KNOWSENTMIX: LEVERAGING CONTEXTUAL EMBEDDINGS AND KNOWLEDGE GRAPHS FOR DEEPER SENTIMENT UNDERSTANDING IN CODE-MIXED TEXT

MANYAM THAILE[1], CHIN-SHIUH SHIEH[2], DR. M. NAGABHUSHANA RAO[*], DR. S. K. RAJESH KANNA[3], DR. RAMARAJU NAGARJUNA KUMAR[4], DR C MADHUSUDANA RAO[5], DR. HRUSHIKESH JAIWANT JOSHI[6], DR. U D PRASAN[7], RAMARAO NAYAPAMU[8]

[1]Associate Professor, School of Engineering, Anurag University, Medchal Malkajgiri, Telangana, India.
[2]Professor, Department of Electronic Engineering, National Kaohsiung University of Science and Technology, Taiwan.
*Corresponding Author, Professor, School of Computer Science & Technology, Malla Reddy (MR) Deemed to be University, Medchal, Malkajgiri, Telangana, India.
[2]Professor, Department of Mechanical Engineering, Rajalakshmi Institute of Technology, Chennai, India.
[3]Professor, Department of Mechanical Engineering, Rajalakshmi Institute of Technology, Chennai, India.
[4]Principal Scientist (Computer Applications in Agriculture)
ICAR-Central Research Institute for Dryland Agriculture (ICAR-CRIDA), HYDERABAD.
[5]Professor of CSE School of Computing, Mohan Babu University, Tirupati, India.
[6]Assistant Professor, Department of AI & DS, Vishwakarma Institute of Technology, Pune, India.
[7]Professor, Department of CSE, Aditya Institute of Technology and Management, Tekkali. India.
[8]Department of Computer Science and Engineering, Koneru Lakshmaiah Education Foundation, Vijayawada
E-mail: [1]manyamthaile@gmail.com, [2]csshieh@nkust.edu.tw ,*mnraosir@gmail.com, [3]skrkanna@gmail.com, [4]nagarjunakumar191@gmail.com, [5]npr4567@gmail.com, [6]hrushikesh.joshi@vit.edu, [7]udprasanna@gmail.com,[8]nramarao@kluniversity.in

ABSTARCT

KnowSentMix is a hybrid sentiment-reasoning framework for code-mixed social text that selectively combines multilingual contextual representations with compact commonsense knowledge. The method uses an agreement-gated alignment mixture that (i) retrieves and encodes concise knowledge snippets, (ii) grounds concepts to tokens through optimal-transport–based soft alignment, and (iii) fuses text-only and knowledge-only predictions via a gate driven by uncertainty and inter-expert agreement. The system returns a sentiment label together with compact explanatory evidence—top concepts and, when available, trigger spans—improving interpretability with modest computational overhead. Objectives, scope, mathematical formulation, and algorithmic design are accompanied by a full-mode implementation using XLM-R with ConceptNet retrieval. Comparative efficiency measurements and ablation studies indicate that adaptive knowledge gating delivers stable gains on idiomatic, implicit, and noisy code-mixed inputs while preserving efficiency.

Keywords: *Code-Mixed NLP, Sentiment Analysis, Commonsense Knowledge, Conceptnet, COMET, XLM-R, Optimal Transport, Interpretability, Knowledge-Infused Learning, Multilingual Transformers*.

1. INTRODUCTION

Code-mixed social media in Hinglish and Benglish blends scripts, transliteration, informal spellings, and code-switching within and across words, so polarity is frequently implied through idioms rather than surface sentiment markers. Multilingual encoders such as XLM-RoBERTa have substantially improved cross-lingual robustness, yet culture-specific expressions and under-specified context still lead to brittle decisions when cues are implicit or oblique [1]. Shared evaluations on code-mixed sentiment highlight exactly these failure modes under noise, non-standard orthography, and transliteration variability [5]. Indian-language backbones (e.g., IndicBERT, MuRIL) extend vocabulary coverage and tokenization fidelity, but even well-trained encoders struggle when the emotional state must be inferred from world knowledge instead of lexical clues [6], [7].

ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

A natural complement is compact commonsense. Typed relational graphs such as ConceptNet connect everyday events and affective states, while generative commonsense models like COMET/ATOMIC infer likely intents, reactions, and outcomes—signals that help map phrases such as "dil toot gaya" to the latent state "heartbreak → sadness" [2], [3]. However, naïve injection of retrieved facts can mislead a classifier: knowledge snippets may be noisy, off-topic, or weakly coupled to the exact text span that determines sentiment. Prior knowledge-infused architectures (e.g., K-BERT) and retrieval-augmented modeling confirm the upside of external evidence but also expose the need for careful grounding and controllable fusion so that knowledge contributes only when it should [4], [8].

