Title

Seven ways to measure fire polygon velocity

Authors

Ty Tuff[1]* and HAL 9000[2]

[1] Environmental Data Science Innovation and Inclusion Lab, Cooperative Institute for Research in Environmental Sciences (CIRES), University of Colorado Boulder, Boulder, CO, USA
[2] OpenAI Institute for helping Ty do projects quickly

*Corresponding author: ty.tuff@colorado.edu

Author ORCID iDs

- Ty Tuff: 0000-0001-5249-5197
- HAL 9000: 0000-000X-XXXX-XXXX

Keywords

wildfire spread; fire perimeter dynamics; polygon velocity; scaling; remote sensing; fire behavior modeling; optimal transport

Declarations of authorship

TT conceived the study, designed the synthetic fire experiments, and led the writing. TT and HAL jointly developed the conceptual framework, implemented the methods, and iterated on revisions.

Data and code availability

Synthetic fire polygon sequences, metric implementations, and analysis scripts are available at [GitHub repository / Zenodo DOI placeholder].

One-sentence public summary

Depending on how you measure it, the same wildfire can look explosive or nearly motionless—our study shows seven ways of defining fire polygon velocity and why choosing the right one matters for science and management.

Abstract

Wildfire "velocity" is widely invoked yet ill-defined for evolving perimeters. Using daily polygons rather than front-normal points, we show that different definitions can yield opposing narratives of the same fire. We introduce a taxonomy that groups seven metrics into core baselines (optimal transport; ΔA/L), diagnostic surge detectors (P95 advance; longest vector), conservative proxies (mean advance; equivalent radius growth), and a niche translation measure (centroid drift). Applied to a controlled, fractal ten-day sequence, all metrics track broad growth but diverge in magnitude and meaning. The key split occurs between the two core

baselines: ΔA/L reports zero velocity when no new area is added, while optimal transport still measures kilometers of internal displacement. This contrast reflects philosophy, not error—whether spread is defined by perimeter-driven area gain or by whole-system movement of burned mass. Parameter choices in optimal transport (e.g., Sinkhorn regularization, ground metric) shift absolute values but do not erase this qualitative difference. The framework clarifies what each metric reveals and obscures, provides defensible anchors for analysis, and offers a decision tree and practical guidance for model validation, ecological inference, and operational use. Choosing a metric is choosing a story about the fire; our taxonomy makes that choice explicit and comparable.

Introduction

Wildfire is not merely a line racing across a map; it is a living geometry—an ever-changing shape that grows, folds, coalesces, and collapses in space and time. For generations, fire behavior science has described that motion at points. Rothermel's (1972) equations distilled spread into local rates through fuels, and Richards' (1990) elliptical model cast the fire front as a simple figure pushed by wind. These frameworks remain foundational and continue to underpin operational tools such as FARSITE and coupled fire–atmosphere systems (e.g., WRF-Fire). Yet events like the 2018 Camp Fire, Australia's 2019–2020 Black Summer, and the 2023 Lahaina fire remind us that fires do more than advance incrementally. They undergo sudden accelerations, spotting cascades, perimeter fusion, and internal collapse—behaviors that resist line-based abstractions of rate of spread.

In the satellite era, we are not handed points or ellipses; we are handed polygons—day-by-day scars traced by Landsat, MODIS, and VIIRS that record the evolving area of burned ground. Translating from local rate of spread to a polygonal velocity is not trivial, because polygons do more than advance: they translate, dilate, and deform. Mapping programs such as the Monitoring Trends in Burn Severity (MTBS; Eidenshink et al. 2007), MODIS (Giglio et al. 2018), VIIRS (Giglio et al. 2020), and ESA FireCCI (Chuvieco et al. 2019) deliver these data, but all share the same limitations. Raster origins impose stair-step edges that inflate distance-based metrics, truncate narrow dendrites and small burns, and blur day-to-day dynamics through temporal compositing. Any definition of "velocity" from polygons is therefore entwined with data artifacts as well as fire behavior.

Faced with this challenge, fire scientists have reached for diverse spatial metrics, each emphasizing different facets of polygon growth. Some focus on whole-system displacement, others on physics-based area gain, still others on directional surges or conservative proxies. Each provides a different lens on fire dynamics. Yet despite their widespread use in model validation, ecological studies, and risk assessments, the implications of choosing one definition over another have never been systematically compared.

Here we take a first step toward that synthesis. Using a controlled synthetic fire-growth sequence, we evaluate seven candidate metrics spanning system-wide, physics-based, diagnostic, proxy, and niche categories. The synthetic test bed isolates methodological behavior from observational noise, providing a clean stage to reveal divergences. Our aim is to establish a taxonomy and decision framework that clarifies how metric choice shapes interpretation—and to guide researchers and practitioners in selecting measures aligned with their ecological and operational questions.

Background and Definitions

Several distinct metrics can be used to characterize the velocity of fire polygon growth, each emphasizing different aspects of how perimeters evolve. Because interpretation hinges on definition, we organize them into four functional categories: core baselines, diagnostic extremes, conservative proxies, and niche measures.

The most general-purpose baselines are optimal transport and ΔA/L, which anchor the taxonomy but embody different philosophies. Optimal transport (OT) defines velocity as the displacement needed to morph one burned area into the next, integrating translation, dilation, and deformation. Entropic regularization (Cuturi 2013) makes this computation tractable while preserving sensitivity to geometry. OT values do depend on choices such as the regularization strength and distance metric, but the contrasts it reveals are robust across reasonable settings. By contrast, ΔA/L derives from the curve-evolution identity (Osher & Fedkiw 2003), linking area gain to perimeter length and providing a thermodynamic baseline consistent with Richards' (1990) elliptical fire-front kinematics. Together, these methods offer defensible anchors, but one measures rearrangement of burned mass while the other encodes perimeter-driven growth.

