MANTECH Publications

International Journal of Data Science and Analytics Innovations
Volume 2, Issue 2, July-December 2026
ISSN: 3139-7638(Online)

Conversational Construction Site Management & Reporting: Extracting Structured Data from Informal Natural-Language Messages

Shreya Jadhav[1], Param Patel[2], Atharva Mangutkar[3]

ABSTRACT

Objective: *Construction site teams report material movement, labour use, expenses, payments and progress in short, informal messages, which are usually converted into structured records by hand. This paper proposes and evaluates a conversational pipeline that turns such messages into validated operational records.*

Methodology: *The system normalizes language, extracts intents and entities, maps them to construction schemas, matches entities against application state, validates the result and assigns a confidence value before a record is created. It was tested on a synthetic benchmark of 300 records in English, Marathi-English code-mixed, transliterated Marathi and informal or noisy text.*

Results: *On the 75-record held-out test set, intent accuracy was 61.33%, entity macro F1 41.47%, entity micro F1 45.01% and exact match 1.33%. On a 12-record multi-action subset with 17 ground-truth intent instances, intent-instance accuracy was 70.59% and action F1 82.35%.*

Limitations: *The evaluation is a laboratory study on synthetic data. The results are not claimed to hold on real construction sites, nor to outperform large language models.*

Originality: *The paper offers a construction-specific pipeline for informal, multilingual and multi-action site messages, together with a benchmark, a configuration ablation and an error analysis.*

Page 76-92 © MANTECH PUBLICATIONS 2026. All Rights Reserved

KEYWORDS: *Construction management; Natural language processing; Information extraction; Conversational systems; Code-mixed language; Entity matching; Intent detection; Marathi.*

INTRODUCTION

1. Background

Building projects depend on communication between site engineers, supervisors, contractors, suppliers, labour agencies and managers, who exchange information mostly through brief informal messages rather than formal documents. A single message may report a delivery, material used or damaged, workers present, an expense paid or an instruction. Before such a message can support project control, it has to be turned into a structured record.

Natural language processing (NLP) is widely applied in construction for document analysis, information extraction and text mining (Wu et al., 2022; Shamshiri et al., 2024). Cho and Lee (2019) showed that daily reports can be collected through a chatbot, later extended to a mobile-messenger setting (Cho et al., 2023), and Xiao et al. (2024) combined language models with computer vision to generate daily reports from site video.

2. Problem statement and motivation

Manual conversion is slow and error-prone. One message may contain several operations, for example receiving, using and damaging the same material. A candidate record that merely matches a text pattern should not be treated as a valid database entry: mandatory fields must be present and operational constraints respected before anything is committed. We therefore separate extraction from the commitment of a transaction, and treat this separation as the core design principle.

The problem is narrower than general conversational assistance. We transform informal site messages into structured transactions that can be checked against the current state of a construction management application, covering materials, labour, ready-mix concrete (RMC), expenses, payments, vendors, reminders and queries.

Page 76-92 © MANTECH PUBLICATIONS 2026. All Rights Reserved

3. Research gap

Prior work shows that language interfaces can support daily reporting (Cho and Lee, 2019; Cho et al., 2023) and that information can be extracted from contracts and construction documents (Al Qady and Kandil, 2010; ul Hassan and Le, 2022; Ren and Zhang, 2022; Ko et al., 2024). Few studies address a construction-specific pipeline that converts informal, multi-action messages into confirmable transactions. Language variation is a further gap: site messages mix English with local languages, transliteration, abbreviations and poor grammar.

4. Objectives and research questions

The aim is to design and evaluate a construction-specific pipeline of language normalization, intent and entity extraction, schema mapping, entity matching, validation and confirmed record creation. The study asks: RQ1, how well can construction intents and entities be extracted from informal messages? RQ2, does performance differ across English, Marathi-English, transliterated Marathi and informal messages? RQ3, does it differ for multi-action messages? RQ4, how do pipeline configurations compare? RQ5, which errors remain?