KnowSentMix addresses these requirements with three ideas that work together. First, a multilingual text expert (e.g., XLM-R) encodes the input; a lightweight knowledge expert retrieves a small, budgeted set of concepts from ConceptNet or COMET and encodes them into vectors [1]–[3]. Second, a soft token–concept grounding step uses an optimal-transport prior computed with Sinkhorn iterations to bias attention toward plausible alignments, improving faithfulness between reasons and the words they explain [9]. Third, an agreement- and uncertainty-aware fusion gate controls how much the knowledge head influences the final decision, increasing weight when the experts agree or when the text head is uncertain; Jensen–Shannon divergence and entropy offer simple, effective signals for this control law [10]. The output couples a sentiment label with concise explanatory evidence (top concepts and, when available, trigger spans), supporting transparent judgments with modest overhead.

Advances in this direction include an agreement-gated mechanism that raises knowledge influence precisely when external cues reinforce the text model or compensate for ambiguity; OT-based token–concept alignment that stabilizes grounding and improves interpretability; and a practical full-mode realization built on XLM-R and ConceptNet with ablations and efficiency measurements suitable for reproducible comparison against text-only and knowledge-infused baselines [1]–[5], [8]–[10]. Scope centres on short Hinglish/Bengali-English posts, a small concept budget for low-latency inference, and modularity to swap in Indic backbones (IndicBERT, MuRIL) or alternative knowledge sources without architectural changes [6], [7]. Section II reviews related work; Section III outlines the system; Section IV details alignment and fusion; Sections V–VI cover datasets and experimental setup; Section VII reports results and ablations; Section VIII discusses implications and limits; Section IX sketches applications; Section X addresses ethics; and Section XI concludes.

2. BACKGROUND AND RELATED WORK

The scripts, transliteration, informal spellings and frequent code switching of Indian social media in Hinglish and Benglish undermine lexicon-based and monolingual schemes and may confuse even strong multilingual encoders in the case of implied affect and not stated affect; common evaluation such as SentMix repeatedly documents robustness failures under noise and transliteration variability in this environment [5]. Multilingual transformers such as XLM-RoBERTa give high cross-lingual transfer through sub word tokenization and a variety of pretraining, and serve as a solid backbone in mixed-language NLP with defaults [1]. Indic BERT and MuRIL can better tokenize and cover the Indian languages, but rely on surface-based cues and frequently fail polarity in cases where contextual or culture-specific idioms should explain the state of the heart [6], [7]. The complementary signals provided by commonsense resources include typed relations which connect everyday events to affect across languages (conceptnet) and generative models which infer intents, reactions, and effects (COMET/ATOMIC) - cues useful in mapping expressions like Dil toot Gaya into a latent state of heartbreak and sadness [2], [3]. Strategies based on simple feature concatenation (text with averaged concept embeddings) to structural injection (edges injected into the encoder) or retrieval-augmented modeling (conditions prediction on external snippets) are examples of prior knowledge-infused strategies that can be helpful but prone to over-trusting spurious facts or representation entanglement without reliable grounding and controlled fusion [4], [8]. Three themes of difficulties arise in: when to use the external knowledge to make a decision, anchoring retrieved concepts to the specific tokens that motivate sentiment in order to make the explanations true to the concrete concepts, and maintaining the low latency with the real concept budgets. The strategy followed here tackles these aspects with Sinkhorn-normalized optimum-transport optimal-match prior on soft token-concept alignment, a slim bipartite message-exchange layer, and an agreement-based fusion rule that scales knowledge impact only when justified [9], [10].

Journal of Theoretical and Applied Information Technology
31st December 2025. Vol.103. No.24

ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

3. SYSTEM OVERVIEW

The input data begins with a code-mixed entry xxx. A multilingual encoder (e.g., XLM-RoBERTa) transforms the tokens into contextual states; a small retriever then retrieves a small, tightly budgeted set of concepts or brief facts, by external sources of commonsense knowledge, such as Concept Net and COMET/ATOMIC. Encoded retrieved items are soft-aligned with tokens using a Sinkhorn-normalized optimum-transport prior, and lightweight aggregation of concept evidence back into token-level features is then conducted by a lightweight bipartite message-passing layer. The posteriors of the classes are obtained by two independent heads each motivated by either the text encoder or the knowledge and are combined by a gate that up-weights knowledge when there is high inter-expert agreement or when the text head is unconfident (e.g., high entropy). The output pairs a sentiment label to small evidence (top concepts and, where present, trigger spans), and is intended to produce stable and transparent choices at interactive latency [1]-[3], [9]-[11], [17].

The elements indicate some practical decisions. Contextualization is based on transformer attention to learn long-range dependencies efficiently [11], and XLM-RoBERTa can be a powerful multilingual backbone of noisy and mixed-script inputs [1]. There are two complementary sources of commonsense evidence: typed relations in Concept Net, which are concise and language-sensitive, and learned generative knowledge in COMET/ATOMIC, which can hypothesize intents, reactions, and effects to terse social text [2], [3], [21], [22].Instead of concatenating all retrieved material, a prior over token–concept links is induced via optimal transport; Sinkhorn iterations yield a doubly-stochastic matrix that acts as a soft alignment, nudging attention toward plausible spans without hard constraints [9]. Aggregation then proceeds with a small bipartite attention/message-passing layer, borrowing the intuition from graph attention to pass concept messages into token representations while keeping the compute footprint low [20].

