

Article

From Context to Aspects: LLM-Based Implicit Aspect Extraction with Paraphrased Input and Knowledge Graph Support

Lujain Abdulrahman Alawwad ^{1,2,*}  and Mohamed El Bachir Menai ¹ 

¹ Department of Computer Science, King Saud University, Riyadh 11451, Saudi Arabia; menai@ksu.edu.sa

² Department of Computing and Informatics, Saudi Electronic University, Riyadh 13323, Saudi Arabia

* Correspondence: l.alawwad@seu.edu.sa

Abstract

While aspect-based sentiment analysis (ABSA) has made significant progress in the identification of explicit opinion targets, the more challenging case of implicit aspects remains insufficiently studied. Implicit aspect extraction is particularly challenging, as it relies on contextual and semantic cues and requires systems to infer what reviewers mean rather than what they state explicitly. A four-component hybrid pipeline is proposed for explicit and implicit aspect extraction, formulating the task as controlled text generation. The pipeline combines (i) a fine-tuned decoder-only large language model as a generative baseline, (ii) an iterative residual generation strategy that recovers multiple aspects through successive masked generation passes, (iii) paraphrase-based input transformation to broaden the contextual signal, and (iv) domain-specific knowledge graphs activated by linguistic signals to infer implicit aspects. The novelty lies not in the individual components themselves but in their principled orchestration and the linguistically motivated gating logic governing the activation of each stage. Extensive experiments are conducted on eight benchmark ABSA datasets spanning both English and Arabic: SemEval-2014, SemEval-2015, SemEval-2016, ACOS, and M-ABSA for English; and SemEval-2016, HAAD, and M-ABSA for Arabic. The proposed solution outperforms strong baseline methods and recent state-of-the-art models on English datasets, with F_1 -scores of 0.8533, 0.713, 0.7859, 0.793, and 0.664, respectively. On Arabic datasets, the best-performing configurations achieve F_1 -scores of 0.7632, 0.4765, and 0.7656 on SemEval-2016, HAAD, and M-ABSA, respectively, with the knowledge-graph component providing consistent and statistically significant gains for implicit aspect identification in both languages. These results demonstrate the effectiveness of generative modeling, iterative generation, paraphrasing, and structured knowledge for aspect extraction and highlight the potential of the proposed approach for implicit aspect identification, in particular for morphologically rich languages such as Arabic, where annotated resources remain scarce.

Keywords: implicit aspect extraction; transformers architecture; transfer learning; text paraphrasing; knowledge graphs



Academic Editor: Isidoros Perikos

Received: 9 May 2026

Revised: 19 June 2026

Accepted: 21 June 2026

Published: 25 June 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

When a reviewer writes “the laptop feels like carrying a brick,” they are not describing the laptop’s material composition—they are complaining about its *weight*, an aspect never mentioned by name. Whereas humans decode such expressions effortlessly through pragmatic inference and world knowledge, automated systems, for the most part, do not. This gap sits at the heart of opinion mining.

The volume of opinion-bearing text produced online—reviews, ratings, and commentary on products, services, and public issues—represents a remarkably rich source of behavioral insight for researchers and practitioners alike. Sentiment analysis (SA) was developed to leverage this resource systematically, with the goal of modeling the subjective orientation expressed in natural language text. Document- and sentence-level SA methods, however, assign a single polarity to an entire text, obscuring cases where a reviewer simultaneously praises one feature and criticizes another. Aspect-based sentiment analysis (ABSA) addresses this limitation by decomposing sentiment analysis into constituent opinion targets: discrete product or service features—such as battery life, performance, food quality, or service—each carrying its own sentiment value.

ABSA subsumes two interdependent sub-tasks: aspect extraction, which identifies the features being evaluated, and aspect sentiment classification, which assigns a polarity label to each identified feature. Within aspect extraction, the distinction between explicit and implicit aspects is of particular consequence. Explicit aspects appear verbatim in the text, as in “The screen is bright.” Implicit aspects, by contrast, must be inferred from contextual and pragmatic cues: “The laptop is heavy” implicates *weight* without naming it. Models lack the world knowledge and pragmatic reasoning that this inference requires, rendering implicit aspect extraction a persistently difficult problem whose prevalence in authentic reviews cannot be overlooked.

The research community has institutionalized its focus on aspect extraction through several benchmark shared tasks, most notably SemEval-2014 Task 4 [1], SemEval-2015 Task 12 [2], and SemEval-2016 Task 5 [3], which established widely adopted evaluation protocols. Each shared task released a companion dataset that has since become a standard benchmark for aspect term extraction (ATE): the SemEval-2014 dataset (<https://alt.qcri.org/semeval2014/task4/> (accessed on 10 June 2026)), the SemEval-2015 dataset (<https://alt.qcri.org/semeval2015/task12/> (accessed on 10 June 2026)), and the SemEval-2016 dataset (<https://alt.qcri.org/semeval2016/task5/> (accessed on 10 June 2026)). Further datasets have since extended benchmark coverage, including ACOS (<https://github.com/NUSTM/ACOS> (accessed on 10 June 2026)) [4] and the recent multilingual benchmark M-ABSA (<https://github.com/swaggy66/M-ABSA> (accessed on 10 June 2026)) [5]. Methodological approaches have evolved in parallel: early work relied on statistical and feature-based machine learning [6–10] before the field shifted toward deep learning [11–15]. More recently, large language models (LLMs) have demonstrated substantial gains across NLP tasks and are increasingly being evaluated for ABSA [16–18].

Several challenges, however, remain unresolved. First, a single sentence may implicate multiple aspects simultaneously: “The place was loud and overpriced” conveys opinions about both *ambience* and *price* without naming either. Second, the same opinion word may index different aspects depending on domain: “fast” implies *service* in a restaurant context but *performance* in a laptop review. Third, annotated corpora that label implicit aspects explicitly remain scarce—particularly for Arabic—constraining both the availability of supervision and the scope of evaluation. These challenges motivate the research questions addressed in this work.

The contributions of this work are as follows:

- We propose an *iterative residual generation* strategy for generative aspect extraction, in which the model performs successive generation passes over a progressively masked input, enabling recovery of aspects missed in single-pass decoding. A paraphrase-based input strategy further broadens the contextual signal available at each pass.
- We introduce *lonely adjective detection* as a linguistically motivated gating signal that selectively activates knowledge-graph inference when opinion-bearing adjectives lack

an explicit syntactic noun head—a surface configuration strongly associated with implicit aspect expression.

- We extend the full pipeline to Arabic, incorporating morphology-aware adaptations at each stage to account for the language’s rich inflectional structure.
- We provide empirical evaluation across eight English and Arabic benchmarks, with ablation studies isolating the contribution of each pipeline component.

The remainder of this paper is organized as follows. Section 2 provides background on text paraphrasing, knowledge graphs, and the challenges of Arabic natural language processing and states the research questions. Section 3 reviews related work on deep learning and hybrid approaches to explicit and implicit aspect extraction. Section 4 presents the proposed four-component pipeline and its Arabic adaptation. Section 5 describes the experimental setup and reports the results, including ablation studies, error and sensitivity analyses, statistical significance tests, comparisons with state-of-the-art methods, and a discussion of limitations. Section 6 concludes this paper and outlines future work.

2. Background and Preliminaries

2.1. Text Paraphrasing

Text paraphrasing—the reformulation of a source expression into alternative wording while preserving its meaning—is a well-established technique in natural language processing, with applications spanning data augmentation, machine translation, and question answering [19]. In the context of aspect extraction, paraphrasing serves a distinct purpose: by introducing syntactic variation while holding semantics constant, paraphrased inputs can surface implicit aspects that would otherwise remain latent under any single syntactic realization.

2.2. Knowledge Graphs

A knowledge graph (KG) represents facts as typed nodes and relations, formalized as triples (h, r, t) , enabling models to reason over entities and their neighborhoods rather than over isolated tokens. For implicit aspect extraction, a KG provides grounded opinion-to-aspect mappings that purely data-driven models cannot reliably infer from distributional statistics.

General-purpose resources fall short for this task. Topic models such as LDA [20] produce soft probabilistic assignments rather than precise mappings; preliminary experiments conducted in this work confirm that their signals consistently degrade extraction performance. WordNet [21] and its Arabic counterpart [22] lack domain-specific opinion-to-aspect coverage. ConceptNet [23] encodes broader commonsense relations and has been applied to sentiment analysis [24,25]; however, its knowledge is domain-general, its Arabic subgraph is sparse, and its undirected relational structure lacks the domain hierarchy required to resolve ambiguous opinion words across domains (e.g., *slow* mapping to SERVICE#GENERAL in restaurants versus LAPTOP#PERFORMANCE in laptop reviews). Domain-specific KGs have demonstrated more consistent ABSA gains [25–28]; the KG constructed in the present work follows that direction, encoding directed, domain-specific opinion-to-aspect mappings derived from training data.

2.3. Arabic NLP and Its Challenges

Arabic presents several NLP challenges due to its extensive inflectional morphology, by which one word may have many surface forms based on gender, number, and syntactic position. Three challenges are of direct relevance to this work.

Orthographic inconsistency. Writers inconsistently interchange letters such as ا with إ with ل, ي with ع, and ة with ه, causing the same term to appear in multiple sur-

face forms [29]. Left unaddressed, this variability causes identical aspect terms to be treated as distinct tokens during model inference, paraphrase comparison, and knowledge graph lookup.

Diacritics. Short vowels alter meaning fundamentally (e.g., *كَتَبَ wrote* vs. *كُتِبَ books*) yet are routinely omitted in informal text, introducing systematic ambiguity that complicates morphological analysis and aspect matching.

Dialectal variation and resource poverty. Arabic social media largely follows dialectal rather than Modern Standard Arabic (MSA) conventions; syntactic analyzers and domain lexicons for dialectal Arabic remain scarce. No large, publicly annotated Arabic dataset for implicit aspect extraction exists, constraining both model training and evaluation.

These challenges directly motivate the Arabic-specific design decisions described in Section 4.2.

2.4. Research Questions

This work addresses the following research questions:

1. How can a generative model be adapted to recover multiple implicit aspects from a single sentence, where each opinion expression must be mapped to its corresponding aspect?
2. How can domain-ambiguous opinion words be resolved to their intended aspect targets across different review domains?
3. To what extent does instruction-tuned generative modeling improve implicit aspect extraction performance relative to discriminative approaches?
4. Does paraphrase-based input diversification yield measurable gains in implicit aspect recall, and under what conditions?

3. Related Work

The year 2014 marked a turning point for the ABSA research community, with the release of benchmark datasets and shared evaluation settings that established the foundation for subsequent research. This review focuses on Arabic ABSA; it emphasizes deep learning, hybrid, and large language model (LLM)-based approaches, which currently dominate aspect extraction research. Non-Arabic studies included here are restricted to recent work targeting implicit aspect extraction. The reviewed studies are grouped into deep learning methods and hybrid approaches that integrate neural models with structured or symbolic components.

3.1. Deep Learning Approaches

The adoption of deep learning marked a significant shift in aspect extraction by enabling models to capture richer contextual representations. In Arabic, early neural architectures applied CNNs [30] and hierarchical BiLSTMs [31], followed by models combining CNNs with LSTMs [32] or GRU-based architectures with bilingual embeddings [11]. More recent work incorporated transformer-based encoders, including multi-task extensions such as MTL-AraBERT, for joint aspect extraction and sentiment classification [33].

Transformer architectures have since become central to Arabic ABSA. AraBERT has been integrated with stacked LSTM–GRU layers and attention mechanisms [12], CRF-based tagging pipelines [34], and sequence-to-sequence models generating aspect–sentiment pairs [35]. Domain-adaptive fine-tuning has also proven effective, as demonstrated by CAMELBERT for noisy and dialectal text [13]. Comparative evaluations consistently show transformer-based models outperforming recurrent neural architectures.

Beyond Arabic, deep learning has driven substantial advances in aspect extraction. BERT-based embeddings have been combined with CNN, BiLSTM, RCNN, and attention mechanisms [36], ensemble BiLSTMs [37], topic-aware CNNs [14,15], and graph convolutional networks [38]. Further refinements include attention-augmented BiLSTMs [39,40], BiLSTM–CRF architectures [41], and domain-adaptive BiLSTMs with multi-head attention [42]. Additional contributions include span-level joint extraction models [43] and constituency-aware architectures [44].

3.2. Deep Learning Approaches for Implicit Aspect Extraction

While most deep learning research focuses on explicit aspects, a smaller body of work directly targets implicit aspect extraction. Soni and Rambola [45] combined recurrent neural networks with semantic similarity measures derived from WordNet and spaCy to map opinion expressions to latent aspects. More recent studies have explored LLMs, with [16] evaluating GPT-3.5 and GPT-4 for aspect category detection and demonstrating that prompting and fine-tuning strategies can capture implicit signals. Similarly, ref. [46] employed BERT in a zero-shot setting to address multiple ABSA sub-tasks, including implicit categories. Data augmentation strategies, such as Word2Vec-based replacements in [47], aimed to enrich training signals, although their impact remained limited.

3.3. Hybrid Approaches

Hybrid approaches integrate neural models with symbolic, statistical, or structural components to enhance robustness. In Arabic, early hybrid work integrated stacked LSTMs with logistic regression for aspect classification [48]. Ontology-driven and neural hybrid designs have also been explored, including BiLSTM–CRF and BiGRU–CRF pipelines combined with fastText or AraBERT embeddings [49–53]. These models leverage recurrent encoders with structured decoders to better capture sequence dependencies. More recent developments include pseudo-labeling with the GPT-4.1-mini model (OpenAI, 2025) prior to transformer fine-tuning [54] and pipelines combining BERTopic clustering, CAMELBERT sentiment analysis, and AraBART summarization [17].

Beyond Arabic, hybrid approaches have similarly evolved. Some combine syntactic rules with BERT embeddings and neural classifiers [55], while others extend encoder–decoder architectures with contrastive learning for structured quadruple extraction [56].

3.4. Hybrid Approaches for Implicit Aspect Extraction

A limited but growing body of hybrid research directly addresses implicit aspect extraction. Ahmed et al. [57] combined topic modeling with BERT to learn implicit aspect-specific representations, while Feng et al. [58] and Lazhar [59] integrated CNNs and association rules to link opinion expressions to hidden aspects. More recent work has shifted toward transformer-based hybrids. Li et al. [60] introduced generative T5-based frameworks augmented with graph neural networks and prompt fusion, surpassing earlier ACOS baselines. Other studies proposed multi-stage pipelines combining BERT embeddings, semantic similarity, and classical classifiers [61] and frequency–syntax–CRF frameworks addressing both explicit and implicit aspects [62]. WoSe [63], a representation-driven hybrid, employed dual-level encoders and CRF decoding to better capture implicit aspects, while ref. [64] explored prompting and synthetic augmentation strategies using large language models.

Most recently, knowledge-enhanced transformer hybrids have advanced implicit aspect detection further. Dubey et al. [28] combined a domain-adaptive knowledge graph with a hierarchical graph transformer, addressing implicit aspects through contrastive learning and semantic role labeling. Ahmed et al. [65] paired instruction-tuned models with auxiliary sentence construction and knowledge-enhanced graphs for implicit aspect

detection. Zhou et al. [66] introduced a generative framework that produces the content of implicit terms directly rather than emitting null placeholders, while Bu et al. [67] extended prompt-based implicit sentiment modeling to multi-lingual settings.