5. Contributions and organization

The contributions are (1) a pipeline that converts site messages into operational transactions; (2) a 300-record synthetic benchmark; (3) an evaluation across languages, splits and multi-action messages; and (4) an ablation and error analysis. Section 2 reviews the literature, Section 3 describes the method, Section 4 the framework, Section 5 the results, Section 6 the discussion, Section 7 limitations and Section 8 concludes.

LITERATURE REVIEW

1. NLP and conversational systems in construction

Earlier construction NLP work concentrated on extracting structured information from documents to support project tracking (Wu et al., 2022; Shamshiri et al., 2024; Al Qady and Kandil, 2010; ul Hassan and Le, 2022; Ren and Zhang, 2022; Ko et al., 2024). More recent work addresses conversational daily reporting: Cho and Lee (2019) built a chatbot for daily report handling, Cho et al. (2023) collected work reports from site workers through a mobile messenger, and Xiao et al. (2024) combined computer vision with ChatGPT to prepare daily reports from site footage. Document-based studies deal with written text rather than short

Page 76-92 © MANTECH PUBLICATIONS 2026. All Rights Reserved

informal messages, and conversational studies deal with collecting or generating reports rather than extracting several operational transactions from field messages.

2. Intent detection and slot filling

Converting a message into a record requires intent classification and slot filling. Goo et al. (2018) proposed a slot-gated model that solves both jointly, and Chen et al. (2019) evaluated a joint BERT-based model on the ATIS and SNIPS datasets. In construction, the relevant slots are domain-specific (project, material, quantity, unit, supplier, labour agency, price, payment detail, expense category), so schema mapping and entity matching are needed in addition to standard extraction. We therefore treat language understanding as a pipeline of normalization, intent recognition, entity extraction, schema mapping, entity matching, validation and processing.

3. Marathi and code-mixed language processing

Marathi appears in digital communication in Devanagari or in Roman transliteration. Joshi (2022) introduced L3Cube-MahaCorpus and MahaBERT, and Litake et al. (2022) released L3Cube-MahaNER, a named-entity dataset. Code mixing adds difficulty through informal spellings and transliterated words. Amin et al. (2024) studied Marathi-English code-mixing for question answering and named-entity recognition and found mixed-language input hard to handle. Gupta et al. (2023) proposed the Hinglish dataset MUTANT, and Sheth et al. (2025) created COMI-LINGUA for Hindi-English code-mixing. For this reason our benchmark includes code-mixed, transliterated and noisy Marathi content.

4. Large language models for structured extraction

Large language models (LLMs) can convert unstructured text into structured representations; Dagdelen et al. (2024) extracted structured records from scientific text without rules for every wording variant. Extraction, however, differs from safely executing a transaction. A construction system must validate entities against known projects, materials, suppliers and labour agencies, check quantities and units, and flag missing information instead of silently committing ambiguous data.

Page 76-92 © MANTECH PUBLICATIONS 2026. All Rights Reserved

5. Synthesis and positioning

The literature provides four foundations: NLP-based extraction in construction, conversational site reporting, joint intent and slot extraction, and language resources for Indian languages. Each covers only part of the problem. Document extraction is not conversational, conversational reporting is limited to daily reports, general slot-filling models are not tied to construction schemas, and language resources are not tied to construction activities. What is missing is an empirical assessment of a construction-specific pipeline that turns short, informal, multilingual site messages into valid operational transactions. We address this gap with a 300-record benchmark.

RESEARCH METHODOLOGY

1. Research design

The study is a systematic evaluation of a deterministic extraction module (Section 4) on a synthetic dataset with predefined splits. Five questions are examined: overall extraction (RQ1, Table II); language and style (RQ2, Table III); multi-action messages (RQ3, Section 5.3); configuration comparison (RQ4, Table II); and remaining failures (RQ5, Section 6.1).

2. Data source and dataset construction

The benchmark (data version V4) contains 300 synthetic construction-site messages with intent and entity labels produced by a dataset-generation program. No real site communication was used. Messages fall into four conditions: English, Marathi-English code-mixed, transliterated Marathi, and informal or noisy text. They are divided into development (135), validation (45), test (75) and challenge (45) splits. The challenge split uses wording not seen in development and includes ambiguous messages that need clarification. Labels follow the generation logic; no human annotation was involved.

