

RedMarx: Teaching AI to Read Marx Dialectically

What is the relation between the socialisation of labour and the rising organic composition of capital?

Let us imagine Mariana, a political economist in São Paulo mandated by the Landless Workers' Movement (MST) to work through precisely this question for the movement's own political education. In addition to working through the canon, she puts the question to an AI tool such as ChatGPT or Gemini. It returns a familiar kind of answer: a definition of the terms, a tidy summary of the relation, and a gesture towards the present. But for a researcher accountable to a movement like the MST, it is not enough.

That failure is not surprising. A serious answer to the question is more complicated than it first appears. It cannot be a definition. It has to trace the development of one of Marx's central insights: that the growing social productivity of cooperative labour is, in the same motion, the process by which that productive power is stripped from the producers and turned against them as capital. A meaningful answer has to hold a contradiction open instead of resolving it into a tidy summary. It also has to remain faithful to what Marx actually wrote, across the *Grundrisse* (1857–1858), the *Economic Manuscript of 1861–63*, and the three volumes of *Capital* (1867–1894), texts he returned to across more than two decades.

A question like Mariana's exposes, in a single stroke, why the Marxist corpus defeats the ordinary AI assistant and what it would take to build one that does not. RedMarx, the system introduced in this article, is built to meet exactly that challenge.

Marxist thought is one of the most consequential intellectual traditions of the past two centuries. Across the world, and especially in research institutions in the Global South, it remains a working framework rather than an academic curiosity. Marxist political economy supplies the categories through which researchers read agrarian crises, extractive economies, financialised states, labour regimes, and the international division of labour.

And yet anyone who has tried to ground a piece of analysis directly in Marxist sources knows the difficulty.

The corpus is vast. The collected works of Marx and Engels alone run to scores of volumes and millions of words, multiplied many times over by Lenin, Luxemburg, Gramsci, Mariátegui, Mao, Nkrumah, Jones, First, Amin, and others. The categories themselves are dialectical: 'value', 'class', 'state', 'imperialism', and many more carry no static definition but acquire their meaning through the relations they describe and the historical moment in which they are deployed. A researcher like Mariana cannot simply quote a sentence from *Capital* (1867–1894) and treat the matter as settled. She must know which formulation to deploy, from which work, under which conditions, and how later writers extended or contested it.

Large language models change the calculation, but only under strict conditions. A researcher no longer needs to search manually through every page of every relevant volume to verify how a concept is treated across the corpus, if, and only if, she has a system that searches the original works faithfully, surfaces the relevant passages, and tracks where every claim came from. The ‘if and only if’ clause is pivotal. A generic chatbot will paraphrase, summarise, and, as is by now well documented, fabricate: one widely cited [study](#) found that 55% of the bibliographic citations generated by ChatGPT-3.5 were entirely fabricated.

Moreover, a meaningful tool must produce dialectical analysis, not reproduce books. It must reason over the corpus and return its own argument, pointing the reader to where each claim can be checked instead of reprinting long passages from the original works. The result must be interpretive: a line of reasoning the researcher can stand behind, never a substitute copy of the source.

These questions deserve more than a generic chatbot can offer, and so Global South Insights (GSI), a research project of Tricontinental: Institute for Social Research, developed RedMarx to meet that need. RedMarx is a research system configured for work in the Marxist tradition. The sections that follow explain the problem it was built to solve, the architecture GSI developed, and the evidence from its use.

1. The Problem RedMarx Is Built to Solve

No serious research tool can be built on the promise that AI will simply ‘read Marx for you’. The difficulty is not only that the corpus is large, but that generic tools fail in predictable ways when they encounter it. They flatten dialectical movement into definitions, blur later interpretation into the primary voice, smooth over political and translational disputes, and mistake retrieval for understanding.

1.1 Why the Marxist Corpus Resists Generic AI

Generic AI fails first by flattening development into definition. The treatment of the state in *The Eighteenth Brumaire of Louis Bonaparte* (1852) is not the treatment in *The Civil War in France* (1871), and neither maps directly onto Lenin’s reformulations in *State and Revolution* (1917). In the *Grundrisse* (1857–1858), ‘contradiction’ carries a Hegelian sense of the motor of development; by Volume I of *Capital* (1867), it has become a structural tendency towards crisis. A useful research tool has to hold these formulations side by side, recognise their historical specificity, and trace how a category develops, instead of collapsing them into a single averaged paraphrase.

