

# Linear Attention Mechanisms in HERO for Long-Form Video Retrieval Efficiency

Assignee Research

June 8, 2026

## Abstract

This report synthesises findings from 12 peer-reviewed papers addressing the following research question: Does replacing the Temporal Transformer in HERO with a linear attention mechanism improve inference throughput on long-form video retrieval benchmarks while maintaining accuracy. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 7.2/10. This report is a machine-generated literature synthesis and does not constitute original research.

## 1 Introduction

This paper examines: T\*: Re-thinking Temporal Search for Long-Form Video Understanding. Research question: Does replacing the Temporal Transformer in HERO with a linear attention mechanism improve inference throughput on long-form video retrieval benchmarks while maintaining accuracy?.

## 2 Methodology

Systematic literature search across multiple databases yielded 12 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 7.2/10.

## 3 Results

12 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 7.2/10.

## 4 Limitations

This report is a machine-generated literature synthesis and does not constitute original research. Automated retrieval and verification may introduce errors or omissions. Review scores reflect automated assessment, not human peer review. Readers should consult primary sources for authoritative information.

## References

- <http://arxiv.org/abs/2505.21439v1>
- <http://arxiv.org/abs/1906.04972v1>
- <http://arxiv.org/abs/2504.02259v3>