3.5. Critical Assessment and Research Gap

Three gaps emerge from this review. First, Arabic ABSA research overwhelmingly targets explicit aspects. The few studies engaging with implicit aspects in Arabic [16,46,47] address category detection or label augmentation rather than proposing a dedicated implicit-to-aspect inference mechanism. Second, even with transformer and LLM-based models, current solutions typically rely on single-pass, model-internal inference, which limits their ability to recover multiple implicit aspects from a single sentence. Third, while KG-augmented inference has been explored in English [28,60,65], and earlier Arabic efforts applied manually engineered ontologies, semantic-similarity measures, and knowledge hierarchies to implicit aspect inference on private corpora [68–72], no prior Arabic work constructed a corpus-derived, domain-specific knowledge graph encoding directed opinion-to-aspect mappings, integrated it with a generative LLM pipeline, or evaluated it on public benchmarks. These gaps motivate the proposed solution.

The proposed solution differs from the closest hybrid systems in three respects. Unlike [28,60], which inject graph knowledge uniformly into every input, KG inference in the present work is gated by linguistically motivated cues—lonely adjectives and verb-phrase clues—so that symbolic reasoning is invoked only where implicit aspect expression is structurally signaled. Unlike [64,65], which augment inputs or instructions prior to single-pass extraction, the proposed iterative residual generation conditions each pass on a progressively masked input, recovering aspects that single-pass decoding systematically misses. Finally, no prior hybrid system extends these mechanisms to Arabic; the corpus-derived Arabic knowledge graph and morphology-aware adaptations constitute the first such extension evaluated on public benchmarks. The novelty therefore lies not in the individual components but in their conditional orchestration and cross-lingual generalization.

4. Methodology

This section presents the proposed four-component pipeline for explicit and implicit aspect extraction. The pipeline comprises a fine-tuned decoder-only baseline model (BL) that provides the generative foundation; an iterative residual generation component (IG) that recovers aspects missed in preceding rounds through successive masked generation; a paraphrase-based input variation module (PARA) that promotes cross-paraphrase agreement and surfaces latent aspects; and a domain-specific knowledge graph (KG) that provides structured opinion-to-aspect mappings for implicit aspect inference. The overall architecture is illustrated in Figure 1.

This section is organized as follows: each component is described in detail, followed by the end-to-end inference procedure that formalizes their interaction, and concluding with the Arabic-specific adaptations and preprocessing steps required to extend the pipeline to Arabic.

4.1. Proposed Solution Components

4.1.1. Baseline Model: Fine-Tuned Generative Model

A decoder-only LLM is fine-tuned for aspect extraction using instruction-based supervised fine-tuning (SFT). Each training instance is formatted as an (**instruction**, **sentence**, **aspect**) triple using the following prompt template:




Figure 1. Overview of the proposed four-component pipeline for explicit and implicit aspect extraction. Arrows indicate the flow of sentence representations. Parallel double lines indicate multiple original and paraphrased inputs passed in parallel to the baseline model.

The instruction field is fixed across all instances, with the language field instantiated per pipeline (English or Arabic):

Act as a Natural Language Processing expert in {language} language. Your task is to extract aspect term(s) that are explicitly mentioned in the given sentence.

- Only return aspect terms that appear verbatim in the sentence.
- Do not infer or invent related words.
- If no aspect is found, return ``none''.

Aspect extraction is framed as text generation: the model produces explicit aspect terms present in the input, or returns none. This fine-tuned model serves as the baseline (BL) component. At inference, the same template is used with the *Aspect* field left empty for the model to predict.

The instruction constraint (“do not infer or invent related words”) is a deliberate hallucination guard; it does not prevent the pipeline from identifying implicit aspects. The BL model receives paraphrased variants of each input sentence (Section 4.1.3), which may render implicit aspects verbatim, allowing the model to extract them while faithfully following its explicit-only instruction. The KG component operates independently of the generative model: it always scans the input sentence for lonely adjectives and verb-phrase clue words drawn from a curated whitelist (Algorithm 1, line 21); if any clue is detected, it maps that clue to a domain-consistent aspect via the knowledge graph. KG activation is based on clue detection alone—not on NULL output from the BL model. The pipeline thus provides two

complementary pathways for implicit aspect recovery: paraphrase-induced verbalization handled by the BL model and clue-gated KG inference running in parallel.

Algorithm 1: End-to-End Inference for Explicit and Implicit Aspect Extraction

Input : s : input sentence; M : fine-tuned generative LLM; G : domain knowledge graph; R : maximum IG rounds; N : number of paraphrases; τ : similarity threshold; $min_consensus$: minimum prediction agreement across paraphrases

Output : A : final set of extracted aspects

```

1  $A \leftarrow \emptyset$  // Final aspect set
2  $C \leftarrow \emptyset$  // Intermediate aspect predictions
// Step 1: Generate paraphrases
3  $P \leftarrow \{s\}$ 
4 for  $i \leftarrow 1$  to  $N$  do
5    $p_i \leftarrow LLM(ParaphrasePrompt(s))$ 
6    $P \leftarrow P \cup \{p_i\}$ 
// Step 2: Iterative generation over original and paraphrases
7 foreach  $x \in P$  do
8    $x' \leftarrow x$ 
9   for  $r \leftarrow 1$  to  $R$  do
10    // Prompt the fine-tuned model to generate a prediction
11     $a \leftarrow M(GeneratePrediction(x'))$ 
12    if  $a = none$  then
13      break
14     $C \leftarrow C \cup \{a\}$ 
15    // Remove the extracted terms from the sentence
16     $x' \leftarrow Remove(x', a)$ 
// Step 3: Paraphrase agreement and semantic filtering
17  $C' \leftarrow \emptyset$ 
18 foreach  $a \in C$  do
19    $sim \leftarrow SemanticSimilarity(a, s)$ 
20    $votes \leftarrow consensus(a, P)$  // Number of paraphrases yielding this prediction
21   if  $sim \geq \tau \wedge votes \geq min\_consensus$  then
22      $C' \leftarrow C' \cup \{a\}$ 
// Step 4: Knowledge-guided implicit aspect inference
23  $L \leftarrow IdentifyClues(s)$  // Search for cues in s
24 if  $L \neq \emptyset$  then
25    $I \leftarrow Algorithm\ A1(s, P, G)$ 
26    $C' \leftarrow C' \cup I$ 
// Step 5: Deduplication and normalization
27  $A \leftarrow NormalizeAndDeduplicate(C')$ 
28 return  $A$ 

```

4.1.2. Iterative Residual Generation

Iterative residual generation decomposes multi-aspect extraction into successive generation rounds. After each round, the extracted term is masked from the input, redirecting the model's attention to the remaining content. Each round targets the residual left by previous rounds, with extraction terminating when the model produces a null output.

This design resolves a structural limitation of beam search, which biases toward the most salient aspect regardless of beam width, as all beams are conditioned on the same unmodified input [73]. Masking extracted terms between rounds eliminates the dominant-aspect attractor at each step, producing a monotonic increase in recall that beam search cannot achieve.

4.1.3. Input Transformation via Paraphrasing

Several paraphrases are generated for every sentence, and each is processed through the same iterative generation procedure. The resulting predictions are combined and fil-

tered, whereby retained aspect terms must either appear verbatim in the original sentence or satisfy a set of validation constraints, including a minimum semantic similarity to the original sentence and sufficient agreement across paraphrases. These consistency checks allow the fine-tuned model to recover latent aspects while suppressing generative drift.

Paraphrases are generated using the GPT-4.1-mini model [74]. For each input sentence, ten paraphrases are requested using the following prompt:

```
``Act as an expert linguist in English. Give 10 different paraphrases to the following text in English, trying to explicitly state every noun described in the sentence. Only include the paraphrased sentence with no additional text before or after.``
```

Ten paraphrases are generated once per sentence to support the sensitivity analysis over the number of paraphrases (list item 4 of Section 5.2.1); the final reported configuration uses $k = 2$ paraphrases per sentence for both English and Arabic, the optimum identified by that analysis.

4.1.4. Knowledge Graph for Implicit Aspect Identification

While the preceding components target explicit aspect extraction, many opinionated sentences convey sentiment without naming an aspect term. In such cases, the fine-tuned generative model may correctly return `NULL`, indicating the absence of lexically grounded aspects. A `NULL` output, however, may correspond either to an objective sentence or to an implicit-aspect sentence where the aspect is expressed indirectly through evaluative clues; distinguishing between these cases requires reasoning beyond direct extraction.

Accordingly, the KG component operates as an independent implicit aspect identification path: for every input sentence, it scans for implicit aspect clues, searching for two categories of cue: *lonely adjectives*, sentiment adjectives that syntactically modify no explicit noun in the sentence; and *verb-phrase clues*, sentiment verbs that signal an evaluative stance without naming an aspect. When at least one such cue is detected, the component maps it to a domain-consistent aspect concept; when no cue is detected, it produces no implicit prediction for that sentence. This clue-driven gating ensures that symbolic inference is invoked only where implicit aspect expression is structurally signaled, independently of the generative component's output.

Lonely Adjective Detection

Lonely adjective detection operates on the dependency parse produced by spaCy (`en_core_web_sm`) for English. A candidate is any token with part-of-speech `ADJ` whose dependency relation is not purely structural (`nummod`, `det`, `advmod`, `compound`) and whose lemma is not in a surface blocklist of non-evaluative adjectives (positional, ordinal, and nationality terms). A candidate adjective a is classified as *lonely* according to its dependency configuration as follows.

(i) Predicative (copular) constructions (`acomp`): the nominal subject of the governing verb is inspected. If the subject is a noun or proper noun (e.g., “The food was excellent”), the aspect is explicit and a is *not* lonely. If the subject is an anaphoric pronoun referring to the reviewed entity (e.g., “it is expensive,” “they are great”), a is lonely; first-person subjects (e.g., “I am satisfied”) are excluded as speaker-referential. Conjoined adjectives inherit the classification of the predicate they attach to.

(ii) Attributive constructions (`amod`): a is lonely only when its head noun is a generic abstract noun carrying no aspect content (e.g., *place*, *thing*, *experience*); any other nominal head renders the aspect explicit.

(iii) **All remaining configurations:** *a* is rejected if its head or any of its dependents is a noun, proper noun, or numeral.

Verb-phrase clue detection complements this mechanism through a two-tier curated lexicon of sentiment-bearing verbs. For Arabic, candidate clues are obtained by matching CAMEL-derived lemmas against the KG clue vocabulary after the preprocessing steps described in Section 4.2.

Domain Inference

The sentence domain (Algorithm A1, line 2) is inferred using the same sentence-embedding model employed throughout the pipeline (SBERT paraphrase-multi-lingual-MiniLM-L12-v2). Each domain vertex is represented by the ℓ_2 -normalized centroid of the embeddings of a small curated domain vocabulary (approximately twenty characteristic terms per domain); the input sentence is embedded with the same encoder, and the domain with maximal cosine similarity to the sentence embedding is selected.

Knowledge Graph Structure

The KG represents domain knowledge using four vertex types: *domains*, *entities*, *aspects*, and *clues*. Clue vertices correspond to opinion-bearing expressions detected in text, while aspect vertices represent semantic aspect concepts, grouped under higher-level entity vertices and ultimately associated with domain vertices. The graph encodes three primary relationship types: *entity is_relevant_to domain*, *aspect belongs_to entity*, and *clue describes_aspect*. This layered structure enables multi-hop reasoning from surface-level clues to domain-specific aspects.

The full formal definition of the KG, relation composition, and algorithmic details are provided in Appendix A.

Knowledge Graph Construction

Both English and Arabic KGs share the same hierarchical structure, constructed by extracting opinion clues from review texts and aligning them to aspects, entities, and domains via training set labels. For English, lexical coverage was expanded with synonyms, antonyms, and related expressions across adjectives, verbs, nouns, and interjections that function as implicit opinion signals (e.g., *clean*, *spotless*, *filthy* → CLEANLINESS). For Arabic, clue nodes group multiple lexical realizations under unified concepts to accommodate nominalized and derivational forms (e.g., *قذارة*, *وسخ*, *قدرة* → CLEANLINESS), and evaluative constructions are captured directly (e.g., *لا يستحق السعر* → PRICES). Figure 2 shows a KG subgraph and Table 1 reports the scale of both KGs. For each dataset, the KG is constructed exclusively from the official training split: clue extraction, aspect–entity–domain alignment, and lexical enrichment all operate on training-split text and labels only. No test-split text or labels are used at any stage of KG construction, ensuring that no test-set information leaks into evaluation.

Table 1. Scale and coverage statistics for the English and Arabic knowledge graphs, reporting the number of vertex types and directed edges per relation type. The Arabic KG encodes all three relation types within a unified edge structure; see note (*).

	Coverage				Edges			Total
	Domains	Entities	Aspects	Clues	<i>Describes_Aspect</i>	<i>Belongs_to_Entity</i>	<i>Is_Relevant_to_Domain</i>	
English KG	7	85	136	1179	2630	301	71	3435
Arabic KG	8	89	3258	3255		5026 *		5026

* In the Arabic KG, clue-to-aspect, aspect-to-entity, and entity-to-domain mappings are encoded within a single unified edge structure rather than as separate edge types, yielding identical counts across all three relation columns.




Figure 2. An illustrative subgraph of the English knowledge graph, showing the four vertex types and their relationships. Domain vertices: green; entity vertices: yellow; aspect vertices: blue; clue vertices: pink. Edges represent the relations *is_relevant_to*, *belongs_to*, and *describes_aspect*, enabling multi-hop inference from surface clues to domain-specific aspects.

4.1.5. End-to-End Inference Procedure

Algorithm 1 presents the end-to-end inference procedure of the proposed solution. Given an input sentence, the pipeline first generates paraphrases to introduce syntactic variation while preserving meaning. Iterative residual generation is then applied jointly to the original sentence and its paraphrases, enabling the fine-tuned model to recover multiple explicit aspects missed under single-pass decoding. The resulting candidate pool is subsequently filtered through semantic similarity and paraphrase consensus constraints to discard spurious predictions. Concurrently, the KG component scans the input for opinion-bearing clues with no syntactic noun head, providing an independent implicit aspect identification path that maps each detected cue to a domain-consistent aspect. Finally, all candidates are normalized, deduplicated, and aggregated into a unified output set.

4.2. Arabic Adaptation

The Arabic adaptation addresses the challenges identified in Section 2.3 through four preprocessing steps applied uniformly across all components.

1. Clitic segmentation:

Aspect clues frequently appear inside affixed tokens (e.g., وبأسعار). CAMEL Tools [75] separates functional prefixes from content-bearing stems (أسعار + ب + → أسعار → lemma سعر), ensuring aspect cues are not obscured by affixation.

2. Orthographic normalization: Segmented tokens undergo hamzah unification (أ، آ، إ → إ), diacritic and tatweel removal, Unicode normalization, and feminine ending (ة) removal from adjectives, ensuring consistent token identity across all pipeline components.

3. Lexical clustering: Semantically related surface forms sharing a common referent (e.g., *مرتفع، غالي، أسعار*) are unified to a single aspect target (e.g., *PRICES*), reducing sparsity in knowledge graph lookups and evaluation matching.
4. Morphological disambiguation: CAMEL-derived lemmas are matched against knowledge graph clue nodes, ensuring that inflected variants across gender and number (*نظيفاً نظيفة نظيف*) map to the same canonical form, increasing clue coverage without manual enumeration of inflected forms.

Beyond preprocessing, the Arabic pipeline shares the English architecture but operates at a language-specific configuration determined by the sensitivity analysis (list item 4 of Section 5.2.2): the paraphrase module is a configurable component whose use is language-dependent, and the recommended Arabic configuration minimizes paraphrase variation, as Arabic morphological ambiguity causes paraphrasing to amplify lexical noise rather than provide semantic reinforcement. The pipeline is thus a single architecture with language-specific operating points rather than two distinct systems.

