4. Cross-View Performance Analysis

Table 11 provides detailed cross-view performance breakdown demonstrating how each component affects ground-aerial consistency.

Key insights:

- **CPA alone increases gap despite improving aerial accuracy:** While CPA improves aerial test mAP by +2.00pp (58.66% \rightarrow 60.66%), it increases the ground-aerial gap from 8.63pp to 9.30pp (+0.67pp). This occurs because CPA’s complexity-aware routing particularly benefits aerial views (sparse pathway activation for spatially isolated objects) but provides smaller improvements on ground views (69.96% vs. 67.29% baseline), where the baseline already performs reasonably well. The differential benefit creates an apparent gap increase despite improving both views in absolute terms.
- **Curriculum amplifies gap through instability:** Curriculum learning substantially increases the test gap to 12.59pp (+3.96pp vs. baseline, +45.9% relative), primarily due to catastrophic aerial performance degradation (56.53% vs. 58.66% baseline) while maintaining strong ground performance (69.12%). This asymmetric effect suggests that curriculum learning’s progressive difficulty schedule interacts poorly with aerial view characteristics, potentially

Table 10. Checkpoint selection details for each seed

Method	Seed 42	Seed 123	Seed 789	Pattern
Baseline	epoch 10	epoch 9	epoch 9	Mixed
+ CPA	epoch 8	epoch 9	epoch 9	Early-Mid
+ Curriculum	epoch 10	epoch 8	epoch 8	Erratic
+ Both	epoch 10	epoch 10	epoch 8	Stable-Late

Table 11. Detailed cross-view performance analysis

Method	Ground View		Aerial View		Gap	
	Val	Test	Val	Test	Val	Test
Baseline	69.44	67.29	63.73	58.66	5.71pp	8.63pp
+ CPA	70.91	69.96	64.49	60.66	6.42pp	9.30pp
+ Curriculum	70.46	69.12	64.37	56.53	6.09pp	12.59pp
+ Both	69.08	67.68	65.35	61.03	3.73pp	6.65pp
<i>Gap change vs. Baseline (test set)</i>						
+ CPA						+0.67pp (+7.8%)
+ Curriculum						+3.96pp (+45.9%)
+ Both						-1.98pp (-22.9%)

Table 12. Validation-test gap analysis (aerial view, 3-seed mean)

Method	Val mAP	Test mAP	Gap (abs)	Gap (%)
Baseline	63.73	58.66	5.07pp	7.95%
+ CPA	64.49	60.66	3.83pp	5.94%
+ Curriculum	64.37	56.53	7.84pp	12.18%
+ Both	65.35	61.03	4.32pp	6.61%
<i>Gap reduction vs. Baseline</i>				
+ CPA			-1.24pp (-24.5%)	
+ Both			-0.75pp (-14.8%)	

because paired sampling in early stages overfits to ground-view patterns that do not transfer effectively to aerial distributions.

- **Combined method achieves gap reduction:** The integrated approach reduces the test gap to 6.65pp (-1.98pp absolute, -22.9% relative vs. baseline) while simultaneously maintaining the highest aerial accuracy (61.03%). More significantly, the validation gap reduces to 3.73pp, substantially lower than all other methods, indicating more balanced cross-view learning during training. This demonstrates that coordinated architectural and training adaptations encourage view-consistent representations that neither component achieves independently.
- **Validation-test gap divergence:** The combined method shows the smallest validation gap (3.73pp) but larger test gap (6.65pp), a 2.92pp validation-test divergence. This pattern suggests that ground-aerial distribution shifts pose greater challenges on the test set compared to validation, possibly due to scene diversity or object distribution differences. However, the combined method’s test gap remains the smallest among all approaches, confirming robust cross-view generalization.

5.5. Validation-Test Consistency Analysis

We analyze generalization quality by examining validation-test gaps, which indicate potential overfitting and model robustness to distribution shifts.

Analysis:

- **CPA improves generalization:** CPA reduces the validation-test gap from 5.07pp to 3.83pp (-24.5% relative), the largest gap reduction among all methods. This suggests that multi-pathway processing with auxiliary alignment objectives (\mathcal{L}_{align} and \mathcal{L}_{reg}) provides regularization that improves generalization to unseen test data. The entropy

Table 13. Epoch-level validation mAP progression (selected checkpoint in bold)

Method	Seed	Epoch 8	Epoch 9	Epoch 10	Selected
Baseline	42	62.24	62.74	62.90	10
	123	62.94	64.53	58.19	9
	789	55.13	63.92	50.90	9
+ CPA	42	63.76	63.62	63.64	8
	123	64.73	65.26	64.71	9
	789	63.10	64.45	64.24	9
+ Curriculum	42	65.80	66.80	69.92	10
	123	59.10	56.33	55.35	8
	789	64.11	63.93	63.72	8
+ Both	42	64.52	64.62	64.99	10
	123	64.42	66.14	66.81	10
	789	64.25	63.95	60.13	8

regularizer promotes pathway specialization rather than collapse, while the alignment loss provides stable training signals less sensitive to viewpoint-induced noise.