Table I gives the distribution by primary intent. The test split is imbalanced, with only one site_expense and two query_expenses records. A 12-record multi-action subset of the test split, whose messages contain more than one intent, holds 17 ground-truth intent instances and is evaluated separately.

Page 76-92 © MANTECH PUBLICATIONS 2026. All Rights Reserved

Table 1: Distribution of Records by Primary Intent

Primary intent	Records (all splits)	Share of dataset	Records in test split
material_received	71	23.7%	16
labour_attendance	34	11.3%	4
material_consumed	33	11.0%	12
query_stock	31	10.3%	12
reminder	29	9.7%	8
material_damaged	28	9.3%	9
payment	26	8.7%	7
query_expenses	26	8.7%	2
rmc_received	12	4.0%	4
site_expense	10	3.3%	1
Total	**300**	**100.0%**	**75**

3. Tools and technologies

The pipeline is integrated with a desktop construction management application built with React 19, TypeScript and Vite, with Electron 34 as the desktop shell and application state in local storage. The pipeline uses a normalization glossary, construction-specific rules and templates, and validation and evaluation utilities. No external language service is used, and the experiments test the pipeline alone.

4. Compared configurations

Six variants are compared (Table II). The rule-based and raw-pattern baselines use simpler extraction. Configurations A and B extract without and with schema validation, respectively. Configuration C enables the remaining pipeline components and yields more entities than B with intent predictions unchanged. The proposed configuration adds the confirmation step, which cannot be exercised in an offline benchmark, so its metrics equal those of C. Likewise A and B are identical because validation does not alter the predictions being scored.

5. Evaluation metrics

Intent accuracy is the share of messages whose intent is correct. Entities are scored with macro and micro precision, recall and F1, and exact match requires a message to be fully

Page 76-92 © MANTECH PUBLICATIONS 2026. All Rights Reserved

correct. For multi-action messages, intent-instance accuracy and action-level precision, recall and F1 are reported.

6. Reproducibility

The fixed 300-record benchmark, defined partitions, deterministic extraction rules and standard metrics allow repeat evaluation. The source code and benchmark are not currently released.

PROPOSED FRAMEWORK

1. System architecture

The framework treats conversational construction reporting as structured data extraction rather than free question answering (Figure 1). A message is first normalized, then intents and entities are extracted, the result is mapped to the application schema and matched against active project entities, and validation and confidence checks are applied before a structured transaction is created.

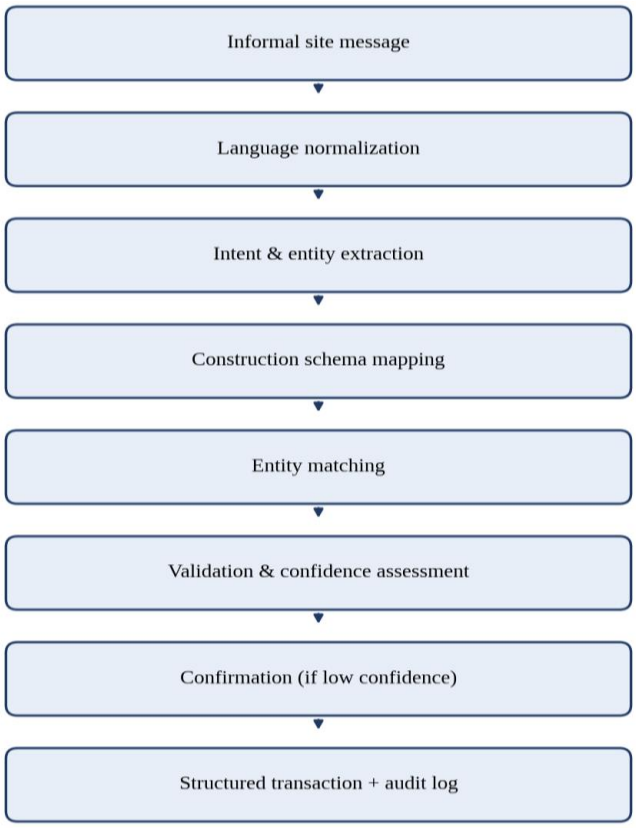


Figure 1: Proposed construction-specific pipeline for converting natural-language messages into structured data.