5. Experimental Study

5.1. Experimental Setup

5.1.1. Implementation Details

A decoder-only large language model, Llama-3-8B [76], was fine-tuned on task-specific datasets using LoRA [77], a parameter-efficient fine-tuning (PEFT) method in which the pre-trained model parameters remain frozen while low-rank adaptation layers are inserted and optimized during training.

Fine-tuning was performed on an NVIDIA A100 SXM4 GPU with 40 GB VRAM, accessed via Google Colaboratory with the High-RAM runtime enabled, which is required at LoRA rank $r = 128$.

The hyperparameter configuration reported in Table 2 follows that of Neveditsin et al. [64], who fine-tuned the same LLM backbone under an identical parameter-efficient fine-tuning setup and reported competitive results, providing a well-validated starting point for the experimental conditions of the present work.

Table 2. Fine -tuning hyperparameters for Llama-3-8B under LoRA-based supervised fine-tuning.

Component	Setting
LoRA rank (r)	128
LoRA alpha	32
LoRA dropout	0.1
Batch size (per device)	8
Gradient accumulation	1
Warmup steps	2
Max training steps	Dataset-dependent *
Learning rate	1×10^{-4}
Optimizer	<code>paged_adamw_32bit</code>
Eval strategy	<code>steps</code> , every 200 steps
Save strategy	<code>steps</code> , every 200 steps
Model selection	<code>load_best_model_at_end = True</code> , <code>metric_for_best_model = eval_loss</code>

* Corresponds to roughly 4–6 epochs depending on dataset size.

Table 3 lists the full software environment and NLP tools used across all experiments.

Table 3. Software environment and NLP tools used across all experiments, including runtime, fine-tuning, paraphrasing, and linguistic processing components.

Component	Version/Details
<i>Runtime</i>	
Operating system	Ubuntu 22.04 LTS
Python	3.12.13 (Colab default)
PyTorch	2.10.0 (CUDA 12.1 build)
CUDA toolkit	12.1 (Colab-managed)
NVIDIA driver	535.x (Colab-managed)
<i>Fine-tuning</i>	
Transformers	4.40.0 (HuggingFace)
PEFT	0.10.0 (HuggingFace)
TRL	0.8.6 (HuggingFace)
BitsAndBytes	0.43.1 (4-bit quantization)
Unsloth	2024.x (LoRA optimization)
Accelerate	0.29.3 (HuggingFace)
<i>Paraphrasing</i>	
OpenAI API	Chat Completions (GPT-4.1-mini)
<i>NLP Processing</i>	
spaCy	3.7.x (en_core_web_sm)
Stanza	1.11.0 (tokenize, pos, lemma, depparse)
CAMEL-Tools	1.5.2 (MLE disambiguator)
SentenceTransformers	paraphrase-multi-lingual-MiniLM-L12-v2

The implementation of the proposed solution, including training, inference, and evaluation scripts, is publicly available at <https://github.com/LujainAlawwad/ImplicitAspectExtraction> (accessed on 10 June 2026).

Computational Cost

Paraphrase generation submits a single API request per sentence, returning ten candidate paraphrases. Across all eight test sets (Tables 4 and 5), the pipeline processes 13,380 test sentences in total (8063 English and 5317 Arabic), yielding 13,380 API requests. At the measured rate of approximately \$0.47 per 1000 sentences (derived from 448,832 input tokens over the English evaluation), the projected total GPT-4.1-mini cost across all datasets is approximately \$6.36. Pipeline inference is dominated by the generative passes: under the final configuration ($k = 2$ paraphrases, up to $R = 3$ iterative rounds), each sentence requires at most $R \times (k + 1) = 9$ generation passes of the fine-tuned Llama-3-8B model, at approximately one second per pass on the A100 GPU. KG inference adds negligible overhead, requiring only a dependency parse and embedding similarity lookups per sentence.

5.1.2. Datasets

The proposed solution is evaluated on both English and Arabic ABSA benchmarks. For English, the evaluation uses the SemEval 2014, 2015, and 2016 shared-task datasets, ACOS, and the English portion of the Multi-lingual ABSA dataset (M-ABSA). The Arabic evaluation covers the SemEval 2016 Arabic Hotel reviews, the Human Annotated Arabic Dataset (HAAD) (<https://github.com/msmadi/HAAD> (accessed on 10 June 2026)) of book reviews, and the Arabic portion of M-ABSA; all datasets are publicly available. Download links for every benchmark are provided in the Data Availability Statement. Tables 4 and 5 summarize review-level statistics for all splits, distinguishing explicit, implicit, mixed, and objective sentences. As Table 5 shows, the HAAD test split contains no implicit-aspect annotations. HAAD is therefore evaluated for explicit extraction only (Metric A)

and is excluded from the KG-based implicit identification evaluation (Metric B) throughout.

Table 4. English ABSA datasets used in the experiments. All counts are at the review level. #Mix.: sentences containing both explicit and implicit aspects.

Dataset	Domain	Split	#Reviews	#Expl.	#Impl.	#Mix.	#Obj.
SE14	Rest.	Train	3041	2021	1020	0	0
	Rest.	Test	800	606	194	0	0
	Lapt.	Train	3045	1488 *	0	0	1557
	Lapt.	Test	800	422	378	0	0
SE15	Rest.	Train	1315	818	302	0	195
	Rest.	Test	684	401	180	0	103
SE16	Rest.	Train	2000	1210	498	0	292
	Rest.	Test	586	419	167	0	0
ACOS	Rest.	Train	1701	1225	476	0	0
	Rest.	Test	583	412	171	0	0
	Lapt.	Train	2485	1710	775	0	0
	Lapt.	Test	816	645	171	0	0
M-ABSA _{en}	Multi	Train	8796	5043	1864	323	1566
	Multi	Test	3794	2564	986	171	73

* 1488 subjective sentences; official labels do not annotate aspect categories.

Table 5. Arabic ABSA datasets used in the experiments. Column definitions follow Table 4.

Dataset	Domain	Split	#Reviews	#Expl.	#Impl.	#Mix.	#Obj.
SE16	Hotel	Train	4776	4445	331	0	0
	Hotel	Test	1225	1029	196	0	0
HAAD	Books	Train	1198	1187	11	0	0
	Books	Test	299	299	0	0	0
M-ABSA _{ar}	Multi	Train	8792	5005	1893	328	1566
	Multi	Test	3793	2504	1039	177	73

5.1.3. Evaluation Metrics

Explicit and implicit aspects require separate evaluation protocols and cannot be collapsed into a single overall score: explicit sentences carry a surface aspect term that the model must identify, whereas implicit sentences carry only an abstract Entity#Aspect category with no surface mention to extract. Two metrics are therefore reported throughout: Metric A for explicit aspect extraction and Metric B for implicit aspect identification via the KG module.

System performance is evaluated using micro-precision (P), micro-recall (R), and micro- F_1 score (F_1), the standard metrics for this task. These metrics are computed globally by aggregating true positives (TP), false positives (FP), and false negatives (FN) over all predictions:

$$P = \frac{TP}{TP + FP}, \quad R = \frac{TP}{TP + FN}, \quad F_1 = \frac{2PR}{P + R}. \quad (1)$$

Metric A: Explicit Aspect Extraction.

For explicit aspect-bearing sentences, predictions are compared against the surface-level aspect field of the gold annotations using micro- $P/R/F_1$. A predicted aspect term is counted as a true positive if it matches a gold aspect term, with a one-to-one matching constraint ensuring each prediction matches at most one gold term and vice versa.

Metric B: Implicit Aspect Identification.

The gold annotations of implicit sentences carry NULL in the aspect field but provide a complete Entity#Aspect pair in the category field (e.g., RESTAURANT#PRICES). Accordingly, KG module predictions are evaluated against the aspect component of the category field. Results are reported with *partial matching*: a prediction is counted as correct if the predicted aspect matches the gold aspect after normalization, accepting token-subset overlaps in addition to exact matches to accommodate lexical variation in compound expressions (e.g., *swimming pool* and *pool*). A strict bipartite matching constraint ensures each prediction is aligned to at most one gold aspect and vice versa.

5.2. Experiment Results

This section presents the experimental results across all English and Arabic datasets. The presentation begins with an ablation study analyzing the contribution of each system component, followed by performance evaluation, sensitivity analysis, error analysis, and comparative results against published state-of-the-art methods.

5.2.1. Proposed Solution Evaluation on English Datasets

1. Ablation Study

Table 6 reports micro- F_1 scores across five English benchmarks: SemEval-2014, SemEval-2015, SemEval-2016, ACOS, and M-ABSA. For ACOS, results are aggregated over both the Restaurant and Laptop subsets.

Table 6. Explicit aspect extraction performance (micro-precision, micro-recall, and micro- F_1) on English datasets. ZS: zero-shot; BL: fine-tuned baseline; BL + IG: baseline with iterative generation; BL + IG + Para: baseline with iterative generation and paraphrase-based input variation.

Method	SE14			SE15			SE16			ACOS			M-ABSA _{en}		
	p	R	F_1	p	R	F_1	p	R	F_1	p	R	F_1	p	R	F_1
ZS	0.3425	0.2322	0.2767	0.2663	0.2037	0.2308	0.3020	0.2160	0.2523	0.3323	0.3323	0.3323	0.2425	0.1701	0.2000
BL	0.8894	0.6030	0.7187	0.7298	0.6110	0.6651	0.8618	0.6491	0.7405	0.8270	0.8270	0.8270	0.7065	0.5990	0.6483
BL + IG	0.8906	0.7864	0.8353	0.7367	0.6776	0.7059	0.8249	0.7995	0.8120	0.7389	0.8541	0.7923	0.6609	0.6353	0.6478
BL + IG + Para	0.8957	0.8148	0.8533	0.7402	0.6878	0.7130	0.7839	0.7879	0.7859	0.7104	0.8969	0.7930	0.6878	0.6415	0.6639

Across all datasets, fine-tuning the base model substantially improves performance over the zero-shot (ZS) setting, confirming the necessity of task adaptation. The strongest absolute performance is observed on SemEval-2014 ($F_1 = 0.8533$), reflecting the relatively constrained and well-structured annotation scheme of this benchmark. Iterative generation (+IG) consistently improves recall over the baseline across all datasets, with the most pronounced gain on SE14 (F_1 : 0.7187 $\rightarrow$ 0.8353). The paraphrase-based component (+Para) achieves the best overall F_1 on all five benchmarks. Its contribution is largest on SE14 and ACOS, while its contribution is more modest on SE15, SE16, and M-ABSA_{en}. On M-ABSA_{en}, the multi-domain nature of the benchmark—spanning restaurant, hotel, laptop, phone, food, sight, and Coursera reviews—presents a broader vocabulary challenge; the system nevertheless achieves $F_1 = 0.6639$, with the largest per-domain gains observed in the hotel ($F_1 = 0.7695$) and laptop ($F_1 = 0.7692$) domains.

2. Performance Evaluation

(a) Explicit vs. Implicit Performance

Table 7 separates performance across explicit and non-explicit subsets; specifically, the non-explicit subset includes implicit, mixed, and objective rows. For implicit rows, the gold label is NULL; a correct NULL prediction counts as a true positive.

Table 7. Explicit extraction vs. null-abstention performance (micro- F_1). Non-explicit rows measure abstention reliability (correct NULL prediction against NULL gold) and do not constitute implicit aspect identification, which is evaluated separately via Metric B (Table 8). Columns are cumulative: BL: baseline; +IG: BL + IG; +Para: BL + IG + Para.

Dataset	Type	BL	+IG	+Para	Δ vs. BL
SE14	Explicit	0.6193	0.7916	0.8245	+0.2052
	Non-Explicit	0.9381	0.9231	0.8906	−0.0475
SE15	Explicit	0.5274	0.6009	0.6224	+0.0950
	Non-Explicit	0.9429	0.9347	0.9309	−0.0120
SE16	Explicit	0.6932	0.7963	0.7739	+0.0807
	Non-Explicit	0.8862	0.8683	0.8299	−0.0563
ACOS-Rest.	Explicit	0.8398	0.7722	0.7743	−0.0655
	Non-Explicit	0.8555	0.8300	0.8415	−0.0140
ACOS-Lap.	Explicit	0.8127	0.8037	0.7937	−0.0190
	Non-Explicit	0.8246	0.8580	0.8290	+0.0044
M-ABSA _{en}	Explicit	0.5718	0.5824	0.5993	+0.0275
	Non-Explicit	0.9392	0.9245	0.9332	−0.0060

Table 8. KG implicit aspect identification performance (partial matching). Any-TP@row: proportion of implicit rows with ≥ 1 correct prediction. SE14 excluded; see text.

Dataset	p	R	F_1	Any-TP@row
SE15	0.7119	0.3077	0.4297	0.3306
SE16	0.7021	0.3708	0.4853	0.3952
ACOS-Rest.	0.6842	0.3757	0.4851	0.3757
ACOS-Lap.	0.4479	0.2515	0.3221	0.2515
M-ABSA _{en}	0.5420	0.2821	0.3710	0.3095

On explicit rows, the pipeline improves consistently from BL to +Para across SemEval datasets, with the largest gain on SE14 ($\Delta F_1 = +0.205$) driven by iterative generation recovering multiple aspects. On ACOS, the baseline achieves the highest precision, while +IG and +Para expand recall at the cost of false positives, consistent with ACOS’s denser multi-aspect annotation. On M-ABSA_{en} explicit rows, all components contribute incrementally ($F_1 : 0.572 \rightarrow 0.582 \rightarrow 0.599$), though overall performance is lower, reflecting the challenge of generalizing across seven heterogeneous domains.

On non-explicit rows, all system components achieve high abstention scores ($F_1 = 0.83\text{--}0.94$). These scores quantify the system’s ability to refrain from generating spurious aspect terms when no surface aspect exists—a prerequisite for cleanly delegating implicit inference to the KG component—and do not measure implicit aspect identification, which is evaluated against ENTITY#ASPECT categories in list item 2 of Section 5.2.1. Counting correct NULL prediction on such rows follows the standard ACOS-style evaluation convention [56,78,79], in which implicit aspects are represented as NULL spans. A degenerate system that always returns NULL would achieve perfect abstention scores on non-explicit rows but $F_1 = 0$ on explicit rows; the abstention measure is therefore meaningful only jointly with the explicit performance reported in the same table. The relevant property is that the proposed system maintains high abstention while simultaneously achieving strong explicit extraction—a non-trivial combination for generative models, whose dominant failure mode in the error analysis is hallucinating aspect terms on aspect-free sentences.

(b) Implicit Aspect Identification via Knowledge Graph

The KG module is evaluated against the category gold annotation (ENTITY# ASPECT format). SemEval-2014 is excluded: its category field is absent in the laptop domain and uses only a generic label (*food, service, anecdotes/miscellaneous, ambience, price*) in the restaurant domain, providing insufficient granularity for meaningful implicit aspect evaluation.

The KG achieves high precision on restaurant and hotel domains ($p = 0.68\text{--}0.71$), confirming reliable predictions when inference is activated. Recall is moderate ($R = 0.31\text{--}0.37$), reflecting the coverage ceiling of lonely-adjective and verb-clue detection: roughly half of implicit rows are suppressed due to no detectable syntactic or lexical cue (Table 9). On ACOS-Laptop, performance drops substantially ($F_1 = 0.32$) as technical aspects such as *operation_performance* and *portability* require domain-specific clues underrepresented in the KG lexicon. On M-ABSA_{en}, the KG achieves $F_1 = 0.37$ overall, with strongest coverage on restaurant ($F_1 = 0.50$) and hotel ($F_1 = 0.41$) subsets and weaker coverage on sight ($F_1 = 0.12$) and Coursera ($F_1 = 0.31$), where implicit sentiment patterns are inadequately represented.