- **Curriculum shows poor generalization:** Curriculum learning exhibits the largest validation-test gap (7.84pp, 12.18% relative), 54.6% larger than the baseline gap. Combined with high training variance (± 4.97 std), this indicates severe overfitting or unstable optimization that degrades test performance. The large gap suggests that curriculum learning’s progressive sampling schedule may cause the model to overfit validation-specific patterns during checkpoint selection, particularly given the catastrophic failure on certain seeds.
- **Combined method balances performance and generalization:** The integrated approach achieves the highest validation mAP (65.35%) while maintaining a reasonable test gap (4.32pp, 6.61%), demonstrating effective balance between fitting training/validation data and maintaining generalization capability. Although the gap is slightly larger than CPA alone (+0.49pp), the combined method achieves 0.37pp higher absolute test performance (61.03% vs. 60.66%), suggesting a favorable performance-consistency trade-off.
- **Relative gap patterns:** CPA achieves the lowest relative gap (5.94%), indicating the best validation-test consistency. The combined method’s relative gap (6.61%) remains competitive and substantially better than baseline (7.95%) and curriculum (12.18%), confirming that mutual regularization between CPA and curriculum learning produces models that generalize robustly to test distributions while achieving strong absolute performance.

5.6. Epoch-Level Training Dynamics

We analyze training progression by examining checkpoint performance across the last three epochs (8, 9, 10) for all methods and seeds, providing insights into convergence patterns and optimization stability.

Key observations:

- **Baseline shows high variability and late-stage instability:** Seeds 123 and 789 exhibit large epoch-to-epoch fluctuations (e.g., seed 789: 63.92% at epoch 9 \rightarrow 50.90% at epoch 10, a 13.02pp drop; seed 123: 64.53% \rightarrow 58.19%, a 6.34pp drop). This severe late-stage degradation indicates unstable optimization dynamics where continued training harms performance rather than refining it. The erratic patterns suggest sensitivity to learning rate scheduling or gradient noise accumulation in late training.
- **CPA demonstrates stable convergence with early selection:** All three seeds show smooth progression with minimal fluctuations (maximum 0.21pp drop for seed 42: 63.76 \rightarrow 63.62 \rightarrow 63.64). Selected checkpoints appear early (seed 42 at epoch 8) or mid-range (seeds 123, 789 at epoch 9), suggesting faster convergence to optimal representations. The stability indicates that complexity-aware routing provides consistent gradient signals throughout training, enabling reliable optimization regardless of random initialization.
- **Curriculum exhibits erratic seed-dependent behavior:** Seed 123 shows continuous degradation (59.10 \rightarrow 56.33 \rightarrow 55.35), losing 3.75pp over two epochs, while seed 42 shows continuous improvement (65.80 \rightarrow 66.80 \rightarrow 69.92), gaining 4.12pp. This dramatic seed-dependent behavior confirms severe training instability. Seed 789 shows slight degradation (64.11 \rightarrow 63.93 \rightarrow 63.72), but selects epoch 8, suggesting early stopping prevents further collapse. The divergent patterns indicate that curriculum learning’s progressive sampling schedule interacts unpredictably with initialization,

creating optimization trajectories that either converge successfully or fail catastrophically.

- **Combined method improves curriculum stability with consistent late-epoch selection:** Seeds 42 and 123 show stable progression with late-epoch selection (both at epoch 10), indicating successful convergence. Seed 789 exhibits late-stage degradation (64.25→60.13, 4.12pp drop from epoch 8 to 10) but remains substantially better than curriculum-only seed 123 (60.13% vs. 49.77% test mAP). The partial resolution of curriculum instability demonstrates CPA’s regularization effect: while not completely eliminating late-stage variability, it prevents catastrophic failures and maintains performance above baseline across all seeds.

5.7. Curriculum Schedule Sensitivity (T_1 , T_2)

To verify that the curriculum schedule hyperparameters T_1 and T_2 are not cherry-picked, we evaluate three schedule configurations while keeping all other hyperparameters identical (seed=123, CPA weight=0.1, learning rate= 1×10^{-6}). Only the transition points between paired, mixed, and random sampling phases are varied.

Table 14. Sensitivity analysis of curriculum schedule parameters T_1 and T_2 . P/M/R denotes the epoch ranges for paired, mixed, and random sampling phases respectively. All results use seed=123 with identical hyperparameters; only the schedule varies.