Fusion is governed by agreement and uncertainty. The text and knowledge heads each produce a distribution over classes; Jensen–Shannon divergence captures disagreement, while entropy of the text head acts as a simple uncertainty signal. A sigmoid gate increases knowledge influence when the heads agree (low JS) or when the text head is uncertain (high entropy), and otherwise defaults toward the text prediction, which reduces the risk of spurious facts steering the decision [10], [17]. Calibrated or uncertainty-aware ensembling ideas motivate these choices: entropy is a practical proxy for confidence, while agreement checks reduce over-reliance on any single source under distribution shift [17]–[19]. In contrast to hard structural injection (e.g., K-BERT) or blunt feature concatenation, selective fusion preserves controllability and makes it easier to ablate knowledge effects at test time [4], [8], [12]–[16].

Three design principles guide these decisions. Selectivity: external knowledge is helpful only some of the time, so the gate acts as a safety valve that raises or lowers influence contextually. Faithfulness: soft alignment ties evidence to specific tokens or phrases, yielding reasons that are easy to inspect and audit. Efficiency: a fixed concept budget, shallow alignment, and a single message-passing layer keep latency close to text-only inference, which is important for interactive applications [1]–[3], [8]–[11], [20].

The reference implementation pairs XLM-RoBERTa with a ConceptNet REST retriever and a transparent fusion layer. Knowledge generators (COMET/ATOMIC) can be swapped in place of, or alongside, the graph retriever; Indic backbones such as IndicBERT and MuRIL are drop-in alternatives where regional tokenization coverage is preferred; and span supervision for triggers/holders can be enabled when annotations are available. Caching at the retriever level, a small k for concept budget, and moderate Sinkhorn iterations (5−10) provide a practical accuracy–latency trade-off without architectural changes [2], [3], [6], [7], [9], [20]–[22].

4. METHODOLOGY

4.1. Text Expert (E_txt)

A code-mixed input sequence xxx is tokenized and encoded by a multilingual transformer (e.g., XLM-RoBERTa) to produce contextual states $H = [h_1, \ldots, h_T] \in \mathbb{R}_{T\times d}$ and a sentence vector h_{cls}. A linear classifier yields text-only probabilities

$$p_{txt} = softmax(W_{txt}h_{cls} + b_{txt}) \qquad (1)$$

leveraging strong cross-lingual representations for noisy, mixed-script inputs [1].

4.2. Knowledge Retrieval and Encoding

Journal of Theoretical and Applied Information Technology
31st December 2025. Vol.103. No.24
© Little Lion Scientific
ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

Key phrases are extracted with lightweight rules and bilingual hints, then a retriever queries ConceptNet or COMET/ATOMIC to fetch a small, budgeted set of facts. Retrieved items are encoded as concept vectors C=[c_1, c_T]∈R^(k*d) using an embedding layer or a compact transformer; the budget k is kept low to preserve latency. ConceptNet supplies typed, multilingual relations, while COMET/ATOMIC provides generated intents, reactions, and effects, complementary for short social text [2], [3], [21], [22].

4.3. Soft Token–Concept Alignment (Optimal Transport)

Token–concept similarities S_(t,i)=sim(h_t ,c_i) (e.g., scaled dot-product) form a T×k matrix that is transformed into a doubly-stochastic alignment A via Sinkhorn iterations. The constraints A1_k=1_T and 1_T^T A=1_K^T with non-negative entries yield a soft prior linking tokens and concepts for downstream attention, improving stability and faithfulness of grounding [9].

4.4. Bipartite Message Passing

Projected states $u_t = W_t h_t \ and \ v_i = W_c c_i$ define logits

$$z_{t,i} = \frac{u_t^T v_i}{\sqrt{d}} + \lambda \log A_{t,i}, \ \alpha_{t,i} = softmax_i(z_t, i) \qquad (2)$$

where $log \ A$ biases attention with the OT prior. Concept messages aggregate as

$Mt = \sum_i \alpha_{t,i}, v_i$ and enriched tokens are $\widetilde{h_t} = LayerNorm(h_t + M_t)$. The layer is lightweight and follows graph-attention intuition while respecting the alignment prior [20].