Table 9. KG inference step distribution across implicit rows (%).

Dataset	STOPPED	INFERRED	NO_MATCH
SE15	54.7	40.4	4.9
SE16	49.1	44.9	6.0
ACOS-Rest.	49.7	45.0	5.2
ACOS-Lap.	54.4	40.4	5.3
M-ABSA _{en}	48.7	41.8	9.5

The recall ceiling of the KG module merits direct discussion. The moderate recall reflects a deliberate precision-first design rather than an architectural deficiency: because KG predictions are merged into the final output without any downstream filtering stage, a permissive inference policy would inject false implicit aspects into every objective or cue-free sentence. The clue-driven gate is precisely what secures high precision when inference fires. The recall ceiling is therefore a lexicon- and cue-coverage problem rather than a structural one, and the evidence supports this view: extending the cue typology with verb-phrase clues—a purely additive change—recovered 23.7% of implicit rows on M-ABSA_{ar} and raised the inferred share on SemEval-2016 Arabic to 53.1% (list item 2 of Section 5.2.2), demonstrating that recall grows with coverage without sacrificing the precision guarantee. Despite its recall ceiling, the KG component remains statistically significant on every eligible dataset ($\Delta F_1 = +0.28$ to $+0.41$, $p < 0.001$; Appendix D) and is the most reliable implicit identification mechanism in the pipeline.

(c) Domain-level Performance

The Restaurant domain benefits most from iterative generation and paraphrasing across SemEval benchmarks (up to $+0.144$ on SE14-Restaurant). SE14-Laptop shows an inversion where +Para outperforms +IG ($0.875 > 0.813$) (Table 10), suggesting paraphrase-based variation is particularly effective for technical vocabulary. SE15-Hotel shows the smallest gain ($\Delta F_1 = +0.021$), consistent with lower KG recall ($R = 0.183$), indicating more restrained language with fewer detectable implicit cues. On ACOS, both domains show slight F_1 decreases from BL to +Para due to precision–recall trade-offs: the baseline predicts conservatively, while +IG and +Para expand coverage at the cost of false

positives. Among M-ABSA_{en} domains, Hotel and Laptop achieve the highest F_1 (0.770 and 0.769); Sight is the most challenging ($F_1 = 0.481$), as tourist attraction reviews use idiomatic and metaphorical language poorly represented in training data, yielding low KG coverage ($R = 0.073$ on implicit Sight rows).

Table 10. Per -domain explicit extraction performance (micro- F_1). Columns are cumulative: BL: baseline; +IG: BL + IG; +Para: BL + IG + Para.

Dataset	Domain	BL	+IG	+Para	Δ vs. BL
SE14	Restaurant	0.6917	0.8528	0.8357	+0.1440
SE14	Laptop	0.7500	0.8132	0.8754	+0.1254
SE15	Restaurant	0.6826	0.7260	0.7421	+0.0595
SE15	Hotel	0.6257	0.6599	0.6466	+0.0209
SE16	Restaurant	0.7405	0.8120	0.7859	+0.0454
ACOS	Restaurant	0.8444	0.7876	0.7919	−0.0525
ACOS	Laptop	0.8148	0.8128	0.7993	−0.0155
M-ABSA _{en}	Hotel	0.7402	0.7451	0.7695	+0.0293
M-ABSA _{en}	Laptop	0.7532	0.7418	0.7692	+0.0160
M-ABSA _{en}	Restaurant	0.7403	0.7412	0.7483	+0.0080
M-ABSA _{en}	Food	0.6451	0.6427	0.6728	+0.0277
M-ABSA _{en}	Phone	0.5668	0.5889	0.5852	+0.0184
M-ABSA _{en}	Coursera	0.6836	0.6724	0.6877	+0.0041
M-ABSA _{en}	Sight	0.4640	0.4556	0.4813	+0.0173

3. Error Analysis

- **Boundary Mismatch.**

Generative models produce aspect spans that are either subspans or superspans of the gold annotation. Table 11 shows representative cases. Boundary errors account for 8–37% of English false positives (highest on SemEval-2014: 37%, $n = 34$). These mismatches represent an annotation convention disagreement rather than a genuine inference failure.

Table 11. Representative boundary identification errors among English false positives, illustrating subspan and superspan mismatches against gold annotations.

Issue	Gold	Predicted
Subspan	product quality	quality
Subspan	Apple Help	Help
Superspan	shipping	two day shipping
Superspan	start up	quick start up

- **Spurious Hallucination.** Spurious predictions sharing no token overlap with any gold term dominate English false positives (84–93% across all English datasets). They are most prevalent in sentences with domain-specific terminology or multi-aspect structures absent from fine-tuning data, suggesting targeted augmentation or constrained decoding as remedies.
- **Taxonomy Mismatch.** The most frequent KG error across all restaurant and hotel datasets is a fine-grained taxonomy mismatch between the KG’s internal aspect vocabulary and the gold annotation scheme. The KG predicts *taste* for sentences expressing food quality, while the gold annotation uses *quality*. For example, the sentence “Mmmmmmmmm so delicious” is predicted as food#taste but annotated as food#quality.

- **KG Coverage: No-Cue Detection (STOPPED rows).** Approximately 49–55% of implicit rows across all datasets are suppressed (STOPPED_no_lonely_adj) because no lonely adjective or verb clue is detected. Examination of these sentences reveals three recurring patterns. First, implicit sentiment expressed through noun phrases alone, with no adjective present (e.g., “Two thumbs up!”, gold: food#quality). Second, figurative expressions carrying no surface-level sentiment marker detectable by the KG (e.g., “You’ll be there for every anniversary...”, gold: restaurant#miscellaneous). Third, sentences in which the implicit aspect is carried entirely by a verb with no accompanying adjective (e.g., “We were charged full price”, gold: service#general). All three cases require discourse-level or pragmatic reasoning beyond lexical KG inference.

4. Sensitivity Analysis

- Iterative Generation Rounds** (Figure 3a): Increasing rounds from 2 to 3 yields the largest F_1 gain (+9.4% recall, −2.3% precision); beyond 3 rounds improvements diminish. The maximum number of iterative generation rounds is set to 3 for English.
- Paraphrase-Aided Parameters** (Figure 3b,c):




Figure 3. Sensitivity analysis: (a) iterative generation rounds; (b) number of paraphrases; (c) semantic similarity threshold τ .

- Number of paraphrases:** Two paraphrases yield the best F_1 ; beyond two, marginal recall gains are offset by precision loss.
- Similarity threshold τ :** Precision improves smoothly with τ ; recall remains nearly constant; F_1 saturates at $\tau \geq 0.70$.
- Candidate retention, consensus, and explicitness:** All three have negligible impact once semantic filtering is active. Full results are reported in Appendix C.1.

5. Statistical Significance Tests

Paired bootstrap resampling [80] ($n = 10,000$ iterations) and McNemar’s test [81] were applied to assess whether observed F_1 differences between pipeline systems reflect genuine improvements rather than sampling variation. Three pairwise comparisons were evaluated: BL vs. BL + IG, BL vs. BL + IG + PARA, and BL + IG vs. BL + IG + PARA. Each comparison was tested across three evaluation views: all rows, aspect-extraction only (rows with at least one non-null gold aspect), and explicit rows only. The KG component was evaluated separately, comparing BL against KG on implicit rows only, using the ASPECT part of the entity#aspect category label as gold. Full numerical results are reported in Appendix D.

BL $\rightarrow$ BL + IG and BL $\rightarrow$ BL + IG + PARA are bootstrap-significant ($p < 0.001$) across all views on SE14, SE15, and SE16, confirming that both IG and the full pipeline deliver genuine gains over the baseline. The marginal step BL + IG $\rightarrow$ BL + IG + PARA is

additionally significant on SE14 ($p < 0.001$, explicit rows), making it the only dataset where every pairwise comparison is bootstrap-confirmed; on SE15, it does not reach significance, indicating partial redundancy between IG and PARA on that dataset. On ACOS, BL is the best system and all additions degrade aggregate performance; bootstrap p -values near 1.0 confirm the degradation is not due to chance, while a significant McNemar result indicates the components shift which sentences are correctly predicted without improving the overall score. The structural properties of ACOS explain why the pipeline additions are unhelpful on this benchmark. ACOS carries the densest multi-aspect annotation of all evaluated corpora, and the fine-tuned baseline already attains recall of 0.83–0.85 together with the highest precision: little residual signal remains for iterative masking to recover, so successive passes surface mostly spurious candidates, and the paraphrase stage compounds this effect by widening the candidate pool. For densely annotated corpora of this kind, the baseline alone is therefore the recommended configuration.

On M-ABSA_{en}, BL + IG alone does not improve significantly over BL by bootstrap, whereas BL → BL + IG + PARA and BL + IG → BL + IG + PARA are both significant ($p \leq 0.003$), suggesting that the paraphrase component drives the improvement and IG is insufficient on its own. The KG component is bootstrap-significant on all four eligible English datasets ($\Delta F_1 = +0.28$ to $+0.37$, $p < 0.001$; Table A8).

5.2.2. Proposed Solution Evaluation on Arabic Datasets

1. Ablation Study

Table 12 reports micro- F_1 on explicit rows across the three Arabic datasets. On M-ABSA_{ar}, iterative generation delivers the largest gain, driven by improved recall as the iterative component recovers multiple aspect mentions missed in single-pass decoding. The paraphrase component moderates precision, with the best result at $F_1 = 0.760$. Notably, +IG alone outperforms the full pipeline, indicating that Arabic paraphrase variants introduce more spurious candidates than in English—a consequence of greater morphological ambiguity reducing paraphrase precision. On HAAD, the pattern is consistent: +IG improves recall most substantially ($R: 0.438 \rightarrow 0.646$), while +Para recovers precision ($p: 0.371 \rightarrow 0.410$), yielding the best $F_1 = 0.477$. On SE16, the baseline achieves the highest precision ($P = 0.795$, $F_1 = 0.763$). Iterative generation expands recall substantially ($R: 0.734 \rightarrow 0.825$) at a notable precision cost, and +Para partially recovers precision while maintaining high recall ($F_1 = 0.734$). The SE16 hotel reviews average 2.05 gold aspects per explicit row, explaining why iterative generation is particularly effective at recovering additional aspects despite increased false positives.

2. Performance Evaluation

(a) Explicit vs. Implicit Performance

Table 13 shows per-type performance on M-ABSA_{ar} under BL + IG + Para. The pipeline achieves $F_1 = 0.760$ on explicit rows and $F_1 = 0.671$ on mixed rows. The implicit-row abstention score ($F_1 = 0.933$) reflects the system's reliable identification of sentences with no surface aspect, consistent with the English results ($F_1 = 0.83$ – 0.94); implicit aspect identification itself is evaluated in list item 2 of Section 5.2.2. Mixed-row performance is lower than explicit-only rows because mixed sentences require the system to simultaneously extract a surface aspect and correctly output NULL for the implicit component—the latter is addressed separately by the KG module.

Table 12. Explicit aspect extraction performance on Arabic datasets. ZS: zero-shot; BL: fine-tuned baseline; BL + IG: baseline with iterative generation; BL + IG + Para: baseline with iterative generation and paraphrase-based input variation.

Method	M-ABSA _{ar}			SE16			HAAD		
	<i>p</i>	<i>R</i>	<i>F₁</i>	<i>p</i>	<i>R</i>	<i>F₁</i>	<i>p</i>	<i>R</i>	<i>F₁</i>
ZS	0.0482	0.0476	0.0479	0.0702	0.0363	0.0478	0.1818	0.0952	0.125
BL	0.7542	0.5907	0.6625	0.7950	0.7339	0.7632	0.4708	0.4378	0.4537
BL + IG	0.8094	0.7264	0.7656	0.6383	0.8249	0.7197	0.3712	0.6462	0.4716
BL + IG + Para	0.7768	0.7441	0.7601	0.6630	0.8210	0.7336	0.4098	0.5692	0.4765

Table 13. Explicit extraction vs. null-abstention performance on Arabic datasets (micro-*F₁*, BL + IG + Para). Implicit rows have NONE gold aspect and measure abstention reliability; they do not constitute implicit aspect identification, which is evaluated separately via Metric B (Table 14). Mixed rows (M-ABSA_{ar} only) contain one explicit and one implicit aspect per sentence. Columns are cumulative: BL: baseline; +IG: BL + IG; +Para: BL + IG + Para.

Dataset	Type	BL	+IG	+Para	Δ vs. BL
M-ABSA _{ar}	Explicit (1356)	0.7542	0.8094	0.7768	+0.0226
	Implicit (718)	0.9392	0.9245	0.9332	−0.0060
	Mixed (104)	0.7586	0.8130	0.6667	−0.0919
SE16	Explicit (1130)	0.7950	0.6383	0.6630	−0.1320
	Implicit (96)	0.8862	0.8683	0.8299	−0.0563
HAAD	Explicit (299)	0.4708	0.3712	0.4098	−0.0610

Table 14. Strategy 1 (cross-lingual English KG) vs. Strategy 2 (Arabic domain-specific KG) for implicit aspect identification.

Dataset	Strategy	<i>p</i>	<i>R</i>	<i>F₁</i>
M-ABSA _{ar}	S1	0.4841	0.0951	0.1590
	S2	0.6060	0.2541	0.3580
SE16	S1	0.4545	0.1429	0.2174
	S2	0.6981	0.3524	0.4684

(b) Implicit Aspect Identification via Knowledge Graph

The Arabic KG is evaluated against ENTITY#ASPECT category annotations independently of explicit extraction. Two strategies are compared: Strategy 1 uses the English KG with multi-lingual embeddings; Strategy 2 uses direct lexical lookup on the Arabic KG. HAAD is excluded as its test split contains no implicit aspects.

Table 14 shows Strategy 2 outperforms Strategy 1 by $\Delta F_1 = 0.211$ on M-ABSA_{ar} and $\Delta F_1 = 0.193$ on SE16, gaining on both precision and recall. The improvement stems from strict domain filtering reducing cross-domain noise, and the Arabic clue lexicon with verb-phrase fallback recovering more implicit rows than cross-lingual embedding similarity alone.

The Arabic KG achieves high precision when inference is activated ($p = 0.61$ – 0.70), comparable to the English KG ($p = 0.68$ – 0.71). Recall is lower ($R = 0.25$ – 0.35) due to high STOPPED rates (Table 15): 52.1% of M-ABSA_{ar} and 42.7% of SE16 implicit rows produce no prediction, reflecting Arabic morphological complexity and the absence of detectable adjective or sentiment verb cues. The verb-phrase fallback contributes substantially: on M-ABSA_{ar}, 23.7% of implicit rows are recovered via sentiment verbs such as أعجبنى وأوصي; on SE16, the INFERRED share reaches 53.1%, reflecting the prevalence of recommendation expressions in hotel reviews.

Table 15. Arabic KG inference step distribution across implicit rows (%).

