

Temporal Zoom Networks: Distance Regression and Continuous Depth for Accurate and Inference-Efficient Action Localization

Ibne Farabi Shihab[‡]

Sanjeda Akter[‡]

Anuj Sharma²

¹Department of Computer Science, Iowa State University

²Department of Civil, Construction & Environmental Engineering, Iowa State University
ishihab@iastate.edu

Abstract

Temporal action localization requires both precise boundary detection and computational efficiency. Current methods apply uniform computation across all temporal positions, wasting resources on easy boundaries while struggling with ambiguous ones. We address this through two complementary innovations: Boundary Distance Regression (BDR), which replaces classification-based boundary detection with signed-distance regression achieving $3.3\text{--}16.7\times$ lower variance; and Adaptive Temporal Refinement (ATR), which allocates transformer depth continuously ($\tau \in [0, 1]$) to concentrate computation near difficult boundaries. On THUMOS14, our method achieves 56.5% mAP@0.7 and 58.2% average mAP@[0.3:0.7] with 151G FLOPs, using 36% fewer FLOPs than ActionFormer++ (55.7% mAP@0.7 at 235G). Compared to uniform baselines, we achieve +2.9% mAP@0.7 (+1.8% avg mAP, 5.4% relative) with 24% fewer FLOPs and 29% lower latency, with particularly strong gains on short actions (+4.2%, 8.6% relative). Training requires $1.29\times$ baseline FLOPs, but this one-time cost is amortized over many inference runs; knowledge distillation further reduces this to $1.1\times$ while retaining 99.5% accuracy. Our contributions include: (i) a theoretically-grounded distance formulation with information-theoretic analysis showing optimal variance scaling; (ii) a continuous depth allocation mechanism avoiding discrete routing complexity; and (iii) consistent improvements across four datasets with gains correlating with boundary heterogeneity.

1. Introduction

Temporal action localization (identifying when actions occur in untrimmed videos) faces a fundamental challenge: boundary detection difficulty varies dramatically. A sharp camera cut may be detectable within a single frame, while a gradual fade creates inherent ambiguity where even human annotators disagree by ± 0.5 seconds [15]. This heterogeneity is critical in applications requiring fine-grained temporal precision, such as traffic surveillance [1, 40] or adverse conditions [42]. Despite this, most methods apply uniform computation across all temporal positions, using the same 6–9 layer transformer at every location [26, 46], wasting resources on easy boundaries while providing insufficient capacity for ambiguous ones.

We present two complementary contributions addressing both precision and efficiency.

First, Boundary Distance Regression (BDR) replaces classification-based boundary detection with signed-distance regression. Classification methods create broad plateaus of ambiguity spanning $W \approx 2\kappa$ frames, with variance scaling by feature smoothness. BDR instead regresses signed distances and extracts zero-crossings, achieving variance $\text{Var}[\hat{b}_{\text{BDR}}] = O(\Delta t^2/T)$ that depends on temporal discretization rather than smoothness (Theorem 2). Empirically, this yields $R \in [0.06, 0.30]$ ($3.3\text{--}16.7\times$ lower variance), with gains exceeding theoretical predictions due to multi-scale accumulation and other factors (Section 4.4). BDR retrofits to existing methods with ~ 50 lines of code, yielding consistent 1.8–3.1% mAP@0.7 improvements (average +2.4%).

Second, Adaptive Temporal Refinement (ATR) allocates transformer depth continuously ($\tau \in [0, 1]$) based on boundary difficulty. A shallow 2-layer transformer identifies uncertain regions, then a deeper 9-layer transformer refines where needed. Unlike discrete routing requiring reinforcement learning [11, 35], ATR uses continuous interpo-

^{*}Equal contribution.

[†]Corresponding author: ishihab@iastate.edu.

lation enabling fully differentiable training with fewer hyperparameters (2 vs 4–7) and less tuning (2h vs 5–12h).

Our contributions are threefold. First, BDR provides a theoretically grounded boundary loss retrofitting to existing methods with minimal code, achieving consistent +2.4% average improvements (4.5–5.8% relative) across BMN, ActionFormer, and TriDet. Second, ATR achieves competitive accuracy (56.5% mAP@0.7, 58.2% avg mAP@[0.3:0.7]) with 36% fewer FLOPs than ActionFormer++ (151G vs 235G), and +2.9% (5.4% relative) over our Uniform-6 baseline with 24% fewer FLOPs and 29% lower latency. Gains are strongest on short actions (<2s: +4.2%, 8.6% relative) where precision matters most. Our primary contribution is efficiency: achieving SOTA-level accuracy at significantly lower computational cost. Third, we demonstrate consistent improvements across four datasets (THUMOS14, FineAction, ActivityNet, Ego4D) with gains correlating with boundary heterogeneity.

Training cost increases modestly ($1.29\times$ FLOPs, $1.62\times$ memory), a practical one-time trade-off for permanent inference gains where models are deployed millions of times. Knowledge distillation further reduces training overhead to $1.1\times$ baseline while retaining 99.5% accuracy (§5).

2. Related Work

Modern temporal action localization methods such as ActionFormer [46] and TriDet [26] rely on fixed-depth multi-scale transformers, so every timestamp receives the same computational budget. This uniform treatment leaves little room to adapt to the heterogeneous difficulty of boundary prediction and keeps boundary heads anchored in classification losses that spread probability mass across wide ambiguous windows. Subsequent work on boundary refinement retains the same classification view; in contrast, we analyze why signed distance regression can achieve Cramér–Rao-level efficiency and show that zero-crossing extraction yields sharper boundaries than level-set or regression-only approaches used in earlier TAL systems.

Efficiency-aware modeling has emerged through adaptive computation and model pruning. Prior adaptive methods [11, 35] route tokens through discrete depths using reinforcement learning, which introduces combinatorial optimization and additional hyperparameters. Outside vision, resource-constrained sequence modeling has recently demonstrated that careful unstructured pruning of Mamba state-space models removes up to 70% of parameters while maintaining accuracy within 3–9% [39], underscoring the importance of amortizing training costs to unlock deployment-level savings. Our formulation folds these threads together: continuous depth allocation $\tau \in [0, 1]$ avoids discrete routing overhead, and boundary distance regression pairs the compute schedule with theoretically grounded precision gains. For a complete literature review

see Appendix B.

3. Method

We present a two-part framework addressing both boundary precision and computational efficiency. The first component, Boundary Distance Regression (BDR), improves localization accuracy through a theoretically-grounded distance formulation. The second component, Adaptive Temporal Refinement (ATR), allocates computation adaptively based on boundary difficulty.

3.1. Problem Formulation

Given a video with T frames and features $\mathbf{F} \in \mathbb{R}^{T \times D}$ from a frozen backbone, our goal is to predict action instances $\{(s_i, e_i, c_i)\}$ with start times s_i , end times e_i , and class labels c_i . The ground truth is $\mathcal{G} = \{(s_i^*, e_i^*, c_i^*)\}$ with boundary set $\mathcal{B}_{\text{GT}}$. Times are measured in frames unless noted: Δt denotes temporal stride, L action duration, b a boundary index, and $d(t) = t - b(t)$ the signed distance to the nearest boundary $b(t) = \arg \min_{b \in \mathcal{B}_{\text{GT}}} |t - b|$. Variances are reported in frames²; see Section 4 for theoretical assumptions (i.i.d. Laplace noise, uniform stride sampling, sufficient capacity, weak temporal dependence).

3.2. Architecture Overview

Our framework processes temporal positions with adaptive depth through four stages (Fig. 1). First, a lightweight 2-layer transformer produces coarse predictions and uncertainty estimates. Second, an MLP predicts continuous refinement depth $\tau_t \in [0, 1]$, controlling how much additional computation each position receives. Third, a deeper 9-layer transformer refines representations where needed. Fourth, final predictions are weighted combinations of shallow and deep outputs based on τ_t , with boundaries extracted via signed-distance regression. Interpolating predictions (logits and boxes) rather than features enables smooth, differentiable depth allocation. Token pruning further reduces computation in low-information regions while preserving full capacity near boundaries; a full FLOPs breakdown appears in Appendix C.5.

3.3. Boundary-Aware Uncertainty Estimation

Accurate depth allocation requires identifying uncertain regions that merit refinement. Generic uncertainty estimation is insufficient because temporal boundary difficulty depends on local context. For example, a sharp camera cut has low intrinsic uncertainty but may have limited context, while a gradual fade has high intrinsic uncertainty but smooth features.

Local context. We compute local context features via a 3-layer transformer on a window $\mathbf{h}_{\text{local}}(t) = \text{Transformer}(\mathbf{F}[t - w : t + w])$ with $w = 3$ frames; sequence edges are handled by reflection padding. Let




Figure 1. Adaptive Temporal Refinement (ATR) architecture. Four stages: (1) a shallow transformer produces coarse predictions and uncertainty; (2) an MLP predicts continuous depth allocation τ_t ; (3) a deep transformer refines difficult regions; (4) residual refinement merges predictions. Boundaries are extracted via signed-distance regression, and token pruning reduces computation in low-information regions.

$\mathbf{h}_t$ denote the shallow encoder token at t . We also compute temporal gradient magnitude $g_t = \|\mathbf{F}[t+1] - \mathbf{F}[t-1]\|_2$ as an explicit signal of boundary sharpness. A lightweight MLP predicts aleatoric uncertainty $\sigma_t^2 = \text{MLP}([\mathbf{h}_{\text{local}}(t); g_t; \mathbf{h}_t]) \in \mathbb{R}^+$. Sharp transitions (large g_t) yield low σ_t^2 , while gradual fades yield high σ_t^2 .

Loss. We train σ_t^2 with heteroscedastic regression [19]:

$$\mathcal{L}_{\text{uncertainty}} = \sum_{t=1}^T \left(\frac{(d(t) - \hat{d}(t))^2}{2\sigma_t^2} + \frac{1}{2} \log \sigma_t^2 \right),$$

where $d(t)$ is the ground-truth signed distance. This aligns uncertainty with empirical error and provides an interpretable signal for depth allocation (see Appendix C.2).

3.4. Continuous Depth Allocation

Using uncertainty as guidance, we allocate computational depth adaptively via continuous interpolation weights $\tau_t \in [0, 1]$ that smoothly blend shallow and deep predictions. This formulation avoids discrete routing or reinforcement learning, enabling stable, fully differentiable optimization.

The shallow encoder has 2 layers and the deep encoder 9 layers. An MLP predicts interpolation weights from shallow features and uncertainty:

$$\tau_t = \sigma(\text{MLP}_{\text{depth}}([\mathbf{h}_{\text{shallow},t}; \sigma_t^2])).$$

Residual refinement combines logits and boxes:

$$\begin{aligned} \text{logits}_t &= f_{\text{cls}}(\text{LayerNorm}(\mathbf{h}_{\text{shallow},t})) + \tau_t r_t^{\text{cls}}, \\ \text{boxes}_t &= f_{\text{box}}(\mathbf{h}_{\text{shallow},t}) + \tau_t r_t^{\text{box}}, \end{aligned}$$

where $r_t^{\text{cls}} = f_{\text{cls}}(\mathbf{h}_{\text{deep},t}) - f_{\text{cls}}(\mathbf{h}_{\text{shallow},t})$ and $r_t^{\text{box}} = f_{\text{box}}(\mathbf{h}_{\text{deep},t}) - f_{\text{box}}(\mathbf{h}_{\text{shallow},t})$. A hysteresis band $\gamma = \pm 0.05$ enforces temporal stability: if $\tau_t \in (\tau_{t-1} - \gamma, \tau_{t-1} + \gamma)$, we set $\tau_t := \tau_{t-1}$, reducing frame-to-frame flips by half (see Table 6).

3.5. Selective Token Processing

Not all temporal positions require equal processing depth. We introduce learned token pruning to skip low-information regions while retaining all boundaries.

For each temporal position t , a small MLP predicts importance $w_t = \sigma(\text{MLP}_{\text{prune}}(\mathbf{h}_{\text{shallow},t})) \in [0, 1]$. During training we use Gumbel-Softmax sampling [16] and a straight-through top- k mask with $k = \lfloor 0.8T \rfloor$. A boundary mask within ± 12 frames of shallow-predicted boundaries overrides pruning ($w_t := 1$), ensuring 100% retention in action regions. At inference, we deterministically keep tokens with w_t above the 80th-percentile threshold. This yields an effective sequence length $T_{\text{eff}} \approx 0.8T$. Sparsity is encouraged via

$$\mathcal{L}_{\text{compute}} = \lambda_c \frac{1}{T} \sum_t \tau_t, \quad \mathcal{L}_{\text{prune}} = \lambda_p \frac{1}{T} \sum_t w_t,$$

with $\lambda_c = 0.05$ and $\lambda_p = 0.01$ selected on the validation mAP-FLOPs Pareto frontier.

3.6. Boundary Distance Regression (BDR): Loss and Extraction

While adaptive depth improves efficiency, precise boundary localization requires a new detection formulation. Classification-based detectors produce multi-modal probabilities with 3–5-frame ambiguity; BDR instead regresses the signed distance to the nearest boundary:

$$d(t) = t - b(t), \quad b(t) = \arg \min_{b \in \mathcal{B}_{\text{GT}}} |t - b|.$$

Here $d(t)$ is negative before the boundary, zero at the boundary, and positive after, with $|\nabla_t d| = 1$ almost everywhere—allowing zero-crossing detection (Fig. 2).

The model outputs $\hat{d}(t) = \text{Linear}(\mathbf{h}_t) \in \mathbb{R}$. The BDR loss combines L1 regression with gradient-magnitude regularization:

$$\mathcal{L}_{\text{BDR}} := \frac{1}{T} \sum_{t=1}^T |d(t) - \hat{d}(t)| + \frac{\alpha}{T-1} \sum_{t=1}^{T-1} (\max\{0, |\hat{d}(t+1) - \hat{d}(t)| - \alpha\})^2$$

with $\alpha=0.1$. Boundaries are extracted as zero-crossings where $\hat{d}(t)\hat{d}(t+1) \leq 0$, filtered by gradient magnitude $|\nabla \hat{d}| \geq \theta_{\text{grad}}=0.5$, and merged using 1D-NMS with window $w_{\text{nms}}=5$. This procedure is robust to $\theta_{\text{grad}} \in [0.3, 0.7]$, causing under 0.4% mAP variation (see Appendix C.6).

Boundary extraction summary. Predict $\hat{d}(t)$; find linear zero-crossings where $\hat{d}(t)\hat{d}(t+1) \leq 0$; keep candidates with $|\nabla \hat{d}| \geq \theta_{\text{grad}}$; apply 1D-NMS with $w_{\text{nms}}=5$ to remove duplicates.

4. Theoretical Analysis

We analyze boundary localization through information-theoretic bounds, showing that BDR achieves superior boundary precision under idealized conditions. Our analysis uses the following notation: Δt for temporal stride (frames), T for the number of positions, κ for feature smoothness (frames), and $W \approx 2\kappa$ for the plateau width; variances are reported in frames².

Assumptions. Our theoretical results hold under the following idealized conditions: (i) i.i.d. Laplace noise, (ii) uniform stride sampling, (iii) sufficient function capacity, and (iv) weak temporal dependence. See Appendix A.3 for detailed assumptions and finite-sample guarantees. **Important limitation:** The i.i.d. noise assumption is violated in practice due to temporal correlations in video features (empirical autocorrelation $\rho \approx 0.4$). However, our analysis in Appendix E.6 shows that variance ratios remain stable (R varies by $< 15\%$) for moderate correlation ($\rho < 0.6$), and empirical validation confirms that the predicted scaling trends hold despite this assumption violation. The theoretical bounds should be interpreted as order-of-magnitude guides rather than exact predictions.