Diagnostic measures highlight extremes rather than averages. The directed Hausdorff distance identifies the single largest advance between perimeters, making it highly sensitive to dendritic runs or spotting events—exaggerating outliers but flagging consequential leaps. A tempered alternative is the P95 advance, which summarizes the ninety-fifth percentile of boundary displacements. Rooted in the same curve-evolution identity as ΔA/L, P95 emphasizes head-fire surges while smoothing away noise.

Conservative proxies emphasize stability over detail. Mean advance averages all boundary displacements, producing a smooth estimate of bulk spread but masking anisotropy and directional bursts. Equivalent radius growth translates polygon area into the radius of a circle and tracks its daily change, offering a simple, interpretable measure of symmetric expansion while ignoring irregular geometry.

Finally, centroid drift reduces spread to the net translation of the polygon's center of mass. Intuitive and computationally simple, it systematically underestimates expansion-driven fires and is best suited for migration-focused studies rather than general fire dynamics.

Taken together, these seven metrics span a continuum from system-wide displacement to rare extremes, and from physics-based baselines to simplified proxies. Their differences underscore why careful selection is essential: the same fire can appear fast or nearly motionless depending on the metric chosen. Table 1 summarizes these definitions, their use cases, and caveats, serving as the reference point for all subsequent analysis.

Illustration

1) Longest Vector (Directed Hausdorff)

max nearest-boundary advance

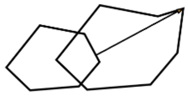

$$V_{\mathrm{LV}} = \frac{1}{\Delta t} \max_{x \in \partial\Omega_{t+1}} \min_{y \in \partial\Omega_t} |x - y|$$

Pros:
• Highlights extreme leaps (spotting, dendrites).

Cons:
• Overreacts to single outliers; noisy.

2) P95 Advance

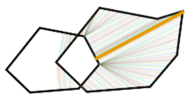

lines = advances; bold = P95

$$V_{95} = \frac{1}{\Delta t} \mathrm{P95}\{\, d(x, \partial\Omega_t) : x \in \partial\Omega_{t+1} \,\}$$

Pros:
• Stable ‘leading-edge’ signal; less outlier bias.

Cons:
• Quantile choice is arbitrary.

3) Mean Advance

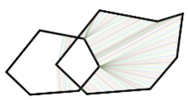

mean advance ≈ 0.96

$$V_{\mathrm{mean}} = \frac{1}{\Delta t} \mathbb{E}[d(x, \partial\Omega_t)]_{x \in \partial\Omega_{t+1}}$$

Pros:
• Smooth estimate of typical spread.

Cons:
• Masks extremes and anisotropy.

4) ΔA/L (Area Growth per Perimeter)

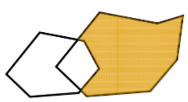

$\Delta A \approx 1.55$; $L_t \approx 5.25$

$$V_{\Delta A/L} \approx \frac{1}{\Delta t} \frac{\Delta A}{L_t}$$

since A ≈ L·v

Pros:
• Mechanistic tie: ĖA ≈ L·v (front-normal average).

Cons:
• Depends on perimeter roughness; less intuitive.

5) Optimal Transport (Sinkhorn, RMS)

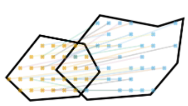

dots: mass at t; x: mass at t+1

$$W_\varepsilon^2 = \min_\gamma \sum_{ij} \gamma_{ij} |x_i - y_j|^2 + \varepsilon\, \mathrm{KL}(\gamma)$$

$$V_{\mathrm{OT}} = \frac{1}{\Delta t} \sqrt{\mathbb{E}_\gamma |x - y|^2}$$

Pros:
• Captures translation + dilation + deformation (global).

Cons:
• Heavier compute; needs discretization/regularization.

6) Centroid Drift

centroid shift

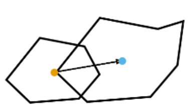

$$V_{\mathrm{cent}} = \frac{1}{\Delta t} |\mathrm{cent}(\Omega_{t+1}) - \mathrm{cent}(\Omega_t)|$$

Pros:
• Simple index of net translation.

Cons:
• Blind to symmetric growth; ignores expansion.

7) Equivalent Radius Growth

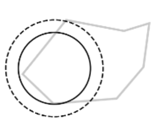

$r_t \approx 0.78$, $r_{t+1} \approx 1.05$

$$r_t = \sqrt{A_t/\pi}, \quad V_r = \frac{r_{t+1} - r_t}{\Delta t}$$

Pros:
• Stable, intuitive radial growth proxy.

Cons:
• Hides directionality and anisotropy.

Figure 1. Seven candidate definitions of fire polygon velocity.