Page 76-92 © MANTECH PUBLICATIONS 2026. All Rights Reserved

2. Language normalization

A normalization layer maps construction terms in the message to a standard internal vocabulary covering materials, actions such as received, used and damaged, and general construction terms. This predefined step reduces lexical variety before pattern matching.

3. Intent and entity extraction

The supported intents are material_received, material_consumed, material_damaged, labour_attendance, rmc_received, site_expense, payment, reminder, query_stock and query_expenses. Quantities, units, project names, materials, organizations, suppliers and dates are recognized using construction dictionaries and templates. A multi-action message is recorded as several action instances rather than a single flattened label.

4. Schema mapping and entity matching

Extracted data is mapped to module schemas. A material receipt includes project, material, quantity, unit and time; a labour record includes project, agency or worker, count and attendance status; an RMC receipt includes project, supplier, quantity, price and delivery details. Extracted entities are then matched to the application state, so shortened or renamed projects, materials, suppliers and agencies resolve to existing records. The ablation in Section 3.4 compares whole configurations and does not isolate the matching module.

5. Validation, confidence and confirmation

Before a transaction is finalized, type-specific checks confirm positive amounts, existence of referenced entities, defined units, no negative inventory and absence of duplicates. Each extraction carries a confidence value, and low-confidence candidates are sent for user confirmation instead of being executed. These are design controls: the offline benchmark does not exercise confirmation, and we do not claim they improve F1. Validated candidates become structured operations that store the original text, intent, entities, confidence, validation status, confirmation state and final action, which supports auditability.

6. Integration with the management application

The application covers planning and estimation, procurement and stock, workforce and site operations, finance and monitoring, and a research dashboard (Figure 2a). The assistant supports ten message types (Figure 2b). In the prototype, messages are typed into an in-

Page 76-92 © MANTECH PUBLICATIONS 2026. All Rights Reserved

application chat that simulates a messaging interface; there was no live WhatsApp Business connection. Each message receives a structured summary of date, project, material, quantity, action, detected intent and confidence (Figure 2c), and is logged as processed or awaiting confirmation (Figure 2d). Because the pipeline writes to shared modules, transactions can be validated against current project, stock and vendor state.

Stock follows

Remaining = Opening + Received − Used − Damaged (1)

where damaged stock is treated as lost. For example, “Aaj ABC site la 120 cement bags ale, 85 use zale, 2 bags damaged” yields receipt, consumption and damage records, each updating the material state.

