

Writing Anima: A Philosophical Writing Counterpart

<https://github.com/Oxford-HAI-Lab/writing-anima>

HAI Lab's Writing Anima is an open-source, decentralised intelligent writing editor that gives perspectival, philosophical criticism on a piece of writing. The backbone of the tool is a concept we call an "Anima", which is an "animation" (in the literal sense of "anima" as "soul") of a particular text or set of texts. Texts (for example, departmental reading lists, tutor recommendations, personal notes or research journals) can be uploaded, and are transformed into an interlocutor that simulates the author of those texts. This interlocutor can be treated as a conversational partner, or as a critical writing assistant. A set of comments on written work can be requested, with a user experience alike MS Word or Google Docs. The user can then iterate their writing in a slow feedback loop, as if a philosopher of the past or present was their personal tutor. The tool is a standalone application and can work with cloud inference infrastructure or, if downloaded onto a machine with sufficient resources, can be used with locally running models of choice (even SOTA models like Deepseek on the right hardware).

Use

- Create a new project with a name and description.
- Feedback is given by "Anima" agents. To add an anima, click "Select Anima" and then click "Manage". The Anima manager will open.
- Click "New", give the anima a name (e.g. "Derek Parfit, Reasons and Persons") and an optional description. Select an embedding model.
- Upload a corpus (.txt or .pdf). The corpus will be chunked, embedded and saved. The Anima is now added.
- On the writing canvas, select an Anima and then select a backend model (for the best results, use an intelligent model class with >128K context, in our testing, e.g. GPT5 or Deepseek 3.2 work well).
- Models, embedding models and database services can be configured in the global app settings menu, including securely adding API keys for cloud inference.

- Add your writing sample and click "Think". The Anima will spend ~3-10 minutes reflecting and serve a series of comments in the right hand "Criticism" column. Criticism references segments in the corpus for further reflection. Click the reference to jump to the relevant section of the corpus.

Philosophical and Pedagogical Motivations

What is it we do when we do philosophy? And the related question: How do we teach someone to do philosophy well? A few things come to mind, which are rather universal across philosophies of different times and cultures. Firstly, philosophical engagement is an activity of “thinking slowly”, an activity of reflective and deliberative thought that cannot be rushed. Secondly, there is something essentially perspectival about philosophy. When philosophising, we do not cultivate or engage with a “view from nowhere”. Instead, we deeply engage with someone’s view again and again, interlocking with different perspectives, until over time we gradually settle on a position of our own. Thirdly, this process inevitably follows the contours of asking well-formed questions and trying to answer them, tracing out a complex, dense and deeply connected trajectory through “erotetic” space. All of this is perhaps why the tutorial system, as practiced in institutions such as Oxford, with its central focus on thinking dialectically, is so powerful a tool in philosophy pedagogy. In the present age, these observations raise the question of whether, and if so how, we can use artificial intelligence to assist and enhance the activity of learning and even doing philosophy.