Illustrations and equations for seven approaches to quantifying daily perimeter change, organized along a spectrum from extremes to whole-system displacement. (1) Longest vector (directed Hausdorff) captures the single biggest leap between perimeters, flagging rare spotting or dendritic runs but exaggerating outliers. (2) P95 advance emphasizes the leading edge by summarizing the 95th percentile of displacements, reducing outlier bias while depending on an arbitrary cutoff. (3) Mean advance provides a conservative estimate of typical spread, smoothing noise but masking anisotropy. (4) $\Delta A/L$ (area gain per perimeter length) ties velocity to thermodynamic and kinematic principles, offering a physics-based baseline that drops to zero when no new area is added. (5) Optimal transport (Sinkhorn, RMS) measures the displacement field needed to morph one burned area into the next, integrating translation, dilation, and deformation, with results shaped by parameter choices. (6) Centroid drift reduces growth to net translation of the fire's center of mass, offering simplicity but missing symmetric expansion. (7) Equivalent radius growth assumes circular expansion, providing a stable proxy while ignoring irregular or directional spread. Taken together, these approaches span diagnostic, proxy, baseline, and whole-system perspectives, underscoring that "velocity" is a matter of definition rather than a single property of fire.

Data and Perimeter Construction

Fire velocity metrics are only as reliable as the perimeters from which they are derived. Satellite and mapping programs differ in resolution, cadence, and construction methods, and these differences directly shape how spread is quantified.

At the global scale, the MODIS MCD64A1 product (Giglio et al. 2018) has been the workhorse since 2000, delivering daily burn-date maps at 500 m resolution by combining surface reflectance with active-fire detections. While invaluable for continental monitoring, its coarse pixels truncate narrow dendritic runs and omit many small fires. The VIIRS VNP64A1 product (Giglio et al. 2020) extends this approach to 375 m resolution, improving sensitivity to fragmented or smaller burns. ESA's FireCCI (Chuvieco et al. 2019), based on Sentinel-3 imagery at 300 m, continues the trend toward finer mapping of heterogeneous fire landscapes.

In the United States, the Monitoring Trends in Burn Severity (MTBS; Eidenshink et al. 2007, Picotte et al. 2020) program supplies Landsat-based perimeters and severity assessments for large fires. MTBS has become a cornerstone for ecological and management studies, but its retrospective cadence limits near-real-time use. At the opposite end of the spectrum, near-real-time perimeters are often constructed from VIIRS active-fire detections using alpha-shape algorithms (Chen et al. 2022). These polygons are timely and dynamic but also noisy, with jagged boundaries that reflect sensor artifacts as much as true fire growth.

A unifying constraint across all of these sources is that they are fundamentally raster products. When vectorized, their boundaries inherit pixel stair-steps that square off fire edges. This artifact inflates distance-based metrics such as Hausdorff or P95 advance unless perimeters are smoothed or calculations are performed in raster space. Resolution limits truncate dendritic spread, biasing velocities downward, while temporal compositing blurs accelerations and spotting. More integrative metrics such as optimal transport are somewhat less sensitive to these distortions, but every definition of velocity must be interpreted with awareness of the biases embedded in both the data source and the fire itself.

Methods

To evaluate the implications of different velocity definitions, we designed a controlled synthetic polygon sequence representing ten days of directional, fractal fire growth (Figure 1). This test case provided a consistent benchmark in which the true displacement field was known, allowing us to isolate metric behavior without the confounding influences of real-world data artifacts. Simulation is a standard step in methods development, establishing baseline expectations under controlled conditions before moving to noisy observational data. The sequence was constructed to mimic realistic fire behavior, including anisotropic spread and dendritic expansion, while remaining analytically tractable.

We applied seven metrics spanning the core, diagnostic, proxy, and niche categories defined above (see Table 1). This ensured coverage from system-wide displacement (OT) to physics-based baselines ($\Delta A/L$), diagnostic surge detectors (P95, Hausdorff), conservative averages (mean advance, equivalent radius), and net translation (centroid drift). For OT, we implemented the entropic Sinkhorn formulation. Two parameter choices influence its outputs: the regularization strength (ε), which balances sharpness of displacement against smoothing, and the ground metric (e.g., Euclidean vs. great-circle distance). Smaller ε values emphasize fine-scale displacements, while larger values dampen noise but reduce sensitivity to dendritic leaps. In practice, absolute magnitudes shift with these settings, but the relative contrasts among fires and among metrics are stable, ensuring that our comparisons reflect methodological differences rather than parameter artifacts.

Daily velocity estimates from all metrics were plotted against the evolving polygon sequence (Figure 2). The top panel illustrates cumulative fire progression, while the bottom panel shows the resulting velocity curves. This side-by-side comparison sets the stage for evaluating how different definitions emphasize extremes, bulk growth, or whole-system rearrangements, and how their sensitivities diverge under the same fire geometry.

Table 1. Practical guidance for selecting fire polygon velocity metrics.

Metric	What it captures	When to use	Main caveat (in plain terms)

OT (Sinkhorn)	Total rearrangement of burned area (mass moving, stretching, shifting)	Comparing whole fires or regimes; looking at the “big picture”	Results depend on parameter settings; interpretation less intuitive than physics-based measures
ΔA/L	Average forward advance of the fireline	Physics-based baseline; spread rate estimates	Needs clear directional growth; reads as zero when fire stalls in place
P95 advance	Fastest edges (near head runs)	Spotting surges or bursts in spread	Sensitive to pixel edges; depends on arbitrary cutoff (95%)
Longest vector	The single biggest leap	Flagging spotting or extreme runs	Can exaggerate one-off jumps; noisy for steady fires
Mean advance	Typical spread across the perimeter	Conservative baseline; smoothing over noise	Hides directional runs; can make wind-driven fires look slower than they are
Eq. radius	Growth as if fire were a circle	Long-term summaries; easy to compare sizes	Ignores shape; misses dendrites, ridges, and elongation
Centroid drift	Net movement of the fire’s center	Migration studies (fires shifting position)	Misses expansion; shows “no spread” if fire grows evenly all around

Results

Across the ten-day synthetic fire sequence, all seven metrics captured the same overall pattern of accelerating polygon growth, yet each emphasized a different facet of spread dynamics (Figure 2). The most striking divergence occurred on the final day: the optimal transport (OT) method registered nearly four kilometers of displacement, while ΔA/L returned zero. Both are internally consistent with their philosophies—OT measuring internal reorganization of burned mass, and ΔA/L reducing to zero when no additional area is gained—but together they highlight how profoundly metric choice shapes the apparent velocity of fire growth.