4.1. Classification Localization Limits

Classification-based boundary detection models $p(t|\text{boundary})$ via smooth probability curves. Near boundaries, feature similarity creates ambiguous plateau regions spanning approximately 2κ frames. This fundamental limitation arises because classification methods must find peaks in probability distributions, which become broad and flat when features are smooth.

Theorem 1 (Classification variance bound). *Let features near the true boundary b^* follow a Gaussian similarity kernel $\mathbf{h}(t) = \phi(t)v$ with $\phi(t) = \exp(-(t - b^*)^2/(2\kappa^2))$ and $v \in \mathbb{R}^D$, $\|v\|_2 = 1$. Let a calibrated classifier*

be $p(t) = \sigma(w^\top \mathbf{h}(t))$ with $\|w\|_2 = 1$. Under regularity conditions ensuring differentiability and identifiability, the Fisher information for b^ satisfies $I_{\text{cls}}(b^*) = \Theta(\kappa^{-1})$, hence $\text{Var}[\hat{b}_{\text{cls}}] = \Omega(\kappa)$ (units: frames²).*

Complete proof in Appendix E.1. The probability plateau spanning $\approx 2\kappa$ frames makes precise peak localization fundamentally difficult when κ is large.

4.2. Distance Regression Precision

Signed distance regression leverages sharp gradient discontinuities at boundaries. Unlike classification, which operates on smooth probability curves, the signed distance field has constant gradient magnitude $|\nabla_t d(t)| = 1$ almost everywhere and crosses zero at b , enabling precise zero-crossing detection. This structural advantage allows distance regression to achieve variance that scales with temporal discretization rather than feature smoothness.

Theorem 2 (BDR Fisher information). *Assume observations of the signed distance satisfy $y_t = d(t) + \varepsilon_t$ on a uniform grid with stride Δt , where $\varepsilon_t \stackrel{\text{i.i.d.}}{\sim} \text{Laplace}(0, b)$ (variance $2b^2$). In a local linearization of $d(t)$ around the boundary b and under standard regularity conditions, the zero-crossing estimator admits Fisher information*

$$I_{\text{BDR}}(b) \geq \frac{T}{4b^2 \Delta t^2} \Rightarrow \text{Var}[\hat{b}_{\text{BDR}}] \leq c \frac{\Delta t^2}{T},$$

for a universal constant c . With L_1 regression under Laplace noise, the zero-crossing estimator approaches the Cramér–Rao lower bound asymptotically. For fixed-video inference (constant T), the variance scales as $O(\Delta t^2)$ and is independent of κ .

Complete proof in Appendix E.2. A finite-sample version with explicit approximation error appears in Appendix E.3.

4.3. Why Classical Bounds Underestimate BDR

We define the empirical variance ratio $R_{\text{emp}} \triangleq \text{MSE}_{\text{BDR}}/\text{MSE}_{\text{cls}}$ computed per-video via blocked bootstrap. Since both estimators are asymptotically unbiased, $R_{\text{emp}} \approx R_{\text{theoretical}}$ for practical purposes. The theoretical ratio from Fisher information analysis provides order-of-magnitude scaling but underestimates BDR’s practical advantage.

Corollary 1 (Naive Fisher bound with action-length averaging). *From Theorems 1 and 2, $\text{Var}[\hat{b}_{\text{cls}}] = \Omega(\kappa) = \Omega(W)$ and $\text{Var}[\hat{b}_{\text{BDR}}] = O(\Delta t^2/T)$, hence for fixed T the ratio $R = \text{MSE}_{\text{BDR}}/\text{MSE}_{\text{cls}} = O(\Delta t^2/W)$. Averaging information across an action of duration L yields*

$$R = C \frac{\Delta t^2}{W^2 \sqrt{L}} \quad (\text{with } C \text{ constant}),$$




Figure 2. BDR vs Classification comparison. BDR produces sharp zero-crossings at boundaries (blue line: distance to start boundary at $t=25$, showing $d(t) = t - 25$ with zero-crossing only at the true boundary) while classification creates fuzzy probability regions (red). The signed distance field $d(t) = t - b(t)$ has constant gradient $|\nabla_t d| = 1$ and clear zero-crossings only at true boundaries, enabling precise localization. End boundaries are detected similarly using distance to the end boundary.

so that, when L is approximately constant across boundaries, $R \propto \Delta t^2 / W^2$ (quadratic in plateau width).

Complete proof in Appendix E.4. Width-stratified analysis (Appendix F.2.9) shows the predicted trend: R is near unity when $W \ll \Delta t$ and decreases as $W/\Delta t$ grows. Empirical ratios (0.06–0.30) are smaller than naive predictions, reflecting violations of idealized assumptions (temporal correlation, heavy tails, capacity limits).

4.4. Why Naive Bounds Fail: Contributing Factors

We identify four factors that contribute to BDR’s practical advantage beyond information-theoretic predictions. These factors partially compound rather than multiply independently, as they exhibit correlations and saturation effects. Multi-scale accumulation provides the dominant advantage, with other factors contributing additively to the residual. The naive Fisher bound predicts $R = O(\Delta t^2 / W^2) \approx 0.64$ for median parameters. Multi-scale accumulation over $L \approx 65$ frames provides approximately $\sqrt{L} \approx 8\times$ additional information, yielding $R \approx 0.64/8 = 0.08$. Heavy-tail correction (Student- t with $\nu = 3$) degrades classification Fisher information by factor $\sim 1.5\times$, yielding $R \approx 0.08/1.5 = 0.053$. The empirical average $R = 0.11$ across all boundary types reflects capacity limitations, calibration degradation near boundaries, and violations of i.i.d. assumptions (temporal correlation $\rho \approx 0.4$). The observed variance ratios $R = 0.06$ – 0.30 align with this decomposition when accounting for heterogeneous L and κ values. See Appendix E.5 for detailed analysis including empirical validation and per-boundary-type breakdown.

The complete training objective combines all components: $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{TAL}} + \lambda_1 \mathcal{L}_{\text{BDR}} + \lambda_2 \mathcal{L}_{\text{uncertainty}} + \lambda_c \mathcal{L}_{\text{compute}} +$

$\lambda_p \mathcal{L}_{\text{prune}}$, where $\mathcal{L}_{\text{TAL}}$ is the standard DETR loss [4] with focal loss [24] for classification and L1 plus GIoU [36] for boxes, $\mathcal{L}_{\text{BDR}}$ is signed distance regression, $\mathcal{L}_{\text{uncertainty}}$ is the calibration loss from Section 3.3, $\mathcal{L}_{\text{compute}}$ encourages shallow processing, and $\mathcal{L}_{\text{prune}}$ encourages token sparsity. We set $\lambda_1 = 1.0$, $\lambda_2 = 0.1$, $\lambda_c = 0.05$, and $\lambda_p = 0.01$ selected on the validation mAP-FLOPs Pareto frontier.

Our continuous formulation enables stable end-to-end gradient flow throughout training without requiring special handling or variance reduction techniques. We measure depth predictor gradient variance as $\sigma_{\nabla}^2 = \frac{1}{|\Theta|} \sum_{\theta \in \Theta} \text{Var}(\partial \mathcal{L} / \partial \theta)$ where Θ are the depth MLP parameters, computed across 10,000 training iterations. ATR achieves $\sigma_{\nabla}^2 = 0.021$, while Gumbel-Softmax routing exhibits $6.8\times$ higher variance ($\sigma_{\nabla}^2 = 0.143$) due to temperature annealing schedules, and reinforcement learning shows $15\times$ higher variance ($\sigma_{\nabla}^2 = 0.318$) from policy gradient stochasticity, averaged across 5 runs. This optimization stability translates directly to faster convergence in practice, with ATR reaching 90% of final validation mAP in 18,000 iterations versus 58,000 iterations required for Gumbel-Softmax routing. The smooth interpolation between shallow and deep predictions avoids discrete decisions during backpropagation, eliminating the need for straight-through estimators or complex variance reduction techniques that would otherwise be required for discrete routing approaches.

We use AdamW [32] with learning rate $1e-4$, weight decay $1e-4$, and cosine schedule [31] over 60K iterations, with training taking 24 hours on $4\times$ A100 GPUs with batch size 32 via gradient accumulation.

5. Experiments

We evaluate ATR/BDR on four TAL benchmarks with matched backbones and rigorous testing. Main text reports the essential results (one main table + Pareto + supporting tables); extended analyses, ablations, and training-cost details are in the appendix.

5.1. Experimental Setup

Datasets: THUMOS14 [15] (413/20, 2.3s avg), ActivityNet-1.3 [13] (20K/200, 36s avg), FineAction [28] (17K/106), Ego4D [10] (3,670h). Backbone: VideoSwin-B [30] pretrained on Kinetics-400 [18], frozen; stride-4 features (768-d). Metrics: mAP@IoU {0.3, 0.5, 0.7}, average mAP@[0.3:0.7] (standard TAL metric), FLOPs via fvcore, latency on a single A100 (bs=1). Baselines: Uniform-6/9, ActionFormer, TriDet; published SOTA references (ActionFormer++, TemporalMaxer) included for context. We observe performance gaps between our reproductions and published results (e.g., ActionFormer++: 55.7% published vs 52.8% reproduced with frozen backbone), reflecting differences in backbones, augmentation

Table 1. Contributing factors to BDR’s empirical advantage.

Factor	Contribution	Mechanism
Multi-scale accumulation	Primary (60–70% of gap)	Information across action span L
Heavy-tailed features	Secondary (15–20%)	Student- t vs Gaussian noise
Capacity efficiency	Secondary (10–15%)	Sharp boundaries easier to fit
Calibration quality	Context-dependent (0–40%)	Varies by boundary type

strategies, hardware, and hyperparameters. To ensure fair comparison, we establish a controlled experimental setting with identical frozen VideoSwin-B backbone, fixed augmentation pipeline, consistent hardware (4×A100), and matched training iterations (60K). All comparisons in this paper (Uniform-6, ATR, ActionFormer, TriDet) use this controlled setting, ensuring fair relative comparisons. Published SOTA numbers are provided for context but reflect different experimental conditions. Stats: paired tests across THUMOS14 videos with Holm–Bonferroni correction; 95% CIs via blocked bootstrap (10k resamples). Reproduction/implementation details in Appx. F.1, C.1.

5.2. Main Results

We present comprehensive THUMOS14 results with 10 seeds and bootstrap CIs in Table 2. We establish our baseline for state-of-the-art performance against ActionFormer++ [47], which at the time of our work’s conception was the leading method. We note the concurrent publication of BRTAL [29], which explores an alternative diffusion-based refinement strategy. As this work was published within three months of the submission deadline and its code was not publicly available, a direct experimental comparison was not feasible. Our work, therefore, focuses on establishing a new SOTA in inference efficiency and high-precision (mAP@0.7) over the established SOTA.

ATR reaches 56.5% mAP@0.7 and 58.2% average mAP@[0.3:0.7] at 151G, achieving accuracy competitive with published SOTA (ActionFormer++: 55.7% mAP@0.7) while requiring 36% fewer inference FLOPs (151G vs 235G) and 29% lower latency (118ms vs ~165ms). While direct statistical comparison with published methods is complicated by different experimental settings (backbones, augmentation, hardware), ATR’s 56.5% [55.8, 57.3] overlaps with ActionFormer++’s 55.7±0.4 [55.3, 56.1], suggesting comparable accuracy under different conditions. The key contribution is efficiency: ATR achieves SOTA-level accuracy at 36% lower computational cost. Compared to our reproduced Uniform-6 baseline (53.6% mAP@0.7, 56.4% avg mAP at 198G FLOPs) under identical settings, ATR achieves +2.9% absolute gain in mAP@0.7 (+1.8% avg mAP, 3.2% relative) with 24% fewer inference FLOPs (151G vs 198G). The improvement is statistically significant: paired tests across 213 videos with Holm–Bonferroni correction yield $p < .01$ for all com-




Figure 3. Pareto on THUMOS14. ATR dominates uniform base-lines across budgets.

parisons (Appx. F.2.1). Recent 2025 methods (e.g., CLIP-AE for unsupervised TAL, TimeLoc for long-form videos) address different problem settings and are not directly comparable; our focus is on efficient fully-supervised TAL with adaptive computation.

The 29% latency reduction (167ms→118ms) exceeds the 24% FLOPs reduction due to memory bandwidth improvements from token pruning and attention’s quadratic scaling: with effective sequence length $T_{\text{eff}} = 0.8T$, attention FLOPs scale as $(0.8)^2 = 0.64$ (36% reduction) while FFN scales as 0.8 (20% reduction), yielding super-linear speedup. Full FLOPs breakdown is in Appx. C.5.

For each test video we average metrics across seeds, then run paired tests per baseline; p -values are corrected via Holm–Bonferroni over 12 comparisons (4 baselines × 3 metrics). We also report blocked bootstrap CIs over videos (10k resamples). Both procedures indicate significant gains ($p < .01$) vs Uniform-6/9, ActionFormer, and TriDet (Appx. F.2.1).

Training FLOPs are measured using fvcare on forward passes in training mode, following standard practice in computer vision. Backward pass computation is implementation-dependent and hardware-specific, making forward-pass measurement the standard for reproducible comparisons. Additionally, mixed-precision training (FP16 activations, FP32 gradients) reduces effective measured FLOPs compared to theoretical FP32 calculations by ap-

Table 2. THUMOS14 test (10 seeds; 95% CI). ATR achieves competitive accuracy (56.5% mAP@0.7, 58.2% avg mAP@[0.3:0.7]) with 36% fewer FLOPs than ActionFormer++ (151G vs 235G). Compared to Uniform-6 baseline, ATR improves by +2.9% mAP@0.7 (+1.8% avg mAP) while *reducing inference* FLOPs by 24% (151G vs 198G). Published SOTA numbers reflect different experimental conditions (backbones, augmentation); all reproduced baselines use identical frozen VideoSwin-B. Recent 2025 methods (e.g., CLIP-AE [25], TimeLoc [5]) focus on different settings (unsupervised, long-form) and are not directly comparable. Our contribution emphasizes efficiency: achieving SOTA-level accuracy at significantly lower computational cost. FlashAttention [8] enabled uniformly.

Method	mAP@0.5 (%)	mAP@0.7 (%)	Avg mAP@[0.3:0.7]	FLOPs (G)	$\mathbb{E}[\tau]$	Latency (ms)	Δ vs SOTA
<i>Published SOTA (reference, different settings):</i>							
TemporalMaxer [27] (ICCV 2023)	58.6 $\pm$ 0.4	54.9 $\pm$ 0.3	-	212	-	-	-
ActionFormer++ [47] (CVPR 2024)	59.8 $\pm$ 0.5	55.7 $\pm$ 0.4	-	235	-	-	-
<i>Reproduced baselines (identical settings):</i>							
ActionFormer	56.8 [56.0, 57.5]	52.8 [52.1, 53.6]	56.4 [55.7, 57.1]	198	-	158	-
TriDet	58.7 [57.9, 59.4]	54.1 [53.4, 54.9]	57.8 [57.1, 58.5]	215	-	173	-
Uniform-6	59.3 [58.6, 60.1]	53.6 [52.9, 54.4]	56.4 [55.7, 57.1]	198	-	167	-
Uniform-9	60.1 [59.3, 60.8]	54.2 [53.5, 55.0]	57.1 [56.4, 57.8]	245	-	192	-
<i>Our method:</i>							
ATR (residual refine)	62.1 [61.4, 62.9]	56.5 [55.8, 57.3]	58.2 [57.5, 58.9]	151	0.16	118	+0.8 vs SOTA
ATR (logit blend)	61.8 [61.1, 62.6]	56.3 [55.6, 57.1]	58.0 [57.3, 58.7]	154	0.16	121	+0.6 vs SOTA

proximately 20–25%. The ratio 152G/198G = 0.77 for Uniform-6 combines mixed-precision (0.8 $\times$) and checkpointing efficiency (0.96 $\times$), yielding the measured values. This methodology is applied uniformly across all methods for fair comparison.

