

DEEPTANS STUDIO: Turning Expert Interventions into Shared Team Knowledge in Agentic Translation Workflows

ZIYANG LIAN, QINGYA ZHANG, HAO WANG*, HUIWEN XIONG, and QI YANG, School of Computer Engineering and Science, Shanghai University, China
 LINGYI MENG, School of Foreign Languages, East China Normal University, China
 XIAOYI GU, School of Foreign Studies, Shanghai University, China
 RUI WANG, Department of Computer Science and Engineering, Shanghai Jiao Tong University, China

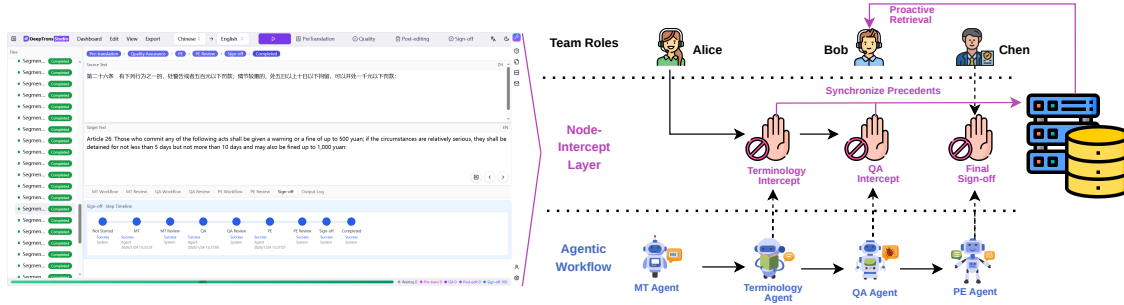


Fig. 1. Overview of DeepTrans Studio. Expert decisions made at terminology, QA, and sign-off intercepts are saved to shared team memory and later retrieved as precedents in teammate workflows. The demo lets attendees experience this flow by correcting a preset legal risk and observing its reuse in another workspace.

Professional translation is often a team-based process: translators, reviewers, and project managers must coordinate terminology, legal force, and accountability across documents. Yet many LLM-based translation tools treat human corrections as isolated edits. Expert decisions made in one segment or by one member are rarely captured as reusable knowledge for the rest of the team. We present DEEPTANS STUDIO, a collaborative translation workspace that lets professionals intercept selected nodes in an agentic translation workflow, review evidence, revise AI outputs, and save approved decisions to a shared team memory. During the demo, attendees will role-play translators and reviewers, resolve preset terminology and legal-modal risks, and see how their decisions are propagated to downstream segments and surfaced in a teammate’s workspace as reusable precedents. The demo illustrates how human interventions

*Corresponding author.

Authors’ Contact Information: Ziyang Lian, 1257115202@shu.edu.cn; Qingya Zhang, zhangqingya@shu.edu.cn; Hao Wang, wang-hao@shu.edu.cn; Huiwen Xiong, hwxiong@shu.edu.cn; Qi Yang, yang-qi@shu.edu.cn, School of Computer Engineering and Science, Shanghai University, Shanghai, China; Lingyi Meng, lingyi@ecnu.edu.cn, School of Foreign Languages, East China Normal University, Shanghai, China; Xiaoyi Gu, xiaoyi-gu@shu.edu.cn, School of Foreign Studies, Shanghai University, Shanghai, China; Rui Wang, wangrui12@sjtu.edu.cn, Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai, China.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.

Manuscript submitted to ACM

Manuscript submitted to ACM

in AI-mediated work can become shared, traceable knowledge rather than one-off corrections. one-off corrections. Code and demo video: <https://github.com/hint-lab/deeptrans-studio>, <https://youtu.be/cNpafhHAEjg>.

CCS Concepts: • **Human-centered computing** → **Interactive systems and tools**; *Collaborative and social computing systems and tools*; • **Computing methodologies** → *Machine translation*.

Additional Key Words and Phrases: Human-AI collaboration, collaborative translation, agentic workflows, team memory, accountability

ACM Reference Format:

Ziyang Lian, Qingya Zhang, Hao Wang, Huiwen Xiong, Qi Yang, Lingyi Meng, Xiaoyi Gu, and Rui Wang. 2026. DEEPTRANS STUDIO: Turning Expert Interventions into Shared Team Knowledge in Agentic Translation Workflows. In *Proceedings of Companion of the 2026 Computer-Supported Cooperative Work and Social Computing (CSCW '26)*. ACM, New York, NY, USA, 4 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 Introduction

Professional translation in high-stakes domains is collaborative work. Translators, reviewers, project managers, and clients must coordinate terminology, legal force, and accountability across people and documents. The central challenge is therefore not only to produce fluent text, but to make translation decisions consistent, visible, and reusable within a team [4, 5].