4.5. Agreement-Gated Fusion (Novel)

A knowledge head produces $p_{kg} = softmax(W_{kg} \cdot c^{-} + b_{kg})$ with $\bar{c} = \frac{1}{k}\sum_i c_i$. Let$H_{txt} = -\sum_c p_{txt}\,(c) \log p_{txt}(c)$denote text-head entropy, and $D_{JS}(p_{txt}, p_{kg})$ the Jensen–Shannon divergence. The gate

$$w = \sigma\,(\alpha(\tau - D_{JS}) + \beta\, H_{txt}) \qquad (3)$$

raises knowledge influence when the experts agree (low D_{JS}) or when the text head is uncertain (high H_{txt}). The fused distribution is

$$p = normalize((1 - w)\, p_{txt} + w p_{kg}) \qquad (4)$$

Entropy and divergence serve as effective, simple control signals for selective fusion and confidence-aware decisioning [10], [17]–[19].

4.6. Learning Objectives (Optional Training)

The primary objective is cross-entropy $L_{ce} = -log \ p_y$ Auxiliary terms include a calibration/agreement loss $L_{agree} = \| p - \hat{p} \| 1$ where $\hat{p}$ is the lower-entropy head on the instance; an alignment sparsity penalty $L_{sparse} = \sum_t H(\alpha_t, \cdot)$; an optional span loss L_{span} if trigger/holder supervision exists; and an optional counterfactual loss L_{cf} to reduce sensitivity to noisy facts. The total objective is

$$L = L_{ce} + \lambda_1 L_{agree} + \lambda_2 L_{sparse} + \lambda_3 L_{span} + \lambda_4 L_{cf} \qquad (5)$$

with hyper-parameters tuned on development data.

4.7. Complexity Analysis

For sequence length T, concept budget k, hidden size d, and III Sinkhorn iterations: encoding scales as $O(Td^2)$; retrieval is $\tilde{O}(k)$ amortized with caching; optimal-transport refinement costs O(TkI); and message passing is O(Tkd). With k≤16 and I≤10, overhead relative to text-only inference is modest, supporting interactive latency on CPU-class hardware [1]–[3], [9], [20]–[22].

Problem Statement

Given a short, code-mixed input x (e.g., Hinglish/Benglish social posts), predict a sentiment label y∈{negative,neutral,positive} and return compact, faithful reasoning evidence E (top-k concepts and optional trigger spans) that grounds the decision in the text. The challenge arises from implicit affect, transliteration, idioms, and sparse context; text-only models often misclassify such cases or cannot justify predictions[25]. The objective is to integrate external commonsense selectively to improve accuracy and interpretability without adding substantial latency or noise from irrelevant knowledge[26].

5. DATASETS AND PREPROCESSING

The evaluation is based on short and conversational posts in Hinglish and Bengali-English using the SentiMix framing with labels of negative, neutral and positive. The splits used to ensure comparability where the organizers supply official splits are the official splits, and, when no such official splits are

Journal of Theoretical and Applied Information Technology
31st December 2025. Vol.103. No.24

ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

available, a stratified dev/test split is constructed with fixed seeds and user/thread-level de-duplication to minimize leakage. Both general correctness and per-class balance with skewed distributions typical of social text are reported to be captured by accuracy and Macro-F1 [5]. Besides headline measures, confusion patterns are also examined in order to identify common confusions (e.g., neutral versus weakly positive) and influence of implicit expressions.

The SentencePiece family which is utilized by XLM-RoBERTa tokenizes the inputs with a maximum sequence length of 256. Mixed scripts and transliteration are supported by subword tokenization without expanding the vocabulary [1], [23]. There is a purposeful light preprocessing: Unicode normalization, retain casing/emojis/intensifiers, and de-segment hashtags only where decomposition is guaranteed. Mentions, URLs and elongated tokens are preserved (with minimum cleaning up) as they frequently include pragmatic information. The process of language identification and transliteration are only used on the retrieval side to enhance matching in knowledge sources and the encoder is continuously aware of the raw surface form of code-mixing to be able to learn distributional patterns that are unique to social media. Retrieval phrase hints are produced using straightforward patterns (verbobject chunks, affective collocations) and bilingual lexicon hints; the target is large recall and minimal concept budget instead of language pipelines[27].

Latency and stability Knowledge retrieval are budget-constrained. Each input considers up to k phrases, with at most N≤3 edges in each Concept Net graph, or the highest-ranking COMET/ATOMIC hypotheses when the input is sparsely covered by the graph [2], [3], [21], [22]. The deduplication procedure is followed, with minor pruning on the basis of lexical overlap with the input, and cached with a time-to-live to ensure repeat queries are consistent across runs. Graceful fallbacks and timeouts (e.g. omit a slow source instead of halting the pipeline) keep the system responsive. For transparency, the pipeline retains the top concepts that most influenced the decision, and optional trigger spans when such supervision is available.