ATR trains both shallow and deep paths: training cost rises to 196G vs 152G FLOPs (1.29 $\times$), 24h vs 18h wall-clock (1.33 $\times$), and 14.9GB vs 9.2GB peak memory (1.62 $\times$) on THUMOS14. This training overhead is justified by permanent inference efficiency gains: in production deployments processing millions of videos, the 29% latency reduction (118ms vs 167ms) and 24% FLOPs savings (151G vs 198G) accumulate to substantial cost savings. For example, processing 1M videos saves $\sim$ 49 GPU-hours at inference, far exceeding the 6-hour training overhead. The 14.9GB training memory requirement remains accessible and is well within the standard for SOTA TAL models. At inference, token pruning and $\mathbb{E}[\tau]=0.16$ yield 151G total (27G post-backbone) vs 198G (74G post-backbone), i.e., 64% lower localization-specific cost; see Appx. C.5.

For practitioners prioritizing training efficiency, knowledge distillation provides an attractive alternative: students retain 99.5% of teacher mAP (56.2 vs 56.5) at $\sim$ 1.1 $\times$ baseline training cost (Appx. F.2.4). We explored conditional backprop for $\tau_t < 0.3$ to reduce training cost, but this lowered accuracy to 55.1% mAP@0.7, so we adopt full backprop for best performance.

5.3. Generalization Analysis

ATR improves across four datasets: THUMOS14 (+2.9%, 5.4% relative), FineAction (+2.7%), ActivityNet (+1.8%), Ego4D (+1.9%); gains correlate with boundary heterogene-

ity (Appx. F.2). Within THUMOS14, short actions (< 2s) gain +4.2% (8.6% relative improvement) with higher refinement ($\mathbb{E}[\tau]=0.24$), demonstrating that ATR provides the most value where boundary precision matters most; long actions (> 10s) gain +0.8% with shallow processing ($\mathbb{E}[\tau]=0.05$). Cross-dataset results (Table 12, Appx.) show consistent improvements across sports (THUMOS14), daily activities (ActivityNet), fine-grained actions (FineAction), and egocentric video (Ego4D), validating domain-agnostic principles.

We ablate each component in Table 18 (Appx. F.2): BDR alone adds +1.3% via sharper boundaries (2.4% relative improvement); continuous depth alone adds +1.5% with 24% fewer FLOPs; combined they yield +2.9% (5.4% relative) and the best Pareto efficiency. Residual refinement outperforms naive logit blending at the same budget (Table 7, Appx. C.3). Sensitivity shows robustness to $\theta_{\text{grad}} \in [0.3, 0.7]$ and stride-2/4; stride-8 harms precision (Table 9). A small hysteresis band ($\gamma = \pm 0.05$) halves frame-to-frame flips without impacting FLOPs (Table 8, Appx. C.9).

Our boundary-aware uncertainty attains 47% lower R-ECE than MC-Dropout (0.076 vs 0.142) at 92% lower FLOPs (151G vs 1,980G, 10 passes), outperforming deep ensembles (990G) at far lower cost (Table 13, Appx.). Standard evidential improves calibration, but boundary-aware contextualization adds a further 22% gain (0.076 vs 0.098). By boundary type, we observe 47% improvement on gradual and 53% on sharp cuts. For selective prediction (abstention by σ_t^2), we maintain 58.1% mAP@0.7 at 80% coverage vs 51.2% for MC-Dropout.

Synthetic data confirms the predicted scaling $R = \mathcal{O}(\Delta t^2/W)$ (log-log slope 1.15 ± 0.08 , $R^2=0.82$).

Table 3. THUMOS14 accuracy–efficiency. Latency on A100 (bs=1). Frozen: VideoSwin-B frozen; E2E: backbone finetuned (10k warmup).

Method	Backbone	mAP@0.5	mAP@0.7	FLOPs (G)	Latency (ms)	Peak Mem (GB)
Uniform-6	Frozen	59.3±0.4	53.6±0.8	198	167	11.8
Uniform-9	Frozen	60.1±0.5	54.2±0.9	245	192	13.0
Early-Exit+Temp (3 heads)	Frozen	61.6±0.4	56.3±0.6	156	128	11.0
MoD-lite (ST routing)	Frozen	61.4±0.5	56.1±0.6	154	130	11.2
ATR (logit blend)	Frozen	61.8±0.4	56.3±0.5	154	121	10.4
ATR (residual refine)	Frozen	62.1±0.4	56.5±0.5	151	118	9.8
Uniform-6	E2E	60.8±0.5	54.7±0.7	276	214	14.6
Early-Exit+Temp (3 heads)	E2E	62.7±0.5	57.0±0.6	221	165	13.9
ATR (residual refine)	E2E	63.2±0.5	57.2±0.6	214	160	13.2

Table 4. Robustness of R (mean over videos) vs. width W on THUMOS14 (lower is better). Trend persists across thresholds and NMS/multi-scale settings.

Setting	$W \leq \Delta t$	$\Delta t < W \leq 2\Delta t$	$W > 2\Delta t$
$\theta_{\text{grad}}=0.3$	0.97	0.61	0.35
$\theta_{\text{grad}}=0.5$	0.94	0.62	0.34
$\theta_{\text{grad}}=0.7$	0.99	0.64	0.36
NMS $w=3$	0.95	0.63	0.35
NMS $w=5$	0.94	0.62	0.34
NMS $w=7$	0.96	0.63	0.35
$w_s=1$	0.96	0.63	0.36
$w_s=1/\sqrt{s}$	0.94	0.62	0.34
$w_s=1/s$	0.95	0.62	0.35

On THUMOS14, BDR outperforms classification with R_{emp} : sharp cuts 0.29, gradual fades 0.06, medium 0.12 (Table 22). The averaged model $R=C \Delta t^2/(W^2 \sqrt{L})$ from Cor. 1 matches within 10%. BDR achieves 43% sharper peaks than focal loss (0.73 vs 0.51) and the lowest Boundary Chamfer Distance (4.8 frames), supporting sub-frame localization. Table 4 shows the width-stratified trend ($R \downarrow$ as $W/\Delta t \uparrow$) is invariant to threshold, NMS window, and multi-scale weights.

To demonstrate BDR’s value independent of ATR, we retrofit it to BMN, ActionFormer, and TriDet without architectural changes. Implementation requires ~ 50 LoC (distance targets, regression head, BDR loss, zero-crossing extraction). Retrofitting yields 1.8–3.1% mAP@0.7 gains (avg +2.4

When to use ATR: actions $\lesssim 5$ s, mixed boundary types (sports/surveillance), or tight budgets (150–200G). When to skip: long-form ($\gtrsim 10$ s), uniform boundaries (scripted), or unconstrained budgets. ATR improves boundary precision for most cases but cannot resolve all ambiguities: from 100 manually inspected failure cases, dense overlaps (32%) create interfering distance fields that deeper transformers cannot disentangle; extreme blur (18%) produces

feature smoothing that prevents precise localization regardless of depth; very gradual transitions (14%) have inherently low gradients below detection thresholds. Where inter-annotator variance exceeds 0.5s (9% of classes), refinement cannot resolve fundamental ground-truth disagreement. ATR’s adaptive refinement helps where additional computation provides signal, but fundamental ambiguities (dense overlaps, extreme blur, annotation disagreement) require architectural innovations beyond depth allocation (Appx. G.1).

6. Discussion and Conclusions

We presented two complementary contributions for temporal action localization that address both precision and efficiency. Boundary Distance Regression (BDR) provides a theoretically-grounded distance-regression formulation with variance ratio $R = O(\Delta t^2/W)$ from basic Fisher bounds, becoming $R = C \cdot (\Delta t^2)/(W^2 \cdot \sqrt{L})$ with action-length averaging (Corollary 1). Empirical ratios (0.06 to 0.30) reveal $3.3\times$ to $16.7\times$ lower variance than classification due to four amplification factors (Table 1). BDR retrofits to existing methods with ~ 50 lines of code, yielding consistent 1.8–3.1% mAP@0.7 improvements (average +2.4%, Table 5).

Adaptive Temporal Refinement (ATR) achieves efficient localization through continuous depth allocation $\tau \in [0, 1]$, avoiding discrete routing complexity. On THUMOS14, ATR establishes new SOTA (56.5% mAP@0.7) with 151G FLOPs, achieving +0.8% absolute gain over the previous best method ActionFormer++ (55.7% mAP@0.7 at 235G FLOPs) while using 36% fewer FLOPs (151G vs 235G) and 29% lower latency (118ms vs ~ 165 ms). Compared to Uniform-6 baseline (53.6% mAP@0.7 at 198G FLOPs), ATR achieves +2.9% absolute gain (5.4% relative improvement) with 24% fewer inference FLOPs (151G vs 198G) and 29% lower latency. Gains scale with boundary heterogeneity across four benchmarks, with particularly strong improvements on challenging short actions (+4.2%, 8.6%

relative). Training overhead ($1.29 \times$ baseline) is modest and amortized over many inference runs; for practitioners prioritizing training efficiency, knowledge distillation provides 99.5% performance retention at $1.1 \times$ baseline training cost, effectively reducing training overhead to $1.08 \times$ compared to the teacher.

Together, BDR and ATR advance temporal action localization through principled design combining information-theoretic analysis and practical efficiency. Our work demonstrates that theoretical insights can translate directly to practical improvements, with consistent gains across diverse datasets and architectures. See Appendix I for limitations and future directions.

References

- [1] Sanjeda Akter, Ibne Farabi Shihab, and Anuj Sharma. Image segmentation with large language models: A survey with perspectives for intelligent transportation systems. *arXiv preprint arXiv:2506.14096*, 2025. 1
- [2] Alexander Amini, Wilko Schwarting, Ava Soleimany, and Daniela Rus. Deep evidential regression. In *NeurIPS*, pages 14927–14937, 2020. 11
- [3] Shyamal Buch, Victor Escorcia, Chuanqi Shen, Bernard Ghanem, and Juan Carlos Nibbles. Sst: Single-stream temporal action proposals. In *CVPR*, pages 2911–2920, 2017. 11
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, pages 213–229, 2020. 5
- [5] Mei Chen, Hao Zhang, Dong Wang, and Yicong Zhou. Time-loc: Unified timestamp localization for long videos. *arXiv preprint arXiv:2503.06526*, 2025. 7
- [6] Xiangde Chen, Bryan M. Williams, Srinivasa R. Vallabhaneni, Gabriela Czanner, Rachel Williams, and Yalin Zheng. Learning active contour models for medical image segmentation. In *CVPR*, pages 11632–11640, 2019. 11
- [7] Xiyang Dai, Yinpeng Chen, Bin Xiao, Dongdong Chen, Mengchen Liu, Lu Yuan, and Lei Zhang. Dynamic head: Unifying object detection heads with attentions. In *CVPR*, pages 7373–7382, 2021. 11
- [8] Tri Dao, Daniel Y. Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. Flashattention: Fast and memory-efficient exact attention with io-awareness. In *NeurIPS*, 2022. 7, 13
- [9] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *Int. Conf. Mach. Learn.*, pages 1050–1059, 2016. 11
- [10] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, Miguel Martin, Tushar Nagarajan, Ilija Radosavovic, Santhosh Kumar Ramakrishnan, Fiona Ryan, Jayant Sharma, Michael Wray, Mengmeng Xu, Eric Zhongcong Xu, Chen Zhao, Siddhant Bansal, Dhruv Batra, Vincent Cartillier, Sean Crane, Tien Do, Morrie Doulaty, Akshay Erapalli, Christoph Feichtenhofer, Adriano Fragomeni, Qichen Fu, Christian Fuegen, Abraham Gebreselasie, Cristina Gonzalez, James Hillis, Xuhua Huang, Yifei Huang, Wenqi Jia, Weslie Khoo, Jachym Kolar, Satwik Kottur, Anurag Kumar, Federico Landini, Chao Li, Yanghao Li, Zhenqiang Li, Karttikeya Mangalam, Raghava Modhugu, Jonathan Munro, Tullie Murrell, Takumi Nishiyasu, Will Price, Paola Ruiz Puentes, Merey Ramazanov, Leda Sari, Kiran Somasundaram, Audrey Southerland, Yusuke Sugano, Ruijie Tao, Minh Vo, Yuchen Wang, Xindi Wu, Takuma Yagi, Yunyi Zhu, Pablo Arbelaez, David Crandall, Dima Damen, Giovanni Maria Farinella, Bernard Ghanem, Vamsi Krishna Ithapu, C. V. Jawahar, Hanbyul Joo, Kris Kitani, Haizhou Li, Richard Newcombe, Aude Oliva, Hyun Soo Park, James M. Rehg, Yoichi Sato, Jianbo Shi, Mike Zheng Shou, Antonio Torralba, Lorenzo Torresani, Mingfei Yan, and Jitendra Malik. Ego4d: Around the world in 3,000 hours of egocentric video. In *CVPR*, pages 18995–19012, 2022. 5, 19
- [11] Alex Graves. Adaptive computation time for recurrent neural networks. In *arXiv preprint arXiv:1603.08983*, 2016. 1, 2, 11
- [12] David Hall, Feras Dayoub, John Skinner, Haoyang Zhang, Dimity Miller, Peter Corke, Gustavo Carneiro, Anelia Angelova, and Niko Sunderhauf. Probabilistic object detection: Definition and evaluation. In *IEEE Winter Conf. Appl. Comput. Vis.*, pages 1031–1040, 2020. 11
- [13] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Nibbles. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*, pages 961–970, 2015. 5, 19
- [14] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q. Weinberger. Deep networks with stochastic depth. In *ECCV*, pages 646–661, 2016. 11
- [15] Haroon Idrees, Amir R. Zamir, Yu-Gang Jiang, Alex Gorbunov, Ivan Laptev, Rahul Sukthankar, and Mubarak Shah. The thumos challenge on action recognition for videos ”in the wild”. *Computer Vision and Image Understanding*, 155:1–23, 2017. 1, 5
- [16] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *ICLR*, 2017. 3
- [17] Steven M. Kay. *Fundamentals of Statistical Signal Processing: Estimation Theory*. Prentice Hall, 1993. 16
- [18] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, Mustafa Suleyman, and Andrew Zisserman. The kinetics human action video dataset. In *arXiv preprint arXiv:1705.06950*, 2017. 5, 13, 25
- [19] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? *NeurIPS*, pages 5574–5584, 2017. 3, 12
- [20] Hoel Kervadec, Jihene Bouchtiba, Christian Desrosiers, Eric Granger, Jose Dolz, and Ismail Ben Ayed. Boundary loss for highly unbalanced segmentation. In *Medical Image Analysis*, page 101851, 2021. 11
- [21] Tengyuan Liang, Tomaso Poggio, Alexander Rakhlin, and James Stokes. Fisher-rao metric, geometry, and complex-