LLM-based translation tools make this coordination harder when they are designed as end-to-end pipelines or one-person chat interfaces [1, 6]. Professionals may correct AI-generated drafts or multi-agent outputs [3, 9], but these corrections often stay inside a local session. As a result, teams repeatedly resolve the same terminology conflicts and legal-modal ambiguities without a shared record, leading to terminological drift, duplicated review work, and accountability gaps [8]. DEEPTRANS STUDIO addresses this by making expert interventions **visible, reusable, and auditable across professional roles**.

To achieve this, the system exposes selected agentic workflow nodes for human interception, balancing automation with expert agency [7]. At terminology, QA, and sign-off points, experts can inspect evidence, revise AI suggestions, and record the rationale for their decisions[2]. Approved interventions are then synchronized into a living team memory, where they guide downstream segments and future work by other members. In this way, the system shifts AI-assisted translation from isolated editing toward shared, traceable team knowledge building.

Our core contributions as a CSCW Demo include:

- An interactive demonstration of **human-interceptable agentic workflows**, allowing professionals to pause selected steps and revise outputs before downstream propagation.
- A **shared team memory mechanism** that turns individual expert interventions into reusable precedents across segments and team members.
- A **role-play demo scenario** showing how translators, reviewers, and project managers coordinate AI-mediated decisions in a high-accountability workflow.

2 System Design: Collaborative Translation Infrastructure

To support team coordination, DEEPTRANS STUDIO implements a stateful dual-loop framework that combines background automation with human-supervised decision points. The system moves beyond “black-box” pipelines through four collaborative modules:

Manuscript submitted to ACM

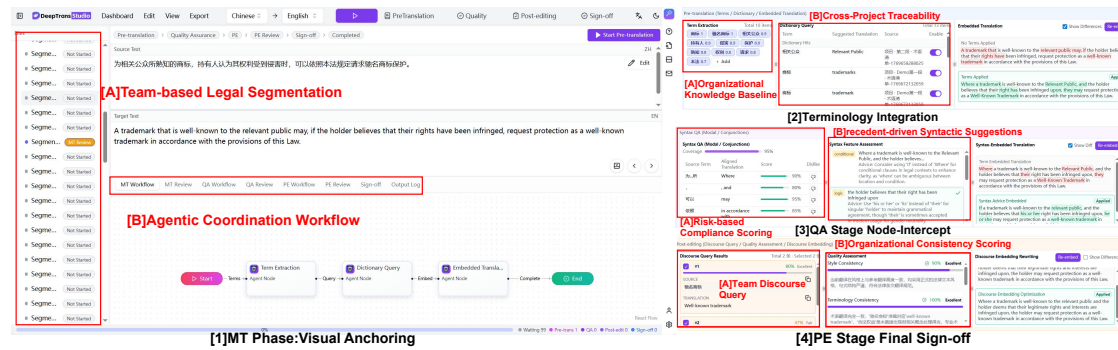


Fig. 2. From individual intervention to team-wide propagation: a translator intercepts an AI decision, reviews evidence, approves a terminology choice, and the system reuses the decision in downstream segments and teammate workflows.

Shared Translation Workspace: Translators, reviewers, and managers work with synchronized document states (e.g., intercepted, signed-off), establishing workflow awareness without reconstructing decisions from local chat histories.

Node-Intercept Layer: Instead of autonomous completion, the pipeline pauses at high-risk nodes (e.g., terminology conflicts, legal-modal ambiguities) to expose intermediate outputs for human review before downstream propagation.

Shared Team Memory: Expert interventions are synchronized into a “Living Dictionary.” These approved decisions become reusable precedents retrieved for teammate workflows, reducing repeated edits and terminological drift.

Accountability Trace: The system logs the provenance of model suggestions, human edits, and reused precedents. This trace allows reviewers to inspect decision chains and verify accountability across roles.

3 DEEPTANS STUDIO in Action: A Collaborative Walkthrough

To illustrate the collaborative workflow, we use a scenario involving a senior translator (Alice), a junior translator (Bob), and a project manager (Chen).

1. MT Phase and Visual Anchoring: Alice uploads a contract. The system preserves key structural cues and initializes a multi-panel workspace (Fig. 2(a)). This helps the team maintain a shared view of nested legal conditions rather than treating translation segments as isolated sentences.

2. Terminology Intercept and Memory Sync: Alice reviews inconsistent terminology through the alignment interface (Fig. 2(b)). After she approves a term, the decision and its rationale are synchronized into the “Living Dictionary,” turning a local judgment into shared team memory.