Take a single question put to two systems: *how does a commodity come to express its value in another?* Google’s NotebookLM, benchmarked against RedMarx, answered with one tidy sentence – in the simple form, ‘the coat serves as a particular equivalent for the linen, reflecting the linen’s value in the coat’s bodily form’ – and moved on. The statement is accurate and inert. The equivalent form has been named and then left as a static fact of exchange.

RedMarx, given the same question, did not stop at naming the equivalent form. It unfolded what Marx [called](#) the three peculiarities of the equivalent form as a chain of reversals: use-value ‘becomes the form of manifestation, the phenomenal form of its opposite, value’; concrete labour ‘becomes the form under which its opposite, abstract human labour, manifests itself’; and private labour ‘takes the form of its opposite, labour directly social in its form’. Each step is a particular, concrete, private determination standing in for its universal, abstract, social opposite – and only by tracing that movement does the ‘enigmatical character of the equivalent form’ come into view, along with its connection to the money-form (*Capital*, Volume I (1867), Chapter 1, Section 3).

One system mentions the equivalent form while the other shows why it is the seed of the money-form. That gap, between a label and a logical chain, is exactly the gap generic tools cannot close on a dialectical corpus.

A second failure follows from a generic tool's tendency to silently blend later commentary into the primary voice, presenting a 1970s reading as though it were Marx's own. For a researcher, that distinction matters as she needs to know whether a category is being formulated by Marx, extended by Lenin, Luxemburg, or Amin, reconstructed through dependency theory, or debated in contemporary scholarship.

Finally, the corpus is politically contested, and much of it exists only in translation, where the choice of words is itself a political act. There are partisan editions, hostile translations, and theoretical traditions that treat the same texts in fundamentally incompatible ways. A research system that smooths over these tensions in the name of producing a confident answer ends up misleading the researcher.

How sharp these failures are became clear the hard way. In GSI's first experiments, the system retrieved the correct passages from the correct texts, and the human experts still rated the answers as shoddy and careless in how they handled Marxist concepts. The problem was not retrieval but interpretation. Pulled out of the surrounding scholarship, an isolated passage of Marx leaves even a capable model reasoning from first principles rather than from the established understanding of what the passage means. It reads the words at their ordinary, plain-English face value, without the technical sense the tradition has given them, and loses the conceptual context that makes the passage intelligible. Solving that turned out to require the rest of the system described below.

1.2 What Researchers Actually Need

In practice, a researcher approaching the Marxist corpus needs four things that generic AI tools do not provide, each taken up in Section 2. The first is faithful retrieval of original passages rather than paraphrase. When a section of an article rests on Marx's argument in the *Grundrisse* (1857–1858), the researcher needs the actual paragraph, in context, with its chapter and section identified. This is the job of MetaRAG retrieval (Section 2.2).

The second is a dialectical, historically aware analytical scaffold, a way to see the conceptual structure of the question before drafting, including which categories are in play, how they relate, which earlier formulations are foundational, and which later developments extend or revise them. This is the work of the knowledge graph (Section 2.1). The third is interpretation anchored in established scholarship rather than reasoning from scratch, so the system reads a passage the way the field reads it instead of reading an isolated chunk without the conceptual context that gives it meaning. This is the work of the research loop and its embedded skills (Section 2.3). The fourth is traceable, section-level citation in the final output, so that any reader or reviewer can follow each claim back to the exact passage that supports it. This is a precondition for serious academic and political work (Section 2.4).

RedMarx is built around these requirements, applying them in proportion to the task. The goal is not to turn every sentence into a forensic fact-check, as one would for a factual report, but to give the researcher a rigorously sourced foundation for deep, dialectical thinking about complex theory.

2. How RedMarx Works

RedMarx takes a research question as input and produces a research report as output. RedMarx is built on MetaRAG – GSI's institutional retrieval-augmented generation infrastructure – configured for work in the Marxist tradition. MetaRAG is 'meta' in a precise sense. Rather than acting as a single research agent, it is the layer that builds research

agents on demand. Faced with a question, it probes the available knowledge, works out what a strong answer would require, and assembles the right combination of search, analysis, and writing strategies for that particular question.