- ity of neural networks. In *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics*, pages 888–896, 2019. 18, 21
- [22] Chuming Lin, Chengming Xu, Donghao Luo, Yabiao Wang, Ying Tai, Chengjie Wang, Jilin Li, Feiyue Huang, and Yanwei Fu. Learning salient boundary feature for anchor-free temporal action localization. In *CVPR*, pages 3320–3329, 2021. 11
- [23] Tianwei Lin, Xu Zhao, Haisheng Su, Chongjing Wang, and Ming Yang. Bmn: Boundary-matching network for temporal action proposal generation. In *ICCV*, pages 3889–3898, 2019. 11, 12
- [24] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollar. Focal loss for dense object detection. In *ICCV*, pages 2980–2988, 2017. 5
- [25] Zhiwei Lin, Yinghao Xu, Jiawei Chen, and Yan Song. Clip-ae: Clip-assisted cross-view audio-visual enhancement for unsupervised temporal action localization. *arXiv preprint arXiv:2505.23524*, 2025. 7
- [26] Shuming Liu, Mengmeng Xu, Chen Zhao, Xu Zhao, and Bernard Ghanem. Tridet: Temporal action detection with relative boundary modeling. In *CVPR*, pages 18857–18866, 2023. 1, 2, 11, 12
- [27] Xiaolong Liu, Yao Hu, Song Bai, and Xiang Bai. Temporalmaxer: Maximize temporal context with only max pooling for temporal action localization. In *ICCV*, pages 18073–18083, 2023. 7
- [28] Yi Liu, Limin Wang, Yali Wang, Xiao Ma, and Yu Qiao. Fineaction: A fine-grained video dataset for temporal action localization. In *arXiv preprint arXiv:2105.11107*, 2022. 5, 19
- [29] Yuxuan Liu, Rui Zhang, Han Li, and Junsong Yuan. Brtal: Boundary refinement for temporal action localization via diffusion. *IEEE TCSVT*, 2025. Early Access. 6
- [30] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. In *CVPR*, pages 3202–3211, 2022. 5
- [31] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. In *ICLR*, 2017. 5
- [32] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019. 5
- [33] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *CVPR*, pages 4460–4470, 2019. 11
- [34] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *CVPR*, pages 165–174, 2019. 11
- [35] David Raposo, Sam Ritter, Blake Richards, Timothy Lillicrap, Peter Conway Humphreys, and Adam Santoro. Mixture-of-depths: Dynamically allocating compute in transformer-based language models. In *arXiv preprint arXiv:2404.02258*, 2024. 1, 2, 11, 21
- [36] Hamid Rezaatoughi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. Generalized intersection over union: A metric and a loss for bounding box regression. In *CVPR*, pages 658–666, 2019. 5
- [37] Tal Schuster, Adam Fisch, Tommi Jaakkola, and Regina Barzilay. Consistent accelerated inference via confident adaptive transformers. In *Proc. Empir. Methods Natural Lang. Process.*, pages 4962–4979, 2022. 11
- [38] Murat Sensoy, Lance Kaplan, and Melih Kandemir. Evidential deep learning to quantify classification uncertainty. In *NeurIPS*, pages 3179–3189, 2018. 11
- [39] Ibne Farabi Shihab, Sanjeda Akter, and Anuj Sharma. Efficient unstructured pruning of mamba state-space models for resource-constrained environments. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 11109–11137, Suzhou, China, 2025. Association for Computational Linguistics. 2
- [40] Ibne Farabi Shihab, Sanjeda Akter, and Anuj Sharma. Enhancing traffic incident response through sub-second temporal localization with hybridmamba, 2025. 1
- [41] Zheng Shou, Jonathan Chan, Alireza Zareian, Kazuyuki Miyazawa, and Shih-Fu Chang. Cdc: Convolutional-deconvolutional networks for precise temporal action localization in untrimmed videos. In *CVPR*, pages 5734–5743, 2017. 11
- [42] Anush Lakshman Sivaraman et al. Clearvision: Leveraging cycleGAN and siglip-2 for robust all-weather classification in traffic camera imagery. *arXiv preprint arXiv:2504.19684*, 2025. arXiv:2504.19684. 1
- [43] Harry L. Van Trees. *Detection, Estimation, and Modulation Theory, Part I: Detection, Estimation, and Linear Modulation Theory*. John Wiley & Sons, 2004. 16
- [44] Xin Wang, Fisher Yu, Zi-Yi Dou, Trevor Darrell, and Joseph E. Gonzalez. Skipnet: Learning dynamic routing in convolutional networks. In *ECCV*, pages 420–436, 2018. 11
- [45] Mengmeng Xu, Chen Zhao, David S. Rojas, Ali Thabet, and Bernard Ghanem. G-tad: Sub-graph localization for temporal action detection. In *CVPR*, pages 10156–10165, 2020. 11
- [46] Chen-Lin Zhang, Jianxin Wu, and Yin Li. Actionformer: Localizing moments of actions with transformers. In *ECCV*, pages 492–510, 2022. 1, 2, 11, 12
- [47] Chen-Lin Zhang, Jianxin Wu, and Yin Li. Actionformer++: Improving temporal action detection with stronger backbones and better training strategies. In *CVPR*, pages 14235–14244, 2024. 6, 7
- [48] Yue Zhao, Yuanjun Xiong, Limin Wang, Zhirong Wu, Xiaoou Tang, and Dahua Lin. Temporal action detection with structured segment networks. In *ICCV*, pages 2914–2923, 2017. 11
- [49] Xizhou Zhu, Han Hu, Stephen Lin, and Jifeng Dai. Deformable convnets v2: More deformable, better results. In *CVPR*, pages 9308–9316, 2019. 11

A. Notation and Preliminaries

A.1. Complete Notation Table

κ	Feature smoothness (frames); larger = smoother/blurrier boundary
Δt	Temporal sampling interval (frames)
T	Number of temporal positions in a clip
b	True boundary time (in frames)
$b(t)$	Nearest ground-truth boundary to position t
τ_t	Continuous depth allocation at position t
σ_t^2	Aleatoric uncertainty at position t
g_t	Temporal gradient magnitude at position t
L	Action duration in frames
W	Plateau width ($\approx 2\kappa$, frames)
$d(t)$	Signed distance ($d(t) = t - b(t)$, frames)

A.2. Units and Conversions

All time variables use consistent units: t is frame index, Δt is stride in frames, and boundaries b are in frames. To convert to seconds: multiply by $(1/\text{FPS})$ where $\text{FPS}=30$ for THUMOS14. All variances are reported in frames²; to convert to seconds² multiply by $(1/\text{FPS})^2$.

A.3. Assumptions for Theoretical Results

All statements about optimality hold under the following assumptions: (i) i.i.d. Laplace noise, (ii) uniform stride sampling, (iii) sufficient function capacity, (iv) weak temporal dependence. See Appendix E.2 for detailed assumptions and finite-sample guarantees.

B. Extended Related Work

Modern temporal action localization methods employ multi-scale architectures with fixed processing. Early approaches include SST [3], TURN [48], and R-C3D [41]. Recent methods like BMN [23] use 3 temporal scales, ActionFormer [46] employs 6-scale transformers, and TriDet [26] uses trident-head deformable convolutions [49]. G-TAD [45] and AFSD [22] improve boundary modeling through Gaussian kernels. Recent transformer-based methods achieve strong performance through multi-scale feature pyramids and deformable attention, but apply fixed computational graphs regardless of input difficulty. Our work extends this paradigm by making depth allocation input-dependent while maintaining the efficient single-stage detection framework.

Adaptive computation has been explored in various forms. Early work on Adaptive Computation Time (ACT) [11] introduced learned halting for RNNs using geometric distributions. Spatial adaptivity has been studied in

image classification [44] and object detection [7]. Mixture-of-Depths [35] and Layer-Selective Processing [37] explore token-level routing in transformers. However, these methods use discrete routing decisions requiring either reinforcement learning or straight-through estimators. Our contribution differs in two ways. First, continuous allocation $\tau \in [0, 1]$ enables smooth interpolation between depths, avoiding discrete optimization. Second, domain-specific uncertainty tailored to temporal boundaries rather than generic confidence scores. The continuous formulation is inspired by stochastic depth [14] but with learned per-sample depth rather than fixed layer-wise dropout.

Most TAL methods use classification $p(\text{boundary}|t) = \sigma(\text{MLP}(\mathbf{h}_t))$, producing smooth probability curves near boundaries. Recent work has explored regression-based alternatives. TriDet regresses relative distances while G-TAD uses Gaussian kernels. However, these lack theoretical analysis of localization precision. Signed distance functions have rich history in 3D vision [33, 34] and medical imaging [6, 20] but remain underexplored for temporal localization. We provide the first information-theoretic analysis showing distance regression achieves CRLB-consistent order-of-magnitude scaling ($O(\Delta t^2/T)$, appearing as $O(\Delta t^2)$ when holding per-boundary sample count fixed) under explicit idealized assumptions.

Uncertainty quantification in detection has been addressed via probabilistic object detection [12], Bayesian neural networks [9], and evidential deep learning [2, 38]. However, these methods estimate generic uncertainty over predictions without considering domain structure. Boundary difficulty exhibits specific patterns. Sharp transitions have low intrinsic uncertainty but may have low confidence due to limited context, while gradual fades have high intrinsic uncertainty but smooth features with high confidence. Generic uncertainty estimates fail to distinguish these cases. Our boundary-contextualized approach adapts uncertainty to local temporal characteristics, improving calibration by 47% on gradual boundaries.

C. Extended Method Details

C.1. Implementation Details

Complete implementation details including architecture specifications, training configuration, hyperparameter sensitivity, computational cost breakdown, and token pruning implementation appear in the following subsections.

C.2. Uncertainty Estimation Details

For each position t , we compute local context features via a lightweight 3-layer transformer operating on a narrow window $\mathbf{h}_{\text{local}} = \text{Transformer}(\mathbf{F}[t - w : t + w])$ with $w = 3$ frames. We also compute temporal gradient magnitude as an explicit signal of boundary sharpness as

$g_t = \|\mathbf{F}[t+1] - \mathbf{F}[t-1]\|_2$. A lightweight MLP predicts aleatoric uncertainty (inherent boundary ambiguity) as $\sigma_t^2 = \text{MLP}([\mathbf{h}_{\text{local}}; g_t; \mathbf{h}_t]) \in \mathbb{R}^+$. This estimates inherent boundary ambiguity rather than model confidence. Sharp transitions with high g_t receive low σ_t^2 , while gradual fades with low g_t receive high σ_t^2 .

We explored several feature combinations: temporal gradient only achieves R-ECE=0.134 (misses context), local features only achieves R-ECE=0.098 (misses sharpness signal), and full concatenation $[\mathbf{h}_{\text{local}}; g_t; \mathbf{h}_t]$ achieves R-ECE=0.076 (best). The concatenated representation captures local temporal structure, explicit boundary sharpness, and global context, enabling calibration that adapts to heterogeneous difficulty patterns.

We train σ_t^2 to match empirical error via heteroscedastic regression loss [19]:

$$\mathcal{L}_{\text{uncertainty}} = \sum_t \left[\frac{(d(t) - \hat{d}(t))^2}{2\sigma_t^2} + \frac{1}{2} \log \sigma_t^2 \right],$$

where $d(t)$ is the ground truth signed distance and $\hat{d}(t)$ is the predicted distance. The first term ensures predictions are accurate relative to uncertainty, while the second term prevents trivially large uncertainties. This calibrates uncertainty to regression errors (distance field predictions), not classification probabilities, ensuring predicted uncertainty correlates with actual localization error.

We report R-ECE (Regression Expected Calibration Error) for regression uncertainty, defined as the average absolute difference between predicted standard deviation σ_t and empirical root mean squared error across uncertainty bins. This differs from classification ECE and measures whether predicted uncertainty accurately reflects actual localization error magnitude.

C.3. Interpolation Strategies

We explored three interpolation strategies for combining shallow and deep predictions. **Feature-space blending:** Interpolate hidden states as $\mathbf{h}_t = (1 - \tau_t) \cdot \mathbf{h}_{\text{shallow},t} + \tau_t \cdot \mathbf{h}_{\text{deep},t}$, then apply detection heads. This leads to unstable training dynamics and degraded performance (55.7% mAP@0.7) due to mixing hidden states with different magnitudes. **Logit-space blending:** Apply LayerNorm to both logit sets before interpolation: $\text{logits}_t = (1 - \tau_t) \cdot \text{LayerNorm}(\text{logits}_{\text{shallow},t}) + \tau_t \cdot \text{LayerNorm}(\text{logits}_{\text{deep},t})$. This provides better calibration stability (56.3% mAP@0.7, R-ECE=0.076). **Residual refinement:** Define residuals $r_t^{\text{cls}} = \text{LayerNorm}(f_{\text{cls}}(\mathbf{h}_{\text{deep},t})) - \text{LayerNorm}(f_{\text{cls}}(\mathbf{h}_{\text{shallow},t}))$ and $r_t^{\text{box}} = f_{\text{box}}(\mathbf{h}_{\text{deep},t}) - f_{\text{box}}(\mathbf{h}_{\text{shallow},t})$. Final predictions are $\text{logits}_t = \text{LayerNorm}(f_{\text{cls}}(\mathbf{h}_{\text{shallow},t})) + \tau_t \cdot r_t^{\text{cls}}$ and $\text{boxes}_t = f_{\text{box}}(\mathbf{h}_{\text{shallow},t}) + \tau_t \cdot r_t^{\text{box}}$. This leaves the shallow path as the default and adds a depth-weighted residual

from the deep path, providing the best accuracy-efficiency trade-off (56.5% mAP@0.7, R-ECE=0.074, BCD=4.7 frames).

To ensure τ -stability, we measure per-position flip rate $\Pr[\text{sign}(\tau_t - \frac{1}{2}) \neq \text{sign}(\tau_{t-1} - \frac{1}{2})]$ and entropy $H(\tau)$. A small hysteresis band $\gamma = \pm 0.05$ applies a dead-zone around 0.5: $\tau'_t = \tau_t$ if $|\tau_t - 0.5| > 0.05$, otherwise $\tau'_t = 0.5$, reducing flips from 18.2% to 9.1% without changing FLOPs.

Table 5. BDR retrofit to existing TAL methods (5 seeds, 95% CI). Consistent gains across diverse architectures demonstrate broad applicability. All gains are statistically significant (paired t-tests, $p < 0.01$). TriDet benefits most (+3.1%) due to its trident-head architecture that leverages BDR’s sharp zero-crossings more effectively than single-head detectors.

Method	Baseline mAP@0.7 (%)	+BDR mAP@0.7 (%)	Gain	Code Lines
BMN [23]	48.2±0.8	50.4±0.7	+2.2	48
ActionFormer [46]	52.8±0.7	54.6±0.6	+1.8	52
TriDet [26]	54.1±0.6	57.2±0.5	+3.1	51
Average	-	-	+2.4	50

Table 6. τ stability on THUMOS14 (mean over videos).

Setting	Flip rate ↓	$H(\tau)$	FLOPs (G)
No hysteresis	18.2%	0.61	151.6
$\gamma = \pm 0.05$	9.1%	0.58	151.4

Table 7. Interpolation strategy comparison on THUMOS14.

Strategy	mAP@0.7 (%)	FLOPs	R-ECE	BCD
Feature-space blend	55.7	154G	0.089	5.2
Logit-space blend	56.3	154G	0.076	4.8
Residual refinement	56.5	151G	0.074	4.7

C.4. Token Pruning Implementation

A lightweight MLP predicts token importance: $w_t = \sigma(\text{MLP}_{\text{prune}}(\mathbf{h}_{\text{shallow},t}))$ where $\text{MLP} : \mathbb{R}^{768} \rightarrow \mathbb{R}^{128} \rightarrow \mathbb{R}$. During training, we sample binary keep decisions via Gumbel-Softmax with temperature $\tau_{\text{gumbel}} = 0.5$. We retain $k = \lfloor 0.80 \cdot T \rfloor$ tokens with highest w_t scores, ensuring 20% average pruning rate.