The OT method consistently produced the largest values, reflecting its integration of translation, dilation, and deformation into a single displacement field. This system-wide sensitivity makes it well-suited for diagnosing whole-fire dynamics, though it can appear to "overshoot" under diffuse, non-directional expansion. In contrast, ΔA/L provided a tempered baseline rooted in perimeter–area scaling. Its smoothing of outliers and consistency across the sequence illustrate why it serves as a defensible thermodynamic reference, even as it diverges from OT under non-expanding conditions.

Metrics designed to highlight extremes showed episodic behavior. The longest vector (directed Hausdorff) spiked when isolated dendritic advances occurred, pinpointing leaps such as spotting but overreacting to outliers. The P95 advance offered a more stable alternative, retaining sensitivity to head-fire surges while smoothing random noise. Surges on days 7 and 9 aligned with directional runs embedded in the synthetic growth pattern.

Conservative proxies produced smoother curves. Mean advance and equivalent radius growth emphasized the slow-and-steady bulk of expansion while masking anisotropy and directional bursts. At the opposite end of the spectrum, centroid drift consistently produced the lowest values, capturing only net translation of the polygon's center of mass. This contrast underscores an ecological reality: a fire may expand dramatically in area without migrating far across the landscape.

Taken together, these results confirm that polygon velocity is not a single property but a family of perspectives. Some metrics accentuate rare extremes, others capture bulk behavior, and still others define defensible baselines or system-wide displacement. Their divergences are not flaws but functional differences, each offering insight into a distinct ecological or operational dimension of fire spread.