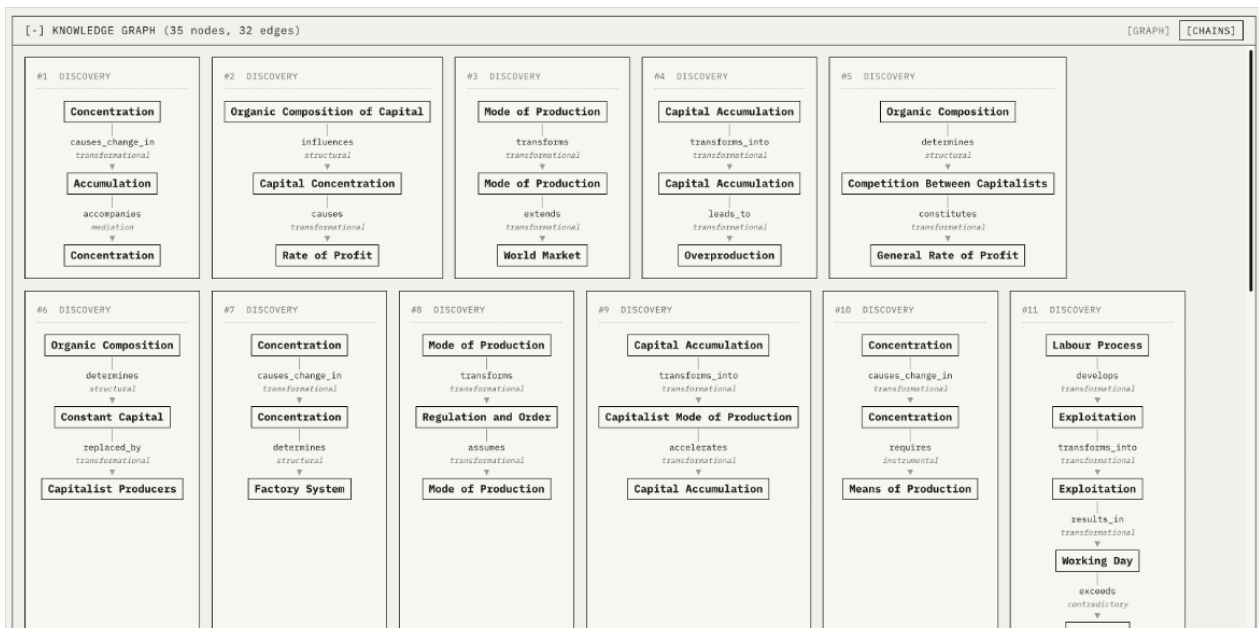
In GSI's [AI for Social Science](#) architecture, MetaRAG is the verified retrieval-and-augmentation substrate on which the institute's research systems are built; RedMarx is what that substrate becomes when a research team points it at a particular body of knowledge and shapes it with that field's own concepts. With MetaRAG's methodology and toolchain, a domain expert who supplies good source material, a sound analytical framework, and a carefully constructed knowledge graph can build a systematic research engine of her own, without writing code. Mariana's question, on the relation between the socialisation of labour and the rising organic composition of capital, offers a useful thread for following what happens in between, including the role human experts play in shaping the system at each step.

2.1 The Knowledge Graph: A Dialectical Skeleton for Analysis

RedMarx's most distinctive feature is its use of a layered conceptual knowledge graph of Marxist political economy, which serves as the analytical skeleton for every research question.

When RedMarx receives a question, it does not begin by searching the corpus but by reasoning over the knowledge graph. The graph encodes the major categories of Marxist political economy – value, labour, capital, class, state, imperialism, mode of production, and others – as a layered, relational structure rather than a set of flat tags. Categories are organised hierarchically from general to specific, dialectically through a concept's opposites and developments, and historically through the reformulation of a category across periods and traditions. These relations are drawn from a controlled vocabulary of social-scientific predicates – 'contradicts', 'mediates', 'reproduces', 'commodifies', 'alienates', 'exploits', 'determines', 'transforms into', and others – so that the graph captures genuinely Marxist relationships rather than generic 'is-related-to' links.