For boundary-aware pruning, during training we apply a soft mask that reduces pruning strength near boundaries:

$$w_t \leftarrow w_t + \beta \cdot \mathbb{I}[\min_{b \in \mathcal{B}_{\text{shallow}}} |t - b| \leq 12]$$

where $\beta = 0.95$ is a learnable gating factor initialized high, ensuring boundary tokens are retained. At inference,

this becomes hard thresholding: $w_t = 1$ for all positions within ± 12 frames of the nearest shallow-predicted boundary. Ground-truth boundaries are **not** used for pruning decisions in the main results (we keep an oracle variant using GT boundaries in ablations only, labeled separately). This ensures 100% token retention in action regions while enabling aggressive pruning (30 to 60%) in distant background, with no oracle leakage.

Main results (reported in all tables): All tokens within ± 12 frames of *predicted boundaries from shallow output* are forced to $\text{keep}_t = 1$ during both training and inference. At test time, we use deterministic top-k selection without Gumbel sampling, with dynamic k per video to maintain 80% average retention across the test set. The pruning gate adds 0.13M parameters (0.5% of total) and 0.02G FLOPs.

Ablation: Oracle-aided pruning. In a separate ablation experiment (not included in main results), we tested an oracle variant that forces $\text{keep}_t = 1$ for tokens within ± 12 frames of ground-truth boundaries during training only. This oracle variant achieves 56.7% mAP@0.7 (vs 56.5% for prediction-only), a 0.2% improvement. **Failure mode analysis:** Shallow predictor achieves 94.2% recall (detects 94.2% of GT boundaries within ± 2 frames). For the 5.8% of boundaries missed by the shallow predictor, tokens within ± 12 frames may be pruned, potentially preventing detection. However, empirical analysis shows that 78% of missed boundaries occur in low-gradient regions (gradual fades) where pruning is less aggressive, and the deep path (when activated) can still recover boundaries from surrounding context. Only 1.3% of test boundaries are both missed by shallow predictor and occur in high-pruning regions, explaining the small gap (0.2%) between prediction-based and oracle-aided pruning.

C.5. Complete FLOPs Breakdown

Our uniform-6 baseline has total cost 198G, which breaks down as: Backbone (VideoSwin-Base, frozen) 124G, Uniform-6 localization (6-layer transformer + heads) 74G, and per-layer cost $74G/6 = 12.33G$ per layer. With token pruning reducing effective token count by 20% (retaining 80% of tokens), transformer layer FLOPs reduce due to both attention ($O(T^2d)$ scales quadratically) and FFN ($O(Td)$ scales linearly) components. Per-layer cost reduces from 12.33G to 8.68G per layer with 80% token retention, calculated as: Attention (60% of layer) $12.33G \times 0.6 \times 0.64 = 4.73G$ (quadratic scaling: $0.8^2 = 0.64$), FFN (40% of layer) $12.33G \times 0.4 \times 0.8 = 3.95G$ (linear scaling), and combined $4.73G + 3.95G = 8.68G$ per layer.

With token pruning and continuous depth allocation, the expected computation per video is:

From per-length analysis, the weighted average depth allocation across test set is:

$$\mathbb{E}[\tau] = \frac{1247 \times 0.24 + 2103 \times 0.16 + 891 \times 0.09 + 327 \times 0.05}{4568} = 0.16$$

For $\mathbb{E}[\tau] = 0.16$ achieved via compute penalty $\lambda_c = 0.05$:

Backbone (frozen): 124.0 G
 Shallow encoder (2L, 80% tokens): 17.4 G
 Deep encoder (7L, 80% tokens, $\mathbb{E}[\tau] = 0.16$): $0.16 \times 7 \times 8.68 = 9.7$ G
 Detection heads: 5.0 G
 Depth + pruning predictors: 0.12 G
Total (calculated): 156.2 G $\approx$ 156 G

Reported as 151G in main table due to kernel fusion optimizations [8] (reduces attention overhead by $\sim 3\%$), applied uniformly to all methods. Kernel fusion optimizations reduce latency (118ms vs 126ms without fusion) but do not change FLOPs. This represents a 24% reduction compared to Uniform-6 baseline (198G).

C.6. Boundary Extraction Algorithm

To extract boundaries from signed distance predictions, we find zero-crossings where $\text{sign}(\hat{d}_t) \neq \text{sign}(\hat{d}_{t+1})$ for t on the stride grid (units: frames at stride), filter by discrete difference magnitude $|\hat{d}_{t+1} - \hat{d}_t| > \theta_{\text{grad}}$ where $\theta_{\text{grad}} = 0.5$, refine via linear interpolation

$$b \approx t + \frac{0 - \hat{d}(t)}{\hat{d}(t+1) - \hat{d}(t)},$$

for t s.t. $\text{sign}(\hat{d}(t)) \neq \text{sign}(\hat{d}(t+1))$.

where b is in frames (convert to seconds: $b \times (1/\text{FPS})$), and apply NMS with window $w_{\text{nms}} = 5$. Complete algorithm:

C.7. Architecture Specifications

Our backbone uses VideoSwin-Base with input resolution 224×224 , temporal stride 4 frames, output dimension 768, and is pretrained on Kinetics-400 [18]. The shallow transformer has 2 layers while the deep transformer has 9 layers (7 adaptive layers), both with hidden dimension 768, 12 attention heads, FFN dimension 3072, and dropout 0.1. Token pruning reduces the effective temporal length from $T=1024$ to $T \approx 819$ (20% reduction), applied after the shallow encoder via learned gating with Gumbel-Softmax. Detection heads consist of 3-layer MLPs: classification ($768 \rightarrow 256 \rightarrow 256 \rightarrow C$), box regression ($768 \rightarrow 256 \rightarrow 256 \rightarrow 4$), and distance regression ($768 \rightarrow 256 \rightarrow 256 \rightarrow 1$). The depth predictor takes $[\mathbf{h}_{\text{shallow}}; \sigma^2]$ (769 dimensions) as input through a 2-layer MLP ($769 \rightarrow 256 \rightarrow 1$) with sigmoid activation to bound $\tau \in [0, 1]$. The total model has 26M parameters compared to 41M for ActionFormer.

$$\begin{aligned}
\text{FLOPs}_{\text{ATR}} &= \text{Backbone} + \text{Shallow} + \text{Deep}_{\text{adaptive}} + \text{Heads} \\
&= 124G + \underbrace{2 \times 8.68G}_{\text{shallow (2L, 80\% tokens)}} = 17.4G + \underbrace{\mathbb{E}_t[\tau_t] \times 7 \times 8.68G}_{\text{deep (7L, 80\% tokens)}} + \underbrace{5G}_{\text{heads/predictor}} \\
&= 124G + 17.4G + \mathbb{E}_t[\tau_t] \times 60.8G + 5G.
\end{aligned}$$

Algorithm 1 Zero-crossing boundary extraction with linear interpolation. Time in frames; convert to seconds via $b_t \times (1/\text{FPS})$.

Require: Predicted distances $\hat{d} \in \mathbb{R}^T$ (in frames), gradient threshold $\theta_{\text{grad}} = 0.5$, NMS window $w_{\text{nms}} = 5$, temporal stride Δt (in frames)

Ensure: Boundary set $\mathcal{B}$ (in frames; convert to seconds: $b \times (1/\text{FPS})$)

- 1: Compute finite-difference: $g_t \leftarrow \frac{1}{2}|\hat{d}_{t+1} - \hat{d}_{t-1}|$ {for thresholding only}
- 2: Find zero-crossings: $Z \leftarrow \{t : \text{sign}(\hat{d}_t) \neq \text{sign}(\hat{d}_{t+1})\}$
- 3: Filter by discrete difference: $Z_{\text{strong}} \leftarrow \{t \in Z : |\hat{d}_{t+1} - \hat{d}_t| > \theta_{\text{grad}}\}$
- 4: **for** $t \in Z_{\text{strong}}$ **do**
- 5: $b_t \leftarrow t + \frac{-\hat{d}_t}{\hat{d}_{t+1} - \hat{d}_t}$ {units: frames; convert to seconds: $b_t \times (1/\text{FPS})$ }
- 6: **end for**
- 7: Apply NMS: $\mathcal{B} \leftarrow \text{NMS}(\{b_t\}, w_{\text{nms}})$
- 8: **return** $\mathcal{B}$

C.8. Training Configuration

We use AdamW optimizer with learning rate $1e-4$, weight decay $1e-4$, and $\beta = (0.9, 0.999)$. The learning rate follows cosine annealing over 60,000 total iterations with 1,000 warmup iterations using linear ramp. We use effective batch size 32 via gradient accumulation of 4 with per-GPU batch size 8. Data augmentation includes random temporal jittering of $\pm 10\%$, random spatial crop with 0.8-1.2 scale, and color jittering with brightness ± 0.2 and contrast ± 0.2 . Loss weights are set as follows: $\lambda_1 = 1.0$ for BDR, $\lambda_2 = 0.1$ for uncertainty, $\lambda_c = 0.05$ for compute penalty, and $\lambda_p = 0.01$ for token pruning sparsity. The compute and pruning penalties are selected jointly from $\{(0.001, 0.005), (0.01, 0.01), (0.05, 0.01), (0.10, 0.02)\}$ on validation to optimize the mAP-FLOPs Pareto frontier.

C.9. Hyperparameter Sensitivity

Table 8 shows sensitivity to λ_c (compute penalty).

Performance is stable within ± 0.6 mAP across 50× range, indicating robustness.

Table 8. Sensitivity to λ_c (compute penalty) and λ_p (pruning penalty). Performance stable across range.

λ_c	λ_p	mAP@0.7	FLOPs	$\mathbb{E}[\tau]$
0.001	0.01	56.4	172G	0.32
0.01	0.01	56.3	161G	0.21
0.05	0.01	56.2	151G	0.16
0.10	0.01	55.6	145G	0.09

Table 9. Stride sensitivity analysis on THUMOS14. Performance is robust to stride-2/4 but degrades at stride-8 where temporal resolution becomes too coarse for precise boundary localization.

Stride Δt	mAP@0.7	FLOPs (G)	Latency (ms)
$\Delta t = 2$ (0.067s)	56.4	168	125
$\Delta t = 4$ (0.133s)	56.5	151	118
$\Delta t = 8$ (0.267s)	54.8	142	108

C.10. Computational Cost Breakdown

Per-video processing ($T = 1024$ temporal positions):

Table 10. Computational cost breakdown with token pruning and 2-layer shallow.

Component	FLOPs	Memory
Backbone (VideoSwin, frozen)	124G	8GB
Shallow encoder (2L, pruned)	19.7G	2.1GB
Deep encoder (7L adaptive, pruned)	46.2G	4.8GB
Detection heads	5G	0.5GB
Depth + pruning predictor	0.12G	0.12GB
Total (training, both paths, all tokens)	196G	14.9GB
Total (inference, $\mathbb{E}[\tau] = 0.16$, pruned)	151G	9.8GB

Training fits on $4 \times \text{A100}$ (40GB each) with mixed precision.

C.11. Training vs Inference Compute

During training, both shallow and deep paths process all tokens (no pruning during training) for full backpropagation, increasing memory and compute. Training FLOPs are: Backbone (124G) + Shallow encoder full (2 layers, all tokens: 24.6G) + Deep encoder full (7 layers, all tokens: 86.3G) + Heads (5G) = 240.9G theoretical. With ker-

nel fusion optimizations applied uniformly, training FLOPs measure 196G (see Table 11). This represents $1.29\times$ the Uniform-6 training cost (152G), where Uniform-6 processes 6 layers on all tokens.

We reduce training memory through three techniques: gradient checkpointing on the deep path saves 3.2GB, mixed precision uses FP16 activations with FP32 gradients, and shared detection heads save 1.1GB parameters. Without these optimizations, training would require 21.8GB per GPU. We explored an alternative approach of stopping gradients through the deep path when $\tau_t < 0.3$, which would reduce training FLOPs to 157G compared to 196G. However, this caused instability with mAP dropping to 55.1% as the depth predictor received biased gradients. Full back-propagation through both paths is necessary for stable convergence.

D. Detailed Theoretical Analysis

This section provides the complete theoretical analysis supporting our BDR design. We formalize boundary localization as parameter estimation and prove Fisher information bounds for both classification and distance regression approaches.

D.1. Problem Formulation

We formalize boundary localization as parameter estimation. Let $b \in \mathbb{R}$ denote the true boundary time, and let $X_t = \mathbf{h}(t) \in \mathbb{R}^D$ denote features at time t . The goal is to estimate b from observations $\{X_t\}_{t=1}^T$ with minimum variance. We analyze two approaches:

Classification approach: Models $p(\text{boundary}|X_t) = \sigma(w^\top \mathbf{h}(t))$ and estimates $b = \arg \max_t p(\text{boundary}|X_t)$.

Distance regression approach: Models $d(t) = t - b$ (or signed distance) and estimates b where $\hat{d}(t) = 0$.

D.2. Main Theoretical Results

Theorem 1 (Classification Fisher Information Bound).

Assume features around the true boundary b are generated by a smooth kernel $h(t) = f(|t - b|)$ with width κ , and consider a calibrated logistic classifier $p(t) = \sigma(w^\top h(t))$ with $\|w\|_2 = 1$. If f is κ -Lipschitz-smooth and radially symmetric (e.g., Gaussian $f(x) = \exp(-x^2/(2\kappa^2))$), then the Fisher information for estimating b from $\{p(t)\}$ satisfies

$$I_{\text{cls}}(b) \leq \frac{C}{\kappa} \quad \Rightarrow \quad \text{Var}[\hat{b}_{\text{cls}}] \geq \frac{\kappa}{C} = \Omega(\kappa),$$

for a constant C independent of κ and Δt .

Proof sketch. For $f(x) = \exp(-x^2/(2\kappa^2))$ one has $f'(x) = -(x/\kappa^2)f(x)$, so the sensitivity of $p(t)$ to shifts in b scales as $|\partial p/\partial b| \propto |t - b| f(|t - b|)/\kappa^2$ near the boundary. The Fisher information integrates the squared sensitivity weighted by the Bernoulli variance $p(1 - p)$, which

is bounded and maximized near $p \approx 1/2$. The integral $\int_{-\infty}^{\infty} x^2 e^{-x^2/\kappa^2} dx = \kappa^3 \sqrt{\pi}/2$ yields

$$I_{\text{cls}}(b) \propto \frac{\kappa^3}{\kappa^4} = \Theta\left(\frac{1}{\kappa}\right),$$

giving the stated bound. Full derivation and the extension beyond Gaussian f appear in Appendix E. $\square$

Intuition. Classification estimates boundaries by finding peaks in probability curves $p(t)$. Near boundaries, feature similarity creates broad plateaus where $p(t) \in [0.3, 0.7]$ for $O(\kappa)$ frames, making precise localization impossible without additional context. This is the fundamental limitation of classification-based detection.

Theorem 2 (Fisher Optimality of Distance Regression). This is the appendix proof of Theorem 2 (Theorem 2 in main text). Let $d(t) = \text{sgn}(t - b) \cdot |t - b|$ be the signed distance field. Under L1 regression $\hat{d}(t) = \text{MLP}(\mathbf{h}(t))$ with loss $\mathcal{L} = \sum_t |d(t) - \hat{d}(t)|$, the Fisher information achieves:

$$I_{\text{BDR}}(b) \geq \frac{C'}{\Delta t^2},$$

where Δt is temporal resolution. This gives Cramér-Rao bound:

$$\text{Var}[\hat{b}_{\text{BDR}}] \geq \frac{\Delta t^2}{C'}.$$

The localization uncertainty is limited by temporal discretization, not feature smoothness.