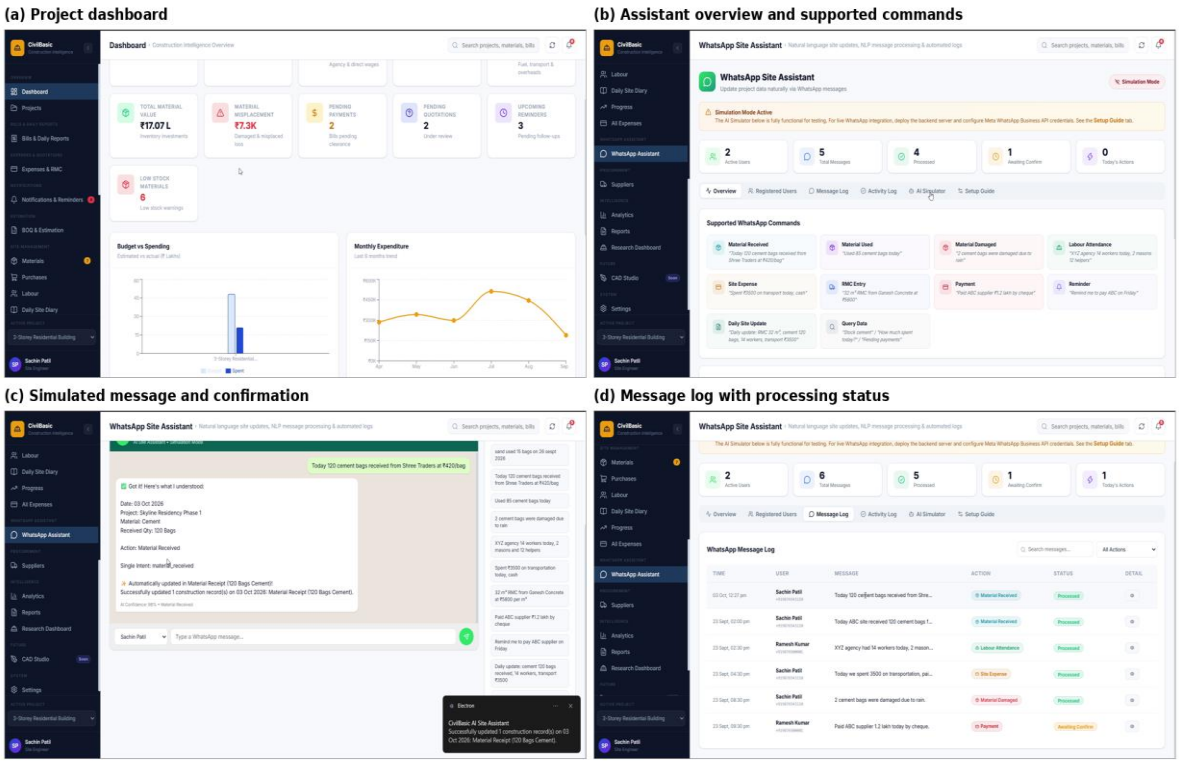


Figure 2: CivilBasic prototype screens: (a) project dashboard; (b) assistant overview and supported message types; (c) simulated message with structured confirmation; (d) message log with processing status.

Page 76-92 © MANTECH PUBLICATIONS 2026. All Rights Reserved

RESULTS

1. Held-out test results

Table II reports results on the held-out test set (N = 75). The proposed system reaches 61.33% intent accuracy, 53.04% entity macro precision, 35.63% macro recall, 41.47% macro F1, 45.01% micro F1 and 1.33% exact match. Its macro F1 exceeds the raw-pattern baseline (20.18%) by 21.29 percentage points, and its micro F1 exceeds the baseline's 22.54% by 22.47 points. Intent accuracy is identical for A, B, C and the proposed configuration, so the gains appear only at entity level.

Table 2: Held-out test results (N = 75)

Variant	Intent accuracy	Entity macro precision	Entity macro recall	Entity macro F1	Entity micro F1	Exact match
Rule-based	49.33%	32.89%	14.74%	19.69%	20.69%	0%
Raw pattern	58.67%	33.64%	15.43%	20.18%	22.54%	0%
A	61.33%	42.24%	21.26%	27.80%	32.87%	0%
B	61.33%	42.24%	21.26%	27.80%	32.87%	0%
C	61.33%	53.04%	35.63%	41.47%	45.01%	1.33%
Proposed	61.33%	53.04%	35.63%	41.47%	45.01%	1.33%

2. Split-wise results

Intent accuracy of the proposed system was 59.26% on development, 46.67% on validation, 61.33% on test and 44.44% on the challenge split. Entity macro F1 was 48.73%, 35.17%, 41.47% and 44.28%, and micro F1 was 51.86%, 37.76%, 45.01% and 49.55%, respectively. The challenge split is therefore harder for intent classification but not for entity extraction. Exact match stayed between 0% and 2.96%. Average processing time was about 0.54–0.57 ms per message, which describes computational cost only and says nothing about deployment accuracy.

Page 76-92 © MANTECH PUBLICATIONS 2026. All Rights Reserved

3. Language-wise and multi-action results

Table III shows language-wise results. English has the highest intent accuracy (87.50%). Marathi-English has lower intent accuracy (57.89%) but the highest entity macro F1 (53.25%). Transliterated Marathi has the lowest intent accuracy (33.33%), and informal messages reach 52.94%.