Dataset	STOPPED	INFERRED	NO_MATCH
M-ABSA _{ar}	52.1	45.1	2.8
SE16	42.7	53.1	4.2

(c) Domain-level Performance (M-ABSA_{ar})

The Hotel domain achieves the highest $F_1 = 0.467$ with $p = 0.778$ (Table 16), reflecting strong alignment between the hotel-focused SemEval-2016 training edges in the Arabic KG and the hotel domain reviews in M-ABSA_{ar}. Restaurant and Food domains perform moderately ($F_1 = 0.391$ and 0.306 , respectively), while Laptop is the weakest ($F_1 = 0.281$), consistent with the English results: laptop implicit aspects require highly specific technical vocabulary—*operation_performance*, *design_features*, *portability*—that is underrepresented in the Arabic KG’s adjective lexicon, which was primarily enriched from hotel and restaurant review corpora.

Table 16. Per -domain KG implicit identification performance on M-ABSA_{ar} implicit rows (Strategy 2).

Domain	p	R	F_1
Hotel	0.7778	0.3333	0.4667
Restaurant	0.6338	0.2830	0.3913
Food	0.5333	0.2149	0.3064
Laptop	0.4792	0.1983	0.2805

(d) KG Performance on Mixed Sentences

Mixed sentences contain both an explicit and an implicit aspect in the same sentence, requiring the solution to handle both tasks simultaneously (Table 17). Explicit extraction achieves $F_1 = 0.678$, a modest decline from $F_1 = 0.777$ on purely explicit rows, suggesting that competing opinion signals in mixed sentences introduce some additional difficulty for aspect boundary identification. Implicit identification achieves $F_1 = 0.367$, closely matching the $F_1 = 0.370$ obtained on purely implicit rows, confirming that the lonely adjective detection mechanism operates independently of whether an explicit aspect is also present. Mixed sentences are therefore a moderately harder case for explicit extraction but not for implicit identification.

Table 17. Evaluation on M-ABSA_{ar} mixed sentences (104 rows). Metric A: full system explicit extraction; Metric B: KG implicit identification (Strategy 2, partial matching).

Metric	Method	p	R	F_1
A (Explicit)	BL + IG + Para + KG	0.6736	0.6831	0.6783
B (Implicit)	S2 KG	0.6071	0.2629	0.3669

3. Error Analysis

- **Boundary Mismatch.**

Boundary errors account for 2–8% of Arabic false positives, where Arabic morphological structure causes the model to either include genitive complements (superspan) or omit clitic-bearing head nouns (subspan).

- **Arabic Annotation Inconsistency.** The Arabic datasets exhibit uneven granularity: fine-grained aspects (e.g., PRICES, CLEANLINESS) coexist with coarse GENERAL annotations that collapse multiple implicit aspects. Table 18 illustrates

cases where the solution correctly identifies specific aspects but is penalized against the coarser gold label. Generic opinion adjectives with no specific aspect noun constitute the largest false positive category on M-ABSA_{ar} (47%, $n = 277$), SemEval-2016 Arabic (24%, $n = 302$), and HAAD (22%, $n = 108$). Nevertheless, GENERAL-label noise remains non-trivial because the boundary between general sentiment and fine-grained aspects is inconsistently drawn in the annotations themselves.

- **KG Coverage: No-Cue Detection (STOPPED rows).** STOPPED rates of 52.1% (M-ABSA_{ar}) and 42.7% (SE16) arise from three structural categories: (i) purely factual or nominal sentences with no sentiment marker (e.g., طلبت ثلاث حزم, gold: food#general), where the aspect is inferred from purchase behavior rather than linguistic cues; (ii) exclamatory expressions (e.g., أووه!) conveying sentiment through pragmatic cues alone; and (iii) comparative constructions (e.g., هذا المكان هو أفضل شيء بعد طبخ أبي), where the comparative adjective modifies a noun phrase and is correctly blocked by the lonely-adjective gate. All three require discourse-level or pragmatic reasoning beyond lexical KG inference.

Table 19 summarizes errors across all English and Arabic datasets. English false positives are dominated by spurious hallucination (84–93%), with boundary errors secondary in the laptop domain. Arabic false positives are more varied: GENERAL-label noise dominates on M-ABSA_{ar}, clitic duplication on HAAD, and spurious hallucination on SE16-Arabic. False negatives are consistent across datasets: 76–85% were never predicted by any stage; 14–24% were correctly identified by the baseline but dropped by later voting stages.

Table 18. Annotation inconsistency in the Arabic SemEval 2016 dataset: the pipeline identifies specific aspects that are collapsed into GENERAL in the gold annotation.

Example	Prediction	Gold Aspect
استقبال سيء ونوادي مزينة	COMFORT, SERVICE	GENERAL
مطلب كل مسافر الراحة وعدم الوفاء بالوعود دون المستوى في التعامل والإستجمام	COMFORT, CLEANLINESS	GENERAL
مزعج جداً للنوم	COMFORT	GENERAL

Table 19. False positive and false negative breakdown across all datasets (explicit rows; full test sets). FP categories: Spur. = spurious hallucination; Bound. = boundary error; Gen. = GENERAL-label noise; Clit. = clitic duplication (Arabic only). FN categories: Never = never predicted; Lost = predicted by baseline but filtered later.

Dataset	n	TP	FP	FN	p	R	FP Breakdown (%)				FN Breakdown (%)	
							Spur.	Bound.	Gen.	Clit.	Never	Lost
SE14 (EN)	1028	1389	91	391	0.939	0.780	63	37	0	—	85	14
SE15 (EN)	605	474	116	324	0.803	0.594	93	7	0	—	85	14
SE16 (EN)	419	478	131	133	0.785	0.782	89	11	0	—	79	19
ACOS Rest. (EN)	412	378	200	34	0.654	0.917	93	7	0	—	76	24
ACOS Lap. (EN)	801	727	267	74	0.731	0.908	84	16	0	—	77	14
MABSA (EN)	2576	2307	951	1654	0.708	0.582	91	8	1	—	78	15
MABSA (AR)	1356	1280	594	439	0.683	0.745	31	8	47	15	78	18
SE16 (AR)	1130	1912	1248	411	0.605	0.823	45	2	24	29	77	18
HAAD (AR)	299	326	487	245	0.401	0.571	36	2	22	40	79	16

4. Sensitivity Analysis

- (a) **Iterative Generation Rounds** (Figure 4a): Unlike English, increasing rounds consistently degrades micro- F_1 in Arabic. Additional rounds surface noisy candidates, causing a sharp precision drop that outweighs recall gains.




Figure 4. Sensitivity analysis of the proposed model on Arabic datasets. (a) Effect of iterative generation rounds. (b) Effect of number of paraphrases. (c) Effect of semantic similarity threshold.

(b) Paraphrase-Aided Parameters:

- i. **Number of paraphrases** (Figure 4b): Adding paraphrases increases recall but causes a monotonic precision and micro- F_1 decline. Unlike English, the optimal Arabic configuration uses no or very few paraphrases, as paraphrasing amplifies lexical noise rather than providing semantic reinforcement.
- ii. **Similarity threshold** (Figure 4c): Raising the threshold improves precision while recall decreases gradually, yielding a monotonic micro- F_1 gain. Conservative semantic filtering is therefore essential for suppressing paraphrase-induced noise in Arabic.

- (c) **Statistical Significance Tests (Arabic).** On M-ABSA_{ar}, BL $\rightarrow$ BL + IG and BL $\rightarrow$ BL + IG + PARA are highly significant across all views ($p < 0.001$), confirming that both components contribute meaningful gains; however, BL+IG $\rightarrow$ BL + IG + PARA is not significant by either test, indicating PARA provides no additional benefit once IG is applied. On SE16_{ar}, both BL $\rightarrow$ BL + IG and BL $\rightarrow$ BL + IG + PARA are associated with statistically significant *degradation* (McNemar $p < 0.001$, bootstrap $p \approx 1.0$), reflecting a precision collapse under generative augmentation in morphologically rich Arabic; BL remains the best system. The only positive comparison on SE16_{ar}, BL + IG $\rightarrow$ BL + IG + PARA ($\Delta F_1 = +0.019$, $p < 0.001$), shows that PARA partially recovers from the degradation introduced by IG but not sufficiently to surpass BL. On HAAD, no comparison reaches bootstrap significance ($p = 0.05\text{--}0.10$), attributable to the small evaluation set ($n = 299$); a significant McNemar result nonetheless confirms that the components alter error patterns. Despite the inconsistent pipeline behavior across Arabic datasets, the KG component is bootstrap-significant on both datasets ($\Delta F_1 = +0.37$ and $+0.41$, $p < 0.001$; Table A10), establishing it as the primary and most reliable mechanism for implicit aspect identification in Arabic.

Taken together, the Arabic results should be interpreted at the component level rather than as a monolithic pipeline, with the sensitivity analysis above determining the appropriate per-dataset configuration. On M-ABSA_{ar}, iterative generation drives large

and statistically significant gains (F_1 : 0.663 $\rightarrow$ 0.766, $p < 0.001$), consistent with its multi-aspect sentence structure. On SemEval-2016 Arabic, whose hotel reviews average 2.05 gold aspects per explicit row, the same mechanism trades precision for recall and the fine-tuned baseline remains the preferred configuration ($F_1 = 0.763$); morphological ambiguity further amplifies paraphrase-induced noise, motivating the conservative settings identified in the sensitivity analysis. On HAAD, the improvement under the full configuration (F_1 : 0.454 $\rightarrow$ 0.477) does not reach bootstrap significance, which is attributed to the small evaluation set ($n = 299$)—McNemar’s test nonetheless confirms that the components systematically alter the error pattern—together with HAAD’s annotation characteristics, where clitic-level boundary disagreements and coarse GENERAL labels constitute the majority of false positives. Across all Arabic datasets, the knowledge-graph component, which carries the implicit-extraction contribution at the core of this work, is uniformly and significantly positive. The full pipeline therefore does not uniformly dominate the baseline on Arabic, and no such claim is made; rather, each component contributes positively under the configuration regime the sensitivity analysis prescribes, and the statistical tests make the boundaries of those regimes explicit.

5.3. Comparative Analysis

The comparative analysis validates the core design objective of the proposed system: robust aspect identification encompassing both explicit and implicit expressions across multiple domains and benchmarks.

A scope distinction must precede any reading of Tables 20 and 21. The Target column separates systems solving two different problems: explicit-only systems are neither required to abstain on implicit-bearing sentences nor penalized for them and can specialize decoding entirely toward surface spans. Systems targeting both aspect types—including the proposed system—accept a structural recall–precision trade-off on the same test sets, since every implicit-bearing sentence is an opportunity for a false positive that explicit-only systems never face. F_1 differences across scopes therefore do not directly measure relative quality; this is the appropriate lens for SemEval-2015 and SemEval-2016, where explicit-only systems report higher F_1 than the proposed solution.

1. English Datasets: Comparison Results

Table 20 validates the proposed solution across both explicit and implicit aspect identification. On SemEval-2014, the solution outperforms [14,82] on Restaurant and surpasses [44,83] on Laptop. The closest competitor targeting both aspect types, ref. [64], reports results slightly below those of the proposed system on both domains. On SemEval-2015, the solution falls short of [83] (explicit only) and [63], reflecting the additional challenge of targeting implicit aspects in a smaller dataset. On SemEval-2016, ref. [62] achieves the highest F_1 ; the gap is partly attributable to differences in evaluation scope and annotation coverage across studies. On ACOS, the solution substantially exceeds [44] in recall (explicit only), demonstrating the effectiveness of iterative and paraphrase-based expansion for increasing aspect coverage. The lower precision relative to explicit-only systems reflects the inherent trade-off of targeting both aspect types.

Two qualifications are necessary when interpreting the comparisons in Table 20. First, the $F_1 = 0.91$ reported by [62] on SemEval-2016 is not measured on the official test set: the authors partition their prepared data into a custom 80/20 train/test split, which also permits sentences from the same review to appear on both sides of the split, and their training injects a manually constructed implicit-aspect seed index. Their score is therefore not directly comparable to results computed on the official

held-out test set. Second, ref. [63] (WoSe) is a discriminative BIO sequence tagger: it can only label tokens present in the text, which by construction precludes both span hallucination—the dominant false-positive source for generative systems (Section 5.2.1, Error Analysis)—and any prediction for sentences whose aspect is not lexically present. Moreover, ref. [63] defines implicit aspects as in-text aspect terms whose category is implied, whereas the present task includes inferring aspects for sentences containing no aspect term at all. Direct F_1 comparison on SemEval-2015 and SemEval-2016 should be read with these scope and protocol differences in mind.

Table 20. Performance comparison with state-of-the-art methods on English benchmark datasets. Target: Explicit = explicit aspects only; Both = explicit and implicit aspects. —: not reported by the study.

Study	Dataset	Domain	Target	p	R	F_1
SemEval-2014						
[44]	SE-14	Lap.	Explicit	—	—	0.8635
[83]	SE-14	Rest.	Explicit	0.8481	0.8619	0.8693
[83]	SE-14	Lap.	Explicit	—	—	0.8299
[82]	SE-14	Rest.	Explicit	0.7987	0.6603	0.7230
[64]	SE-14	Lap.	Both	0.91	0.72	0.81
[64]	SE-14	Rest.	Both	0.88	0.77	0.82
[15]	SE-14	Combined	Explicit	—	—	0.8200
[14]	SE-14	Rest.	Explicit	—	—	0.7596
[14]	SE-14	Lap.	Explicit	—	—	0.5174
Ours	SE-14	Rest.	Both	0.8909	0.7869	0.8357
Ours	SE-14	Lap.	Both	0.9014	0.8508	0.8754
SemEval-2015						
[83]	SE-15	Rest.	Explicit	—	—	0.7551
[63]	SE-15	Rest.	Both	—	—	0.7286
Ours	SE-15	Rest.	Both	0.7402	0.6878	0.7130
SemEval-2016						
[44]	SE-16	Rest.	Explicit	—	—	0.8163
[62]	SE-16	Rest.	Both	0.912	0.8953	0.91
[63]	SE-16	Rest.	Both	—	—	0.8113
Ours	SE-16	Rest.	Both	0.7839	0.7879	0.7859
ACOS						
[44]	ACOS	Rest.	Explicit	0.8054	0.8208	0.8163
[44]	ACOS	Lap.	Explicit	0.8635	0.8635	0.8635
Ours	ACOS	Rest.	Both	0.7085	0.8974	0.7919
Ours	ACOS	Lap.	Both	0.7193	0.8992	0.7993

2. Arabic Datasets: Comparison Results

Table 21 compares the proposed solution against state-of-the-art Arabic methods. On SE16, explicit-only systems such as [33] report higher F_1 but are trained and evaluated on explicit aspects only, with no requirement to handle implicit sentences. The proposed system achieves higher recall, confirming that iterative generation effectively surfaces multiple aspects per sentence, at the cost of additional false positives. Against category-detection systems [34,47], the system is competitive while additionally identifying implicit aspects. On HAAD, the system exceeds the LLM-based category-detection approach of [16]. Strong recall ($R = 0.569$) reflects the robustness of iterative generation, while the lower precision ($p = 0.410$) stems from paraphrase-induced noise and clitic-level boundary errors, which together dominate the false-positive profile on this dataset (Table 19).

Table 21. Performance comparison with state-of-the-art methods on Arabic benchmark datasets. Target: Explicit = explicit aspects only; Both = explicit and implicit aspects. Subtask: ATE = Aspect Term Extraction; ACD = Aspect Category Detection. —: not reported.

