

# CAT-GPT: A Skopos-Driven, LLM-Based Computer-Assisted Translation Tool

**Paşa Abdullah Bayramoğlu**

Üsküdar University / Istanbul

pasa.bayramoglu@uskudar.edu.tr

## Abstract

This paper introduces CAT-GPT, an innovative Computer-Assisted Translation (CAT) tool designed to address context-awareness and terminological consistency challenges often encountered in standard CAT workflows. Grounded in Skopos theory and powered by a Large Language Model (LLM) backend, CAT-GPT integrates context-sensitive segmentation, automatically generated and adjustable translation instructions, and an advanced machine translation component. Comparative observations with a widely used CAT tool (RWS Trados Studio) suggest that CAT-GPT reduces post-editing effort and improves text-level coherence, especially in specialized or domain-specific scenarios.

## 1 Introduction

CAT tools form a cornerstone of modern translation workflows, providing features such as translation memories, terminology management, and built-in machine translation (O'Brien et al., 2017). However, many of these systems rely on sentence-level segmentation without robust methods for maintaining broader context (Läubli et al., 2020). Consequently, translations can become fragmented, leading to increased post-editing and potential inconsistencies in specialized content (Kappus & Ehrensberger-Dow, 2020). Furthermore, texts requiring high terminological precision and clear functional alignment—such as legal or technical documentation—can suffer when each sentence is treated in isolation.

To address these gaps, I present CAT-GPT, a tool that combines GPT-4o with context-sensitive segmentation and user-defined instructions

grounded in Skopos theory. By allowing translators to specify functional goals and revise guidelines throughout the process, CAT-GPT aligns the final product with the intended communicative purpose (Vermeer, 2014).

## 2 Product Description

### 2.1 Key Features

CAT-GPT employs context-sensitive segmentation that determines segment boundaries by analyzing linguistic structure, discursive flow, and paragraph-level cues rather than relying solely on sentence breaks. This design ensures that long or syntactically dense sentences—commonly found in highly regulated documents—remain coherent, minimizing the risk of fragmenting essential information (Läubli et al., 2020). A specialized prompt-based routine merges semicolon-ended lines, preserves bullet-list integrity, and avoids superficial breaks, reflecting the document's actual structure and communicative logic.

Before translation begins, the system automatically generates a set of translation instructions incorporating user preferences on style, terminology, and overall communicative goals. Crucially, these instructions remain active throughout the workflow, so if the translator later updates stylistic or terminological choices, subsequent segments are re-translated accordingly. Rooted in Skopos theory (Vermeer, 2014), the instructions can be updated at any point, giving translators the flexibility to adjust as project needs evolve.

Once the instructions are finalized, an LLM-based engine (GPT-4o) references them to provide on-demand machine translation suggestions. This approach allows domain-specific terminology to be applied consistently from one segment to another, a common challenge in texts where certain expressions, roles, or designations recur. By continuously aligning the LLM's output with both

the text’s purpose and the user’s evolving instructions, CAT-GPT aims to reduce repetitive corrections, streamline the revision process, and produce more coherent target texts (Vieira et al., 2023).

## 2.2 User Interface

Developed in PyQt5, the CAT-GPT interface presents source–target segments with real-time status indicators (e.g., “Not Translated,” “Draft,” “Approved”) and machine translation suggestions. Figure 1 shows a partial view of the editor, where translators can merge or split segments, edit or refine their instructions, and track overall progress.

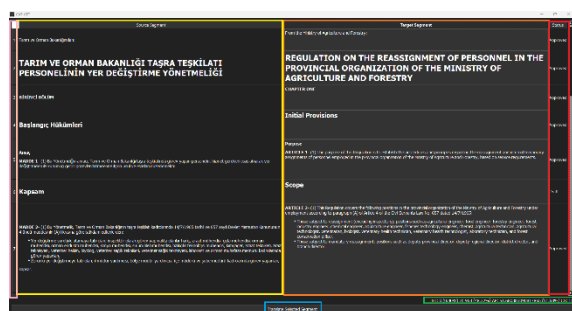


Figure 1: A partial screenshot of CAT-GPT’s editor. Segment numbers are highlighted in pink, source text segments are marked in yellow, target text segments are indicated in orange, segment statuses are displayed in red, the ‘Translate Segment’ button is marked in blue, and the character count for each segment’s status is shown in green.

## 2.3 Pricing, Licensing, and Availability

CAT-GPT will be released as open-source software on GitHub once development is complete. The tool itself incurs no license fee or subscription, and users only pay for GPT-4o API usage directly to OpenAI on a “pay as you go” basis.

## 3 Comparison with Existing CAT Tools

A mainstream sentence-based CAT tool (RWS Trados Studio 2022) was observed alongside CAT-GPT for a project that contained repeated references to specific articles and multiple official titles across several paragraphs. The sentence-based system often broke up closely related items, forcing minor inconsistencies to accumulate whenever a role or article name recurred (Kappus & Ehrensberger-Dow, 2020; O’Brien et al., 2017). For instance, designations might appear in slightly varied translations across different sentences, requiring corrections each time.

CAT-GPT’s paragraph-level segmentation, by contrast, preserved the logical structure of these references, enabling them to be rendered consistently each time they appeared. Once the translator or reviewer introduced updated guidelines—for example, a new stylistic approach to referencing articles—CAT-GPT immediately integrated these changes into subsequent machine translation output. As a result, the text maintained a uniform presentation of repeated terms from one paragraph to another, reducing editing passes and aligning with the text’s overall communicative objectives (Läubli et al., 2020). Future work will expand testing against other LLM-based CAT tools.

## 4 Conclusion

By merging GPT-4o with context-sensitive segmentation and Skopos-focused instructions, CAT-GPT addresses key deficiencies in conventional CAT workflows. Early outcomes suggest that it reduces post-editing demands, enhances terminological consistency, and preserves the text’s communicative purpose. Future development plans include scaling up to larger projects, expanding language support, and refining the interface to meet varied professional and academic needs.

## References

- Kappus, M., & Ehrensberger-Dow, M. (2020). The ergonomics of translation tools: Understanding when less is actually more. *The Interpreter and Translator Trainer*, 14(4), 386–404.
- Läubli, S., Simianer, P., Wuebker, J., Kovacs, G., Sennrich, R., & Green, S. (2020). The impact of text presentation on translator performance. *Target*.
- O’Brien, S., Ehrensberger-Dow, M., Connolly, M., & Hasler, M. (2017). Irritating CAT tool features that matter to translators. *HERMES–Journal of Language and Communication in Business*, 56, 145–162.
- Vermeer, H. J. (2014). The priority of purpose (Skopos theory). In *Towards a general theory of translational action* (pp. 85–94). Routledge.
- Vieira, N. R., Zelenka, N. R., Youdale, R. L., Zhang, X., & Carl, M. (2023). Translating science fiction in a CAT tool: Machine translation and segmentation settings. *Translation & Interpreting*, 15(1), 216–235.