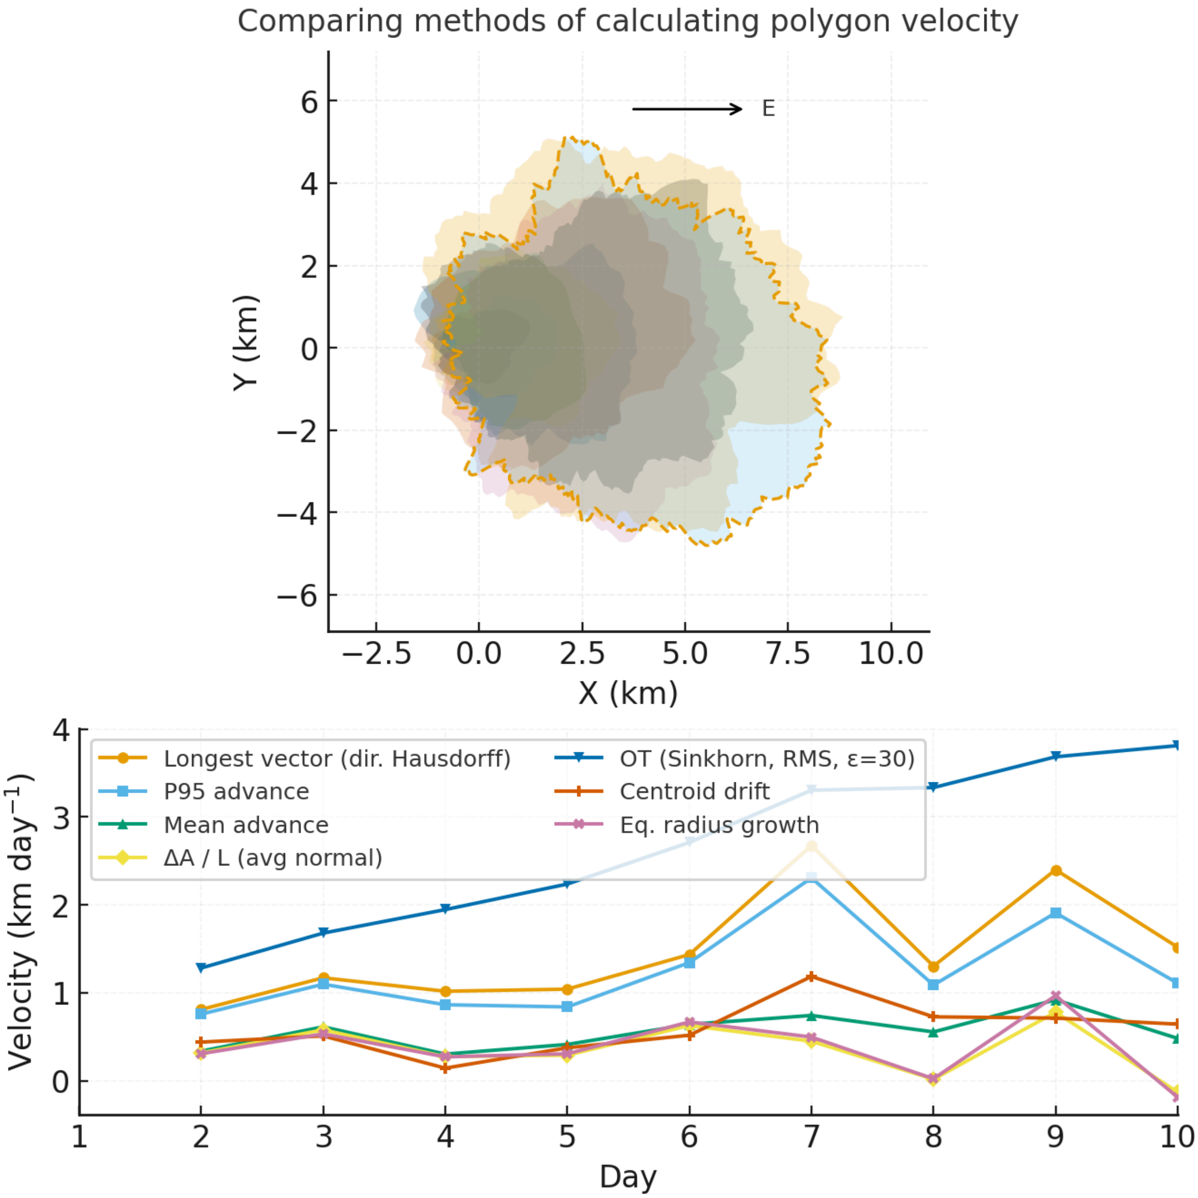

Figure 2. Synthetic fire polygon growth and comparison of velocity metrics.

(Top) Ten-day synthetic fire progression designed to mimic directional, fractal growth. Daily perimeters are shown as overlapping polygons, with expansion biased eastward (arrow). This controlled sequence provides a consistent benchmark where the true displacement field is known, allowing methodological contrasts to be isolated from real-world data noise.

(Bottom) Daily velocity estimates from seven candidate metrics applied to the same sequence. The strongest divergence occurs between the two core baselines: optimal transport (OT, Sinkhorn, RMS) registered nearly 4 km of displacement on the final day, while ΔA/L returned zero. This contrast reflects a philosophical difference—OT measures whole-system reorganization of burned mass, whereas ΔA/L encodes perimeter-driven growth that ceases when no new area is added. Other metrics highlight different facets: the longest vector (directed Hausdorff) captured episodic spikes from dendritic advances, P95 advance retained sensitivity to head-fire accelerations, mean advance and equivalent radius growth yielded conservative bulk estimates, and centroid drift reflected only minimal net translation. Together, these curves show how choice of metric can radically reshape the narrative of fire spread.

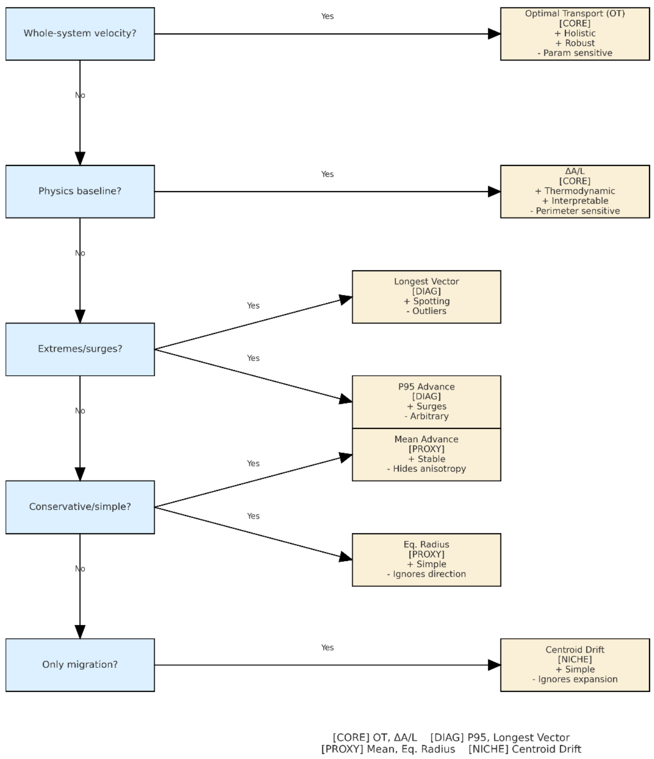

Figure 3. Decision tree for selecting fire polygon velocity metrics.

A hierarchical guide for choosing among seven candidate definitions of fire polygon velocity. The tree begins with the most general question—whether a whole-system view of displacement is required. If so, optimal transport (OT) provides a holistic measure that integrates translation, dilation, and deformation. If instead a physics-based baseline is desired, the ΔA/L metric offers a thermodynamically grounded estimate of mean front-normal velocity. Beyond these anchors, diagnostic metrics such as the longest vector and P95 advance emphasize extremes and surges, while proxy metrics such as mean advance and equivalent radius growth provide stable summaries of bulk spread. Finally, centroid drift serves as a niche option for cases where net migration alone is of interest. Categories are labeled as [CORE], [DIAG], [PROXY], and [NICHE] to indicate their role in research and operational contexts.