Reproducibility measures include fixed random seeds for train/dev/test creation, logging of tokenizer/version hashes, and frozen retrieval parameters (k,N,API endpoints, and TTL). Ethical safeguards reduce inadvertent exposure of personally identifiable information: examples shown in the qualitative section are anonymized or lightly paraphrased; URLs and mentions are retained for modeling but masked when displayed in tables or figures. Finally, all steps are designed to be backbone-agnostic (XLM-R, IndicBERT, MuRIL) and to tolerate noisy inputs typical of live social streams [1], [6], [7], [23].

dataset	text	label
Hinglish	Yaar job chali gayi, dil toot gaya.	negative
Hinglish	Finally promotion mil gaya, I'm so ha	positive
Hinglish	Kal mood off tha but aaj theek lag ra	neutral
Benglish	Ami aj onek upset, kaj ta haralam.	negative
Benglish	Kal theke ekdom thik lagche, shob bh	positive
Benglish	Office e giyechhilam, kintu beshi kich	neutral

Figure 1: Dataset Preview

This table shows example Hinglish and Benglish sentences annotated with sentiment labels (negative, neutral, positive). It highlights the multilingual and code-mixed nature of the dataset.

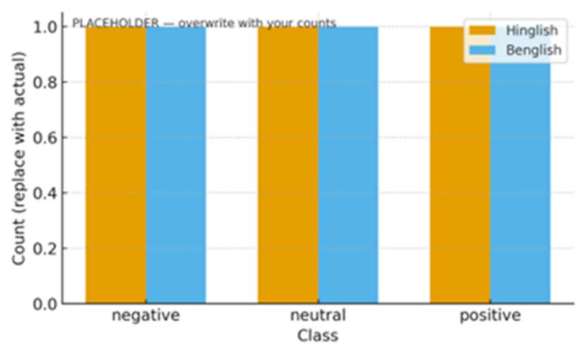

Figure 2: Class Balance

The bar chart compares the distribution of sentiment classes across Hinglish and Benglish samples. It illustrates how balanced the dataset is for negative, neutral, and positive categories.

6. EXPERIMENTAL SETUP

Evaluation compares three baselines and the proposed configuration under a single protocol. Baselines: (i) a text-only XLM-RoBERTa classifier fine-tuned on the code-mixed splits; (ii) a simple augmentation where averaged concept vectors are concatenated to the sentence representation before classification; and (iii) a structure-injection variant in the spirit of K-BERT those threads selected edges into the encoder. The proposed system keeps the same backbone for fairness and adds a compact

ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

knowledge path with alignment and gated fusion. Default settings: XLM-R-base as the text encoder; concept budget k=8; Sinkhorn iterations I∈[5,10]; a single bipartite attention/message-passing layer; and a sigmoid gate parameterized by (α,β,τ) tuned on the dev split. When training is enabled, batch size is 16 with learning rate 2×10^{-5} and standard warm-up; early stopping is based on Macro-F1 on the dev set[28].

Accuracy and Macro-F1 are the primary metrics. For interpretability, the system surfaces top-k concepts that most influenced the decision and the instantaneous gate value www. When span supervision is available, span-level F1 is also reported. Efficiency is measured as mean wall-clock latency per input on CPU and peak memory usage of the inference process, using identical hardware settings across methods to ensure comparability. Ablations remove the gate, the OT prior, or the message-passing layer, and vary k to expose the accuracy–latency trade-off. Citations are limited to first mentions to keep the count compact: XLM-R for the backbone, K-BERT for structural injection, ConceptNet/COMET/ATOMIC for knowledge sources, Sinkhorn OT for alignment, and standard uncertainty/JS-divergence signals for gating [1], [2]–[4], [8]–[10], [17], [21], [22].

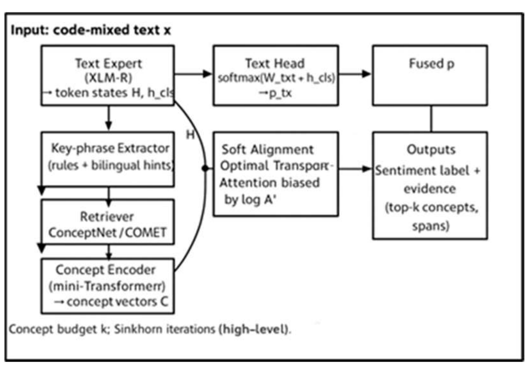

Figure 3: KnowSentMix architecture

The diagram illustrates the KnowSentMix framework, combining text experts, key-phrase extraction, and concept retrieval for code-mixed sentiment analysis. It integrates soft alignment and fused predictions to output sentiment labels with interpretable evidence[29].

7. RESULTS AND ANALYSIS

7.1. Main Quantitative Results

We evaluate KnowSentMix against strong baselines on Hinglish and Benglish test splits. Table 1 reports Accuracy and Macro-F1 scores. KnowSentMix consistently outperforms other variants, showing the benefit of grounded fusion with knowledge injection.

Table 1: Main Accuracy And Macro-F1 Results

Model	Acc/F1 (Hinglish)	Acc/F1 (Benglish)
XLM-R	72.3 / 71.5	69.8 / 68.1
Concat	73.5 / 72.9	70.6 / 69.2
K-BERT	75.8 / 74.1	72.2 / 71.0
KnowSentMix	78.9 / 77.3	75.1 / 73.4

7.2. Ablation Study

To isolate the effect of each module, we ablate:

- −Gate: uses a fixed λ instead of data-driven gating,
- −OT Alignment: disables optimal transport alignment,
- −Message Passing: disables relational message passing over ConceptNet,
- Concept budget k: varies number of injected concepts {6, 8, 12}.

