# Language for Intelligent Machines: A Prospectus

**K. Eric Drexler**
Future of Humanity Institute, University of Oxford
eric.drexler@oxfordmartin.ox.ac.uk

September 21, 2021

**Abstract**

To achieve human-like intellectual competence, machines must be fully literate, able not only to learn by reading, but to write things worth retaining as contributions to collective knowledge. Literate machines can and should employ learned, machine-native representations that go beyond NL syntax and words to provide a graph-structured, embedding-based, quasilinguistic medium that is more expressive and computationally tractable than natural language. Progress toward this objective would extend ongoing research in representation learning, reasoning, and language processing, and could potentially complement or replace today's opaque, error-prone "foundation models" with systems that are more capable, interpretable, scalable, and epistemically reliable. Potential applications and implications are extraordinarily broad.

The statements in the abstract above raise several questions that are explored below:

- Why must intellectually competent machines be literate?
- Why should literate machines employ machine-native representations?
- How can machine-native literacy support compelling applications?
- What are some actionable directions for research?

This brief document motivates the content of a more comprehensive document[1] that explores prospects for quasilinguistic neural representations in greater depth and breadth.

## Why Must Intellectually Competent Machines Be Literate?

Among humans, intellectual competence requires full literacy, including competence in reading, writing, and reasoning about content. At a societal level, the ability to read, reason, and fix ideas in writing has enabled the cumulative growth of knowledge.

1. Drexler (2021); 175 pages, 12 figures, 371 references.

Despite their remarkable linguistic skills, today's large language models lack full literacy in this sense—although they can learn from vast corpora and produce fluent, plausible, fact-rich text, their outputs neither build nor communicate reliable knowledge. To perform at the level of intellectually competent humans, machines must read, assess, and write useful content with cumulative results. Like humans, intellectually competent machines must be *literate.*

## Why Should Literate Machines Employ *Machine-Native* Representations?

Neither encodings of articulated sounds (native to humans) nor sequences of vector embeddings (native to machines) are good models for machine-native language. Machine representations can, however, *combine embeddings with graphs* to enable the learning and use of highly expressive quasilinguistic neural representations (QNRs).[2]

As a medium for writing, QNRs can surpass the expressive capacity of natural languages.[3] To facilitate reading, QNRs can disentangle semantic representations.[4] In training, QNR-oriented models can be end-to-end differentiable.[5] In inference, QNRs can provide a tractable medium for reasoning.[6] NL content is abundant, and QNRs are suitable for building corpora at scale;[7] large QNR corpora can include—and refine, integrate, and enrich—content translated from NL.

The expressive capacity of a linguistic system emerges from its components and structure:[8] In NL, these are words

2. There is no sharp distinction between QNRs and current vector/graph representations; arguably, QNRs are already in use.

3. Graphs generalize syntax, while continuous vector embeddings generalize words and can describe entities (*e.g.,* images) in ways that words and phrases cannot.

4. Rich embeddings and explicit coreference can improve local compositionality by reducing non-local semantic dependencies.

5. Graph representations can be made differentiable; continuous vector embeddings can be trivially differentiable.

6. Potential advantages include disentangled representations with emergent lattice structure.

7. Stored representations can be compact: Integers can designate embeddings in discretized vocabularies, while graph-valued functions can efficiently encode common syntactic structures.

8. There is no need here to understand how this actually works.

and syntax; in QNRs, embeddings and graphs. Vector embeddings are strictly more expressive than NL words (they can designate NL words, and more), while directed graphs are strictly more expressive than NL syntax (they can represent NL syntax, and more). It follows that QNRs can be strictly more expressive than natural language.

To explore and support a stronger claim—that QNRs can, in practice, *strongly outperform* natural languages—requires a wider-ranging discussion. Aspects discussed in (Drexler 2021) include:

- The potential expressive scope of (quasilinguistic) vector/graph representations
- The strengths and weaknesses of natural languages and formal representations
- Continuity of quasilinguistic representations across informal, semi-formal, and formal domains
- Mechanisms for knowledge integration through approximate unification on soft semantic lattices
- Applications of multi-task NL → QNR → NL training in quasilinguistic representation learning
- Semantic-level, QNR-domain cloze tasks for learning inference on high-level abstractions.
- Roles for NL → QNR translation and QNR-domain reasoning in building and applying QNR corpora.
- Potential applications of current neural ML components as building blocks for QNR-oriented systems

## How Can Machine-Native Literacy Support Compelling Applications?

QNR-based systems could be applied to translate, refine, and extend knowledge corpora at internet scale,[9] while efficient billion-scale similarity search in semantic spaces can support semantic search—in effect, associative memory—likewise at scale.[10] Prospective advantages in cost, scalability, interpretability and epistemic quality position QNR-based systems to complement or displace today's opaque "foundation models" (Bommasani et al. 2021) at the frontiers of machine learning.

Supported by QNR-facilitated inference, large corpora, and fluent conditional language models, QNR-based systems could be applied to a wide range of tasks. Improved language comprehension and reasoning could be used to improve question answering, writing assistance, content moderation, and search. Higher-level goals include organizing, refining, extending, and applying knowledge in areas that range from mathematics to medicine. Further potential applications include the implementation of agents capable of high-level, well-informed reasoning regarding actions, consequences, and human preferences (Russell 2019).

9. Some relevant comparisons: OpenAI's GPT-3 trained on ~300 billion words; Google's translation service reads, processes, and writes >30 trillion words per year; see Drexler (2021).

10. *E.g.,* Alibaba's Taobao Marketplace employs a low-latency, billion-scale recommender system based on similarity search in a 160-dimension semantic space; cited in Drexler (2021).

## Actionable directions for research

The broad vision outlined here aligns with both current research directions in neural machine learning and long-standing aspirations for machine intelligence. From a more concrete perspective, prospective developments are natural extensions of architectures and methodologies familiar to neural ML practitioners, and align with objectives measured by familiar benchmarks.

**Some directions in outline:**

- Advances in heterogeneous graph-oriented neural networks and Transformers could be applied to develop architectures with an inductive bias toward quasilinguistic representations.
- Training quasilinguistic systems with *semantic-level* loss functions could improve the transparency and epistemic quality of foundation models.
- QNR-oriented models pretrained with auxiliary lattice loss functions could potentially improve graph network representations and neural reasoning.
- NL → QNR translation, aided by QNR → QNR reasoning, could be applied to build and refine scalable QNR corpora.
- Reasoning mechanisms aided by access to QNR corpora could be applied to a range of multi- and cross-domain tasks, both linguistic and non-linguistic.

The considerations outlined above suggest that capabilities supported by QNR corpora and computation—which can enable machine-based literacy and reasoning—could help the ML community advance machine intelligence toward human-like intellectual competence and apply that competence to help solve problems that matter to the world. Applicable work is already underway.

## References

Bommasani, Rishi, Drew A Hudson, Ehsan Adeli, Russ Altman, et al. 2021. "On the Opportunities and Risks of Foundation Models," arXiv: 2108.07258 [cs.LG].

Drexler, K. Eric. 2021. *QNRs: Toward Language for Intelligent Machines.* Technical Report 2021-3. Future of Humanity Institute, University of Oxford. https://www.fhi.ox.ac.uk/QNRs.

Russell, Stuart. 2019. *Human compatible: Artificial intelligence and the problem of control.* Penguin.