Table 3: Language-wise test results for the proposed system

Language/style	N	Intent accuracy	Entity macro F1
English	24	87.50%	38.75%
Marathi-English	19	57.89%	53.25%
Marathi-transliterated	15	33.33%	37.91%
Informal/noisy	17	52.94%	35.29%

On the 12-record multi-action test subset (17 ground-truth intent instances), 12 instances were identified correctly, giving 70.59% intent-instance accuracy. At action level there were 14 true positives, 3 false positives and 3 false negatives, so precision, recall and action F1 were each 82.35%. This score covers only the multi-action subset and is not comparable to the full-test-set F1. On the other splits, intent-instance accuracy and action F1 were 63.64% and 81.82% (development, 20 records), 57.14% and 76.19% (validation, 8 records), and 46.67% and 72.22% (challenge, 10 records).

DISCUSSION

1. Error analysis

The system misclassified 29 of the 75 test records: six payment, five material_damaged, four material_consumed, four query_stock, four rmc_received, three reminder, two query_expenses and one material_received. All labour_attendance and site_expense records were correct. Most errors share a few causes: material_received was assigned to 20 records, five payment messages were also labelled labour_attendance, and in three cases rmc_received was predicted with or instead of the correct intent.

Examples illustrate the problem. “What are the pending expenses at Ivy Gate site” (R000116) was classified as site_expense instead of query_expenses. “Paid Rs 25000 to Shinde agency

Page 76-92 © MANTECH PUBLICATIONS 2026. All Rights Reserved

via bank transfer for work at Maple Residency" (R000167) was labelled labour_attendance together with payment. "Lakeview Complex site mein 45 kg steel reinforcement use hua column casting ke liye" (R000132) was labelled material_received instead of material_consumed, and the Marathi stock query "Green Valley site par 20mm aggregate ka kitna stock shillak hai" (R000212) was also labelled material_received. Recognizing construction entities is therefore not enough; the action must also be identified correctly.

2. Configuration ablation

Configuration C improves entity macro F1 from 27.80% to 41.47% and micro F1 from 32.87% to 45.01% over B, with intent accuracy unchanged at 61.33%. Per-intent entity F1 is higher for C than B in 8 of 10 intent classes, with the largest gains for query_stock (0.00 to 0.35), rmc_received (0.36 to 0.55) and labour_attendance (0.59 to 0.74). Because whole configurations are compared, the gain cannot be attributed to a single module. C and the proposed configuration perform identically, so no benefit can be assigned to validation or confirmation.

3. Language and noise variation

Intent accuracy is far higher for English than for transliterated Marathi, whereas Marathi-English reaches a reasonable entity macro F1 despite lower intent accuracy (Table III). Lexical diversity, spelling, transliteration, class distribution or the coverage of the deterministic patterns could explain this, but the study does not separate these factors.

4. Comparison with previous research and implications

Earlier conversational studies (Cho and Lee, 2019; Cho et al., 2023; Xiao et al., 2024) did not extract multiple operational transactions from informal messages, so numerical comparison is not possible and we do not claim to outperform them. The weaker intent detection for Marathi-English and transliterated messages is consistent with the difficulty of mixed-language input reported by Amin et al. (2024). Errors in which correct entities are attached to the wrong action support modelling intent jointly with slots (Goo et al., 2018; Chen et al., 2019). The lexical-trigger errors suggest comparing against LLM-based extraction (Dagdelen et al., 2024), which we did not test.

Page 76-92 © MANTECH PUBLICATIONS 2026. All Rights Reserved

The findings suggest that a deterministic construction-specific pipeline can recover part of the entity content of informal messages. Given the many intent errors and the 1.33% exact match, however, automatic transaction confirmation is not yet feasible.