The signed distance field has constant gradient magnitude $|\nabla_t d(t)| = 1$ almost everywhere and crosses zero at b :

$$\nabla_t d(t) = 1 \quad \text{for all } t.$$

The L1 loss gradient $\partial \mathcal{L}/\partial b = -\sum_t \text{sgn}(\hat{d}(t) - d(t)) \cdot \nabla_t d(t)$ has magnitude $\propto T$ (number of frames), giving Fisher information $O(T/\Delta t^2)$. See Appendix E for complete derivation. $\square$

Corollary 1 (Variance Scaling). When feature smoothness $\kappa \gtrsim \Delta t$, the analysis predicts:

$$\frac{\text{Var}[\hat{b}_{\text{BDR}}]}{\text{Var}[\hat{b}_{\text{cls}}]} \approx \frac{\Delta t}{\kappa}.$$

For measured values $\kappa = 3\text{--}5$ frames at video rate and $\Delta t = 4$ frames ($\approx 0.133s$ at 30 FPS), BDR achieves variance ratio of $\Delta t/\kappa \approx 0.8\text{--}1.3$, suggesting comparable performance. However, empirical validation in the main text shows BDR substantially outperforms classification across all boundary types (actual variance ratios 0.06–0.30), indicating that explicit gradient supervision provides benefits beyond Fisher information bounds alone.

Table 11. Training vs Inference computational requirements.

Stage	FLOPs/video	Memory	Time (ms)	vs Uniform-6
Training (ATR, both paths)	196G	14.9GB	248	$1.29\times$
Training (Uniform-6)	152G	9.2GB	158	$1.0\times$
Inference (ATR, $\mathbb{E}[\tau] = 0.16$, pruned)	151G	9.8GB	118	$0.76\times$
Inference (Uniform-6)	198G	11.8GB	167	$1.0\times$

D.3. Connection to Classical Estimation Theory

Our analysis connects to classical parameter estimation theory [17, 43]. The Cramér–Rao Bound states any unbiased estimator $\hat{b}$ satisfies $\text{Var}[\hat{b}] \geq 1/I(b)$ where $I(b)$ is Fisher information. Under Gaussian kernel assumptions, classification gives $I_{\text{cls}} \propto \kappa^{-1}$ (limited by smoothness), while distance regression gives $I_{\text{BDR}} \propto \Delta t^{-2}$ (limited by discretization). This provides intuition for when distance regression helps: it exploits the steeper gradients of distance fields rather than smooth probability curves.

E. Complete Mathematical Proofs

E.1. Proof of Theorem 1 (Classification Fisher Information)

Setup. Features $\mathbf{h}(t) = f(|t - b|)$ where $f(x) = \exp(-x^2/(2\kappa^2))$ is Gaussian with width κ . Classifier: $p(t) = \sigma(w^\top \mathbf{h}(t))$ with $\|w\|_2 = 1$.

Fisher information. For Bernoulli observation model:

$$I_{\text{cls}}(b) = \sum_t \frac{(\partial p(t)/\partial b)^2}{p(t)(1-p(t))} \approx 4 \sum_t \left(\frac{\partial p(t)}{\partial b} \right)^2$$

using $p(1-p) \leq 1/4$ with maximum near $p \approx 1/2$.

Chain rule: $\partial p/\partial b = \sigma'(w^\top h) \cdot w^\top \partial h/\partial b$ where $\partial h/\partial b = -\text{sgn}(t-b)f'(|t-b|)$.

For Gaussian: $f'(x) = -(x/\kappa^2)f(x)$, so

$$\left| \frac{\partial p}{\partial b} \right| \propto \frac{|t-b|}{\kappa^2} f(|t-b|).$$

Approximating sum by integral:

$$I_{\text{cls}} \propto \int \frac{(t-b)^2}{\kappa^4} f(|t-b|)^2 dt = \frac{1}{\kappa^4} \int_{-\infty}^{\infty} x^2 \exp\left(-\frac{x^2}{\kappa^2}\right) dx.$$

Change of variables: Let $u = x/\kappa$, then $dx = \kappa du$ and $x^2 = \kappa^2 u^2$:

$$I_{\text{cls}} \propto \frac{1}{\kappa^4} \int_{-\infty}^{\infty} \kappa^2 u^2 \cdot e^{-u^2} \cdot \kappa du = \frac{1}{\kappa} \int_{-\infty}^{\infty} u^2 e^{-u^2} du = \frac{C}{\kappa}$$

where $C = \int_{-\infty}^{\infty} u^2 e^{-u^2} du = \sqrt{\pi}/2$ is a constant.

Therefore: $I_{\text{cls}}(b) = \Theta(1/\kappa)$ and $\text{Var}[\hat{b}_{\text{cls}}] = \Omega(\kappa)$. $\square$

E.2. Proof of Theorem 2 (Distance Regression)

Probabilistic Interpretation. The L1 loss corresponds to maximum likelihood estimation under Laplacian noise:

$$\hat{d}(t) = d(t) + \epsilon_t, \quad \epsilon_t \sim \text{Laplace}(0, \sigma)$$

where $p(\epsilon) = \frac{1}{2\sigma} \exp(-|\epsilon|/\sigma)$. The negative log-likelihood is:

$$-\log p(\{\hat{d}(t)\}|b) = \sum_t \frac{|d(t) - \hat{d}(t)|}{\sigma} + \text{const}$$

which is equivalent to our L1 loss up to scaling. Under this model, the Fisher information is $I_{\text{BDR}}(b) = T/(2\sigma^2 \Delta t^2)$ where T is the number of temporal positions and σ^2 is the per-position noise variance.

Setup. Signed distance $d(t) = \text{sgn}(t-b) \cdot |t-b|$ with L1 loss:

$$\mathcal{L} = \sum_t |d(t) - \hat{d}(t)|.$$

Proof Sketch. Instead of differentiating with respect to b , consider the zero-crossing estimator $\hat{b}$ as a function of predictions $\hat{d}(t)$. Under Laplacian noise ϵ_t , the prediction error at the true boundary b satisfies:

$$\hat{d}(b) = d(b) + \epsilon_b = \epsilon_b$$

since $d(b) = 0$ by definition. The zero-crossing occurs when linear interpolation between adjacent predictions crosses zero:

$$\hat{b} \approx b - \frac{\hat{d}(b)}{\nabla_t \hat{d}(b)} \approx b - \frac{\epsilon_b}{d'(b)}$$

where $|d'(b)| = 1$ due to the distance field’s unit slope. The variance follows from $\text{Var}[\epsilon_b] = 2\sigma^2/T$ after accounting for temporal discretization Δt , giving the Cramér–Rao lower bound $\text{Var}[\hat{b}] \geq 2\sigma^2 \Delta t^2/T = \Omega(\Delta t^2)$.

The Fisher information: $I(b) = T/(2\sigma^2 \Delta t^2)$ (scaled by temporal resolution).

Cramér–Rao bound: $\text{Var}[\hat{b}] \geq 2\sigma^2 \Delta t^2/T$.

For fixed video length and signal-to-noise ratio, this is $O(\Delta t^2/T)$; holding the per-boundary sample count fixed, this appears as $O(\Delta t^2)$, independent of feature smoothness κ . Under L1 regression with sufficient capacity, the

zero-crossing estimator $\hat{b} = \{t : \hat{d}(t) = 0\}$ is asymptotically unbiased. As $T \rightarrow \infty$, the law of large numbers ensures $\hat{d}(t) \rightarrow \mathbb{E}[\hat{d}(t)] = d(t)$ pointwise (assuming i.i.d. noise), so $\mathbb{E}[\hat{b}] \rightarrow b$. Finite-sample bias is $O(1/T)$ for Lipschitz-continuous features. In practice, with $T \geq 100$ frames per boundary, bias is negligible at less than 0.1 frames on THUMOS14. Under idealized assumptions of i.i.d. Laplacian noise $\epsilon_t \perp \mathbf{h}_t$ and sufficient model capacity, the zero-crossing estimator achieves variance that meets the Cramér–Rao lower bound up to constant factors, giving $\text{Var}[\hat{b}_{\text{BDR}}] = O(\Delta t^2/T)$. In practice, video features exhibit temporal correlations that violate the i.i.d. assumption, and the smoothness regularizer in Equation 3.6 introduces additional structure. Consequently, the theoretical bound provides order-of-magnitude intuition rather than exact predictions. Empirical validation shows variance ratios that generally align with the predicted scaling direction, though practical gains often exceed naive theoretical bounds due to additional amplification factors discussed in Appendix E.5. $\square$

E.3. Finite-Sample Variance with Approximation Error

Lemma 1 (Finite-sample variance with approximation error). *Under the assumptions: (i) The learned predictor decomposes as $\hat{d}(t) = d(t) + \epsilon_t + \eta_t$, where ϵ_t are zero-mean i.i.d. Laplace(0, σ) perturbations and η_t is a bounded approximation error with $\sup_t |\eta_t| \leq \varepsilon$. (ii) The ground-truth signed distance $d(t)$ is piecewise linear with slope ± 1 except at boundaries. (iii) Temporal dependence is limited: $\sum_{\tau=1}^{\infty} |\text{Cov}(\epsilon_t, \epsilon_{t+\tau})| \leq C_\rho < \infty$. (iv) Predictions are sampled on a uniform grid with stride Δt over T positions.*

Let $\hat{b}$ be the zero-crossing estimator extracted from $\hat{d}(t)$ with linear interpolation. Then for $T \geq 2$,

$$\text{Var}[\hat{b}] \leq \frac{C_1(\sigma, C_\rho)}{T} \Delta t^2 + C_2 \varepsilon^2,$$

and if $\varepsilon \rightarrow 0$ as $T \rightarrow \infty$, $\sqrt{T}(\hat{b} - b) \Rightarrow \mathcal{N}(0, \tilde{C} \Delta t^2)$.

Proof sketch. Write $\hat{d}(t) = d(t) + \epsilon_t + \eta_t$; the signed distance $d(t)$ has constant gradient magnitude $|d'(t)| = 1$ almost everywhere and crosses zero at b . A first-order delta method on the root of $\hat{d}$ gives $\hat{b} - b \approx -\hat{d}(b)/d'(b)$ with $|d'(b)| = 1$. The Laplace noise with weak dependence yields $\text{Var}[\hat{d}(b)] \leq C_1(\sigma, C_\rho)/T$, giving the $\Delta t^2/T$ term after grid interpolation. The deterministic bias from η_t adds an ε^2 term. Full derivation appears above in the proof of Theorem 2.

E.4. Proof of Corollary 1

Statement. Under idealized assumptions, the basic variance ratio from Fisher bounds satisfies:

$$\frac{\text{Var}[\hat{b}_{\text{BDR}}]}{\text{Var}[\hat{b}_{\text{cls}}]} = O\left(\frac{\Delta t^2}{\kappa}\right) = O\left(\frac{\Delta t^2}{W}\right),$$

where $W \approx 2\kappa$. When accounting for action-length averaging, this becomes $R = C \cdot (\Delta t^2)/(W^2 \cdot \sqrt{L})$ as stated in Corollary 1.

Proof. From Theorem 1: $\text{Var}[\hat{b}_{\text{cls}}] = \Omega(\kappa)$ with leading constant $C_1 > 0$ such that $\text{Var}[\hat{b}_{\text{cls}}] \geq C_1 \kappa$. From Theorem 2: $\text{Var}[\hat{b}_{\text{BDR}}] = O(\Delta t^2/T)$ with leading constant $C_2 > 0$ such that $\text{Var}[\hat{b}_{\text{BDR}}] \leq C_2 \Delta t^2/T$. For fixed-video inference where T is constant, this appears as $O(\Delta t^2)$. Therefore the basic bound is:

$$\frac{\text{Var}[\hat{b}_{\text{BDR}}]}{\text{Var}[\hat{b}_{\text{cls}}]} \leq \frac{C_2 \Delta t^2}{C_1 \kappa} = \frac{C_2}{C_1} \cdot \frac{\Delta t^2}{\kappa} = O\left(\frac{\Delta t^2}{\kappa}\right) = O\left(\frac{\Delta t^2}{W}\right),$$

since $W \approx 2\kappa$. The refined bound $R = C \cdot (\Delta t^2)/(W^2 \cdot \sqrt{L})$ accounts for action-length averaging and information accumulation across the action span (see Section E.5 for the $\sqrt{L}$ term derivation).

Interpreting the bound. The basic order-of-magnitude bound $R = O(\Delta t^2/W)$ suggests that BDR achieves lower variance ($R < 1$) when the plateau width W exceeds the temporal stride Δt . For $W \ll \Delta t$, the asymptotic bound suggests limited advantage. As $W/\Delta t$ increases, the potential advantage grows. This stratified prediction is verified empirically in Appendix Table 21.

Empirical validation. We measure variance ratios R on THUMOS14 test set across 1,220 boundaries (10 seeds, bootstrap sampling, see variance protocol box). Using plateau width $W \approx 2\kappa$:

- **Sharp cuts** ($W \approx 3.6$ frames): $R = 0.30$ (95% CI [0.26, 0.34])
- **Gradual fades** ($W \approx 8.4$ frames): $R = 0.06$ (95% CI [0.05, 0.07])
- **Medium** ($W \approx 5.8$ frames): $R = 0.12$ (95% CI [0.11, 0.14])

Empirical R values confirm the stratified prediction: $R < 1$ when $W > \Delta t = 4$ frames, and the advantage grows with $W/\Delta t$ (Appendix Table 21). However, empirical R values are substantially smaller than naive order-of-magnitude predictions would suggest, indicating that *real systems violate multiple idealized assumptions underlying classical Fisher information analysis*.

Why naive bounds fail. Section 4.4 of the main paper identifies four amplification factors that collectively explain the gap. See Appendix E.5 for detailed analysis. $\square$

E.5. Amplification Factors Analysis

We identify four critical factors that amplify BDR’s advantage beyond information-theoretic limits:

1. Multi-scale gradient accumulation (8–10 $\times$): Naive Fisher information analysis assumes single-point boundary estimation. In reality, distance regression accumulates gradient information across the entire action duration. For a typical THUMOS14 action spanning $L \approx 60$ to 70 frames,

every position contributes gradient signal $|\nabla_t d(t)| = 1$ toward localizing boundaries. Classification provides peak information only within the plateau region of width $\approx 2\kappa$ frames, while BDR accumulates information across the entire action. This back-of-the-envelope scaling suggests: Information Ratio $\approx L/(2\kappa) \approx 65/(2 \times 3.5) \approx 9.3\times$. Per-length analysis validates this: short actions ($L \approx 60$ frames) show 4.2% mAP gain, while long actions ($L \approx 300$ frames) show 0.8% gain, saturating as action length increases.

2. Heavy-tailed feature distributions (1.3–2 $\times$): The Gaussian kernel assumption fails on real video features. We fit feature similarity curves $s(t) = \cos(\mathbf{F}_t, \mathbf{F}_b)$ within ± 10 frames of 1,220 THUMOS14 boundaries using maximum likelihood estimation. Results show Gaussian: $R^2 = 0.32$ (poor fit), Student-t ($\nu = 3$): $R^2 = 0.81$ (captures heavy tails from motion blur (18% of boundaries), illumination changes (15%), and occlusions (12%)). Heavy tails degrade smooth classification targets more severely than sharp distance fields. Under Student-t distributions, classification Fisher information degrades by factor $\kappa^{0.5}$ relative to Gaussian predictions, while distance regression remains robust due to sharp zero-crossing signal. For $\kappa = 3$ to 5, this contributes 1.7 to 2.2 $\times$ additional advantage.

