TOWARD AI-ASSISTED CIVIL QUANTITY TAKEOFF and Earthwork Estimation

A Position Paper and Proposed Framework for Excavation-Focused Preconstruction Workflows

Muhammad Haris Saleem, EIT
Independent Researcher and Civil Engineering Professional
M.S. in Civil Engineering, Louisiana Tech University | Former Field/Office Engineer, Kiewit | Pursuing PE Licensure

Abstract: Quantity takeoff remains a foundational but labor-intensive step in civil preconstruction, particularly for excavation, earthwork, and sitework scopes derived from 2D plan sheets and surface-based reports. Current digital workflows already support measurement capture and exportable quantity data through platforms such as Bluebeam Revu, Trimble Business Center, Autodesk Takeoff, and PlanSwift, yet estimators still spend significant effort interpreting legends, reconciling labels, applying assumptions, and converting measurements into estimate-ready line items. This paper proposes a human-in-the-loop framework for AI-assisted civil quantity takeoff focused on excavation and earthwork estimation. Instead of replacing estimator judgment, the framework positions artificial intelligence as a layer for document understanding, quantity normalization, error flagging, and cost-item mapping on top of existing takeoff workflows. The paper defines the technical architecture, data strategy, implementation path, and validation roadmap for an applied civil-engineering initiative and implementation direction.

Keywords: civil engineering; quantity takeoff; excavation; earthwork estimation; construction AI; Bluebeam; Trimble Business Center; preconstruction automation

1. Introduction

Civil preconstruction depends heavily on quantity takeoff because estimating accuracy begins with the quality of the measured quantities. In current practice, commercial tools already support digital measurement and quantity generation: Bluebeam enables PDF-based takeoff markups, Markups List export to Excel, and Quantity Link connections to spreadsheets; Trimble Business Center generates an Earthwork Summary in Excel specifically for takeoff estimation; Autodesk Takeoff combines 2D and 3D quantity workflows; and PlanSwift supports digital takeoff, formulas, and spreadsheet-linked estimating. These capabilities show that the industry has already moved well beyond paper workflows, but they also show that the last mile of estimation still depends on human interpretation and manual data organization.

The unresolved problem is not the absence of measurement software. The unresolved problem is the continued dependence on estimator labor for reading legends, reconciling inconsistent labels, applying depth assumptions, assigning pay items, checking units, and translating takeoff outputs into estimate-ready line items. The practical opportunity for AI is therefore not to replace the estimator with a black-box model, but to reduce repetitive cognitive and administrative work around an already structured workflow.

This paper frames that opportunity as a near-term applied research and industry-oriented initiative. The proposed starting point is narrow by design: excavation and earthwork quantity takeoff. This scope is attractive because it is commercially meaningful, technically tractable, and already supported by exportable digital workflows in Bluebeam Revu and Trimble Business Center.

2. Industry context and motivation

Estimators in U.S. civil and general construction workflows often work across several software ecosystems rather than within a single fully unified platform. Bluebeam is commonly used for PDF-based measurements, markup organization, data export, and spreadsheet linkage. Trimble Business Center is widely used for surface comparison, cut/fill analysis, and earthwork reporting. Autodesk Takeoff markets integrated 2D and 3D quantity workflows, while PlanSwift emphasizes digital takeoff, formulas, and estimate-oriented exports. The existence of these products confirms that the market values speed, traceability, and integration between measurements and cost workflows.

At the same time, these tools do not eliminate estimator interpretation. Legends still need to be understood. Measurement objects still need to be grouped and normalized. Notes still need to be read. Average-depth assumptions still need to be applied. Earthwork summaries still need to be translated into pricing logic. The friction is therefore concentrated in the interpretation layer between measurement capture and final estimate assembly.

That gap is exactly where AI can be introduced responsibly. Large language models and document-intelligence pipelines are now capable of extracting structured information from drawings, notes, legends, and tabular files. When combined with deterministic quantity formulas and estimator review, they can support a practical civil-engineering copilot rather than an untrustworthy autonomous estimator.

3. Problem statement

The core research and product question is: how can AI reduce manual effort in civil quantity takeoff and earthwork estimation without undermining reliability, traceability, and professional judgment?

This question matters because estimate quality is shaped not only by mathematical formulas but also by consistency, review burden, and decision latency. A useful system must therefore reduce effort while preserving transparency. In other words, success is not simply a higher model accuracy score. Success is a reduction in estimator handling time, fewer preventable errors, stronger auditability, and easier conversion of measured quantities into estimate-ready cost items.

4. Proposed framework

The proposed framework is a human-in-the-loop AI workflow for excavation-focused civil takeoff. It contains five functional layers. First, a file-ingestion layer accepts plan PDFs, legend sheets, general notes, Bluebeam CSV or XML exports, TBC earthwork reports, and estimate templates. Second, a document-intelligence layer extracts sheet titles, discipline labels, unit clues, legend items, and depth-related notes from uploaded files. Third, a quantity-normalization layer standardizes inconsistent shorthand such as 'EXC', 'common exc', 'undercut', or 'pad cut' into controlled internal categories.

Fourth, a rules-and-pricing layer applies deterministic formulas and explicit business logic. Area-to-volume conversion, stripping thickness, shrink/swell adjustments, and pay-item mapping should remain auditable and rules-based. Fifth, a review-and-QA layer flags missing depth values, conflicting units, duplicate entries, and uncertain item mappings for estimator approval. In this architecture, AI interprets and organizes while transparent rules compute and price.

