From Translation Tool to Epistemic Technology: Generative AI and the Changing Nature of Knowledge Translation Across Languages

Aleksandar Trklja
University of Innsbruck

ECSPM Symposium – Saarbrücken, 2026

Translation and the Internationalisation of Higher Education

- Translation enables global circulation of research
- Supports collaboration across linguistic communities
- Provides access to academic knowledge across languages
- Increasingly mediated by new technologies
  - Example: *A European Joint Master programme in Digital Society, Social Innovation and Global Citizenship* (**University of Naples Federico II** (Italy), the **University of Innsbruck** (Austria), and **Palacký University Olomouc** (Czechia)

GenAI Translation as Infrastructure

- GenAI translation embedded in digital research workflows
- Used in academic publishing, teaching and research communication
  - Shapes how arguments and knowledge travel across languages
  - May influence discourse patterns and knowledge structures

GenAI as Epistemic Technology

- Operates on representational content
- Generates linguistic representations
- Performs epistemic operations (prediction, inference)
- Participates in knowledge production and circulation

The Epistemic Gap in genAI Translation

- Revision (human–human interaction):
  - Authors can explain reasoning behind expressions
  - Epistemic accountability exists

- Post-editing (human–machine interaction):
- Machine cannot explain translation decisions
  - Statistical patterns drive output
  - Creates an epistemic gap

Post-Editing as Repair

- Post-editing is not simply revision
- Human intervenes to repair machine output
- Reconstructs intended meaning and context
- Restores coherence and adequacy

Empirical Illustration: ECB Translation Experiments

- Dataset: European Central Bank texts
- Translation directions: German → English, French → English
- Systems tested:
  - Baseline MarianOpusMT
  - Fine-tuned model trained on domain corpus

Finetuning in translation

- What is it?
- How does it work?

What data is used?

- The fine-tuning relied on:
- 100,000 sentence pairs for EN-DE, FR-EN
  - 95,000 for training
  - 5,000 for validation
- This is a medium-sized dataset, sufficient for adapting a model to a specific domain.

What data is used?

- The 95,000 training examples are **pairs of sentences**:
  - source sentence (e.g. English)
  - target sentence (e.g. German human translation)
- Example:
  - Input: “The agreement is legally binding.”
  - Target: „Die Vereinbarung ist rechtlich bindend.“
- During training:
  - the model produces its own translation
  - this output is compared to the **human translation**
  - the difference (loss) is used to adjust the model

What are the 5,000 validation sentences?

- Process:
  - The model translates these sentences
  - The output is compared to the human translation
  - A validation loss is computed
- But:
  - **Not used to update the model**
  - Used only to **evaluate performance during training**

Why do we need validation data?

- Because otherwise we cannot tell whether the model:
  - is actually improving

  or

  - is just memorizing the training examples

Where and how was the training performed?

- The process was executed on:
  - a high-performance computing system (LEO5 cluster)
  - using an NVIDIA A100 GPU
- This allows fast numerical computation.
- As a result:
  - about 266 sentence pairs were processed per second
  - the entire process took approximately 18 minutes

What is loss?

- Loss is a numerical measure of error.
  - High loss: the model's output differs strongly from the correct translation (=human translation)
  - Low loss: the output is close to the correct translation
- Training aims to minimize this value.
- 5.0 → very different output
- 0.2 → quite similar
- 0.1 → very close

How does the learning process evolve?

- **Initial phase: rapid adaptation**
- At the beginning:
  - the model performs poorly on the new domain
  - the loss is high (above 5.0)
- After exposure to domain-specific examples:
  - the loss drops sharply (below 0.2)

How does the learning process evolve?

- **Middle phase: stabilization**
- After the initial improvement:
  - changes become smaller
  - the model refines its behaviour
- This includes:
  - improving terminology
  - aligning stylistic conventions
  - increasing consistency

How does the learning process evolve?

- **Final phase: convergence**
- At the end:
  - the loss stabilizes around 0.11–0.13
  - further improvement is minimal
- This indicates that the model has reached a stable configuration.

How do we know the model did not just memorize?

- To test this, we use validation data:
  - sentences not used during training
- In our case:
  - validation loss decreases steadily
- This shows:
  - the model improves on unseen data
  - it has learned general patterns of the domain

What is fine-tuning?

- Fine-tuning is a process in which an existing model, understood as a mathematical function with many parameters, is adapted to a specific domain by adjusting those parameters using example data, so that the model produces outputs that better reflect the patterns of that domain.