For Mariana's question, the graph yields a skeleton before any text is fetched: socialisation of labour stands as the real social content, while organic composition of capital – the ratio of a firm's investment in machinery and materials to its investment in labour – is its value-form expression. The two are linked through cooperation, machinery, and the division of labour to the concentration and centralisation of capital, and onward to the tendency of the rate of profit to fall. That structure is the analytical plan, the dialectically and historically aware scaffold on which the rest of the research hangs.



Discovery chains RedMarx traverses for Mariana's question, from a knowledge graph of 35 concepts and 32 relations

The graph, a carefully constructed analytical map rather than a dictionary, is built and maintained by GSI's domain experts to reflect the actual structure of the tradition. This is where expert labour is most concentrated, and two features deserve emphasis. The first is historical disambiguation: each concept carries a time-period annotation – a category traced '1844 → 1857 → 1867', for instance – so the graph can distinguish the early formulation from the mature one instead of averaging them. The second is sense-splitting, when the same term appears in two distinct theoretical settings with little overlap in meaning. In such cases, the graph splits it into separate senses rather than fusing them, letting the system respect the dialectical movement of Marx's vocabulary instead of erasing it. The impact of this expert curation was evident when GSI tried to extract concepts from a given source document. When guided by a generic, default ontology, the extraction pipeline drew only a handful of concepts while the purpose-built Marxist ontology helped draw several hundred concepts from the very same text.

This step matters because it makes the system's analytical assumptions visible, allowing Mariana to inspect the skeleton before any retrieval happens and adjust it if a category is missing or misplaced. The graph is a working object, not a black box.

2.2 MetaRAG Retrieval: Going Back to the Original Texts

Once the skeleton is in place, RedMarx hands each of its nodes – each concept, each claim – to MetaRAG, configured here against a curated corpus of Marxist works. That corpus is itself an artefact of careful expert preparation. The source volumes are digitised through a two-pass vision-language reading process, with one model reading each scanned page and a second refining and cross-checking it. This allows tables, formulae, and the original printed page numbers to survive intact, so that every passage can later be cited back to the book and page it came from.

The instruction to MetaRAG is straightforward. For every concept and every position identified in the skeleton, locate the original passages in the corpus and return them with their chapter and section identifiers. Do not paraphrase. Do not summarise. The downstream analysis depends on having the real text in front of the system, not a synthesised approximation.

A second, subtler commitment operates here too. When RedMarx moves from finding a passage to analysing it, it does not reason from the isolated snippet. Chunks do the work of finding – they are how the system locates the relevant passage – but they are not the unit it then reasons over. Having used the chunk as a pointer, the system discards it and ingests the entire surrounding source document – the full chapter rather than the fragment – so that the author's overarching argument frames the reading.¹ This is the practical answer to the loss of conceptual context described in Section 1.1, since meaning is no longer artificially confined to the unit that happened to match the query.

This is where the difference between RedMarx and a generic chatbot becomes most visible. A generic system will produce fluent prose about Marx's theory of ground rent without ever consulting Volume III of *Capital* (1894). RedMarx is constructed to refuse this shortcut. The retrieval step is the foundation, and everything downstream is built on direct engagement with the sources.

2.3 The Research Loop: Reading the Way a Scholar Reads

Faithful retrieval is necessary but not sufficient. As GSI's early experiments showed, a system can hold the right passages and still interpret them carelessly. RedMarx therefore wraps retrieval in an iterative research loop – plan, harvest, analyse,

¹ Full-document ingestion processes a large volume of text for each answer, and so requires a model with a one-million-token context window. This currently restricts RedMarx's full configuration to frontier models; the leaner configuration described in Section 3 runs on smaller models at significantly lower cost.

synthesise – that works through a question the way a researcher like Mariana would: scoping it, gathering evidence, assessing whether the evidence is sufficient, and only then drafting.

Two things give this loop its rigour. The first is an embedded expert method. How RedMarx investigates a Marxist question – how to navigate the corpus, how to disambiguate a contested concept, how to read a passage against the consolidated scholarship rather than from first principles – is not hard-coded. It lives in editable instruction documents, distilled from the primary and secondary literature through retrieval work carried out with domain experts. These documents capture how an experienced scholar interrogates a Marx text. Because they are plain documents, a domain expert can revise the system’s method directly by adding an interpretive framework, which will take effect immediately.