LIMITATIONS AND FUTURE DIRECTIONS

1. Limitations

The benchmark consists of 300 synthetic records with automatically generated labels, not field messages or human annotation. Records and labels come from the same generator, so results may be optimistic, and the language use is narrow. With 75 test records and a single exact match, results are noisy, and no confidence intervals are reported. The imbalanced split makes per-class and macro-averaged metrics unreliable. The system is purely deterministic and was not compared with other systems or LLMs, and no usability study was conducted. The ablation compares whole configurations rather than individual components. Code and data are not publicly available.

2. Unresolved challenges and future work

The main challenge is action identification: inappropriate lexical triggers and rule conflicts create false intents such as material_received. Entity recall (35.63%) is lower than precision (53.04%), exact match remains low, and transliterated Marathi has the weakest intent accuracy. The effects of entity matching, validation and confirmation on record quality are still untested. Future work should use real anonymized site messages, independent human annotation and larger multilingual datasets, compare against stronger baselines and LLM-based extraction, isolate individual components, and run usability studies.

CONCLUSION

This paper developed a construction-specific conversational framework that turns informal site messages into structured operational data through language normalization, intent and entity extraction, schema mapping, entity matching, validation, confidence-based confirmation and transaction creation within a construction management application.

On a 75-record test set the system reached 61.33% intent accuracy, 41.47% entity macro F1, 45.01% entity micro F1 and 1.33% exact match. On the multi-action subset, intent-instance

Page 76-92 © MANTECH PUBLICATIONS 2026. All Rights Reserved

accuracy was 70.59% and action F1 82.35%. It outperformed configurations A and B on entity-level metrics, although the ablation does not isolate individual components.

The results show that structured extraction from informal construction messages is feasible, but that the task remains difficult for multilingual, transliterated and noisy messages. This is a controlled prototype evaluation on synthetic data and is not evidence of field readiness. Future work will add usability evaluation and comparison with language-model-based extraction, using real anonymized site data, independent annotation and larger multilingual benchmarks.

DECLARATIONS

Conflict of interest: The authors declare no financial or personal interests that could have influenced the work reported in this paper.

Ethics statement: This study did not involve humans, animals or identifiable data. The dataset consists of synthetically generated messages and contains no real site communication, so ethical approval and consent were not applicable.

Funding: The authors received no funding for this research.

REFERENCES

1. Al Qady, M. and Kandil, A. (2010), “Concept relation extraction from construction documents using natural language processing”, Journal of Construction Engineering and Management, Vol. 136 No. 3, pp. 294–302, doi: 10.1061/(ASCE)CO.1943-7862.0000131.
2. Amin, D., Govilkar, S., Kulkarni, S., Vyawahare, M., Chavan, S. and Pandey, P. (2024), “Exploring Marathi-English code-mixing: comprehensive analysis of NLP applications (QA and NER)”, International Journal of Intelligent Systems and Applications in Engineering, Vol. 12 No. 18s, pp. 420–427.
3. Chen, Q., Zhuo, Z. and Wang, W. (2019), “BERT for joint intent classification and slot filling”, arXiv preprint arXiv:1902.10909, doi: 10.48550/arXiv.1902.10909.