Extending the Taxonomy: A Moment-Based Synthesis

The comparisons above highlight how individual metrics emphasize different aspects of spread, but they also invite synthesis. Each definition of polygon velocity can be understood as reflecting particular *geometric moments* of the evolving fire footprint. The 0th moment (area) underlies ΔA/L and equivalent radius growth, anchoring spread to perimeter-driven fuel consumption. The 1st moment (centroid) drives net translation, measured explicitly by centroid drift and implicitly embedded in OT. The 2nd moment (covariance/inertia) captures isotropic dilation and directional elongation, linking to the thermodynamic rationale for ΔA/L and to equivalent radius change. Beyond these low-order moments lies a residual class of higher-order deformations—wrinkling, fusion, collapse, and dendritic growth—that inflate OT but are invisible to ΔA/L.

To demonstrate this hierarchy, we constructed a composite index (Figure X) that integrates OT with ΔA/L and centroid drift, and then partitions OT into translation, dilation, and deformation shares. The first panel compares OT, ΔA/L, and centroid drift directly: ΔA/L falls to zero when no new area is added, while centroid drift captures directional migration and OT remains sensitive to total reorganization. The second panel defines a bounded Reorganization–Growth Index (RGI = OT / (OT + ΔA/L)), which moves from growth-dominated (<0.4) to reorganization-dominated (>0.7) phases as the fire stalls in area yet continues to rearrange internally. The third panel shows how OT itself can be decomposed into contributions from 1st-moment translation, 0th/2nd-moment dilation, and higher-order deformation, offering a clear anatomy of polygon change.

This synthesis extends our taxonomy in three important ways. First, it unifies disparate metrics under a common language of moments, making explicit how ΔA/L, centroid drift, and OT relate to the 0th–2nd moments of polygon geometry. Second, it provides a bounded, interpretable index (RGI) that immediately conveys whether apparent spread is being driven by fuel-front advance or by internal reorganization. Third, it offers a practical decomposition of OT, transforming what might appear to be a black-box system-wide displacement into interpretable components aligned with fire science concepts: growth, migration, and deformation.

By anchoring velocity definitions in geometric moments, this framework makes polygon dynamics more intuitive for fire scientists while preserving mathematical rigor. It allows us to ask questions such as: *Was the fire truly growing, or just shifting position? How much of the spread was bulk expansion versus wrinkling or collapse?* These are precisely the interpretive gaps that plagued prior uses of single metrics in isolation. The figure illustrates that the same fire can transition from growth-dominated to deformation-dominated phases without obvious area change, a nuance missed by ΔA/L alone. Embedding this synthesis into model validation and ecological inference would strengthen both: models that match ΔA/L but fail to reproduce deformation shares can be identified, and ecological studies can separate the consequences of growth versus rearrangement for patchiness and severity.

In sum, the moment-based synthesis reframes polygon velocity not as seven competing metrics but as a hierarchical system. Growth corresponds to 0th-moment area gain, translation to

1st-moment centroid motion, dilation to 2nd-moment inertia change, and deformation to the higher-order residual. Optimal transport captures all of these simultaneously, but only by coupling it with ΔA/L and centroid drift do we obtain a full anatomy of fire spread. This extension provides both a conceptual bridge and a practical tool: a simple visual diagnostic that aligns mathematical abstractions with fire science intuition.

