MusIML Workshop @ NeurIPS 2026 · Paper #89

From Moses to Multi-Agent Systems: A Paradigm Shift in Machine Translation

Mohd Rihan Yawar · Mohammad Areeb Uddin · Shahab Saquib Sohail · Sohrab Ali

Jamia Hamdard, New Delhi, India

Two ways to tell the history of MT

Architecture view

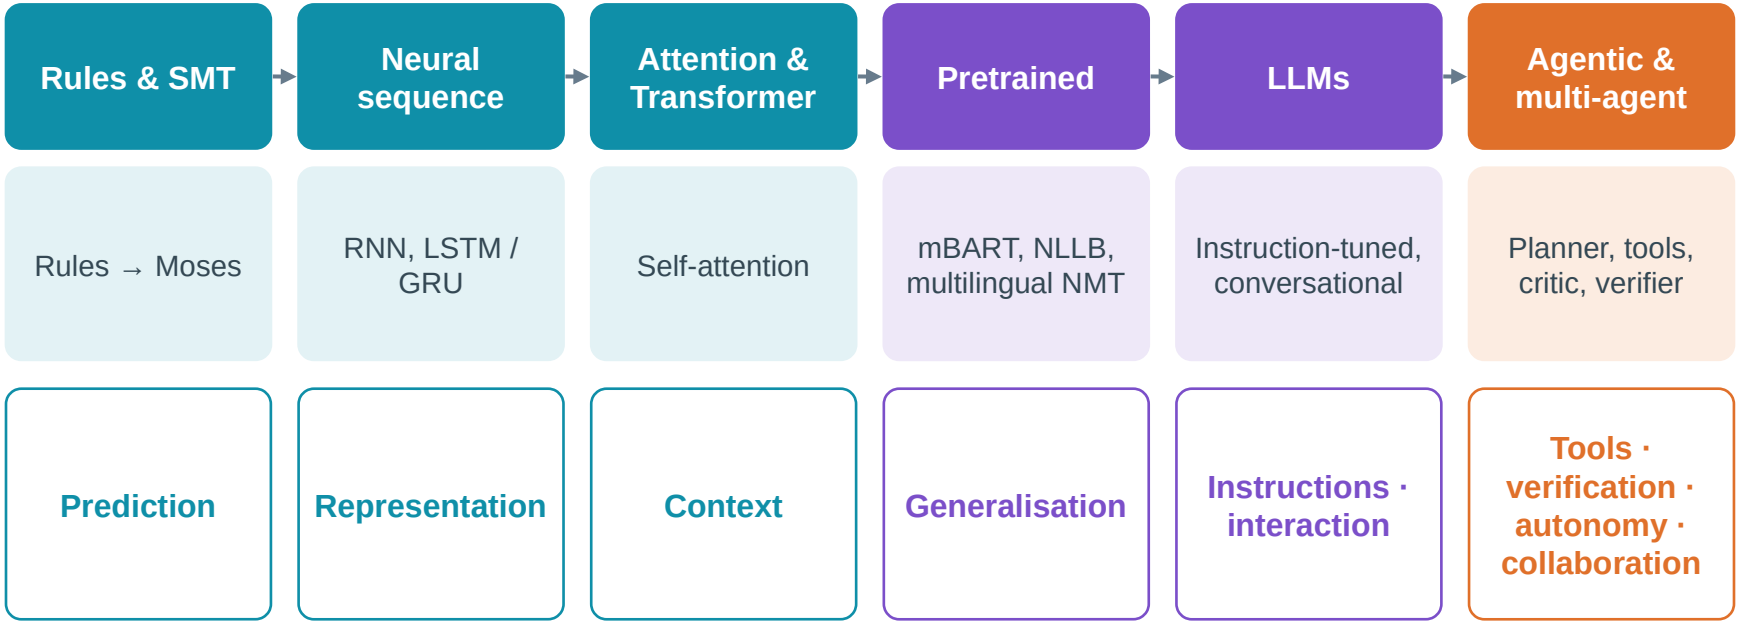
- Rules → SMT → RNN → Transformer → LLM → agents
- Says what replaced what
- Implies each step is simply better

Capability view (ours)

- What could systems newly do?
- Separates the model from the procedure around it
- Makes trade-offs visible: quality vs cost and complexity

