Online Resource 1: Experimental records

Compact Detection of Rebar Intersections and Tying States: Computational and Cross-Acquisition Trade-offs

Yuanxun Zheng, Qiling Guo, Xiangyang Chen, Junwei Shi, Yanyun Li, Yihan Chen

Correspondence: Yuanxun Zheng, Zhengzhou University; yxzheng@zzu.edu.cn

The International Journal of Advanced Manufacturing Technology

All 21 fixed-budget training runs and eight device configurations are complete. Results below are bound to the recorded acquisition manifests and checkpoint hashes.

1. Image pools and fixed acquisition boundaries

Pool	Image records	Bound targets	Unbound targets
Legacy archive	830	3286	3197
16 December 2025	159	1328	1324
17 December 2025	375	2017	1975

Sixteen later images lack valid annotations and are excluded by the frozen pre-inference rule. The legacy archive remains together in training. Capture dates define acquisition groups; they do not establish counts or identities of independent physical meshes.

Three annotators prepared the existing annotations using LabelImg, and two reviewers checked them. Labels for visually unclear wire states or occluded intersections were determined during review. The revised experiments retain these existing two-class box labels.

Holdout 2025-12-17 manifest SHA-256:

`7cf452dfaa6d34655baf1837bb2b346a6976046c9fd04f2a103cd2752fb83ef3`

Holdout 2025-12-16 manifest SHA-256:

`985d2ed0ad1cab427e2f3aa6930b8905db6cea6f65c5449fc02105166d0c46fa`

2. Training and inference settings

Setting	Value
Training input / duration	640 x 640 / 80 epochs
Physical batch / nominal batch	4 / 64, gradient accumulation
Initialization / precision	Random initialization / FP32; AMP disabled
Optimizer	SGD; initial LR 0.01; final LR factor 0.01; momentum 0.937; weight decay 0.0005
Warm-up	3 epochs; momentum 0.8; bias LR 0.1
Seeds / workers	0, 1, 2 for primary models / 0 data-loader workers
Augmentation	HSV 0.015/0.7/0.4; translation 0.1; scale 0.5; horizontal flip 0.5; mosaic 1.0
Disabled augmentation	Rotation, shear, perspective, vertical flip, mixup and copy-paste
Mosaic closing / checkpoint	No closing window / final epoch last.pt
AP evaluation	Batch 1; longest side 640; stride-aligned rectangular letterbox; FP32; confidence 0.001; NMS IoU 0.7; max 300 detections
Operational matching	Confidence 0.25; class-consistent one-to-one matching at IoU 0.50
Runtime input / software	640 x 640; unfused PyTorch eager; Python 3.9.21; PyTorch 2.0.0+cu117; Ultralytics 8.3.141
Hardware	GTX 1650 Ti 4 GB; Intel Core i7-10750H; Windows

The framework validation entry aliases the training pool for fixed-epoch operation. These validation values are not independent test results and do not select model weights. AP50-95 averages IoU 0.50 through 0.95 at increments of 0.05. Framework P/R at a test-optimized F1 point are not used for the operational P/R below.

3. Configuration keys and measured resource costs

Key	Configuration	Parameters	GFLOPs
baseline	YOLO11n	2,590,230	6.374838
full_two	Two-branch SATE-YOLO	1,838,830	4.163640
full_three	Three-branch SATE-YOLO	1,838,930	4.163923
budget_match	Parameter-matched YOLO11	1,813,650	4.875501
shuffle_dual	ShuffleNetV2-DualConv adaptation	2,583,434	6.656701
kernel_unweighted	All-ones kernel template	1,838,830	4.163640
minus_a	Restore standard downsampling	2,318,062	5.186001
minus_e	Restore standard detection head	2,109,294	5.355320
minus_w	Restore standard C3k2	1,840,734	4.163213
minus_t	Remove attention	1,838,630	4.161223

Operations use the PyTorch profiler at batch 1 and 640 x 640, including functional weighted convolutions. A multiply-accumulate counts as two operations where supported by the profiler. Normalization, pooling, activations and memory movement are not completely represented. The matched width is 26/128, selected by parameter proximity on a fixed grid without detection scores.