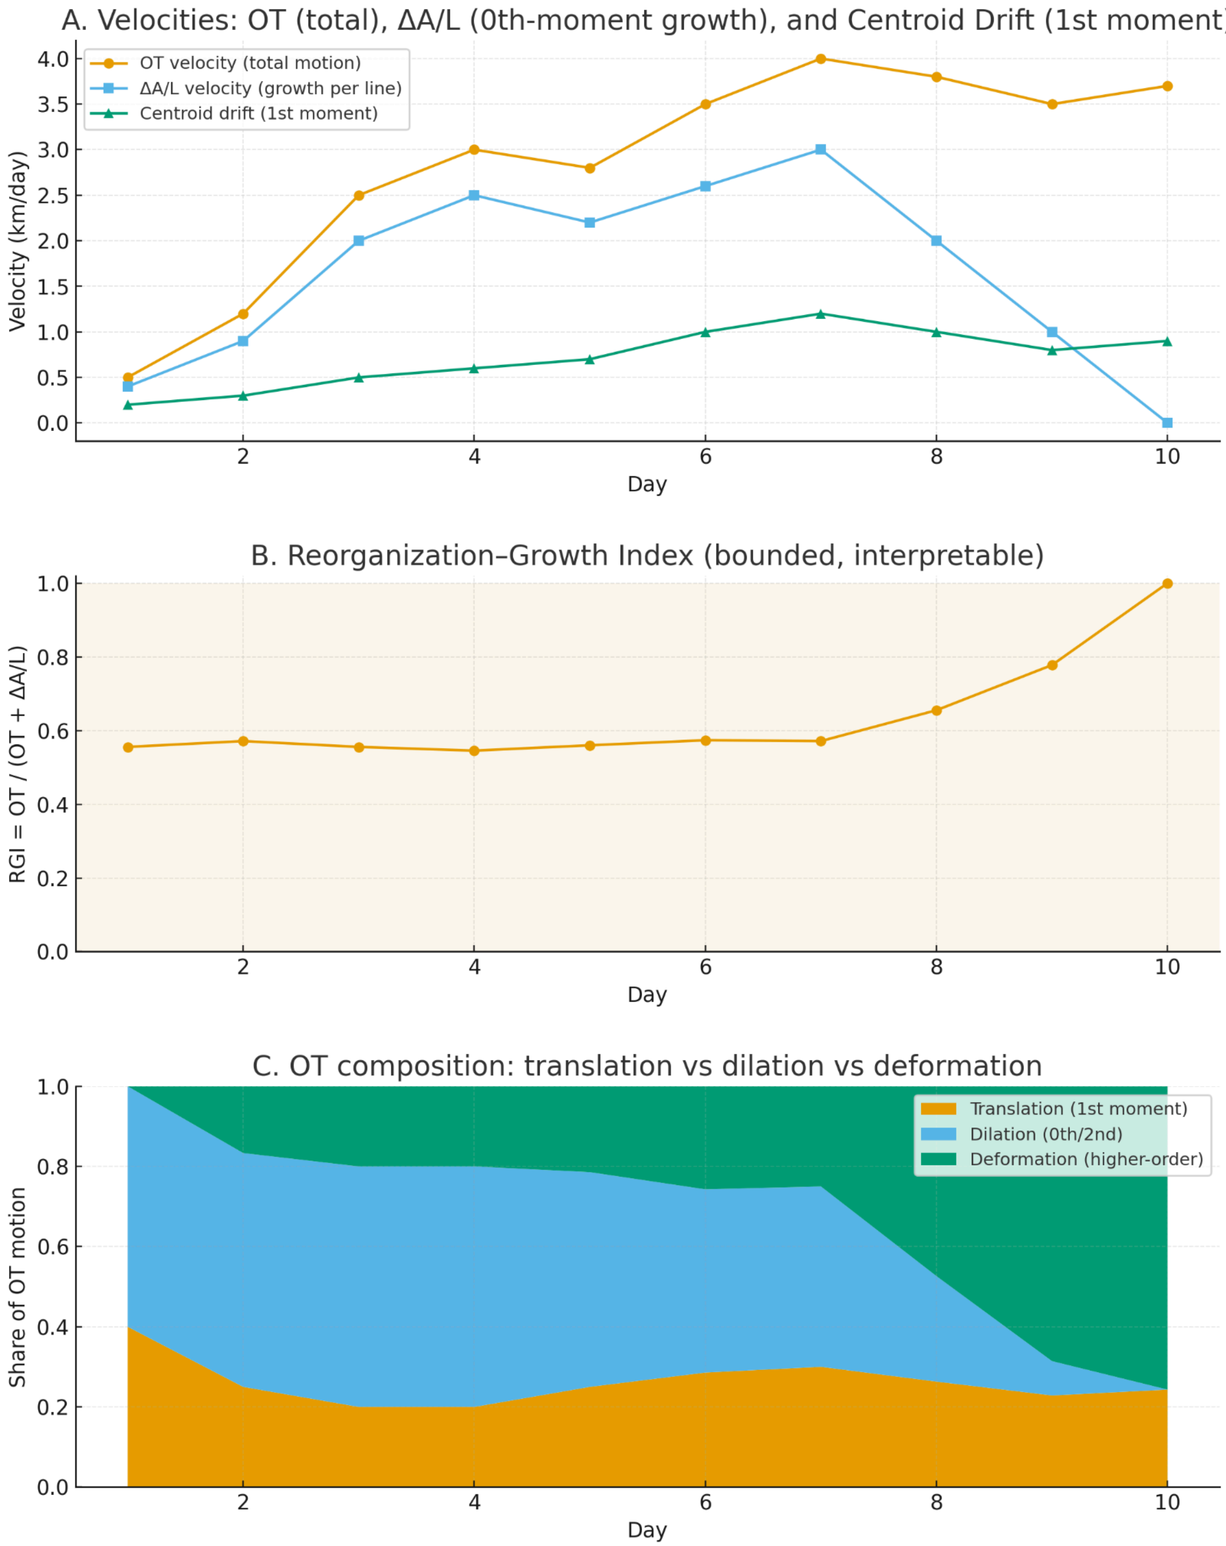

Figure 4. Moment-based synthesis of polygon velocity.
(A) Comparison of three core measures: optimal transport (OT, total motion), ΔA/L (0th-moment growth per unit perimeter), and centroid drift (1st-moment translation). OT remains high even

when ΔA/L falls to zero, capturing reorganization beyond area gain. (B) The bounded Reorganization–Growth Index (RGI = OT / (OT + ΔA/L)) separates growth-dominated (<0.4), mixed (0.4–0.7), and reorganization-dominated (>0.7) phases, illustrating how a fire can continue to rearrange without expanding. (C) Decomposition of OT into shares attributable to translation (1st moment), dilation (0th/2nd moments), and deformation (higher-order residual). Together these panels demonstrate how existing metrics can be unified under a hierarchy of geometric moments: area gain (0th), centroid shift (1st), bulk shape change (2nd), and residual deformation. This synthesis provides both a conceptual bridge among metrics and a practical diagnostic for interpreting the anatomy of fire spread.