Page 76-92 © MANTECH PUBLICATIONS 2026. All Rights Reserved

4. Cho, J. and Lee, G. (2019), "A chatbot system for construction daily report information management", in Proceedings of the 36th International Symposium on Automation and Robotics in Construction (ISARC), pp. 429–437, doi: 10.22260/ISARC2019/0058.
5. Cho, J., Lee, G., Song, T. and Jeong, H.D. (2023), "Chatbot-engaged construction daily work report using mobile messenger", Automation in Construction, Vol. 154, Article 105007, doi: 10.1016/j.autcon.2023.105007.
6. Dagdelen, J., Dunn, A., Lee, S., Walker, N., Rosen, A.S., Ceder, G., Persson, K.A. and Jain, A. (2024), "Structured information extraction from scientific text with large language models", Nature Communications, Vol. 15, Article 1418, doi: 10.1038/s41467-024-45563-x.
7. Goo, C.-W., Gao, G., Hsu, Y.-K., Huo, C.-L., Chen, T.-C., Hsu, K.-W. and Chen, Y.-N. (2018), "Slot-gated modeling for joint slot filling and intent prediction", in Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, New Orleans, LA, pp. 753–757, doi: 10.18653/v1/N18-2118.
8. Gupta, R., Srivastava, V. and Singh, M. (2023), "MUTANT: a multi-sentential code-mixed Hinglish dataset", in Findings of the Association for Computational Linguistics: EACL 2023, Dubrovnik, pp. 744–753, doi: 10.18653/v1/2023.findings-eacl.56.
9. Joshi, R. (2022), "L3Cube-MahaCorpus and MahaBERT: Marathi monolingual corpus, Marathi BERT language models, and resources", in Proceedings of the WILDRE-6 Workshop, 13th Language Resources and Evaluation Conference (LREC 2022), Marseille, pp. 97–101.
10. Ko, T., Shrestha, R. and Lee, J. (2024), "Pro-active allocation of project requirements through natural language processing (NLP) and project information system integration", in Construction Research Congress 2024, pp. 1308–1316, doi: 10.1061/9780784485262.133.
11. Litake, O., Sabane, M.R., Patil, P.S., Ranade, A.A. and Joshi, R. (2022), "L3Cube-MahaNER: a Marathi named entity recognition dataset and BERT models", in Proceedings of the WILDRE-6 Workshop, 13th Language Resources and Evaluation Conference (LREC 2022), Marseille, pp. 29–34, doi: 10.48550/arXiv.2204.06029.

Page 76-92 © MANTECH PUBLICATIONS 2026. All Rights Reserved

12. Ren, R. and Zhang, J. (2022), "Construction procedural information extraction from textual sources to support scheduling", in Construction Research Congress 2022, pp. 330–339, doi: 10.1061/9780784483978.035.
13. Shamshiri, A., Ryu, K.R. and Park, J.Y. (2024), "Text mining and natural language processing in construction", Automation in Construction, Vol. 158, Article 105200, doi: 10.1016/j.autcon.2023.105200.
14. Sheth, R., Beniwal, H. and Singh, M. (2025), "COMI-LINGUA: expert annotated large-scale dataset for multitask NLP in Hindi-English code-mixing", in Findings of the Association for Computational Linguistics: EMNLP 2025, Suzhou, pp. 7973–7992, doi: 10.18653/v1/2025.findings-emnlp.422.
15. ul Hassan, F. and Le, T. (2022), "Extraction of activities information from construction contracts using natural language processing (NLP) methods to support scheduling", in Construction Research Congress 2022, pp. 773–781, doi: 10.1061/9780784483961.081.
16. Wu, C., Li, X., Guo, Y., Wang, J., Ren, Z., Wang, M. and Yang, Z. (2022), "Natural language processing for smart construction: current status and future directions", Automation in Construction, Vol. 134, Article 104059, doi: 10.1016/j.autcon.2021.104059.
17. Xiao, B., Wang, Y., Zhang, Y., Chen, C. and Darko, A. (2024), "Automated daily report generation from construction videos using ChatGPT and computer vision", Automation in Construction, Vol. 168, Article 105874, doi: 10.1016/j.autcon.2024.105874.

Page 76-92 © MANTECH PUBLICATIONS 2026. All Rights Reserved

***Author for Correspondence**

Param Patel

E-mail: pp9361682@gmail.com

[1]Student, Department of Computer Science and Engineering (Data Science), D. Y. Patil College of Engineering and Technology, Kolhapur, Maharashtra, India

[2,3]Student, Department of Computer Science and Engineering (Data Science), D. Y. Patil Technical Campus, Talsande, Maharashtra, India

Received Date: October 21, 2026

Accepted Date: October 31, 2026

Published Date: November 03, 2026

Citation: Shreya Jadhav, Param Patel, Atharva Mangutkar**.** Conversational Construction Site Management & Reporting: Extracting Structured Data from Informal Natural-Language Messages. International Journal of Data Science and Analytics Innovations. 2026; 2(2): 76-92p.

Page 76-92 © MANTECH PUBLICATIONS 2026. All Rights Reserved