The ShuffleNetV2-DualConv model is the explicitly described local adaptation of Wang et al. (Applied Sciences 2025, DOI 10.3390/app152412898), not an official author reproduction. The weighted bottleneck, ADown, Triplet Attention and grouped head are attributed in the main manuscript.

4. Completed holdout evaluations

All values are percentages. AP50 and AP50-95 average both classes. Unbound P and R use the fixed operational threshold. Asterisk marks a resumed trajectory.

Holdout	Key	Seed	AP50	AP50-95	U P	U R
16 Dec	baseline	0	25.42	7.91	44.80	37.08
16 Dec	full_two*	0	23.25	7.33	41.63	35.12
16 Dec	shuffle_dual	0	25.44	7.75	42.65	36.18
16 Dec	baseline	1	23.04	7.18	37.92	37.46
16 Dec	full_two	1	25.00	7.80	44.13	38.07
16 Dec	baseline*	2	25.87	7.90	41.79	33.84
16 Dec	full_two	2	18.98	5.85	37.04	24.92
17 Dec	baseline	0	10.46	3.88	13.46	8.00
17 Dec	full_two	0	12.52	4.51	9.50	3.65
17 Dec	full_three	0	14.84	5.50	9.27	5.11
17 Dec	budget_match*	0	10.65	3.86	8.34	3.54
17 Dec	shuffle_dual	0	6.54	2.27	5.83	2.53
17 Dec	kernel_unweighted	0	7.01	2.30	6.39	4.15

Holdout	Key	Seed	AP50	AP50-95	U P	U R
17 Dec	minus_a	0	10.43	3.82	11.05	5.72
17 Dec	minus_e	0	13.56	5.36	11.34	6.99
17 Dec	minus_w	0	12.54	4.78	10.91	5.06
17 Dec	minus_t	0	17.21	7.03	17.43	6.33
17 Dec	baseline	1	10.33	3.77	11.49	4.91
17 Dec	full_two	1	15.55	5.79	5.71	2.23
17 Dec	baseline	2	9.26	3.23	7.93	5.01
17 Dec	full_two	2	11.59	4.02	6.70	4.46

The parameter-matched control resumed after epoch 69; the two-branch reverse-fold seed-zero run resumed after epoch 10; the YOLO11n reverse-fold seed-two run resumed after epoch 79. EMA and optimizer state were restored. Random-number and data-loader states were not archived, so these are identified resumed trajectories.

Primary model comparisons are summarized over seeds only after both models have completed all three seeds for the acquisition fold. Seed variation does not estimate variation between independent meshes.

5. Primary comparisons over paired training seeds

Holdout 2025-12-17: mean +/- sample SD over seeds 0, 1 and 2. All values are percentages.

Model	AP50	AP50-95	Bound AP	Unbound AP	U P	U R
baseline	10.02 +/- 0.66	3.62 +/- 0.35	9.34 +/- 1.00	10.69 +/- 1.83	10.96 +/- 2.80	5.97 +/- 1.75
full_two	13.22 +/- 2.07	4.77 +/- 0.91	13.55 +/- 2.50	12.89 +/- 1.66	7.30 +/- 1.97	3.44 +/- 1.13

SATE minus YOLO11n (pp)	Seed 0	Seed 1	Seed 2	Mean	SD
mAP50	+2.06	+5.23	+2.33	+3.21	1.75
mAP50_95	+0.63	+2.02	+0.80	+1.15	0.76
bound_AP50	+4.67	+6.04	+1.93	+4.21	2.09
unbound_AP50	-0.54	+4.41	+2.73	+2.20	2.52
unbound_precision_025	-3.96	-5.79	-1.24	-3.66	2.29
unbound_recall_025	-4.35	-2.68	-0.56	-2.53	1.90

Holdout 2025-12-16: mean +/- sample SD over seeds 0, 1 and 2. All values are percentages.