Study	Dataset	Target	Subtask	p	R	F_1
SemEval-2016 Arabic						
[33]	SE-16	Explicit	ATE	0.8096	0.7970	0.8032
[34]	SE-16	Explicit	ACD	—	—	0.7100
[47]	SE-16	Explicit	ACD	—	—	0.6610
Ours	SE-16	Both	ATE	0.6630	0.8210	0.7336
HAAD						
[16]	HAAD	Explicit	ACD	—	—	0.3200
[18]	HAAD	Explicit	ACD	—	—	0.4830
[46]	HAAD	Explicit	ATE	—	—	0.3447
[47]	HAAD	Explicit	ATE	—	—	0.661
Ours	HAAD	Both	ATE	0.4098	0.5692	0.4765

5.4. Limitations

The error analysis (Section 5.2.1 item 3 and Section 5.2.2 item 3) reveals four system-level limitations that transcend individual datasets.

1. Pragmatic implicit aspects. The lonely adjective gate and verb-clue fallback address surface-level implicit signals but cannot handle sentiment expressed through purchase behavior, exclamatory fragments, or figurative language. Resolving these cases requires discourse-level or pragmatic reasoning beyond the scope of lexical KG inference.
2. Boundary identification. Generative decoding produces subspans and superspans of gold aspect terms, particularly in morphologically complex Arabic constructions. Constrained decoding or span-extraction post-processing would mitigate this failure mode.
3. Arabic annotation noise. The inconsistent granularity of Arabic gold annotations—where both fine-grained and coarse GENERAL labels appear for structurally similar sentences—introduces irreducible noise that system design alone cannot resolve. Standardizing Arabic ABSA annotation guidelines is a prerequisite for progress on this front.
4. Domain vocabulary coverage and KG recall. The KG’s clue lexicon was built primarily from hotel, restaurant, and food corpora; the Laptop, Coursera, and Sight domains are consequently underrepresented, leading to systematic false negatives on technical and non-restaurant implicit aspects. This coverage limitation is the principal driver of the KG module’s moderate recall: roughly half of implicit rows carry no clue the current lexicon can detect. Because the recall ceiling is a coverage problem rather than an architectural one, three additive mitigations follow directly: (i) domain-adaptive lexicon expansion from larger Arabic and English corpora; (ii) additional cue classes beyond lonely adjectives and sentiment verbs, in particular noun-phrase and exclamatory cues identified in the error analysis; and (iii) embedding-based clue matching as a back-off when exact lexical lookup fails. The verb-phrase fallback already added to the Arabic pipeline demonstrates the viability of this path, recovering 23.7% of implicit rows on M-ABSA_{ar} without degrading precision.

6. Conclusions

This work presented a multi-component pipeline for implicit aspect extraction, combining a fine-tuned generative backbone, iterative residual generation, paraphrase-based

input transformation, and a domain-specific knowledge graph governed by linguistically motivated gating. The central finding is that no single component is sufficient: the KG component provides gains precisely where the generative backbone has no surface term to extract—sentences with isolated opinion adjectives and no syntactic noun head—while iterative generation recovers multi-aspect sentences that single-pass decoding misses. Ablation results confirm that this complementarity holds across both language families.

Experiments across eight datasets demonstrate competitive performance, surpassing established baselines and recent state-of-the-art methods on six of eight benchmarks. On Arabic, the gains are configuration-dependent: iterative generation drives substantial and statistically significant improvements on M-ABSA_{ar}, the fine-tuned baseline remains the strongest configuration on SemEval-2016 Arabic, and the knowledge-graph component is the most consistent contributor, significant on all eligible datasets in both languages. On Arabic, primary evidence for implicit aspect identification comes from SemEval-2016 Hotels and M-ABSA_{ar}; HAAD is evaluated for explicit extraction only, as its annotation scheme does not mark implicit aspects. The English–Arabic performance gap reflects greater annotation consistency, stronger boundary alignment, and lower morphological complexity in English.

A consistent finding across both languages is that strict span-based evaluation penalizes semantically correct predictions that differ lexically or in boundary extent from gold annotations, systematically underestimating generative model performance. Semantic equivalence metrics are a prerequisite for fairly benchmarking generative ABSA systems, and their development constitutes a natural next step for the field.

Two directions follow directly from these results. First, extending the Arabic KG to improve morphological coverage across dialects and derivational forms, increasing robustness beyond the MSA-dominant training corpora used here. Second, investigating cross-lingual transfer to bootstrap implicit aspect extraction in lower-resource Arabic domains where annotated data remain scarce.

Author Contributions: Conceptualization, L.A.A. and M.E.B.M.; methodology, L.A.A.; software, L.A.A.; validation, L.A.A. and M.E.B.M.; formal analysis, L.A.A.; investigation, L.A.A.; resources, L.A.A. and M.E.B.M.; data curation, L.A.A.; writing—original draft preparation, L.A.A.; writing—review and editing, L.A.A. and M.E.B.M.; visualization, L.A.A.; supervision, M.E.B.M.; project administration, M.E.B.M. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable. This study does not involve human participants, human data, or animal subjects. All data used are publicly available benchmark datasets for natural language processing research.

Informed Consent Statement: Not applicable. This study uses only publicly available, pre-existing datasets and does not involve direct interaction with human subjects.

Data Availability Statement: The datasets analyzed in this study are all publicly available. SemEval 2014 Task 4 data are available at <https://alt.qcri.org/semeval2014/task4/>, accessed on 10 June 2026. SemEval 2015 Task 12 data are available at <https://alt.qcri.org/semeval2015/task12/>, accessed on 10 June 2026. SemEval 2016 Task 5 data are available at <https://alt.qcri.org/semeval2016/task5/>, accessed on 10 June 2026. The ACOS dataset is available at <https://github.com/NUSTM/ACOS>, accessed on 10 June 2026. The M-ABSA dataset is available at <https://github.com/swaggy66/M-ABSA>, accessed on 10 June 2026. The HAAD dataset is available at <https://github.com/msmadi/HAAD>, accessed on 10 June 2026. The implementation code for the proposed solution, including training, inference, and evaluation scripts, is publicly available at <https://github.com/LujainAlawwad/ImplicitAspectExtraction>, accessed on 10 June 2026.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

ABSA	Aspect-Based Sentiment Analysis
ATE	Aspect Term Extraction
BL	Baseline component
IG	Iterative Generation component
PARA	Paraphrase-based input variation component
KG	Knowledge Graph
LLM	Large Language Model
LoRA	Low-Rank Adaptation
MSA	Modern Standard Arabic
NLP	Natural Language Processing
SA	Sentiment Analysis

Appendix A. Formal Notation

- **Graph Definition.** Let $G = (V, E)$ be a labeled multi-graph with four vertex types:

$$V = V_d \cup V_e \cup V_a \cup V_c,$$

representing *domains*, *entities*, *aspects*, and *clues*, respectively, connected by three typed relations:

$$R_{c \rightarrow a} \subseteq V_c \times V_a \text{ (clue describes aspect),}$$

$$R_{a \rightarrow e} \subseteq V_a \times V_e \text{ (aspect belongs to entity),}$$

$$R_{e \rightarrow d} \subseteq V_e \times V_d \text{ (entity relevant to domain).}$$

- **Candidate Retrieval.** Given a matched clue node $c \in V_c$, aspect candidates are retrieved as:

$$C = N_G^+(c; R_{c \rightarrow a}) = \{u \in V_a \mid (c, u) \in R_{c \rightarrow a}\}.$$

- **Domain-Consistent Filtering.** Domain consistency is enforced via relation composition:

$$R_{a \rightarrow d} = R_{a \rightarrow e} \circ R_{e \rightarrow d},$$

yielding the domain-filtered candidate set for inferred domain $d \in V_d$:

$$C_d = \{a \in C \mid (a, d) \in R_{a \rightarrow d}\}.$$

- **Scoring and Selection.** Each $a \in C_d$ receives a relevance score combining clue–aspect and sentence–aspect similarity:

$$\text{score}(a) = \text{Sim}(\text{Embed}(c), \text{Embed}(a)) + \text{Sim}(\text{Embed}(s), \text{Embed}(a)), \quad (\text{A1})$$

where $\text{Embed}(\cdot)$ denotes SBERT embeddings and $\text{Sim}(\cdot, \cdot)$ cosine similarity. The implicit aspect set is then:

$$I = \arg \text{topk}_{a \in C_d, \text{score}(a) \geq \tau_{\text{asp}}} \text{score}(a).$$

Appendix B. Algorithm

Algorithm A1 operationalizes the formal definitions above. Lonely adjectives are detected and mapped to clue nodes V_c ; the sentence domain is inferred by selecting the most semantically similar domain vertex. For each clue, candidates are retrieved via $R_{c \rightarrow a}$, filtered for domain consistency through $R_{a \rightarrow d}$, and scored via Equation (A1); if no candidate

survives filtering, a *general* fallback is assigned. Top-ranked aspects above τ_{asp} are normalized, deduplicated, and returned as the implicit aspect set I .

Algorithm A1: KG-Based Implicit Aspect Extraction

Input : s : input sentence; P : paraphrased sentences; $G = (V, E)$: labeled digraph; $D \subseteq V$: domain vertices

Output : $I \subseteq V_a$: implicit aspect set

Variables : $L \subseteq V_c$: matched clue vertices; $C, C_d \subseteq V_a$: candidate and filtered aspects; $d \in D$: inferred domain

```

1  $L \leftarrow \text{IdentifyClues}(\{s\} \cup P) \cap V_c$  // Detect lonely adjectives and map to clue nodes
2  $d \leftarrow \arg \max_{d' \in D} \text{Sim}(\phi(s), \phi(d'))$  // Infer sentence domain
3  $I \leftarrow \emptyset$ 
4 foreach  $c \in L$  do
5    $C \leftarrow \{a \in V_a \mid (c, a) \in R_{c \rightarrow a}\}$  // Retrieve aspect candidates
6    $C_d \leftarrow \{a \in C \mid \exists e : (a, e) \in R_{a \rightarrow e} \wedge (e, d) \in R_{e \rightarrow d}\}$  // Domain-consistent filtering
7   if  $C_d = \emptyset$  then  $C_d \leftarrow \{\text{general}\}$ 
8   foreach  $a \in C_d$  do
9      $\text{score}(a) \leftarrow \text{Sim}(c, a) + \text{Sim}(s, a)$  // Equation (A1)
10   $I \leftarrow I \cup \arg \text{topk}_{a \in C_d, \text{score}(a) \geq \tau_{asp}} \text{score}(a)$ 
11  $I \leftarrow \text{FilterAndNormalize}(I)$ 
12 return  $I$ 

```

Appendix C. Sensitivity Analysis

All sensitivity experiments use the same gold-annotated test splits as the main evaluation, varying only inference-time parameters with no retraining. Each parameter is adjusted independently while all others are fixed at their defaults. Performance is measured using micro- P , R , and F_1 as the main evaluation.

Appendix C.1. Sensitivity Analysis on English Datasets

Tables A1–A3 report the numerical results. The English results are largely stable across hyperparameters; semantic similarity thresholding and iterative generation depth show the clearest monotonic effects, while remaining parameters introduce only marginal variation.

Table A1. Paraphrases sensitivity (English datasets).

Setting	p	R	F_1
Original ($k = 0$)	0.9082	0.7542	0.8241
$k = 2$	0.9042	0.7602	0.8260
$k = 4$	0.8994	0.7619	0.8250
$k = 6$	0.9016	0.7610	0.8254
$k = 8$	0.8995	0.7623	0.8252
$k = 10$	0.8970	0.7602	0.8229

Table A2. Similarity threshold sensitivity (English datasets).

Thr.	p	R	F_1
0.50	0.8264	0.7686	0.7965
0.55	0.8493	0.7640	0.8044
0.58	0.8601	0.7631	0.8087
0.60	0.8720	0.7623	0.8135
0.62	0.8764	0.7602	0.8142
0.65	0.8835	0.7585	0.8162
0.70	0.8931	0.7572	0.8195
0.75	0.9024	0.7559	0.8227
0.80	0.9069	0.7551	0.8240

Table A3. IG rounds sensitivity (English datasets).

Rounds	p	R	F_1
2	0.9278	0.5822	0.7154
3	0.9180	0.6453	0.7579
4	0.9089	0.6682	0.7702
5	0.9053	0.6763	0.7742

Appendix C.2. Sensitivity Analysis on Arabic Dataset

Tables A4–A6 summarize the numerical sensitivity results for the Arabic dataset. In contrast to English, Arabic sensitivity trends exhibit stronger trade-offs between precision and recall. Paraphrase quantity and iterative generation depth significantly increase recall but introduce substantial noise.

Table A4. Paraphrases sensitivity (Arabic datasets).

Setting	p	R	F_1
$k = 0$	0.6526	0.7296	0.6890
$k = 2$	0.6223	0.7427	0.6772
$k = 4$	0.6138	0.7493	0.6749
$k = 6$	0.6137	0.7502	0.6751
$k = 8$	0.6085	0.7533	0.6732
$k = 10$	0.6091	0.7515	0.6728

Table A5. Similarity threshold sensitivity (Arabic datasets).

Thr.	p	R	F_1
0.50	0.5871	0.7674	0.6653
0.55	0.5944	0.7617	0.6677
0.58	0.5994	0.7573	0.6691
0.60	0.6030	0.7529	0.6697
0.62	0.6088	0.7507	0.6723
0.65	0.6171	0.7471	0.6759
0.70	0.6269	0.7388	0.6782
0.75	0.6383	0.7335	0.6826
0.80	0.6461	0.7304	0.6857

Table A6. IG rounds sensitivity (Arabic datasets).

Rounds	p	R	F_1
2	0.6512	0.7300	0.6884
3	0.5926	0.7643	0.6676
4	0.5680	0.7898	0.6608
5	0.5510	0.7977	0.6518

Appendix D. Statistical Significance Tests

Table A7 reports significance results for English datasets, and Table A8 reports the KG component significance on implicit rows. SemEval2014 is excluded from KG evaluation owing to absent category annotations.

Table A9 reports significance results for Arabic datasets, and Table A10 reports the KG component on implicit rows. HAAD is excluded from KG evaluation owing to absent category annotations.

Table A7. Statistical significance of pipeline components on English datasets. Bootstrap p (Boot.) and McNemar p (McN.) reported for three evaluation views: all rows (All), aspect-extraction only (Asp.), explicit only (Exp.). F_1 scores and ΔF_1 shown for the explicit-only view. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, — not significant. Negative ΔF_1 indicates degradation relative to system A.

Dataset	Comparison	All		Asp.		Exp.		Exp. F_1		
		Boot.	McN.	Boot.	McN.	Boot.	McN.	F_1 (A)	F_1 (B)	ΔF_1
SE14 ($n = 1028$)	BL $\rightarrow$ BL + IG	***	***	***	***	***	***	0.6202	0.7921	+0.1719
	BL $\rightarrow$ BL + IG + PARA	***	***	***	***	***	***	0.6202	0.8228	+0.2026
	BL + IG $\rightarrow$ BL + IG + PARA	*	—	***	—	***	—	0.7921	0.8228	+0.0306
SE15 ($n = 605$)	BL $\rightarrow$ BL + IG	***	*	***	**	***	**	0.5256	0.6005	+0.0749
	BL $\rightarrow$ BL + IG + PARA	***	*	***	**	***	**	0.5256	0.6207	+0.0950
	BL + IG $\rightarrow$ BL + IG + PARA	—	—	—	—	—	—	0.6005	0.6207	+0.0202
SE16 ($n = 419$)	BL $\rightarrow$ BL + IG	***	*	***	**	***	**	0.6822	0.7880	+0.1058
	BL $\rightarrow$ BL + IG + PARA	***	—	***	—	***	—	0.6822	0.7628	+0.0807
	BL + IG $\rightarrow$ BL + IG + PARA	—	**	—	**	—	**	0.7880	0.7628	−0.0251
ACOS ($n = 1213$)	BL $\rightarrow$ BL + IG	—	***	—	***	—	***	0.8175	0.7870	−0.0305
	BL $\rightarrow$ BL + IG + PARA	—	***	—	***	—	***	0.8175	0.7824	−0.0352
	BL + IG $\rightarrow$ BL + IG + PARA	—	***	—	***	—	***	0.7870	0.7824	−0.0046
M-ABSA _{en} ($n = 2576$)	BL $\rightarrow$ BL + IG	—	**	—	*	—	*	0.5719	0.5824	+0.0106
	BL $\rightarrow$ BL + IG + PARA	**	—	***	—	***	—	0.5719	0.5993	+0.0274
	BL + IG $\rightarrow$ BL + IG + PARA	***	**	**	**	**	**	0.5824	0.5993	+0.0169

Table A8. Statistical significance of the KG component on English implicit rows. Gold = ASPECT part of entity#aspect. SE14 excluded (no category annotations). *** $p < 0.001$.