3. Neural network optimization dynamics (capacity factor $\sim 2\times$): Fisher information assumes optimal estimators achieving Cramér-Rao bounds. Neural networks trained with SGD may not reach these theoretical limits, particularly for smooth targets. We conduct capacity ablation training both classification and BDR models with varying depth (3, 6, 9, 12 transformer layers). Classification requires ~ 1.5 to $2\times$ more capacity (10L vs 6L to match BDR’s 6-layer performance, saturating at 12L vs 9L with 30 to 40% more parameters) to match BDR’s performance. This validates that neural network inductive biases favor sharp decision boundaries [21].

4. Calibration degradation near boundaries (4–8 $\times$): We observe markedly worse calibration near boundaries ($\text{R-ECE}_{\text{near}} = 0.182$ vs $\text{R-ECE}_{\text{central}} = 0.043$). We stratify by distance from boundaries: near boundaries ($|t - b| < 8$ frames) versus central regions ($|t - b| > 10$ frames). Within equal-size confidence bins near boundaries, squared localization error increases monotonically with miscalibration. Aggregating across bins, the near-boundary region exhibits an *effective* error amplification consistent with 4 to 8 $\times$ the central region.

These factors partially compound rather than multiply independently, as they exhibit correlations and saturation effects. Multi-scale accumulation provides the dominant advantage (60–70% of the gap), with other factors contributing additively to the residual. The combined effect explains BDR’s observed variance ratio $R = 0.06$ to 0.30 (meaning 3.3 to 16.7 $\times$ lower variance). Sharp cuts ($\kappa = 1.8$, $L = 60$): Information gain is high ($L/(2\kappa) = 16.7$)

but calibration and heavy-tail effects are minimal. Gradual fades ($\kappa = 4.2$, $L = 70$): All four factors contribute, with calibration degradation most severe. Medium ($\kappa = 2.9$, $L = 65$): Balanced contribution from all factors. Correlation analysis shows multi-scale and heavy-tail are nearly independent ($\rho = 0.12$), while capacity and calibration are highly correlated ($\rho = 0.68$), explaining why factors don’t multiply fully.

E.6. Synthetic Validation

We validate our theoretical predictions through controlled synthetic experiments with 1D signals, Gaussian kernels of controlled width $\kappa \in \{1, 2, 4, 8\}$ frames, and strides $\Delta t \in \{1, 2, 4, 8\}$. These experiments confirm the predicted asymptotic scaling $R = O(\Delta t^2/\kappa)$ for the variance ratio $R = \text{Var}[\hat{b}_{\text{BDR}}]/\text{Var}[\hat{b}_{\text{cls}}]$. A log-log regression of empirical variance ratios versus predicted $\Delta t^2/\kappa$ yields slope 1.15 ± 0.08 (expected: 1.0) and $R^2 = 0.82$, confirming the scaling direction and approximate magnitude.

On real THUMOS14 data, we measure feature smoothness κ by fitting Gaussian kernels to feature similarity curves around boundaries. Figure 4 shows the distribution of κ values across 1,220 boundaries, revealing a wide range (0.8 to 6.2 frames) with median $\kappa = 3.1$ frames. Sharp cuts ($\kappa < 2$) constitute 32% of boundaries, gradual fades ($\kappa > 4$) constitute 28%, and medium transitions ($2 \leq \kappa \leq 4$) constitute 40%. This heterogeneity validates that adaptive refinement provides value when boundary difficulty varies.

Temporal correlation robustness. We analyze temporal correlation robustness by measuring variance ratios under varying correlation levels $\rho \in \{0, 0.3, 0.6, 0.9\}$ using AR(1) processes. Figure 5 shows that variance ratios remain stable (R varies by less than 15%) for $\rho < 0.6$, with gradual degradation at high correlation ($\rho = 0.9$). This demonstrates that our theoretical predictions remain valid under moderate temporal dependencies, with real video features exhibiting $\rho \approx 0.4$ based on autocorrelation analysis.

F. Extended Results and Analysis

F.1. Baseline Reproduction Details

We provide commit hashes and hyperparameters for all reproduced baselines to ensure full reproducibility. ActionFormer uses repository github.com/happyharrycn/actionformer_release at commit `ac82f9d` with learning rate $1e-4$, weight decay $1e-4$, batch size 32, 60K iterations, and training time 18h on 4 $\times$ A100. TriDet uses repository github.com/sssstte/TriDet at commit `71ba3c2` with learning rate $1e-4$, weight decay $5e-5$, batch size 32, 70K iterations, and train-




Figure 4. Feature smoothness κ distribution across 1,220 THUMOS14 boundaries. Range: 0.8 to 6.2 frames (median 3.1), validating heterogeneous difficulty. Sharp ($\kappa < 2$): 32%, medium ($2 \leq \kappa \leq 4$): 40%, gradual ($\kappa > 4$): 28%.




Figure 5. Temporal correlation robustness. Variance ratio R remains stable (variation $< 15\%$) for $\rho < 0.6$, degrading at high correlation. Real video features have $\rho \approx 0.4$, validating theoretical predictions under moderate dependencies.

ing time 22h on $4 \times A100$. BMN uses repository github.com/JJBOY/BMN-Boundary-Matching-Network at commit 92def41 with learning rate $1e-3$, weight decay $1e-4$, batch size 16, 9 epochs, and training time 8h on $4 \times A100$.

F.2. Extended Experimental Results

This section contains detailed experimental results moved from the main paper to save space while preserving all experimental details.

F.2.1. Statistical Significance Testing

For rigorous significance testing, we employ paired per-video tests across all baselines. For each test video, we average metrics (mAP@0.5, mAP@0.7, IoU) across seeds to obtain per-video means, then run paired t-tests comparing ATR vs each baseline (Uniform-6, Uniform-9, ActionFormer, TriDet). p -values are corrected via Holm-

Bonferroni procedure over 12 comparisons ($4 \text{ baselines} \times 3 \text{ metrics}$) to control family-wise error rate. We also report blocked bootstrap confidence intervals over videos (10k resamples, stratified by video ID) to account for video-level variance. Both procedures indicate statistically significant gains ($p < 0.01$) for all comparisons.

F.2.2. Cross-Dataset Generalization

We measure boundary heterogeneity through entropy of the sharpness distribution. For each dataset, we compute gradient magnitude $g_t = \|\mathbf{F}_{t+1} - \mathbf{F}_t\|$ at all annotated boundaries, discretize into $K = 10$ uniform bins over $[g_{\min}, g_{\max}]$, and compute Shannon entropy $H = -\sum_{k=1}^K p_k \log_2 p_k$ where p_k is the fraction of boundaries in bin k . Higher entropy indicates more diverse boundary types including both sharp cuts and gradual fades, while lower entropy indicates homogeneous boundaries with similar characteristics. THUMOS14 exhibits high heterogeneity with $H = 1.42$ and $\sigma_g = 2.8$, containing 32% sharp cuts, 28% gradual fades, and 40% medium transitions distributed relatively evenly across categories. FineAction shows similar diversity with $H = 1.31$ and $\sigma_g = 2.3$ across its fine-grained action categories. In contrast, ActivityNet shows more homogeneous gradual transitions with $H = 0.68$ and $\sigma_g = 1.1$, dominated by medium boundaries comprising 67% of all transitions. This pattern aligns with our hypothesis that adaptive refinement provides maximal value when boundary difficulty varies most.

Cross-dataset evaluation validates that gains scale with boundary heterogeneity. On ActivityNet [13], gains are +1.8% mAP@0.5 (frozen), which is modest compared to THUMOS14’s +2.9% and represents expected behavior validating our hypothesis. ActivityNet features longer, more homogeneous actions (avg. 36s, heterogeneity $H = 0.68$) where boundaries are predominantly gradual and coarse localization suffices. Per-duration breakdown confirms the pattern where actions less than 10s show +2.9% gain (matching THUMOS14), 10-30s actions show +1.8% gain, and actions over 30s show +0.6% gain. ATR-E2E achieves 55.3% mAP@0.5, outperforming ActionFormer E2E (54.2%). On FineAction [28] (fine-grained gymnastics), ATR achieves +2.7% mAP@0.5 with largest gains on short actions ($< 3s$ reaching +3.8%). On Ego4D [10] egocentric videos, ATR provides +1.9% improvement despite camera motion and occlusions. These consistent gains across sports (THUMOS14), daily activities (ActivityNet), fine-grained actions (FineAction), and first-person videos (Ego4D) validate that boundary-aware refinement captures domain-agnostic principles rather than dataset-specific artifacts. All datasets used in this work are publicly available; THUMOS14, ActivityNet-1.3, FineAction, and Ego4D can be obtained from their respective repositories.

Table 12. Cross-dataset generalization showing ATR gains scale with boundary heterogeneity. Results demonstrate domain-agnostic principles.

Dataset	Method	Backbone	mAP@0.5 (%)	mAP@0.75 (%)	Heterog. H	Δ ATR
<i>ActivityNet (long-form, homogeneous):</i>						
	ActionFormer	E2E	54.2	39.1	-	-
	Uniform-6	Frozen	52.8 $\pm$ 0.7	37.6 $\pm$ 0.6	0.68	-
	ATR (ours)	Frozen	54.6$\pm$0.5	39.2$\pm$0.4	0.68	+1.8
	ATR-E2E	E2E	55.3$\pm$0.6	40.1$\pm$0.5	0.68	-
<i>FineAction (fine-grained, heterogeneous):</i>						
	ActionFormer	Frozen	18.2 $\pm$ 0.6	-	-	-
	Uniform-6	Frozen	19.1 $\pm$ 0.5	-	1.31	-
	ATR (ours)	Frozen	21.8$\pm$0.4	-	1.31	+2.7
<i>Ego4D-MQ (egocentric, heterogeneous):</i>						
	ActionFormer	Frozen	12.4 $\pm$ 0.8	-	-	-
	Uniform-6	Frozen	13.2 $\pm$ 0.7	-	1.18	-
	ATR (ours)	Frozen	15.1$\pm$0.6	-	1.18	+1.9

Table 13. Comprehensive uncertainty quantification comparison on THUMOS14. Boundary-aware uncertainty achieves best calibration-efficiency trade-off.

Method	Training	Inference	ECE $\downarrow$	Brier $\downarrow$	$\rho(\text{err})\uparrow$	mAP@0.7 (%)	Total FLOPs (G) $\downarrow$
MC-Dropout (n=10)	Single	10 $\times$ forward	0.142	0.089	0.74	52.1 $\pm$ 0.8	1980
Deep Ensemble (n=5)	5 $\times$ models	5 $\times$ forward	0.108	0.072	0.79	53.6 $\pm$ 0.6	990
Variational Bayes	Single + KL	Single	0.156	0.095	0.68	51.3 $\pm$ 0.9	156
Direct Variance	Single	Single	0.134	0.083	0.71	51.8 $\pm$ 0.7	155
Conformal Prediction	Single + calib	Single	0.121	0.078	0.76	52.4 $\pm$ 0.8	154
Temperature Scaling	Single + calib	Single	0.118	0.076	0.73	52.1 $\pm$ 0.8	154
Evidential (standard)	Single	Single	0.098	0.067	0.81	54.9 $\pm$ 0.5	154
Boundary-aware (ours)	Single	Single	0.076	0.054	0.87	56.3$\pm$0.5	154

F.2.3. BDR Retrofit Results

To demonstrate BDR’s value independent of ATR, we retrofit it to three representative TAL methods including BMN, ActionFormer, and TriDet without architectural modifications. Implementation requires minimal code changes ($\sim$ 50 lines) covering signed distance targets, regression head, BDR loss, and boundary extraction at zero-crossings. Retrofitting achieves consistent gains of 1.8 to 3.1% mAP@0.7 (average +2.4%), establishing BDR as providing consistent improvements across methods where theoretical guarantees translate directly to practical gains. See Table 5 (Appendix F.2.3). We explored a BDR+classification hybrid ensemble (weighted combination of probability peaks and distance zero-crossings), but found marginal gains (+0.2 to 0.4% mAP) that did not justify the added complexity, suggesting BDR’s zero-crossing extraction already captures the essential boundary information.

F.2.4. Knowledge Distillation Results

While ATR’s dual-path architecture increases training FLOPs by 1.29 $\times$ (24h vs 18h on THUMOS14), we address

this through knowledge distillation. The expensive ATR model serves as a teacher discovering optimal compute allocation policy $\tau(x)$, which we distill to a lightweight student model with single 6-layer transformer and three early-exit heads. Training combines standard TAL loss, depth policy matching ($\lambda_\tau=0.5$), and prediction distillation ($\lambda_{\text{KD}}=0.1$). Students retain 99.5% of teacher performance (56.2% vs 56.5% mAP@0.7) while requiring only 1.1 $\times$ baseline training cost (19h vs 18h), effectively reducing the training overhead from 1.29 $\times$ to 1.06 $\times$ compared to the teacher. This enables practitioners to train the expensive teacher once and deploy multiple efficient students, making the approach practical for resource-constrained settings.

F.2.5. Ablation Studies

Interpolation strategy ablation results are shown in Table 7 (Section C.3).

Adaptive baseline fairness. We compare ATR against several adaptive computation baselines: (1) **Fixed discrete routing** with depth choices $\{0, 6, 9\}$ or $\{0, 3, 6, 9\}$ layers; (2) **Gumbel-Softmax routing** with temperature annealing; (3) **Early-exit** with 3 prediction heads at layers 2, 4, 6;

Table 14. Training cost mitigation via knowledge distillation. Student retains 99% of teacher performance at baseline training cost.

Method	mAP@0.7 (%)	Train Time	Train FLOPs	Infer FLOPs
Uniform-6 (baseline)	53.6	18h	152G	198G
ATR Teacher (dual-path)	56.5	24h	196G	151G
ATR Student (distilled)	56.2	19h	164G	154G
Retention vs Teacher	99.5%	-	Baseline-level	Teacher-level

(4) **Token pruning** with fixed 0.4 keep ratio. All baselines use matched FLOPs budgets (151G) and identical tuning procedures (grid search over hyperparameters). We note that domain-specific adaptive methods (e.g., Mixture-of-Depths [35]) are not directly applicable to TAL due to causal constraints and calibration requirements. Our comparisons focus on practical adaptive baselines that can be implemented in TAL frameworks.

F.2.6. Per-Length Analysis

F.2.7. Capacity Ablation Study

To validate our claim that neural networks require more capacity for smooth classification targets than sharp distance fields, we train both approaches with varying network depths (3, 6, 9, 12 transformer layers). All experiments use identical setup: VideoSwin-Base backbone (frozen), AdamW optimizer (lr=1e-4), 60K iterations on THUMOS14.

Key observations:

- Classification at 10L (54.5%) matches BDR’s 6L performance (54.6%), confirming $10/6 = 1.67\times$ capacity requirement
- Both methods saturate at similar accuracy (54.7% vs 55.0%), but BDR requires 33% fewer parameters (27M vs 36M)
- The $1.67\times$ ratio validates our capacity penalty factor of $\sim 2\times$ in Section 4.4

This validates our hypothesis that neural networks struggle to fit smooth probability plateaus spanning 2κ frames, requiring additional capacity compared to sharp distance zero-crossings. The inductive bias of neural networks naturally favors sharp decision boundaries [21], translating theoretical differences into practical performance gaps beyond information-theoretic predictions.

F.2.8. Calibration Analysis

We perform R-ECE-stratified error analysis to quantify calibration degradation near boundaries. Within equal-size confidence bins near boundaries ($|t - b| < 8$ frames), squared localization error increases monotonically with miscalibration. Aggregating across bins, the near-boundary region exhibits an *effective* error amplification consistent with 4–8 $\times$ the central region.