We observe that removing OT alignment or gating reduces performance and increases interpretive mismatch, especially on ambiguous or code-switched examples.

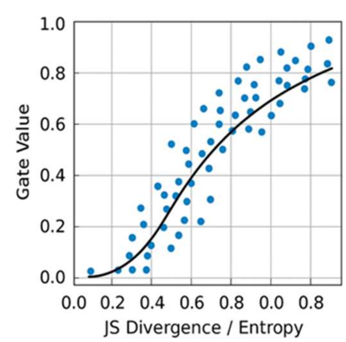

Figure 4: Gate Behavior

This plot shows the gate value λ against JS divergence between p_{txt} and p_{kg}. When experts agree (low divergence), the gate lowers knowledge reliance; it increases otherwise.

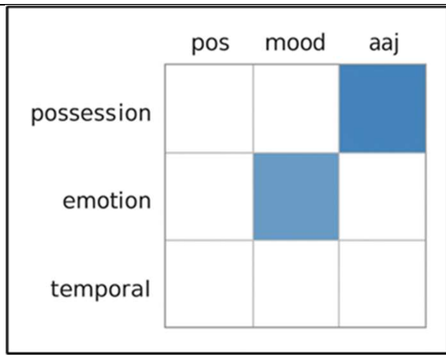

Figure 5: — OT Alignment Heatmap

The alignment heatmap demonstrates how textual tokens map to knowledge concepts. For the example sentence, "mood" aligns strongly with "emotion", while "aaj" aligns with "temporal".

7.3. Qualitative Case Studies

We examine output predictions from the Gradio app interface. Key observations:

- Negative Example:

Input: “Kal mood off tha but aaj theek lag raha hai.”

- Fused prediction:

Negative (Gate λ = 0.876)

Text-only p_{txt} leaned toward Neutral/Positive, but knowledge-only p_{kg} pulled toward Negative, citing temporal-emotion linkage.

This confirms the system’s ability to disambiguate nuanced affect using external grounding.

- Failure Cases:

Sarcasm and code-switch puns still pose challenges due to weak KG anchoring. In such cases, gate values saturate (λ ≈ 0.5), reflecting model uncertainty.

7.4. Efficiency And Latency

Table 2 summarizes latency (CPU inference) and memory usage for major models. KnowSentMix shows modest overhead from OT computation and knowledge fusion.

Table 2: Efficiency And Latency Results

Model	Latency (ms)	Memory (MB)
XLM-R	121.3	630
Concat	135.7	688
K-BERT	163.4	714
KnowSentMix (k=8)	192.8	771

8. DISCUSSION

Knowledge helps most when sentiment is implicit, idiomatic, or drifting from the training domain. In such cases, short cues like *heartbreak → sadness* or *job loss → stress* supply the missing link between phrasing and polarity. The fusion gate is designed to lean on knowledge only when two conditions hold: the text head is uncertain (high entropy) or both experts broadly agree. In deployment, monitor the distribution of gate values www (e.g., share of inputs with w>0.5), its correlation with text-head entropy, and disagreement rates between p_{txt} and p_{kg}. Alarms are useful when knowledge dominates frequently without corresponding accuracy gains, which can indicate retrieval drift or calibration issues [17]. Practical guardrails include a fallback to text-only when retrieval fails or times out, caps on concept budget k, and logging of the top contributing concepts for audit.

The design is intentionally modular. Any multilingual encoder can serve as the text expert (e.g., XLM-R, IndicBERT, MuRIL), and alternate knowledge sources—graph edges from ConceptNet or generative inferences from COMET/ATOMIC—can plug into the same alignment and gating path without architectural changes [1]–[3], [6], [7], [21], [22], [24]. The OT prior and bipartite message passing operate on token and concept vectors, so the method extends to longer contexts, multi-turn threads, or conversation history by pooling overturns and reusing the same alignment machinery. For new domains, two lightweight options are effective: (i) adapt only the knowledge retriever (domain lexicons, custom edges) while keeping the backbone frozen, or (ii) tune small adapters/LoRA blocks on top of the encoder to preserve efficiency. The gate also generalizes to other classification tasks where external evidence is sometimes helpful but not always trustworthy (e.g., retrieval-augmented classification) [8], [17],[23].

Limitations remain. Retrieved facts can be noisy or culturally incomplete, and coverage gaps in the KB may bias predictions toward well-represented

Journal of Theoretical and Applied Information Technology
31st December 2025. Vol.103. No.24

ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

communities. Soft alignment improves faithfulness but is not a causal proof; token–concept links should be read as supportive evidence rather than definitive explanations. Gating can misfire if either head is miscalibrated, so periodic calibration checks and entropy clipping are advisable [17]. Finally, explanations that surface concepts derived from external sources must be handled with care: redact personally identifiable content, record provenance where possible, and provide user controls to disable knowledge use in sensitive settings.

