

# Ensemble Size Variations and Relation Extraction Consistency in Product Knowledge Graphs

Assignee Research

June 13, 2026

## Abstract

The rapid expansion of e-commerce platforms generates vast amounts of unstructured product data, creating significant challenges for information retrieval, recommendation systems, and data analytics. Knowledge Graphs (KGs) offer a structured, interpretable format to organize such data, yet constructing product-specific KGs remains a complex and manual process. This paper introduces a fully automated, AI agent-driven framework for constructing product knowledge graphs directly from unstructured product descriptions. Leveraging Large Language Models (LLMs), our method operates in three stages us

## 1 Introduction

This paper examines: AI Agent-Driven Framework for Automated Product Knowledge Graph Construction in E-Commerce. Research question: What is the impact of varying the size of the AI agent ensemble on the consistency of extracted relations in product knowledge graphs, measured by F1-score on a held-out test set of unstructured product data?.

## 2 Methodology

Systematic literature search across multiple databases yielded 9 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

9 papers retrieved. 18 claims extracted; 15 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.



## 5 Extracted Claims

Claim	Verified	Confidence
The framework was evaluated on a dataset of 291 product descriptions for the category 'air conditioners' from a retail s	✓	0.19
The evaluation methodology includes three dimensions: Ontology Coverage, Ontology Quality, and Knowledge Graph Populatio	×	0.14
Ontology Coverage is assessed by measuring the number of extracted classes, attributes, and relationships.	✓	0.20
Ontology Quality is evaluated via expert annotation focusing on coherence, generality, usability, redundancy, inconsiste	✓	0.18
Knowledge Graph Population is evaluated by measuring the number of generated RDF triples and the proportion of instantia	✓	0.21
The constructed ontology comprises 42 classes.	✓	0.15
The constructed ontology comprises 69 properties, consisting of 20 data attributes and 49 object relationships.	✓	0.20
Manual review found the schema to be modular, hierarchically structured, and readily extendable.	✓	0.15
The qualitative assessment reported minimal redundancy and no significant inconsistencies in the ontology.	×	0.13
The framework successfully processed 282 out of 291 product descriptions.	✓	0.19
Processing failures occurred on 9 instances (3%) due to invalid RDF outputs.	×	0.11
The resulting knowledge graph contains 7,459 RDF triples.	✓	0.20
The knowledge graph achieved 97.1% coverage of the properties defined in the ontology.	✓	0.16
The framework consists of three main stages: ontology creation and expansion, ontology refinement, and knowledge graph p	✓	0.19
Each stage of the framework is orchestrated through a modular pipeline of LLM-powered agents.	✓	0.18
The Ontology Creation agent extracts initial elements and organizes them into RDF/Turtle format with rdfs:domain, rdfs:r	✓	0.19
Ontology expansion proceeds iteratively by presenting additional product descriptions to the agent to generalize beyond	✓	0.17
During expansion, the agent integrates new classes or properties without removing existing elements.	✓	0.17

## References

- <http://arxiv.org/abs/2404.17000v1>
- <http://arxiv.org/abs/2605.05991v1>
- <http://arxiv.org/abs/2511.11017v1>