# Summary

**CivicTwin** creates transportation digital twins, high-fidelity, physics-grounded, editable models of real streets and corridors, directly from street-view video and images. Unlike pipelines that rely on expensive, non-scalable multi-view capture (especially drone/aerial), CivicTwin eliminates overhead dependencies instead uses generative video models and visual-geometric foundation models to reconstruct and fill unobserved regions with plausible, structured geometry. The result is a simulator-ready twin that preserves measured detail where data exist, fills gaps where they don't, and scales across corridors and cities with dramatically ten times lower cost.

**Why this matters.** AVs, delivery robots, and automated traffic management require massive, safe, repeatable testing, particularly for rare, safety-critical scenarios that are hard to observe or stage on real roads. Today's best simulators still struggle with the authoring burden, physics realism, and city-scale coverage needed for safety cases and operations planning.

**CivicTwin's capability** rests on five integrated elements: (1) object-aware reconstruction from street-view data guided by foundation-model segmentation; (2) generative inpainting to handle occlusions and sparse viewpoints while maintaining lane/sidewalk/curb topology and material priors; (3) physics-property inference (materials, contact, friction) with export to standard formats; (4) a GPU-parallel runtime with multimodal sensor simulation and behaviorally realistic agents; and (5) scenario editing for long-tail risk coverage with continuous updates as the network changes.

**Performance targets** include compiling a 5-10 km corridor in <24 hours on a single modern GPU node; validating physics with sensor-sim divergence ≤5% and friction error ≤20-25%; delivering ≥10× faster closed-loop training; and achieving ≥50% safety-metric improvement on rare events, thereby enabling ≥70% fewer on-road validation miles and ≥60% lower closed-course test costs for a target AV feature.

**CivicTwin is purpose-built** to accelerate ADS/ADAS (automated driving or advanced driver assist systems) and traffic management development, validation, and safety-case generation. The twin enables closed-loop simulation with real-world data, systematic rare-event generation (e.g., collisions/surrogates and perception robustness under weather/lighting). Secondary applications, such as micromobility and ground-support robotics, benefit from the same tooling.

We envision a 36-month effort advancing from corridor-scale twins (Phase I) to district-scale pilots (Phase II) and operator deployments (Phase III). Core deliverables center on ADS/ADAS: an open compiler, a runtime with multimodal sensor simulation and behaviorally realistic agents, an AV-targeted scenario library for long-tail scenarios, and standards-aligned Verification and validation artifacts to support safety cases. *CivicTwin* directly supports ARPA-I's Knowledge and Optimization priorities by providing dynamically updated digital twins and massively benchmarking for autonomy, reducing cost and risk while accelerating innovation in U.S. transportation.