Dataset	N_{impl}	F_1 (BL)	F_1 (KG)	ΔF_1	Boot.	McN.
SE15 _{en}	245	0.0000	0.3218	+0.3218	***	***
SE16 _{en}	167	0.0000	0.3739	+0.3739	***	***
ACOS _{en}	344	0.0058	0.2885	+0.2827	***	***
M-ABSA _{en}	913	0.0010	0.3016	+0.3005	***	***

Table A9. Statistical significance of pipeline components on Arabic datasets. ** $p < 0.01$, *** $p < 0.001$, — not significant. Negative ΔF_1 indicates degradation relative to system A.

Dataset	Comparison	All		Asp.		Exp.		Exp. F_1		
		Boot.	McN.	Boot.	McN.	Boot.	McN.	F_1 (A)	F_1 (B)	ΔF_1
M-ABSA _{ar} ($n = 1356$)	BL $\rightarrow$ BL + IG	***	***	***	***	***	***	0.6559	0.7744	+0.1185
	BL $\rightarrow$ BL + IG + PARA	***	***	***	***	***	**	0.6559	0.7642	+0.1083
	BL + IG $\rightarrow$ BL + IG + PARA	—	—	—	—	—	—	0.7744	0.7642	−0.0103
SE16 _{ar} ($n = 1130$)	BL $\rightarrow$ BL + IG	—	***	—	***	—	***	0.7676	0.7156	−0.0521
	BL $\rightarrow$ BL + IG + PARA	—	***	—	***	—	***	0.7676	0.7343	−0.0333
	BL + IG $\rightarrow$ BL + IG + PARA	***	—	***	—	***	—	0.7156	0.7343	+0.0187
HAAD ($n = 299$)	BL $\rightarrow$ BL + IG	—	***	—	***	—	***	0.4465	0.4720	+0.0254
	BL $\rightarrow$ BL + IG + PARA	—	***	—	***	—	***	0.4465	0.4743	+0.0278
	BL + IG $\rightarrow$ BL + IG + PARA	—	—	—	—	—	—	0.4720	0.4743	+0.0024

Table A10. Statistical significance of the KG component on Arabic implicit rows. Gold = ASPECT part of entity#aspect. HAAD excluded (no category annotations). *** $p < 0.001$.

Dataset	N_{impl}	F_1 (BL)	F_1 (KG)	ΔF_1	Boot.	McN.
SE16 _{ar}	96	0.0000	0.4099	+0.4099	***	***
M-ABSA _{ar}	718	0.0000	0.3725	+0.3725	***	***

References

1. Pontiki, M.; Galanis, D.; Pavlopoulos, J.; Papageorgiou, H.; Androutsopoulos, I.; Manandhar, S. SemEval-2014 Task 4: Aspect Based Sentiment Analysis. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*; Nakov, P., Zesch, T., Eds.; Association for Computational Linguistics: Stroudsburg, PA, USA, 2014; pp. 27–35. <https://doi.org/10.3115/v1/S14-2004>.
2. Pontiki, M.; Galanis, D.; Papageorgiou, H.; Manandhar, S.; Androutsopoulos, I. SemEval-2015 Task 12: Aspect Based Sentiment Analysis. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*; Nakov, P., Zesch, T., Cer, D., Jurgens, D., Eds.; Association for Computational Linguistics: Stroudsburg, PA, USA, 2015; pp. 486–495. <https://doi.org/10.18653/v1/S15-2082>.

3. Pontiki, M.; Galanis, D.; Papageorgiou, H.; Androutsopoulos, I.; Manandhar, S.; Al-Smadi, M.; Al-Ayyoub, M.; Zhao, Y.; Qin, B.; De Clercq, O.; et al. SemEval-2016 Task 5: Aspect Based Sentiment Analysis. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*; Bethard, S., Carpuat, M., Cer, D., Jurgens, D., Nakov, P., Zesch, T., Eds.; Association for Computational Linguistics: Stroudsburg, PA, USA, 2016; pp. 19–30. <https://doi.org/10.18653/v1/S16-1002>.
4. Cai, H.; Xia, R.; Yu, J. Aspect-Category-Opinion-Sentiment Quadruple Extraction with Implicit Aspects and Opinions. In *Proceedings of the ACL-IJCNLP 2021—59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, Proceedings of the Conference*; Association for Computational Linguistics: Stroudsburg, PA, USA, 2021; pp. 340–350. <https://doi.org/10.18653/V1/2021.ACL-LONG.29>.
5. Wu, C.; Ma, B.; Liu, Y.; Zhang, Z.; Deng, N.; Li, Y.; Chen, B.; Zhang, Y.; Xue, Y.; Plank, B. M-ABSA: A Multilingual Dataset for Aspect-Based Sentiment Analysis. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*; Christodoulopoulos, C., Chakraborty, T., Rose, C., Peng, V., Eds.; Association for Computational Linguistics: Stroudsburg, PA, USA, 2025; pp. 2530–2557. <https://doi.org/10.18653/v1/2025.emnlp-main.128>.
6. Karagoz, P.; Kama, B.; Ozturk, M.; Toroslu, I.H.; Canturk, D. A framework for aspect based sentiment analysis on Turkish informal texts. *J. Intell. Inf. Syst.* **2019**, *53*, 431–451. <https://doi.org/10.1007/S10844-019-00565-W>.
7. Nandhini, M.D.S.; Pradeep, G. A Hybrid Co-occurrence and Ranking-based Approach for Detection of Implicit Aspects in Aspect-Based Sentiment Analysis. *SN Comput. Sci.* **2020**, *1*, 128. <https://doi.org/10.1007/S42979-020-00138-7>.
8. Al-Ayyoub, M.; Gigieh, A.; Al-Qwaqenah, A.; Al-Kabi, M.; Talafha, B.; Alsmadi, I. Aspect-Based Sentiment Analysis of Arabic Laptop Reviews. In *Proceedings of the 18th International Arab Conference on Information Technology (ACIT)*, Hammamet, Tunisia, 22–24 December 2017.
9. Allassaf, M.; Qamar, A.M. Aspect-Based Sentiment Analysis of Arabic Tweets in the Education Sector Using a Hybrid Feature Selection Method. In *Proceedings of the 2020 14th International Conference on Innovations in Information Technology (IIT)*; IEEE: New York, NY, USA, 2020; pp. 178–185. <https://doi.org/10.1109/IIT50501.2020.9299026>.
10. Sana, T.; Boujelben, I.; Jamoussi, S.; Benayed, Y. A hybrid method for Arabic aspect-based sentiment analysis. *Int. J. Hybrid Intell. Syst.* **2020**, *16*, 1–12. <https://doi.org/10.3233/HIS-200285>.
11. Al-Smadi, M.; Hammad, M.M.; Al-Zboon, S.A.; Al-Tawalbeh, S.; Cambria, E. Gated Recurrent Unit with Multilingual Universal Sentence Encoder for Arabic Aspect-Based Sentiment Analysis. *Knowl.-Based Syst.* **2021**, *261*, 107540. <https://doi.org/10.1016/j.knosys.2021.107540>.
12. Alghamdi, A.; Taileb, M.; Almani, N. Stacked LSTM-GRU with Double Attention model for Arabic Aspect-Based Sentiment Analysis. In *Proceedings of the 2025 2nd International Conference on Advanced Innovations in Smart Cities (ICAISC)*; IEEE: New York, NY, USA, 2025. <https://doi.org/10.1109/ICAISC64594.2025.10959531>.
13. Aljomah, F.; Aldhafeeri, L.; Alfadel, M.; Alshahrani, S.; Abbas, Q.; Alhumoud, S. Enhancing Arabic Sentiment Analysis with Pre-Trained CAMELBERT: A Case Study on Noisy Texts. *Comput. Mater. Contin.* **2025**, *84*, 5318–5335. <https://doi.org/10.32604/cmc.2025.062478>.
14. Kalra, S. The Aspect Extraction using Topic-aware dynamic Convolutional Neural Network. In *Proceedings of the 2023 3rd International Conference on Smart Generation Computing, Communication and Networking (SMART GENCON)*; IEEE: New York, NY, USA, 2023; pp. 1–7. <https://doi.org/10.1109/SMARTGENCON60755.2023.10442593>.
15. Zhang, Z.; Rao, Y.; Lai, H.; Wang, J.; Yin, J. TADC: A Topic-Aware Dynamic Convolutional Neural Network for Aspect Extraction. *IEEE Trans. Neural Netw. Learn. Syst.* **2023**, *34*, 3912–3924. <https://doi.org/10.1109/TNNLS.2021.3119026>.
16. Youssef, L.; Elhoussaine, Z.; Soufiane, N.; Nouredine, M. Enhancing Arabic aspect category detection using large language models (LLMs). *Results Eng.* **2025**, *26*, 105049. <https://doi.org/10.1016/j.rineng.2025.105049>.
17. Alotaibi, A.; Nadeem, F. An unsupervised integrated framework for Arabic aspect-based sentiment analysis and abstractive text summarization of traffic services using transformer models. *Smart Cities* **2025**, *8*, 62. <https://doi.org/10.3390/smartcities8020062>.
18. AlNasser, S.; AlMuhaideb, S. Listening to Patients: Advanced Arabic Aspect-Based Sentiment Analysis Using Transformer Models Towards Better Healthcare. *Big Data Cogn. Comput.* **2024**, *8*, 156. <https://doi.org/10.3390/bdcc8110156>.
19. Gadag, A.; Sagar, B.M. A review on different methods of paraphrasing. In *Proceedings of the 2016 International Conference on Electrical, Electronics, Communication, Computer and Optimization Techniques (ICEECOT)*; IEEE: New York, NY, USA, 2016; pp. 188–191. <https://doi.org/10.1109/ICEECOT.2016.7955212>.
20. Titov, I.; McDonald, R. Modeling Online Reviews with Multi-grain Topic Models. In *Proceedings of the 17th International Conference on World Wide Web*; Association for Computing Machinery: New York, NY, USA, 2008; pp. 111–120. <https://doi.org/10.1145/1367497.1367513>.
21. Miller, G.A. WordNet: A lexical database for English. *Commun. ACM* **1995**, *38*, 39–41. <https://doi.org/10.1145/219717.219748>.
22. Elkateb, S.; Black, W.; Rodríguez, H.; Alkhalifa, M.; Vossen, P.; Pease, A.; Fellbaum, C. Building a WordNet for Arabic. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC'06)*, Genoa, Italy, 22–28 May 2006.
23. Speer, R.; Chin, J.; Havasi, C. ConceptNet 5.5: An Open Multilingual Graph of General Knowledge. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*; PKP Publishing: Burnaby, BC, Canada, 2017; pp. 4444–4451.

24. Liao, J.; Wang, M.; Chen, X.; Wang, S.; Zhang, K. Dynamic Commonsense Knowledge Fused Method for Chinese Implicit Sentiment Analysis. *Inf. Process. Manag.* **2022**, *59*, 102934. <https://doi.org/10.1016/j.ipm.2022.102934>.
25. Xu, M.; Wang, D.; Feng, S.; Yang, Z.; Zhang, Y. KC-ISA: An Implicit Sentiment Analysis Model Combining Knowledge Enhancement and Context Features. In *Proceedings of the 29th International Conference on Computational Linguistics*; International Committee on Computational Linguistics: New York, NY, USA, 2022; pp. 6906–6915.
26. Liang, B.; Su, H.; Gui, L.; Cambria, E.; Xu, R. Aspect-based sentiment analysis via affective knowledge enhanced graph convolutional networks. *Knowl.-Based Syst.* **2022**, *235*, 107643.
27. Teo, A.; Wang, Z.; Pen, H.; Subagdja, B.; Ho, S.B.; Quek, B.K. Knowledge Graph enhanced Aspect-Based Sentiment Analysis Incorporating External Knowledge. In *Proceedings of the IEEE International Conference on Data Mining Workshops (ICDMW)*; IEEE: New York, NY, USA, 2023; pp. 791–798.
28. Dubey, G.; Dubey, A.K.; Kaur, K.; Raj, G.; Kumar, P. Adaptive Contextual Memory Graph Transformer with Domain-Adaptive Knowledge Graph for Aspect-Based Sentiment Analysis. *Expert Syst. Appl.* **2025**, *278*, 127300. <https://doi.org/10.1016/j.eswa.2025.127300>.
29. Shaalan, K.; Siddiqui, S.; Alkhatib, M.; Monem, A.A. Chapter 3: Challenges in Arabic Natural Language Processing. In *Computational Linguistics, Speech and Image Processing for Arabic Language*; Series on Language Processing, Pattern Recognition, and Intelligent Systems; World Scientific Publishing: Singapore, 2018; pp. 59–83. https://doi.org/10.1142/9789813229396_0003.
30. Ruder, S.; Ghaffari, P.; Breslin, J.G. INSIGHT-1 at SemEval-2016 Task 5: Deep Learning for Multilingual Aspect-based Sentiment Analysis. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*; Bethard, S., Carpuat, M., Cer, D., Jurgens, D., Nakov, P., Zesch, T., Eds.; Association for Computational Linguistics: Stroudsburg, PA, USA, 2016; pp. 330–336. <https://doi.org/10.18653/v1/S16-1053>.
31. Ruder, S.; Ghaffari, P.; Breslin, J.G. A Hierarchical Model of Reviews for Aspect-based Sentiment Analysis. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*; Su, J., Duh, K., Carreras, X., Eds.; Association for Computational Linguistics: Stroudsburg, PA, USA, 2016; pp. 999–1005. <https://doi.org/10.18653/v1/D16-1103>.
32. Al-Dabet, S.; Tedmori, S.; AL-Smadi, M. Enhancing Arabic aspect-based sentiment analysis using deep learning models. *Comput. Speech Lang.* **2021**, *69*, 101224. <https://doi.org/https://doi.org/10.1016/j.csl.2021.101224>.
33. Fadel, A.; Saleh, M.; Salama, R.; Abulnaja, O. MTL-AraBERT: An Enhanced Multi-Task Learning Model for Arabic Aspect-Based Sentiment Analysis. *Computers* **2024**, *13*, 98. <https://doi.org/10.3390/computers13040098>.
34. Youssef, L.; ELhoussaine, Z. Arabic Aspect Category Detection using traditional neural networks and ARBERT. In *Proceedings of the 2024 11th International Conference on Wireless Networks and Mobile Communications (WINCOM)*; IEEE: New York, NY, USA, 2024. <https://doi.org/10.1109/WINCOM62286.2024.10657409>.
35. Chouikhi, H.; Jarray, F.; Alsuhaibani, M. A Sequence-to-Sequence Neural Network for Joint Aspect Term Extraction and Aspect Term Sentiment Classification Tasks. In *Proceedings of the 15th International Conference on Agents and Artificial Intelligence, INSTICC*; SciTePress: Setubal, Portugal, 2023; pp. 117–123. <https://doi.org/10.5220/0011620500003393>.
36. Wang, L.; Yao, C.; Li, X.; Yu, X. BERT-based implicit aspect extraction. In *Proceedings of the 2021 IEEE 3rd International Conference on Civil Aviation Safety and Information Technology, ICCASIT 2021*; IEEE: New York, NY, USA, 2021; pp. 758–761. <https://doi.org/10.1109/ICCASIT53235.2021.9633578>.
37. Busst, M.M.A.; Anbananthen, K.S.M.; Kannan, S.; Krishnan, J.; Subbiah, S. Ensemble BiLSTM A Novel Approach for Aspect Extraction From Online Text. *IEEE Access* **2024**, *12*, 3528–3539.
38. Zheng, Y.; Li, X. You Only Read Once: Constituency-Oriented Relational Graph Convolutional Network for Multi-Aspect Multi-Sentiment Classification; AAAI'24/IAAI'24/EAAI'24; AAAI Press: Washington, DC, USA, 2024. <https://doi.org/10.1609/aaai.v38i17.29945>.
39. Chauhan, G.S.; Nahta, R.; Meena, Y.K. Improved Deep Learning Model for Aspect Words Extraction. In *Proceedings of the 2023 IEEE World Conference on Applied Intelligence and Computing (AIC)*; IEEE: New York, NY, USA, 2023; pp. 286–292. <https://doi.org/10.1109/AIC57670.2023.10263853>.
40. Chauhan, G.S.; Nahta, R.; Meena, Y.K. A Review Level Aspect Extraction in Sentiment Analysis. In *Proceedings of the 2024 2nd International Conference on Computer, Communication and Control (IC4)*; IEEE: New York, NY, USA, 2024; pp. 1–7. <https://doi.org/10.1109/IC457434.2024.10486257>.
41. Jayanthi, S.; Arumugam, S.S. Refining Sentiment Analysis: Explicit Aspect Extraction with Diverse Datasets and Advanced Models. In *Proceedings of the 2024 Second International Conference on Intelligent Cyber Physical Systems and Internet of Things (ICoICI)*; IEEE: New York, NY, USA, 2024; pp. 182–187. <https://doi.org/10.1109/ICOICI62503.2024.10696470>.
42. Wang, T. Domain Adaptive English Aspect Word Extraction Method Based On Bidirectional Long And Short-term Memory Network And Multi-head Attention Mechanism. *J. Appl. Sci. Eng.* **2025**, *28*, 2661–2669. [https://doi.org/10.6180/jase.202512_28\(12\).0013](https://doi.org/10.6180/jase.202512_28(12).0013).
43. Huang, Z.; Li, L.; Li, Z.; Zhang, H.; Yan, J. Bidirectional interaction and inference strategy joint approach for span-level aspect sentiment triplet extraction. *Int. J. Mach. Learn. Cybern.* **2025**, *16*, 7621–7636. <https://doi.org/10.1007/s13042-025-02675-0>.