5. Initial application focus

The recommended first product scope is intentionally narrow. It should focus on excavation area summaries from Bluebeam, average-depth excavation volumes, topsoil stripping quantities, TBC earthwork summaries for cut/fill intake, and estimate-ready grouping by pay item. This narrower scope is preferable to a broad promise of 'AI for all civil engineering' because it aligns with a real estimator pain point and can be validated against current practice with limited data.

A typical use case would involve a building pad or parking undercut workflow. The estimator measures the region in Bluebeam or imports a TBC Earthwork Summary. The AI layer then reads the labels, checks whether the applicable depth assumption exists, maps the item to a standard pay item, and prepares an estimate-ready line item for review. The estimator remains the final decision-maker, but the system absorbs the repetitive digital labor.

6. Data strategy and pilot development path

A credible applied research and implementation pathway requires a benchmark dataset rather than an abstract idea. The initial target dataset should include roughly 20 to 30 Bluebeam export packages, at least 10 TBC earthwork reports, matching estimate sheets, and the associated plan, legend, and notes pages. Each project bundle should preserve both the raw inputs and the reviewed benchmark outputs so the system can be tested against a traceable reference.

The most practical way to collect this material is through a combination of the founder's own practice files, de-identified sample files from industry contacts, and pilot partners who are willing to share redacted takeoff packages and workflow feedback. These pilot partners are not necessarily employees. They can be friendly estimators, preconstruction engineers, earthwork subcontractors, or small contractors who agree to act as design partners and early testers. Their value lies in providing representative files and validating whether the output actually fits estimating practice.

7. Significance for industry and academia

For industry, the proposed direction is attractive because it does not require estimators to abandon familiar software. It builds on the structured byproducts that existing tools already generate. That lowers adoption friction and makes early deployment more plausible. For academia, the topic opens a useful research space at the intersection of civil engineering, construction informatics, document intelligence, quantity reasoning, and human-in-the-loop AI.

It also creates a credible publication sequence. A position paper can define the framework and motivation now. A later empirical paper can report pilot results, data quality findings, estimator time savings, or quantity-variance analysis once enough benchmark projects have been collected.

8. Conclusion

The path toward AI automation in civil engineering should begin with problems that are both technically tractable and operationally valuable. Excavation and earthwork quantity takeoff satisfy that condition. Existing tools already digitize measurements and generate exportable summaries, but they do not eliminate the estimator's burden of interpretation, normalization, and quality control.

A human-in-the-loop AI framework offers a more credible near-term direction. AI can read, organize, classify, and flag. Deterministic rules can compute and price. The estimator can retain final control. The resulting system is not a replacement for civil-engineering judgment; it is an attempt to reduce the repetitive digital effort that surrounds it.

9. Practical implementation snapshot

For an early industry-focused initiative, the most practical build order is shown below. The goal is to establish a credible prototype first, then layer in increasing levels of intelligence and validation.

Phase	Primary input	System function	Expected output
1	Bluebeam CSV / XML TBC Excel	Parse quantities, normalize labels, preserve source traceability	Clean takeoff table and estimate-ready export
2	Plan PDFs Legends General notes	Extract sheet metadata, legend items, units, and note-based assumptions	Editable project context and review flags
3	Normalized quantities + rules	Apply area/volume conversions, pay-item mapping, and pricing logic	Line-item estimate with assumptions log
4	Pilot project bundles	Benchmark against reviewed quantities and estimator feedback	Variance report, usability findings, publication-ready evidence

References

1. Bluebeam. (2026). Construction quantity takeoff and estimating software; Markups List export; Quantity Link; and custom columns documentation. Bluebeam product and support resources.

2. Trimble. (2026). Run an Earthwork Summary report. Trimble Business Center help documentation.

3. Trimble. (2026). Analyze the balance of cut/fill volumes for a job site. Trimble Business Center help documentation.

4. Autodesk. (2026). Construction Takeoff software and integrated 2D/3D quantity workflows. Autodesk Construction Cloud product resources.

5. Autodesk. (2026, March 24). Construction takeoffs: A complete how-to guide. Autodesk Construction Blog.

6. Autodesk. (2026, January). Autodesk Construction Cloud releases and Takeoff updates. Autodesk release notes.

7. PlanSwift. (2026). Takeoff software for construction estimating. PlanSwift product resources.

8. PlanSwift. (2026). Takeoff and estimating software features. PlanSwift product resources.

9. Elsevier. (2026). Automation in Construction: aims and scope. ScienceDirect journal homepage.

10. Elsevier. (2026). Guide for authors - Automation in Construction. ScienceDirect author guidance.

11. ASCE. (2026). Author Center: journals overview and submission guidance. American Society of Civil Engineers.

12. ASCE. (2026). Journal of Computing in Civil Engineering: journal information and editorial scope. ASCE Library.

13. ASCE. (2026). Journal of Construction Engineering and Management: journal information and editorial scope. ASCE Library.

14. Kang, T., et al. (2024). Earthwork Network Architecture (ENA): research toward automated quantity takeoff for earthwork operations.

15. Alathamneh, S., et al. (2024). BIM-based quantity takeoff: current state and future perspectives.

16. Valinejadshoubi, M., et al. (2024). Automated system for high-accuracy quantity takeoff using BIM.

17. Jafary, P., et al. (2025). AI-augmented construction cost estimation.

18. Research article. (2024). Towards fully automated processing and analysis of construction drawings.

19. Kalasapudi, V. S., et al. (2025). The overlooked frontier of AI in construction.

20. Frontiers in Built Environment article. (2025). Human-in-the-loop and applied AI directions in construction workflows.