3. Modal Intercept and Accountability: The system triggers a *node-intercept* at a legally ambiguous modal phrase, such as *shall* versus *may* (Fig. 2(c)). Alice inspects the surfaced evidence, revises the AI output, and records the decision before it propagates to downstream segments.

4. Team Reuse and Final Sign-off: When Bob later encounters a similar clause, DEEPTANS STUDIO retrieves Alice’s precedent and displays it in his workspace, reducing duplicated research and aligning team decisions. Chen then uses the dashboard to inspect the **Accountability Trace** and verify the decision chain before final sign-off (Fig. 2(d)).

4 Demo Setup and Preliminary Feedback

Interactive Demo: At CSCW 2026, attendees will take part in a 5-minute role-play. First, they choose one of two preset risks in a legal contract: a terminology conflict or a legal-modal ambiguity. Next, DEEPTANS STUDIO triggers a *node-intercept*, and attendees act as reviewers by approving, revising, or rejecting the AI proposal. The system then switches to a teammate view and shows how the approved intervention is retrieved as a precedent for a later segment. This setup demonstrates that attendees are not merely correcting an AI output; they are shaping how the team handles similar decisions afterward.

Preliminary Feedback: Formative walkthroughs with 12 professionals suggest that DEEPTANS STUDIO may help shift corrections from isolated edits to reusable team decisions. Participants especially valued the ability to inspect why a segment was interrupted, record the rationale for a correction, and make that correction visible to teammates. One senior translator (P7) noted: “I no longer just clean machine messes; my intercepts now anchor consistency for all members, allowing individual expertise to act as a shared knowledge asset.”

5 Conclusion and Future Work

DEEPTANS STUDIO demonstrates how expert interventions in agentic translation workflows can be recorded as shared, traceable, and reusable team knowledge. The CSCW demo will let attendees experience how node-level intervention, team memory, and accountability tracing can support collaborative work around AI-generated outputs. Future work will extend this infrastructure to other high-stakes domains where expert judgment must be coordinated across teams.

Acknowledgments

The authors used Gemini to improve the grammar and clarity of author-written text. All system design, evaluation descriptions, and claims were written and verified by the authors, who reviewed and edited all generated suggestions and take full responsibility for the final content.

References

- [1] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N Bennett, Kori Inkpen, et al. 2019. Guidelines for human-AI interaction. In *Proceedings of the 2019 chi conference on human factors in computing systems*. 1–13.
- [2] Michael S Bernstein, Greg Little, Robert C Miller, Björn Hartmann, Mark S Ackerman, David R Karger, David Crowell, and Katrina Panovich. 2010. Soylent: a word processor with a crowd inside. In *Proceedings of the 23rd annual ACM symposium on User interface software and technology*. 313–322.
- [3] Eleftheria Briakou, Jiaming Luo, Colin Cherry, and Markus Freitag. 2024. Translating step-by-step: Decomposing the translation process for improved translation quality of long-form texts. In *Proceedings of the Ninth Conference on Machine Translation*. 1301–1317.
- [4] Paul Dourish and Victoria Bellotti. 1992. Awareness and coordination in shared workspaces. In *Proceedings of the 1992 ACM conference on Computer-supported cooperative work*. 107–114.
- [5] Markus Freitag, George Foster, David Grangier, Viresh Ratnakar, Qijun Tan, and Wolfgang Macherey. 2021. Experts, errors, and context: A large-scale study of human evaluation for machine translation. *Transactions of the Association for Computational Linguistics* 9 (2021), 1460–1474.
- [6] Zhiwei He, Tian Liang, Wenxiang Jiao, Zhuosheng Zhang, Yujiu Yang, Rui Wang, Zhaopeng Tu, Shuming Shi, and Xing Wang. 2024. Exploring human-like translation strategy with large language models. *Transactions of the Association for Computational Linguistics* 12 (2024), 229–246.
- [7] Jeffrey Heer. 2019. Agency plus automation: Designing artificial intelligence into interactive systems. *Proceedings of the National Academy of Sciences* 116, 6 (2019), 1844–1850.
- [8] Katherine Thai, Marzena Karpinska, Kalpesh Krishna, Bill Ray, Moira Inghilleri, John Wieting, and Mohit Iyyer. 2022. Exploring document-level literary machine translation with parallel paragraphs from world literature. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. 9882–9902.
- [9] Minghao Wu, Jiahao Xu, and Longyue Wang. 2024. TransAgents: Build your translation company with language agents. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. 131–141.