Model	AP50	AP50-95	Bound AP	Unbound AP	U P	U R
baseline	24.78 +/- 1.52	7.66 +/- 0.42	24.03 +/- 2.54	25.52 +/- 2.48	41.50 +/- 3.45	36.13 +/- 1.99
full_two	22.41 +/- 3.10	6.99 +/- 1.02	20.82 +/- 2.52	24.00 +/- 3.76	40.93 +/- 3.60	32.70 +/- 6.90

SATE minus YOLO11n (pp)	Seed 0	Seed 1	Seed 2	Mean	SD
mAP50	-2.16	+1.96	-6.89	-2.37	4.43
mAP50_95	-0.58	+0.62	-2.05	-0.67	1.34
bound_AP50	-1.62	+0.71	-8.71	-3.21	4.90
unbound_AP50	-2.71	+3.21	-5.07	-1.52	4.26
unbound_precision_025	-3.17	+6.21	-4.75	-0.57	5.93
unbound_recall_025	-1.96	+0.60	-8.91	-3.42	4.92

These SDs describe training-seed variation within fixed acquisition groups, not uncertainty across independent physical meshes. AP integrates detection ranking; U P and U R use confidence 0.25 and one-to-one matching at IoU 0.50. All single-seed controls remain separate.

6. Checkpoint and implementation provenance

2025-12-16 / baseline / seed 0

`eaebb8e9f1b7eb1074e509b4bf80b4acca5d2a9aeff99eca27df18b620feeb83`

2025-12-16 / full_two / seed 0

`38c0aaad25150a72677cbcc6c3d141747fd240943fd0549dcca6e1cb0755eeb1`

2025-12-16 / shuffle_dual / seed 0

`4e50aa59fa3c54638d1511e40c69576de310d6ab958ea3b7bf379a195536fa91`

2025-12-16 / baseline / seed 1

`799cfe526e4cd54a93cdbee819cdffef054db9fb4fc0565b215e494b494913d1`

2025-12-16 / full_two / seed 1

`cd8f7cae5ccc110b049df9b61c186dfcc4e0b7fa1669ed4ea6586f6ce288c63c`

2025-12-16 / baseline / seed 2

`6510bac7b666f86ae1a412a1826ed90814a3c36fd9e16735ef23a5689846e226`

2025-12-16 / full_two / seed 2

`01338f6030a74e5718b8f6b1b974266f966eee3b3494fcfc5a461f6729569899`

2025-12-17 / baseline / seed 0

`11aa80171a364e229e5bcee641a3e52a13e5b0c1da696a27f7a354efd603a6c7`

2025-12-17 / budget_match / seed 0

`65cee71797c61562cb8a580b81656f3013d9a4a321b06cae9f3f8f8f797d8f23`

2025-12-17 / full_three / seed 0

`18345945d7ae57d01991a11ec4e0c4a3feba829b61ba6a63b2ca0dc420e2ec26`

2025-12-17 / full_two / seed 0

`a9bbf399977a00c2ddf3d6dbfea0b970304521f025bd2064b0ee9c21a1cbf17f`

2025-12-17 / kernel_unweighted / seed 0

`91d5ee297e801c2e6952c4f6190e284ba83fa8fb8107441d13e3fe3cc67f3726`

2025-12-17 / minus_a / seed 0

`7c8f72b9090889afc7d98e8a20626dcc18378829d2ce3807aa7f027ed60e88d2`

2025-12-17 / minus_e / seed 0

`c475c3db132e03864179311126ccb687fe4e712df4578a6075f72dd4491b5138`

2025-12-17 / minus_t / seed 0

`c6115b8c74a3870fa659556b506e7a084ab1cca137a510a7a99696d57cb2d21a`

2025-12-17 / minus_w / seed 0

`bf3d25e004ff6a114c79151b876402b669e54ece17ccf67f1417d87180cd7542`

2025-12-17 / shuffle_dual / seed 0

`21ccb9b2168b8df2bb4e4df8fd721850000c01d84223e2ba53043f7652768126`

2025-12-17 / baseline / seed 1

`a2d0555bb6d3f5d3a72cb75a0bfd2631e08fed6b45369d2161e76ffd10f4de4d`

2025-12-17 / full_two / seed 1

`530ac69eed1a1e5e67ce2729acd18ccd5f3204a6706f0aee1664573858c02599`

2025-12-17 / baseline / seed 2

`af93076059dcc2a27bb5eea6a7ded8d8a435cc5a1fd1dd5e94904edea0b1e6ea`

2025-12-17 / full_two / seed 2

`6004104a66d7aa6da5516356e689390e5af8f6390bd5232b56bcb9c36026ac51`

sate_training_runtime.py SHA-256:

`67a95882e575c6bab4e7f3189be4725944682579f3644d66d640f5a6d9471324`

wang_domain_control.py SHA-256:

`4565a987c294b8c7354d1ec2d1f7d838e4295f22cf1cdb916c7db7aeccd24f96`

Kernel folding replaces 22 weighted convolutions with ordinary convolutions after multiplying their learned weights by the fixed template. Three FP32 input trials produced a decoded-output maximum absolute difference of 0.0 in the checked runtime. The all-ones control preserves normalization and activation placement; restoring standard C3k2 changes that layout as well.