44. Quan, X.; Min, Z.; Li, K.; Yang, Y. Compound Aspect Extraction by Augmentation and Constituency Lattice. *IEEE Trans. Affect. Comput.* **2023**, *14*, 2323–2335. <https://doi.org/10.1109/TAFFC.2022.3161683>.
45. Soni, P.K.; Rambola, R.K. Deep Learning, WordNet, and spaCy based Hybrid Method for Detection of Implicit Aspects for Sentiment Analysis. In *Proceedings of the 2021 International Conference on Intelligent Technologies, CONIT 2021*; IEEE: New York, NY, USA, 2021. <https://doi.org/10.1109/CONIT51480.2021.9498372>.
46. Hammad, M.; AbuEnnab, N.; Al-Refai, M. An aspect-based sentiment analysis model for Arabic game reviews based on hybrid transformers models. *Neural Comput. Appl.* **2025**, *37*, 10309–10331. <https://doi.org/10.1007/s00521-025-11032-9>.
47. Almasre, M.A. Enhance the Aspect Category Detection in Arabic Language using AraBERT and Text Augmentation. In *Proceedings of the 2022 Fifth National Conference of Saudi Computers Colleges (NCCC)*; IEEE: New York, NY, USA, 2022; pp. 7–10. <https://doi.org/10.1109/NCCC57165.2022.10067648>.
48. Tamchyna, A.; Veselovská, K. UFAL at SemEval-2016 Task 5 - Recurrent Neural Networks for Sentence Classification. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*; Association for Computational Linguistics: Stroudsburg, PA, USA, 2016; pp. 367–371. <https://doi.org/10.18653/v1/S16-1059>.
49. Al-Smadi, M.; Talafha, B.; Al-Ayyoub, M.; Jararweh, Y. Using long short-term memory deep neural networks for aspect-based sentiment analysis of Arabic reviews. *Int. J. Mach. Learn. Cybern.* **2019**, *10*, 2163–2175. <https://doi.org/10.1007/s13042-018-0799-4>.
50. Abdelgwad, M.M.; Soliman, T.H.A.; Taloba, A.I.; Farghaly, M.F. Arabic aspect based sentiment analysis using bidirectional GRU based models. *J. King Saud Univ.—Comput. Inf. Sci.* **2021**, *34*, 6652–6662. <https://doi.org/10.1016/j.jksuci.2021.08.030>.
51. Fadel, A.S.; Saleh, M.E.; Abulnaja, O.A. Arabic Aspect Extraction Based on Stacked Contextualized Embedding with Deep Learning. *IEEE Access* **2022**, *10*, 30526–30535. <https://doi.org/10.1109/ACCESS.2022.3159252>.
52. Bensoltane, R.; Zaki, T. Towards Arabic aspect-based sentiment analysis: A transfer learning-based approach. *Soc. Netw. Anal. Min.* **2022**, *12*, 7. <https://doi.org/10.1007/s13278-021-00794-4>.
53. Hamam, H.; Rehman, A.U.; Othman, M.T.B.; Chouikhi, H.; Alsuhaibani, M.; Jarray, F. BERT-Based Joint Model for Aspect Term Extraction and Aspect Polarity Detection in Arabic Text. *Electronics* **2023**, *12*, 515. <https://doi.org/10.3390/ELECTRONICS12030515>.
54. Alamoudi, E.; Solaiman, E. EHSAN: Leveraging ChatGPT in a Hybrid Framework for Arabic Aspect-Based Sentiment Analysis in Healthcare. In *Proceedings of the International Conference on Intelligent Data Engineering and Automated Systems (IDEAS) 2025*; Preprint; Zenodo: Geneva, Switzerland, 2025. <https://doi.org/10.5281/zenodo.15418860>.
55. Thi, T.L.; Tran, T.K.; Phan, T.T. Deep Learning Using Context Vectors to Identify Implicit Aspects. *IEEE Access* **2023**, *11*, 39385–39393. <https://doi.org/10.1109/ACCESS.2023.3268243>.
56. Xiong, H.; Yan, Z.; Wu, C.; Lu, G.; Pang, S.; Xue, Y.; Cai, Q. BART-based contrastive and retrospective network for aspect-category-opinion-sentiment quadruple extraction. *Int. J. Mach. Learn. Cybern.* **2023**, *14*, 3243–3255. <https://doi.org/10.1007/S13042-023-01831-8/METRICS>.
57. Ahmed, M.; Pan, S.; Su, J.; Cao, X.; Zhang, W.; Wen, B.; Liu, Y. BERT-ASC: Implicit Aspect Representation Learning through Auxiliary-Sentence Construction for Sentiment Analysis. *arXiv* **2022**, arXiv:2203.11702. <https://doi.org/10.48550/arxiv.2203.11702>.
58. Feng, J.; Cai, S.; Ma, X. Enhanced sentiment labeling and implicit aspect identification by integration of deep convolution neural network and sequential algorithm. *Clust. Comput.* **2019**, *22*, 5839–5857. <https://doi.org/10.1007/S10586-017-1626-5/METRICS>.
59. Lazhar, F. Implicit feature identification for opinion mining. *Int. J. Bus. Inf. Syst.* **2019**, *30*, 13–30. <https://doi.org/10.1504/IJBIS.2019.097042>.
60. Li, X.; Wang, X.; Yao, C.; Li, Y. Graph-enhanced implicit aspect-level sentiment analysis based on multi-prompt fusion. *Sci. Rep.* **2025**, *15*, 17460. <https://doi.org/10.1038/s41598-025-02609-4>.
61. Harjo, B.; Abdullah, R. Attention-Based Sentence Extraction for Aspect-Based Sentiment Analysis with Implicit Aspect Cases in Hotel Reviews Using Machine Learning Algorithm, Semantic Similarity, and BERT. *Int. J. Comput. Commun. Control* **2023**, *18*, 450–468.
62. Kabir, M.M.; Othman, Z.A.; Yaakub, M.R. A Hybrid Frequency Based, Syntax, and Conditional Random Field Method for Implicit and Explicit Aspect Extraction. *IEEE Access* **2024**, *12*, 72361–72373. <https://doi.org/10.1109/ACCESS.2024.3403479>.
63. Agathangelou, P.; Katakis, I.; Kasnesis, P. Word- and Sentence-Level Representations for Implicit Aspect Extraction. *IEEE Trans. Comput. Soc. Syst.* **2024**, *11*, 5935–5948. <https://doi.org/10.1109/TCSS.2024.3391833>.
64. Neveditsin, N.; Lingras, P.; Mago, V.K. From Annotation to Adaptation: Metrics, Synthetic Data, and Aspect Extraction for Aspect-Based Sentiment Analysis with Large Language Models. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 4: Student Research Workshop)*; Association for Computational Linguistics: Stroudsburg, PA, USA, 2025; pp. 142–161. <https://doi.org/10.18653/v1/2025.naacl-srw.14>.
65. Ahmed, K.; Nadeem, M.I.; Wang, G.; Zuo, F.; Han, Z. Instruction-tuned ABSA with auxiliary sentences and knowledge-enhanced graphs for implicit aspect detection. *Expert Syst. Appl.* **2025**, *289*, 128284. <https://doi.org/10.1016/j.eswa.2025.128284>.

66. Zhou, J.; Li, F.; Su, F.; Aziz, K.; Yuan, J.; Teng, C.; Ji, D. Generative implicit opinion mining with term correlation prompts. *Inf. Sci.* **2026**, *730*, 122869. <https://doi.org/10.1016/j.ins.2025.122869>.
67. Bu, K.; Liu, Y.; Wang, W.; Cao, Z. M-PIN: Multilingual prompt-based implicit sentiment analysis network. *Appl. Soft Comput.* **2026**, *195*, 114945. <https://doi.org/10.1016/j.asoc.2026.114945>.
68. Behdenna, S.; Barigou, F.; Belalem, G. Towards semantic aspect-based sentiment analysis for Arabic reviews. *Int. J. Inf. Syst. Serv. Sect.* **2020**, *12*, 1–13. <https://doi.org/10.4018/IJISSS.2020100101>.
69. Behdenna, S.; Barigou, F.; Belalem, G. *Performance Evaluation of an Implicit Aspect Extraction Approach in Arabic Reviews*; Lecture Notes in Networks and Systems; Springer: Berlin/Heidelberg, Germany, 2022; Volume 318, pp. 445–456. https://doi.org/10.1007/978-981-16-5689-7_39.
70. Eldin, S.S.; Mohammed, A.; Eldin, A.S.; Hefny, H. An enhanced opinion retrieval approach via implicit feature identification. *J. Intell. Inf. Syst.* **2021**, *57*, 101–126. <https://doi.org/10.1007/S10844-020-00622-9>.
71. Mir, J.; Mahmood, A.; Khatoon, S. Multi-level knowledge engineering approach for mapping implicit aspects to explicit aspects. *Comput. Mater. Contin.* **2022**, *70*, 3491–3509. <https://doi.org/10.32604/CMC.2022.019952>.
72. Khabour, S.M.; Al-Radaideh, Q.A.; Mustafa, D. A New Ontology-Based Method for Arabic Sentiment Analysis. *Big Data Cogn. Comput.* **2022**, *6*, 48. <https://doi.org/10.3390/BDCC6020048>.
73. Wiseman, S.; Rush, A.M. Sequence-to-Sequence Learning as Beam-Search Optimization. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*; Association for Computational Linguistics: Stroudsburg, PA, USA, 2016; pp. 1296–1306. <https://doi.org/10.18653/v1/D16-1137>.
74. OpenAI. Introducing GPT-4.1 in the API. Available online: <https://openai.com/index/gpt-4-1/> (accessed on 10 June 2026).
75. Obeid, O.; Zalmout, N.; Khalifa, S.; Taji, D.; Oudah, M.; Alhafni, B.; Inoue, G.; Eryani, F.; Erdmann, A.; Habash, N. CAMEL Tools: An Open Source Python Toolkit for Arabic Natural Language Processing. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*; Calzolari, N., Béchet, F., Blache, P., Choukri, K., Cieri, C., Declerck, T., Goggi, S., Isahara, H., Maegaard, B., Mariani, J., et al., Eds.; European Language Resources Association: Paris, France, 2020; pp. 7022–7032.
76. Grattafiori, A.; Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Vaughan, A.; et al. The Llama 3 Herd of Models. *arXiv* **2024**, arXiv:2407.21783. <https://doi.org/10.48550/arXiv.2407.21783>.
77. Hu, E.J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W. LoRA: Low-Rank Adaptation of Large Language Models. *arXiv* **2022**, arXiv:2106.09685. <https://doi.org/10.48550/arXiv.2106.09685>.
78. Zhou, J.; Yang, H.; He, Y.; Mou, H.; Yang, J. A Unified One-Step Solution for Aspect Sentiment Quad Prediction. In *Proceedings of the Findings of the Association for Computational Linguistics: ACL 2023*; Association for Computational Linguistics: Stroudsburg, PA, USA, 2023; pp. 12249–12265. <https://doi.org/10.18653/v1/2023.findings-acl.777>.
79. Xu, X.; Zhang, J.D.; Xiong, L.; Liu, Z. iACOS: Advancing Implicit Sentiment Extraction with Informative and Adaptive Negative Examples. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*; Association for Computational Linguistics: Stroudsburg, PA, USA, 2024. <https://doi.org/10.18653/v1/2024.naacl-long.241>.
80. Berg-Kirkpatrick, T.; Burkett, D.; Klein, D. An Empirical Investigation of Statistical Significance in NLP. In *Proceedings of the EMNLP*; Association for Computational Linguistics: Stroudsburg, PA, USA, 2012; pp. 995–1005.
81. McNemar, Q. Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika* **1947**, *12*, 153–157.
82. Haq, B.; Daudpota, S.; Ali, I.; Kastrati, Z.; Noor, W. A Semi-Supervised Approach for Aspect Category Detection and Aspect Term Extraction from Opinionated Text. *Comput. Mater. Contin.* **2023**, *77*, 115–137.
83. Chen, M.; Mao, Y.; Hua, Q.; Wu, J. Aspect Extraction from E-Commerce Reviews Based on Word–Word Relationship Classification. In *Proceedings of the 2023 7th Asian Conference on Artificial Intelligence Technology (ACAIT)*; IEEE: New York, NY, USA, 2023; pp. 226–231. <https://doi.org/10.1109/ACAIT60137.2023.10528431>.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.