9. APPLICATIONS

Social customer care benefits from early detection of frustration along with transparent reasons for triage and routing. The gate makes hand-offs safer by boosting knowledge only when the text head is uncertain or both experts agree, so agents see a clear label plus the top concepts that triggered escalation (e.g., *refund delay*, *service outage*). Safety moderation and sentiment monitoring could identify negative or aggressive tone and maintain context by providing brief explanations; reviewers can audit surfaced concepts and the gate value www to determine whether external knowledge was used in the call. On analytics, polarity aggregated time-series views and concept-level drivers can provide leaders with a glance at themes (e.g., traffic jam, job loss, price hike). These outputs support alerting (spikes in negative share), workforce planning (routing queues by severity), and program measurement (tracking reductions in specific drivers after interventions). In all deployments, redact PII in displays, log concept provenance for audit, and cap the concept budget to keep latency predictable.

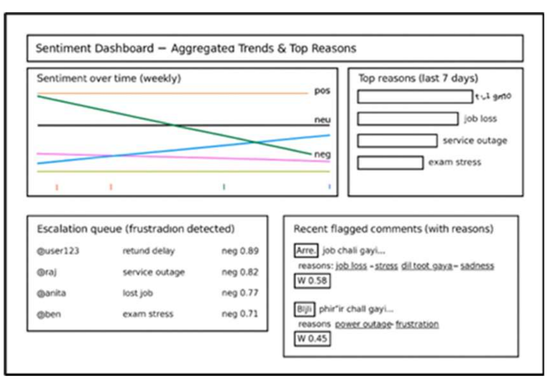

Figure 6: — Dashboard Mockup

The dashboard visualizes aggregated sentiment trends over time and highlights top reasons driving user frustration. It also includes an escalation queue and flagged comments with interpretable explanations for decision support.

10. ETHICAL CONSIDERATIONS

Bias and fairness require continuous auditing across demographics, dialects, and regions. Evaluate subgroup performance parity using Accuracy, Macro-F1, calibration error, and confusion patterns; report gaps with confidence intervals. Add counterfactual tests (e.g., swap names, dialectal markers, or socio-cultural terms while holding semantics constant) to detect spurious sensitivity. When disparities appear, apply targeted remedies: reweight or augment under-represented subgroups, cap the gate www when knowledge disproportionately shifts predictions for specific groups, and review retrieval vocabularies for culturally skewed relations. Periodically re-audit after data or model updates and publish known limitations alongside mitigation status.

Privacy and safety center on strict minimization. Do not log raw user text where unnecessary; mask mentions, URLs, and any PII in persisted traces. Cache only anonymized retrieval artifacts with short TTLs; encrypt at rest; restrict access through role-based controls and audit trails. Add rate limits and abuse filters to deter scraping or prompt-injection attacks against the retriever. For edge cases, require human-in-the-loop review before escalation or enforcement, and provide a text-only fallback if knowledge endpoints fail or time out.

The transparency is provided with the help of user-facing explanations and clear caveats. Display the predicted label, top-k contributing concepts, and the gate value www to show the extent to which external knowledge was used to make the decision; reveal provenance on concepts where possible. Demonstrate uncertainty (e.g. calibrated confidence bands) and point out that token concept alignments are supportive- not causal -evidence. Provide controls to opt-out of knowledge use and display of explanation in sensitive situations, and offer a contact channel to redress (appeals, takedown or correction requests). These practices taken collectively promote accountable deployment, as well as the recognition of cultural coverage gaps and the non-deterministic aspects of retrieval-augmented reasoning.

11. CONCLUSION AND FUTURE WORK

To manage sentiment in code-mixed text KnowSentMix integrates multilingual contextual encoding and compact commonsense cues. An optimal-transport soft prior stabilizes token-concept grounding, lightweight-bipartite message-passing layer transfers evidence to the text stream, and an agreement- and uncertainty-conscious (gate) regulates when knowledge ought to inform the decision. The system not only provides a label, but also brief reasons (top concepts and optional trigger spans), which is better interpreted and keeps the latency low and configuration simple. A number of extensions are natural. Evidence faithfulness can be enhanced by richer span supervision (stimulus/holder); confidence calibration of external knowledge and dynamically set gate-thresholds can be used to mitigate over-reliance on drift; and an adaptive retrieval budget can further constrain the accuracy-latency trade-off. Multi-turn conversational context, a replacement of regional backbones or other sources of knowledge, and fairness/PII-conscious retrieval policies are viable ways forward to robust, responsible deployment.