7. Device timings and memory

Latency starts with a decoded BGR frame in host memory and ends with native-image-coordinate detections in host memory. It includes letterboxing, tensor conversion, H2D transfer, prediction and decoding, NMS, coordinate scaling and D2H transfer. File/JPEG decoding, acquisition, robot actuation and calibrated spatial reconstruction are outside the boundary.

Batch one, FP32, unfused PyTorch eager on GTX 1650 Ti. Twenty warm-up calls precede each measurement boundary; each of three rounds records 100 model-forward calls and 90 pipeline calls on nine equally spaced indices from the frozen 17 December test list. Ranges below span round-level statistics, not confidence intervals. Models are timed serially; no concurrent research training runs during timing.

Configuration	Forward P50 ms	Forward P95 ms	Pipeline P50 ms	Pipeline P95 ms	Peak MiB
YOLO11n	12.66-13.98	16.14-18.49	17.11-19.93	20.33-32.67	81.99
Matched-width YOLO11	11.64-11.75	14.41-14.82	16.80-17.68	19.79-24.96	75.26
ShuffleNetV2-DualConv adaptation	13.48-15.30	16.86-19.73	18.55-19.84	23.43-24.59	103.88
SATE, two branches	11.18-11.35	13.24-15.32	16.39-16.80	18.66-20.21	73.62
SATE, folded	10.39-12.19	13.32-15.99	15.44-15.55	18.39-18.65	73.60
Three branches, folded	11.50-12.40	15.01-15.80	16.57-18.18	19.71-22.04	73.60
Without attention	11.17-11.67	13.32-14.68	16.91-17.35	19.03-20.04	78.65
Without attention, folded	10.93-13.52	13.12-15.42	16.81-17.06	19.25-22.84	73.59

Configuration	Above 33.33 ms (%)	Above 16.67 ms (%)	Above 11.11 ms (%)
YOLO11n	0.00-4.44	71.11-87.78	100.00-100.00
Matched-width YOLO11	0.00-0.00	53.33-76.67	100.00-100.00
ShuffleNetV2-DualConv adaptation	0.00-0.00	98.89-100.00	100.00-100.00
SATE, two branches	0.00-0.00	37.78-53.33	100.00-100.00
SATE, folded	0.00-0.00	12.22-21.11	100.00-100.00
Three branches, folded	0.00-0.00	47.78-86.67	100.00-100.00
Without attention	0.00-0.00	63.33-77.78	100.00-100.00
Without attention, folded	0.00-1.11	58.89-65.56	100.00-100.00

Peak memory is the maximum PyTorch CUDA tensor allocation including model and transient inference tensors; it is not total VRAM or an embedded-device measurement. The 30/60/90 Hz periods are sensitivity scenarios declared before this timing set. Exceedance counts do not establish an application deadline, a scheduling guarantee or worst-case execution time. Attention-removal timing was added after the seed-zero ablation and is a resource/runtime comparison, not new independent accuracy evidence.

Online Resource 2 contains all 4,560 raw forward and pipeline timing samples. Quantiles use linear interpolation between the ordered observations, as in NumPy percentile. Every timing configuration uses the seed-zero checkpoint listed in Section 6.

results.json SHA-256:

`ba2fc61a051c9103d399ae1a477f419ee6267b7eca46adb0ad60e6cd81049d34`

attention_control.json SHA-256:

`a70761255471bc61db83485a1190799ec6a0c8674078945236a7505a1f873f70`

8. Fixed qualitative examples

Figure 5 uses the lower-median image of each frozen test list. The same 1024-pixel crop is centered on the annotated intersection nearest the image center for every model. Predictions come from archived full-image evaluations at confidence 0.25; the crop is an illustration and does not define an additional metric. Selection geometry and source hashes are retained in the audit record.