

DCSR-LLM - Toolkit for Research at UNIL

Large language models are attracting growing interest across research fields, but many academic uses require more than a simple chatbot interface. Researchers often need to compare models, test them on specific tasks, extract structured information from documents, or adapt them to a domain-specific workflow. For these needs, reproducibility, local control, and transparent experimentation matter as much as convenience.

[dcsr-llm](#) was developed with that reality in mind. It is a command-line toolkit designed to support research workflows with large language models in a more controlled and reproducible way. Rather than focusing only on conversational use, it brings together several core functions in a single framework: inspecting models before use, downloading and running them locally, generating predictions, benchmarking results, extracting structured data from text corpora, fine-tuning models, and exporting them for other environments.

A typical dcsr-llm workflow

From model selection to evaluation, extraction, adaptation, and export.



For UNIL researchers, the value is practical. The tool is designed to work on local machines as well as on UNIL-supported GPU environments such as the Curnagl and Urblauna clusters. This makes it possible to move beyond isolated prompting and toward more systematic workflows. A team can, for example, inspect whether a model is compatible with its infrastructure, benchmark several models on the same question set, extract targeted variables from a document collection, or fine-tune an instruction model for a specialized task or terminology.

Several use cases are especially relevant in a research context. One is model selection: before downloading large files, researchers can inspect a model and estimate whether it is suitable for their hardware and intended workflow. Another is evaluation: instead of relying on impressions, researchers can benchmark baseline, quantized, or fine-tuned models on the same dataset and compare results consistently. A third is structured extraction: dcsr-llm can transform unstructured text into validated JSON outputs, with evidence tracking and review mechanisms that are useful for corpus-based work. For more advanced projects, the toolkit also supports fine-tuning existing

instruct models to better match a domain, style, or task protocol.

A key strength of dcsr-llm is that it treats LLM use as a research workflow rather than a one-off interaction. Configurations, saved artifacts, and explicit processing steps help support reproducibility and make experiments easier to document, rerun, and compare. This is particularly important in academic settings, where results need to be traceable and methods need to remain understandable.

dcsr-llm is best understood as a technical research tool rather than a one-click application. It does not replace critical judgment, and model outputs still need to be checked and validated. But for researchers who want a more rigorous and flexible way to work with LLMs, it offers a strong foundation.

UNIL members who would like to learn more, try the tool, or provide feedback can visit the [dcsr-llm repository](#) or contact us at helpdesk@unil.ch with subject: DCSR-LLM.

Repository: <https://git.dcsr.unil.ch/Scientific-Computing/dcsr-llm>

Créé 2026-03-18 19:02:57 UTC par Philippe Jacquet

Mis à jour 2026-03-19 07:39:13 UTC par Philippe Jacquet