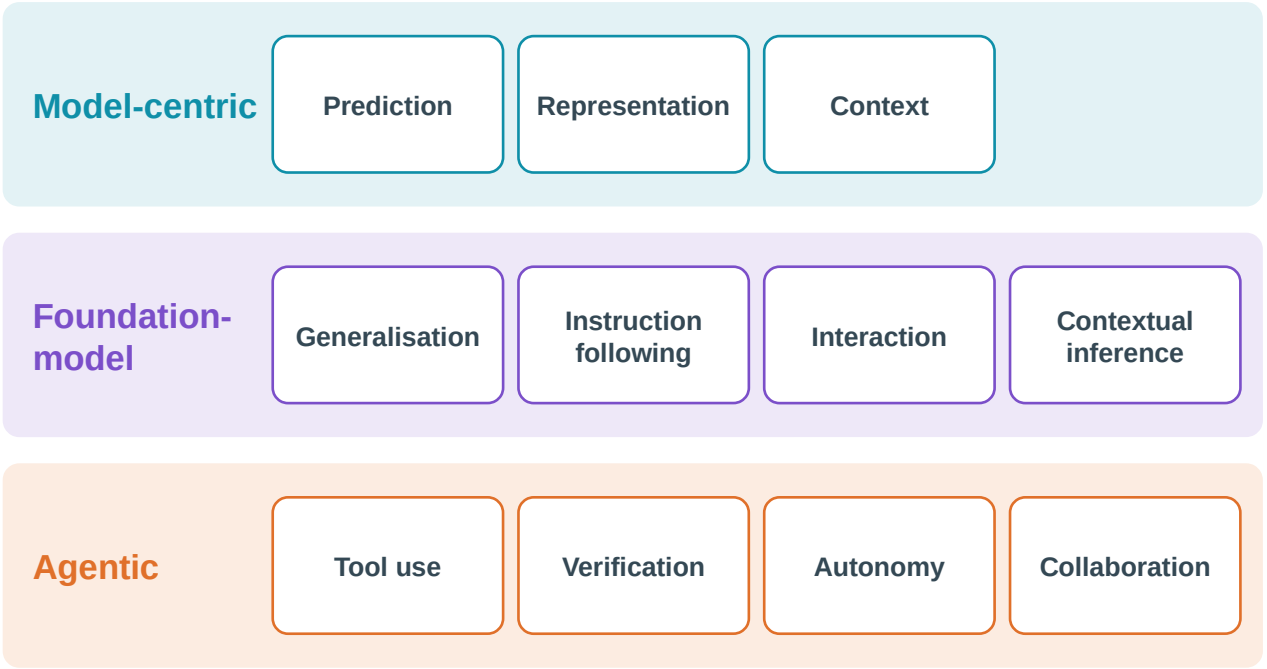
WMT24: human references were still among the best systems for 7 of 11 language pairs. MT is not solved.

Six eras, six new capabilities



Capability gained (boxed) ■ Model-centric ■ Foundation-model ■ Agentic

The Translation Intelligence Continuum



Cumulative
Later systems still predict and represent

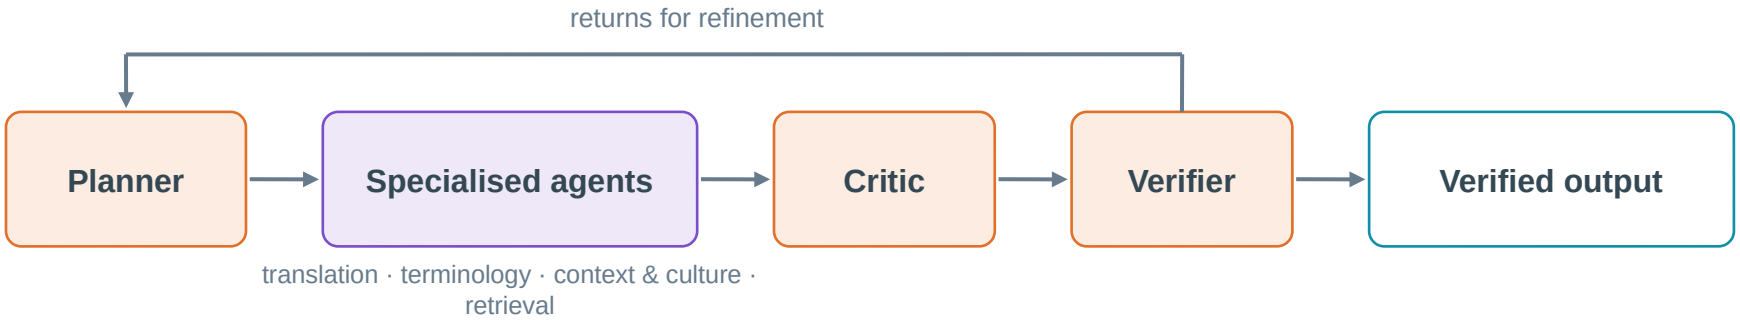
Not architectural
Interaction and verification need no new network

Not a quality ranking
A far-end system can translate worse

From one pass to a procedure

LLM translation $y \sim P(I, x, C)$
one pass: instruction, source, context

Agentic translation $y = \Phi(x; \pi, T, V)$
policy π, tools T, verifier V that can hand control back



+18.8–25.7% average gains **at ~3.2×** processing time (Lv et al., 2025)

TAKEAWAY

More autonomy widens what translation systems can do, and every widening brings new fragility.

Workflow evaluation
Measure planning, retrieval and verification, not just outputs

Cost-aware design
Gains come with multiplied runtime

Cultural grounding
Retrieve the knowledge low-resource text needs

Human + agents
Translators as participants, not post-editors

Limitations: a conceptual synthesis with no new experiments; the multi-agent architecture is not implemented or evaluated.