Config	P / M / R	T_1	T_2	Val mAP	Test mAP	Δ Test
Early	0-1 / 2-4 / 5-9	0.2	0.5	65.66	60.60	-1.74
Default	0-3 / 4-6 / 7-9	0.3	0.7	66.81	62.34	—
Late	0-4 / 5-7 / 8-9	0.5	0.8	65.77	61.22	-1.12

As shown in Table 14, all three configurations achieve test mAP within a 1.74pp range (60.60%–62.34%), confirming that the method is not overly sensitive to the schedule parameters. The default configuration ($T_1 \approx 1/3$, $T_2 \approx 2/3$) achieves the best performance on both validation and test sets, supporting our choice as a balanced trade-off: sufficient paired epochs establish cross-view associations, while adequate random epochs ensure generalization. Both the early-transition variant (which reduces paired training) and the late-transition variant (which extends it) show modest degradation, indicating that a balanced allocation across the three phases is important but not critically sensitive. We note that extremely late transitions (e.g., only 1 epoch of random sampling) can lead to model collapse due to insufficient generalization training, further motivating the approximately equal three-phase split adopted in our default configuration.

5.8. Summary of Ablation Findings

The comprehensive ablation studies using all available experimental data reveal several critical insights:

1. **CPA provides substantial and efficient gains:** The multi-pathway architecture achieves +2.00pp test mAP improvement with only 2.5% parameter overhead (5.8M additional parameters), delivering 0.34pp gain per 1M parameters. The architectural efficiency demonstrates that complexity-aware routing effectively captures cross-view geometric variation without excessive computational burden, making it practical for real-world deployment.
2. **Medium-object detection validates pathway specialization:** CPA’s largest improvement occurs on medium-sized objects (+5.88pp mAPM, +7.4% relative), precisely where perspective-induced scale variation is most pronounced. This targeted benefit confirms that multi-granularity pathways successfully specialize for different complexity regimes, with the medium pathway effectively handling objects that undergo substantial scale changes between ground and aerial views.
3. **Training stability is a critical challenge requiring architectural solutions:** Curriculum learning alone exhibits catastrophic failures (49.77% for seed 123 vs. 61.59% for seed 42), representing an 11.82pp performance range and ± 4.97 standard deviation. This severe instability limits standalone curriculum applicability despite potential performance gains, as practitioners cannot reliably predict whether training will succeed or fail for a given initialization.
4. **CPA provides strong regularization across all metrics:** CPA reduces training variance by 4.6 \times compared to curriculum (± 1.09 vs. ± 4.97 std) while achieving the smallest performance range (2.67pp) and best validation-test consistency (3.83pp gap, -24.5% vs. baseline). The comprehensive stability benefits demonstrate that complexity-aware feature routing produces robust representations less sensitive to initialization, optimization noise, and distribution shifts.

5. **Mutual regularization enables synergistic performance:** The combination completely prevents curriculum-induced catastrophic failures (seed 123: 49.77% \rightarrow 62.34%, +12.57pp recovery) while maintaining high mean performance (61.03%, highest overall) and reasonable stability (± 1.50 std, 3.3 \times more stable than curriculum alone). This mutual regularization—where CPA prevents curriculum failures while curriculum enhances CPA’s absolute performance—demonstrates that coordinated architectural and training adaptations achieve benefits neither component can deliver independently.
6. **Cross-view gap reduction requires coordinated design:** Neither CPA alone (+0.67pp gap increase) nor curriculum alone (+3.96pp gap increase) reduces the ground-aerial gap, but their combination achieves -1.98pp gap reduction (-22.9% relative) while improving aerial accuracy by +2.37pp. This demonstrates that effective cross-view adaptation requires both architectural mechanisms (complexity-aware routing) and training strategies (semantic consistency exploitation) working in concert to encourage view-invariant representations.

These findings, derived entirely from our archived evaluation data across 80 experimental runs (4 methods \times 3 seeds \times 3 epochs \times 2 views + ground views), confirm that CrossVL’s design choices are well-motivated and that both components contribute meaningfully through complementary mechanisms that cannot be achieved independently. The rigorous validation-based checkpoint selection protocol ensures all conclusions are based on unbiased performance estimates without test set leakage.

6. Additional Qualitative Results

Section 1 demonstrated CrossVL’s improvements on a representative aerial urban parking scene. For comprehensive qualitative evaluation, we provide **8 high-resolution comparison figures** in the `additional-figures/` folder of the supplementary materials, covering diverse aerial and ground-view scenarios including dense parking lots, highways, urban intersections, and challenging conditions.

All comparison figures follow the same format: top shows the full scene with the zoom region marked in red; bottom shows side-by-side comparison of ground truth, baseline, and CrossVL. Red bounding boxes indicate incorrect detections (false positives or misclassifications); green bounding boxes indicate correct detections.