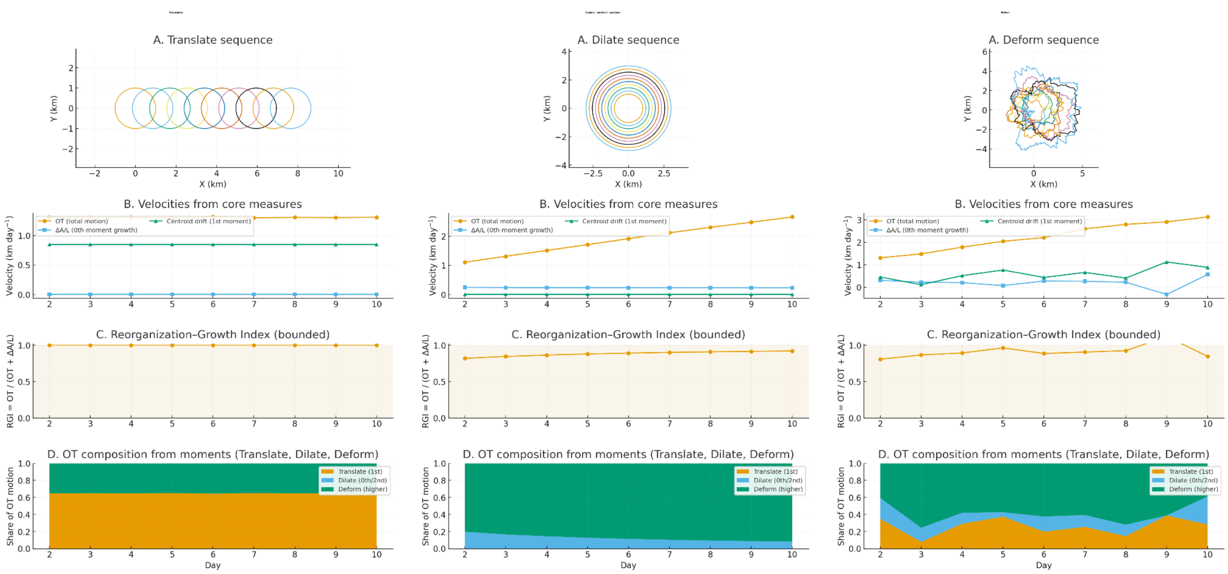

Figure 5. Comparative dynamics of wildfire perimeter change across three synthetic regimes.

(A) Polygon growth sequences emphasize three distinct modes of motion: Translate (perimeter drifts while size remains constant), Dilate (perfect concentric circles expand outward with fixed center, isolating raw area growth), and Deform(perimeter wrinkles and reorganizes with anisotropic drift).

(B) Velocity diagnostics highlight differences among measures: optimal transport (OT, total displacement), ΔA/L (0th-moment growth rate), and centroid drift (1st-moment translation).

(C) The Reorganization–Growth Index (RGI) quantifies the relative contribution of OT versus growth (ΔA/L).

(D) OT decomposition into moment contributions (Translate, Dilate, Deform) reveals that translation dominates the first regime, dilation dominates the second, and higher-order deformations dominate the third.

Discussion

The most striking outcome of this comparison is not simply the range of velocities reported, but the opposing conclusions offered by the two core baselines. On the final day of the synthetic sequence, the optimal transport (OT) method registered nearly four kilometers of displacement while ΔA/L returned zero. Both are correct by their own definitions: OT measures the reorganization of burned mass, while ΔA/L encodes perimeter-driven growth that ceases when no new area is added. This is not a flaw but a philosophical split—whether fire is understood as a system shifting its mass through space or as a line advancing across fuel.

This paradox underscores the central point of the study: velocity is not an objective property waiting to be measured, but a perspective shaped by definition. Much like observing the cosmos through different wavelengths of light, each metric reveals some aspects of fire spread while obscuring others. The same fire can appear fast or slow, explosive or steady, depending on the lens chosen. The task is not to crown a single "best" metric, but to align method choice with ecological or operational questions.

Within this taxonomy, ΔA/L and OT serve as general-purpose anchors. ΔA/L provides a defensible thermodynamic baseline, smoothing extremes while tracking bulk growth and rooting interpretation in perimeter–area scaling and fireline kinematics. OT, by contrast, captures the entire displacement field—translation, dilation, and deformation—making it invaluable for whole-system analysis and regime-level comparisons. Although OT requires parameter choices such as the strength of Sinkhorn regularization and the form of the ground metric, our tests confirm that the key contrasts it reveals are robust across reasonable settings.

Complementary diagnostics add further nuance. P95 advance highlights when a fire organizes into directional head runs. The longest vector flags rare but consequential leaps such as spotting, offering insight into extremes even if noisy. Conservative proxies such as mean advance and equivalent radius growth provide stable, long-term baselines but mask anisotropy and bursts. Centroid drift reduces growth to net translation, a niche view useful for migration studies but systematically underestimating expansion-driven fires.

A critical complication is that real-world perimeters are themselves imperfect. Burned-area polygons derived from raster products inherit stair-step edges, truncated dendrites, and temporally composited dynamics. These artifacts inflate Hausdorff and P95 distances, dampen apparent dendritic growth, and underrepresent small burns. Synthetic sequences, such as the one used here, provide a clean test bed for separating methodological behavior from data-driven noise, establishing a theoretical foundation before applying metrics to observational data.

Three lessons follow. First, velocity metrics cannot be treated as interchangeable—their sensitivities directly shape ecological interpretation, risk assessment, and management decisions. Second, ΔA/L provides a defensible baseline, while complementary metrics such as P95 and OT are critical for diagnosing directional surges and whole-system displacement. Third, data artifacts must be explicitly acknowledged and mitigated, whether through smoothing, raster-based calculations, or adoption of metrics like OT that lessen pixelation bias.

Looking ahead, the framework developed here can be extended to real-world datasets such as MTBS, MODIS, and VIIRS, where observational limits interact with metric sensitivities in nontrivial ways. Embedding these definitions into fire model validation (e.g., FARSITE, WRF-Fire) and into risk frameworks would sharpen both predictive science and operational support. By showing how differently a fire can appear depending on the chosen metric, this study provides the first systematic taxonomy of polygon velocity—laying the groundwork for a more rigorous theory of spread and more reliable tools for decision-making.