Observations from Experiments

- A comparison of the performance of the four models:
  - MarianOpusMT Baseline (domain-specific)
  - MarianOpusMT Finetuned
  - M2M100 Baseline (domain-general)
  - M2M100 Finetuned

Automatic evaluation metrics

- When a model produces a translation, we need a way to measure how good it is.
- We do this by comparing:
  - the model's translation
  - a human reference translation
- Automatic metrics assign a number to this comparison.
- We have used:
  - BLEU
  - COMET

Automatic evaluation metrics

- BLEU measures surface similarity
  - how many words or short sequences (phrases) in the model output also appear in the human translation
- Human:
  - „Die Vereinbarung ist rechtlich bindend.“
- Model:
  - „Die Vereinbarung ist rechtlich verpflichtend.“
- If many words overlap, then the BLEU score will be high.

Automatic evaluation metrics

- COMET is more advanced.
- It evaluates:
  - whether the meaning of the translation is correct
  - whether the sentence is fluent
  - whether it fits the context

Automatic evaluation metrics

- MarianOpusMT Baseline: BLEU 49.8, COMET 0.8988
- M2M100 Baseline: BLEU 35.88, COMET 0.8628

- MarianOpusMT is clearly stronger than M2M100 (baseline) because
  - BLEU difference: ~14 points
  - COMET difference: ~0.036
- MarianOpusMT produces translations that are:
  - closer in wording to the reference (BLEU)
  - also slightly better in meaning and fluency (COMET)

MarianOpusMT: What changes after fine-tuning?

- Global lexical behaviour
- MarianOpusMT-baseline is already very close to human translation
- After fine-tuning the model remains close, but
  - shows a slight increase in divergence, especially in mid-frequency vocabulary
  - the model introduces new **lexical choices** that are not present in the human translation.
- Fine-tuning does not improve global lexical alignment and may slightly disturb an already well-calibrated system.

M2M100: What changes after fine-tuning?

- **Global lexical behaviour**
- M2M100-baseline is **far from human translation**
- After fine-tuning:
  - Still **substantially divergent**
  - No convergence to MarianOpusMT-level alignment
  - Fine-tuning introduces many **new lexical items**, but not aligned with human usage.

Core insights

- Fine-tuning behaves differently depending on the starting point:
  - If the model is already well aligned (MarianOpusMT): the fine-tuning process might **distort lexical balance**
  - If the model is weakly aligned (M2M100): the fine-tuning process **does not necessarily correct structural divergence**

Translators as Epistemic Mediators

- Mediate between machine output and human interpretation
- Preserve conceptual nuance and contextual meaning
- Safeguard diversity of perspectives
- Maintain epistemic responsibility in multilingual communication

Empirical Illustration: ECB Translation Experiments

- Dataset: European Central Bank texts
- Translation directions: German → English, French → English
- Systems tested:
  - Baseline MarianMT
  - Fine-tuned model trained on domain corpus

Observations from Experiments

- Fine-tuning improves:
  - Terminology accuracy
  - Lexical consistency
  - Domain-specific phrasing
- But also reinforces:
  - Institutional discourse patterns
  - Standardised phrasing

Epistemic Implications

- GenAI systems reproduce the epistemic structure of training data
- Corpus design influences translation behaviour
- Training data shapes discourse patterns and terminology
- AI systems generate epistemic content

Multilingualism and genAI

- GenAI can support multilingualism by:
- Expanding access to research
- Supporting multilingual teaching environments

- But may also:
  - Standardise discourse
  - Reinforce dominant linguistic norms ..

Implications for Higher Education Policy

- Universities must address:
  - Responsible use of AI translation tools
  - Transparency of AI-generated content
  - Protection of linguistic and epistemic diversity

Translator Education in the genAI Era

- Future translators need:
  - GenAI literacy
  - Understanding of training data effects
  - Skills in human-AI collaboration
  - Ability to critically evaluate AI output

But also…

- **Learning Through Experimentation**
- Students should engage in:
  - Corpus building
  - Model fine-tuning
  - Translation evaluation experiments
- Aim: explainable genAI for students who
  - should develop explainability-oriented AI literacy.
  - understand how epistemic technologies work

Human-Controlled genAI Translation

- Goal: Human-controlled genAI ecosystems
  - Transparent training data
  - Human post-editing and oversight
  - Critical genAI literacy in education

- Translators act as epistemic mediators

Conclusion

- GenAI transforms translation into
  - human–machine epistemic collaboration
- Translators safeguard:
  - contextual adequacy
  - epistemic diversity
  - responsible multilingual knowledge circulation

Open questions

- How can universities and other institutions integrate genAI translation
  - while preserving epistemic diversity
  - in multilingual academic communication?

Thank you for your attention!