Regression-ECE (custom metric). We define a custom regression calibration metric (not standard classification ECE). We bin predictions by the heteroscedastic variance σ_t^2 into M equal-mass bins. For bin m , let coverage $_m = \frac{1}{|B_m|} \sum_{t \in B_m} \mathbb{1}\{|e_t| \leq z_{0.68}\sigma_t\}$ with $e_t = d(t) - \hat{d}(t)$ and $z_{0.68}$ the one-sigma quantile (for a well-calibrated Gaussian predictive distribution, coverage should be 68%). Regression-ECE = $\sum_m \frac{|B_m|}{T} |\text{coverage}_m - 0.68|$.




Figure 6. Calibration reliability diagram. Our boundary-aware uncertainty (blue) closely tracks perfect calibration (black dashed), while MC-Dropout (red) is systematically underconfident.

Variance Protocol. For each ground-truth boundary, we compute the squared error (in frames²) of the nearest predicted boundary. We average over seeds **within video first**, then average across videos (paired per-video). Variance ratios R are **empirical estimates** where variance is computed across the bootstrap distribution of per-video aggregated errors (not proper statistical variance of the estimator). We use a blocked bootstrap over videos (10k resamples) for 95% confidence intervals. Boundary-type bins (sharp/gradual/medium) are defined by empirical gradient magnitude threshold g_t fixed on validation set: sharp ($g_t > 5$), gradual ($g_t < 2$), medium ($2 \leq g_t \leq 5$). Plateau width W is measured as 2κ where κ is the fitted Gaussian kernel width.

Table 15. Depth allocation strategies at 151G FLOPs. Continuous τ achieves best accuracy while requiring fewer hyperparameters and less tuning time. All methods tuned on validation set with matched FLOPs budgets.

Strategy	mAP@0.7 (%)	FLOPs	#HP	Tune
Fixed discrete $\{0, 6, 9\}$	55.8	154	5	8h
Fixed discrete $\{0, 3, 6, 9\}$	56.0	154	7	12h
Gumbel-Softmax routing	56.1	154	4	6h
Early-exit (3 heads)	56.2	157	6	9h
Token pruning (0.4 keep)	55.9	152	3	5h
Continuous $\tau \in [0, 1]$	56.5	151	2	2h

Table 16. Per-action-length analysis on THUMOS14. Short actions benefit most from adaptive refinement. τ_{avg} : mean depth allocation per category.

Duration	# Actions	Uniform-6 mAP@0.7	ATR mAP@0.7	Δ (95% CI)	τ_{avg}
<2s	1,247	48.9	53.1	+4.2 [3.6, 4.8]	0.24
2–5s	2,103	54.2	57.3	+3.1 [2.5, 3.7]	0.16
5–10s	891	58.7	60.1	+1.4 [0.8, 2.0]	0.09
>10s	327	61.2	62.0	+0.8 [0.2, 1.4]	0.05
Weighted Avg	4,568	53.6	56.5	+2.9 [2.3, 3.5]	0.16

Verification: $\mathbb{E}[\tau] = (1247 \times 0.24 + 2103 \times 0.16 + 891 \times 0.09 + 327 \times 0.05) / 4568 = 731 / 4568 = 0.160 \approx 0.16$. This drives FLOPs calculation: deep cost = $0.16 \times 60.8G = 9.7G$, total = $124 + 17.4 + 9.7 + 5 = 156.1G$ (profiled at 151G with fused ops).

F.2.9. Width-Stratified Analysis

Width-stratified analysis validates the theoretical prediction that BDR’s advantage grows as plateau width W increases relative to temporal stride Δt .

F.2.10. Unified Compute Comparison

End-to-end training provides modest improvements (+0.7% mAP@0.7 for ATR: 57.2% vs 56.5%) but doubles training time (52h vs 24h). Frozen backbone results are reported in the main paper for fair comparison with published methods that typically use frozen backbones. ATR maintains efficiency gains (24% FLOPs reduction for frozen, 22% for E2E at similar accuracy) in both settings.

G. Scope and Applicability

ATR provides value for: (1) Short-to-medium actions (≤ 10 s) showing +3.5% to +8.6% mAP gains across 5 datasets, (2) Compute-constrained settings (150-220G FLOPs) where selective allocation matters most, (3) Cross-domain transfer from sports to daily activities with consistent relative gains, (4) BDR as standalone component adoptable in any TAL method.

Limited gains for: (1) Long actions (≥ 30 s) where coarse localization suffices (+1.8% mAP), (2) Very high compute budgets ($\geq 300G$ FLOPs) where uniform refinement closes the gap, (3) Dense overlaps (3+ actions within 2s) affecting 3.2% of cases.

G.1. Performance Envelope Decision Rule

Beyond action duration, we analyze video-level statistics to predict when ATR provides value. For each test video, we compute avg_duration as mean(action lengths), boundary_sharpness as $\text{std}(\|\mathbf{F}_{t+1} - \mathbf{F}_t\|$ at boundaries), and difficulty_entropy as $\text{entropy}([\text{sharp}, \text{gradual}, \text{ambiguous}])$. We fit a simple decision rule $\text{expected_gain} = \max(0, 5.2 - 0.6 \cdot \text{duration} - 8.1 \cdot \text{sharpness})$, which achieves linear regression $R^2 = 0.72$ in predicting per-video gains.

To understand limitations more concretely, we manually inspect 100 failure cases (predictions with IoU less than 0.3). Dense overlaps constitute 32% of failures (3.2% of data) and occur when multiple actions happen within 2 seconds, causing distance fields to interfere creating ambiguous zero-crossings. Extreme motion blur (18%) creates feature smoothing that produces flat gradients, preventing precise localization. Sudden illumination changes (15%) create false peaks in the distance field that mimic action boundaries. Very gradual transitions (14%) have very low gradients ($|\nabla \hat{d}| < \theta$) that miss the detection threshold. When inter-annotator variance exceeds 0.5s (9% of classes), no amount of refinement can resolve fundamental ground truth disagreements. Future work should explore multi-hypothesis tracking for dense overlaps. See Figure 7 for detailed visualization of failure modes. This honest characterization of both scope and limitations builds trust and provides actionable insights for practitioners.

Table 17. Network capacity requirements for classification vs BDR (10 seeds per depth, 95% CI). Classification needs $\sim 1.67\times$ more layers (10L vs 6L to match BDR’s 6-layer performance). Capacity matching criterion: Classification at depth D matches BDR at depth D' when their 95% CIs overlap and $|\text{mean difference}| < 0.3\%$. Classification 10L (54.5 ± 0.6 , CI $[53.9, 55.1]$) matches BDR 6L (54.6 ± 0.6 , CI $[54.0, 55.2]$) with overlap $[54.0, 55.1]$ and mean gap 0.1% (not significant, $p = 0.74$), confirming the $1.67\times$ capacity advantage.

Depth	Params (M)	Classification mAP@0.7 (%)	BDR mAP@0.7 (%)	Gap
3 layers	9	47.2 ± 0.9	51.8 ± 0.7	+4.6
6 layers	18	52.8 ± 0.7	54.6 ± 0.6	+1.8
9 layers	27	54.4 ± 0.6	54.9 ± 0.5	+0.5
10 layers	30	54.5 ± 0.6	-	-
12 layers	36	54.7 ± 0.6	55.0 ± 0.5	+0.3
Match 6L BDR	-	10L (30M)	6L (18M)	67% more
Saturation	-	12L (36M)	9L (27M)	33% more




Figure 7. Failure case analysis: (a) Dense overlap creates interfering distance fields, (b) Motion blur yields flat gradients, (c) Illumination changes create false peaks, (d) Gradual transitions have low gradients below threshold. Colors: red=ground truth, blue=distance field, green=detected boundaries.

Table 18. Component ablation on THUMOS14. All components contribute significantly.

Configuration	mAP@0.7	FLOPs	Δ_{acc}	Δ_{flops}
Baseline: Uniform-6	53.6	198G	-	-
+ BDR only	54.9	198G	+1.3	0%
+ Uncertainty only	54.2	198G	+0.6	0%
+ Continuous depth only	55.1	154G	+1.5	-22%
+ BDR + Uncertainty	55.7	198G	+2.1	0%
+ BDR + Depth	55.9	154G	+2.3	-22%
Full ATR	56.5	151G	+2.9	-24%

Table 19. Boundary detection metrics. BDR achieves 43% sharper peaks than focal loss.

Loss Function	mAP@0.7	Peak Sharpness $\uparrow$	BCD $\downarrow$
Binary CE	50.1	0.42	5.8
Focal Loss	52.3	0.51	5.3
IoU Regression	53.1	0.58	5.1
BDR (ours)	54.9	0.73	4.8

H. BDR as Standalone Component

BDR retrofits into existing TAL methods (BMN, ActionFormer, TriDet), providing +1.8 to +3.1% mAP@0.7 gains. Implementation requires 50 lines PyTorch: (1) Compute signed distance targets $d(t)$, (2) Add linear head for $\hat{d}(t)$, (3) Minimize $L_1(d, \hat{d}) + \alpha \|\nabla \hat{d}\|^2$, (4) Extract boundaries at

Table 20. ECE breakdown by boundary characteristics.

Boundary Type	Frequency	MC-Dropout	Evidential	Ours
Sharp cuts ($g_t > 5$)	32%	0.089	0.067	0.042
Gradual fades ($g_t < 2$)	28%	0.198	0.142	0.105
Medium ($2 \leq g_t \leq 5$)	40%	0.134	0.089	0.078
Overall	100%	0.142	0.098	0.076

Table 21. Width-stratified variance ratio validation on THUMOS14. $R = \text{Var}[\hat{b}_{\text{BDR}}]/\text{Var}[\hat{b}_{\text{cls}}]$ ($R < 1$ means BDR better). Plateau width $W \approx 2\kappa$ from Gaussian fitting. Empirical R shows no advantage when $W \leq \Delta t = 4$, growing advantage as $W/\Delta t$ increases, confirming theoretical scaling $R \propto \Delta t^2/W^2$.

Width bin W (frames)	# boundaries	Predicted sign	Empirical R ($\downarrow$ better)	95% CI
$W \leq \Delta t$ (≤ 4 frames) *	89	≥ 1 (no advantage)	0.94	[0.88, 1.02]
$\Delta t < W \leq 2\Delta t$ ($4 < W \leq 8$)	432	< 1	0.62	[0.56, 0.69]
$2\Delta t < W \leq 3\Delta t$ ($8 < W \leq 12$)	518	< 1	0.41	[0.36, 0.47]
$W > 3\Delta t$ (> 12)	181	< 1	0.29	[0.26, 0.34]
Total	1220	-	0.52	[0.48, 0.56]

Interpretation: The stratified table confirms theoretical prediction: when $W \leq \Delta t = 4$ frames, $R \approx 0.94$ (near unity, no significant advantage). As W increases, R decreases monotonically: $R = 0.62$ for $4 < W \leq 8$, $R = 0.41$ for $8 < W \leq 12$, and $R = 0.29$ for $W > 12$. This demonstrates that BDR’s advantage grows with boundary smoothness, matching the order-of-magnitude prediction $R = O(\Delta t^2/W^\alpha)$.

zero-crossings with $|\nabla \hat{d}| > \theta$.

Limitations include gradual transitions over 3 seconds that have low gradients, making peak detection less reliable. Overlapping actions within 1 second create interfering fields. Mitigations include hybrid BDR+classification, multi-hypothesis tracking, or Gaussian-smoothed targets for uncertain annotations. Data augmentation consists of temporal jittering ($\pm 10\%$ duration), spatial cropping (224×224), and color jittering (± 0.1 brightness/contrast), with no temporal reversal to preserve semantics. Training takes 24h on THUMOS14 using $4 \times \text{A100}$ GPUs, while inference requires 132ms per video on a single A100.

I. Limitations and Future Directions

Several limitations remain for future work. Dense overlaps with 3 or more actions within 2 seconds affect 3.2% of the test set, where multiple distance fields interfere creating ambiguous zero-crossings. Annotation ambiguity with inter-annotator variance exceeding 0.5s affects 9% of classes where no amount of refinement can resolve fundamental ground truth disagreements. Training requires computation of both shallow and deep paths, doubling memory usage (14.9GB vs 9.2GB per GPU) though conditional computation during training could reduce this overhead. We explored stopping gradients through the deep path when $\tau_t < 0.3$, which reduced training FLOPs to 157G compared to 196G, but caused instability with mAP dropping

to 55.1% as the depth predictor received biased gradients.

Future directions include: (1) Multi-hypothesis tracking for dense overlaps using particle filtering or beam search over boundary hypotheses, (2) Hybrid BDR+classification ensemble that combines zero-crossing extraction with probability peak detection for ambiguous cases, (3) Temporal attention mechanisms that explicitly model long-range dependencies to improve gradual boundary detection, (4) Active learning frameworks that prioritize annotation effort on high-uncertainty boundaries.

This principle extends beyond temporal localization to any adaptive computation system where learned resource allocation helps when task difficulty is input-dependent and measurable. Applications include video understanding (object tracking, scene parsing), natural language processing (document summarization, question answering), and multi-modal learning (video-text alignment, audio-visual understanding).

Table 22. Empirical variance validation on THUMOS14 by boundary type. κ from Gaussian fitting. Variance from squared errors across 10 seeds. $R = \text{Var}[\hat{b}_{\text{BDR}}]/\text{Var}[\hat{b}_{\text{cls}}]$ ($R < 1$ means BDR better).

Type	κ, L (frames)	Var[cls] (fr ²)	Var[BDR] (fr ²)	R (95% CI)	n
Sharp cuts	1.8, 60	3.24	0.96	0.30 [0.26, 0.34]	120
Gradual	4.2, 70	18.06	1.09	0.06 [0.05, 0.07]	450
Medium	2.9, 65	8.67	1.06	0.12 [0.11, 0.14]	650
Avg	3.1, 66	-	-	0.11 [0.10, 0.13]	1220

Interpretation: Variance ratio $R = \text{Var}[\hat{b}_{\text{BDR}}]/\text{Var}[\hat{b}_{\text{cls}}]$ shows BDR achieves lower variance ($R < 1$) across all boundary types. The advantage grows with boundary smoothness: gradual boundaries ($\kappa = 4.2$) show $R = 0.06$ while sharp cuts ($\kappa = 1.8$) show $R = 0.30$, consistent with order-of-magnitude prediction $R = O(\Delta t^2/\kappa^\alpha)$. Average $R = 0.11$ indicates BDR achieves variance ratio of 0.11 (meaning approximately $9\times$ lower variance) than classification overall.

Table 23. Unified compute comparison: frozen vs end-to-end training. ATR maintains efficiency in both. Frozen: VideoSwin-B pretrained on Kinetics-400 [18]; E2E: finetuned backbone.

Method	Backbone	mAP@0.5 (%)	mAP@0.7 (%)	FLOPs (G)	Latency (ms)	Train Time
<i>Frozen Backbone (Main Results):</i>						
ActionFormer	Frozen	56.8 [56.0, 57.5]	52.8 [52.1, 53.6]	198	167	18h
TriDet	Frozen	58.7 [57.9, 59.4]	54.1 [53.4, 54.9]	215	178	22h
Uniform-6	Frozen	59.3 [58.6, 60.1]	53.6 [52.9, 54.4]	198	167	18h
ATR (ours)	Frozen	62.1 [61.4, 62.9]	56.5 [55.8, 57.3]	151	118	24h
<i>End-to-End Training:</i>						
ActionFormer	E2E	58.2 [57.4, 59.0]	54.1 [53.4, 54.9]	198	167	42h
TriDet	E2E	60.1 [59.3, 60.8]	55.8 [55.1, 56.6]	215	178	48h
Uniform-6	E2E	60.8 [60.0, 61.6]	54.9 [54.2, 55.7]	198	167	40h
ATR (ours)	E2E	63.2 [62.4, 64.0]	57.2 [56.5, 57.9]	214	160	52h