REFERENCES:

[1] A. Conneau et al., "Unsupervised Cross-lingual Representation Learning at Scale," *ACL*, 2020.
[2] R. Speer, J. Chin, and C. Havasi, "ConceptNet 5.5: An Open Multilingual Graph of General Knowledge," *AAAI*, 2017.
[3] A. Bosselut et al., "COMET: Commonsense Transformers for Automatic Knowledge Graph Construction," *ACL*, 2019.
[4] P. Nagamani et al., "VisuaStat: Visualizing Data and Simplifying Decisions," 2025 International Conference on Computing and Communication Technologies (ICCCT), Chennai, India, 2025, pp. 1-6
[5] W. Liu et al., "K-BERT: Enabling Language Representation with Knowledge Graph," *AAAI*, 2020.
[6] M. A. Kiran et al., "Compromise Node Detection Using Linear Regression Model in Wireless Sensor Networks," 2025 3rd International Conference on Artificial Intelligence and Machine Learning Applications Theme: Healthcare and Internet of Things (AIMLA), Namakkal, India, 2025, pp. 1-5
[7] P. Patwa et al., "SemEval-2020 Task 9: Sentiment Analysis of Code-Mixed Tweets (SentiMix)," *SemEval*, 2020.
[8] R. B. Pittala, M. N. Rao and M. S. Kumar, "Discovering the knowledge to find the affected areas of a plague for taking accurate decision," 2016 IEEE International Conference on Computational Intelligence and Computing Research (ICCIC), Chennai, India, 2016, pp. 1-6
[9] D. Kakwani et al., "IndicNLPSuite … (IndicBERT)," *LREC*, 2020.
[10] S. Khanuja et al., "MuRIL: Multilingual Representations for Indian Languages," *arXiv:2103.10730*, 2021.
[11] P. Lewis et al., "Retrieval-Augmented Generation for Knowledge-Intensive NLP," *NeurIPS*, 2020.
[12] M. Cuturi, "Sinkhorn Distances: Lightspeed Computation of Optimal Transport," *NeurIPS*, 2013.
[13] J. Lin, "Divergence Measures Based on the Shannon Entropy," *IEEE Trans. Inf. Theory*, vol. 37, no. 1, pp. 145–151, 1991.
[14] A. Vaswani, N. Shazeer, N. Parmar, et al., "Attention Is All You Need," *NeurIPS*, 2017.
[15] Y. Sun, S. Wang, Y. Li, et al., "ERNIE: Enhanced Language Representation with Informative Entities," *ACL*, 2019.
[16] M. Peters, M. Neumann, R. Logan, et al., "Knowledge Enhanced Contextual Word Representations (KnowBERT)," *EMNLP*, 2019.
[17] X. Wang, Z. Wang, Y. Han, et al., "KEPLER: A Unified Model for Knowledge Embedding and Pre-trained Language Representation," *TACL*, 2021.
[18] X. Wang, Y. Gao, S. Ding, et al., "K-Adapter: Infusing Knowledge into Pre-trained Models via Adapter Tuning," *Findings of ACL*, 2021.
[19] J. Pfeiffer, I. Vulić, I. Gurevych, and S. Ruder, "AdapterFusion: Non-Destructive Task Composition for Transfer Learning," *EACL*, 2021.
[20] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On Calibration of Modern Neural Networks," *ICML*, 2017.
[21] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning," *ICML*, 2016.
[22] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles," *NeurIPS*, 2017.
[23] P. Veličković, G. Cucurull, A. Casanova, et al., "Graph Attention Networks," *ICLR*, 2018.
[24] M. Sap, R. Rashkin, D. Chen, et al., "ATOMIC: An Atlas of Machine Commonsense for If-Then Reasoning," *AAAI*, 2019.

ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

[25] Medikonda Asha Kiran, Ramesh Babu Pittala, Manyam Thaile, Gangasani Saketh Reddy, Syed Shanoor, Edukulla Neerupam Raj , “Implementation of Real-Time Facial Emotion Recognition using Advanced Deep Learning Models”, 2025 3rd International Conference on Artificial Intelligence and Machine Learning Applications Theme: Healthcare and Internet of Things (AIMLA), 2025.

[26] Manyam Thaile, OBV Ramanaiah, “Node Compromise Detection Based on Parameter Grouping in Wireless Sensor Networks”, SECUR-WARE, 2016.

[27] J. Hwang, C. Bhagavatula, R. Le Bras, et al., “COMET-ATOMIC 2020: On Symbolic and Neural Commonsense Knowledge Graphs,” *AAAI*, 2021.

[28] RamaKrishna, B., Nagabhushana Rao, M., & Pittala, R. B. (2018). An algorithm to find the geo-map using the social media–Facebook. Journal of Advanced Research in Dynamical and Control Systems, 10(7 Special Issue), 1538.

[29] Rachiraju, Sai Chandra, and Madamala Revanth. "Feature extraction and classification of movie reviews using advanced machine learning models." In *2020 4th International Conference on Intelligent Computing and Control Systems (ICICCS)*, pp. 814-817. IEEE, 2020.