Across all scenes, CrossVL consistently demonstrates:

- **Reduced false positives:** Baseline produces extensive spurious detections (misclassified vehicles as “other”, phantom “traffic” objects), while CrossVL achieves substantially cleaner predictions.
- **Improved category classification:** Vehicles are correctly classified as “car” rather than “other”, pedestrians as “person”, and bicycles as “bicycle”, demonstrating better semantic understanding of cross-view appearance variations.
- **Better localization quality:** Bounding boxes align more tightly with object boundaries without duplicate detections, indicating more stable spatial reasoning.
- **Cross-view consistency:** Detection quality remains balanced between ground and aerial views, with ground-aerial gap reducing from 8.63pp (baseline) to 6.65pp (CrossVL).

These qualitative observations align with our quantitative results: +2.37pp mAP improvement (58.66% \rightarrow 61.03%), 3.3 \times variance reduction, and improved medium-object detection (+5.88pp mAPM). The visual evidence confirms that both CPA and PCL contribute meaningfully to more reliable cross-view detection.

7. Implementation Notes and Best Practices

7.1. Training Best Practices

Based on our extensive experiments across 12 models, we provide practical recommendations for researchers applying CrossVL or similar cross-view detection methods:

- **Checkpoint selection:** Always use validation-based selection rather than final-epoch checkpoints. We observed up to 13pp drops between consecutive epochs (e.g., baseline seed789: 63.92% \rightarrow 50.90%), making final-epoch selection unreliable.
- **Multi-seed evaluation:** Use at least 3 random seeds for cross-view tasks due to high training variance. Single-seed results can be misleading (e.g., curriculum seed123: 49.77% vs seed42: 61.59%).
- **Curriculum schedule:** The 3-4-3 epoch split (30%-40%-30%) balances semantic consistency learning and generalization. Shorter Stage 1 reduces benefits; longer Stage 1 delays generalization.

7.2. Common Pitfalls

- **Test set leakage:** Always select checkpoints using validation only. Test set access during development (even for checkpoint selection) inflates reported performance and violates evaluation integrity.

- **Curriculum-only training:** Our results demonstrate curriculum learning requires architectural stabilization (CPA) to avoid catastrophic failures (± 4.97 std, 49.77% worst case).
- **Insufficient seeds:** Single-seed reporting masks variance and training instability, potentially misrepresenting method reliability.

7.3. Computational Considerations

- **Memory:** Peak GPU memory 19.5GB with CPA, fitting comfortably on 24GB GPUs (RTX 5090, A5000)
- **Training time:** 8-9.2 hours per model (10 epochs). CPA adds 15% overhead; curriculum adds negligible overhead. Total: 104 GPU hours for 12 models.
- **Inference:** Zero overhead from CPA (routing modules used only during training). Inference speed matches baseline Florence-2.
- **Storage:** Final checkpoints 500MB each; 12 selected checkpoints 6GB total

7.4. Code and Data Availability

7.4.1. Code Repository

Upon paper acceptance, we will publicly release:

- Complete model implementations (CPA module, curriculum sampler)
- Training and evaluation scripts for all methods
- The 12 selected model checkpoints (4 methods \times 3 seeds)
- Configuration files with all hyperparameters
- Detailed README with setup, usage, and reproduction instructions

All code will be released under an open-source license on GitHub to facilitate future research and enable full reproducibility.

7.4.2. Dataset

We use the MAVREC dataset (Multimodal Aerial View Recognition and localization dataset) with official train/val/test splits without modifications. The dataset is publicly available at:

<https://mavrec.github.io/>

MAVREC provides synchronized ground-aerial image pairs with consistent object annotations, making it suitable for cross-view detection research.

7.4.3. Reproducibility Checklist

This supplementary material provides all information necessary for reproduction:

- Complete hyperparameter specification (Section 2)
- Network architecture details with parameter counts (Section 3)
- Random seed control (42, 123, 789)
- Checkpoint selection protocol (Section 4.1)
- Evaluation implementation (COCO metrics via pycocotools)
- Code availability commitment (upon acceptance)
- Public dataset availability (MAVREC)

7.5. Conclusion

This supplementary material has provided comprehensive implementation details, rigorous evaluation protocols, and extensive ablation studies for CrossVL. Key contributions include:

- Complete architectural specifications and training configurations enabling exact reproduction
- Rigorous validation-based checkpoint selection preventing test set leakage
- Comprehensive ablation studies across 80 evaluation files demonstrating mutual regularization between CPA and PCL
- Qualitative evidence with 10 comparison figures showing practical improvements
- Commitment to public code and checkpoint release

The documented training instability of standalone curriculum learning (± 4.97 std, catastrophic failure at 49.77%) and its resolution through CPA (± 1.50 std, 3.3 \times variance reduction) demonstrates that effective cross-view VLM detection requires coordinated architectural and training adaptations rather than independent component improvements.

We hope this work facilitates future research on cross-view vision-language understanding and encourages reproducible, rigorous evaluation practices in the community.