This is also how the team corrected an early failure in which the system merely stacked retrieved points in a flat list, like a student reciting facts, instead of building the logical relations between them. Teaching the system, through these documents, to write as serious scholarship demands turned a stack of retrieved points into a genuine argument.

The second is a human checkpoint. After the system has scoped the question and surveyed the available evidence, it pauses for the researcher to review and approve the direction before committing to full analysis. The human is not a spectator at the end of an automated process but is in the loop at the decisive moment. Work in progress is preserved at every step, so the path from question to report stays auditable rather than opaque.

2.4 Synthesis: Faithful to the Text, Footnoted to the Section

The final stage takes the conceptual skeleton, the retrieved primary passages, and the researcher’s question, and produces a synthesised research report.

The synthesis is bound by a strict requirement that every substantive claim that rests on a source must carry a footnote that identifies the source down to the chapter and section. Because each retrieved passage arrives already tagged with its origin, citations are born with their sources rather than reconstructed afterwards. This removes the risk of the fabricated or mismatched citation that plagues generic tools. A single footnote in a RedMarx report resolves not merely to a book but through its full structure – Book → Part → Chapter → Section → Subsection – so that ‘Marx on the form of value’ becomes, concretely, a pointer to Section 3, ‘The Form of Value or Exchange Value’, in Volume I of *Capital* (1867). The reader of a RedMarx report, whether a colleague, a reviewer, or a critic, can follow each substantive claim back to the precise passage in the original work that grounds it.

For academic work, this satisfies the rigour that serious scholarship requires. For political work, it matters because claims about what Marx, Lenin, or Gramsci argued must be checkable, particularly when those claims are contested. A report whose footnotes resolve to chapter and section is a report that can be defended.

Traceability serves thought; it is not an end in itself. The footnotes exist so that a researcher can build on the sources with confidence, shifting energy from the hunt for passages to argument, judgement, and the dialectical work that is the point. On Mariana’s question, for instance, citation is only the starting point. What matters is the argument RedMarx is able to sustain: that the rising social productivity of labour is the very process through which labour is subordinated to capital. Where Marx’s texts leave that tension genuinely open, the system marks it as such, with sources on both sides, rather than forcing the corpus into a resolution it does not contain. That is the register RedMarx is built for.

3. Current Status: Evidence from Use

RedMarx is being trialled by GSI researchers, who have so far run it on 16 substantive questions of the kind reproduced throughout this article, across 29 logged sessions – five of those questions selected for the controlled four-system benchmark reported below, the other 11 developed as full research reports – and judged the results to be strong.

That judgement is not merely impressionistic. In a controlled benchmark, five expert-selected questions on Marxist political economy were scored by human evaluators against a six-dimension rubric – comprehensiveness, representation of contradictions, terminological precision, dialectical exposition, critical reflexivity, and analytic depth. Together, these produced 118 structured pass/fail checks per system.

RedMarx's full configuration scored 118 out of 118, or 100%. A leaner configuration of the same system, running on the substantially cheaper Kimi K2.5 rather than a frontier model, scored 97.5%. Claude Opus 4.7 with no retrieval at all scored 83.1%, and NotebookLM scored 50.8%. The decisive gap appeared not in surface vocabulary, where every system could name the categories, but in the dimensions that test Marxist understanding: surfacing internal contradictions, maintaining critical and historical reflexivity, and sustaining analytic depth. There, the weaker tools produced clean taxonomies that sanded off Marx's critical edge. NotebookLM's result should be read as a mismatch of purpose rather than a simple failure: it is built to answer in seconds, where RedMarx is built for analyses that take 10 to 15 minutes. One difference did not need a rubric to detect: across the same five questions, the two retrieval systems produced 216 references resolving to a named chapter or section, and the two systems without retrieval produced none at all. The benchmark matters because it tests the difference between summarising Marxist vocabulary and reasoning dialectically with it.

These claims can be checked against the answers themselves. Four questions are published in full – each with RedMarx's answer, the knowledge graph and discovery chains behind it, and a general AI tool's answer to the same question – in [GSI's GitHub repository](#).

The benchmark was run on a restricted, canonical slice of Marx on which even an unaided model performed respectably. That performance is expected to fall away – and the value of retrieval and the knowledge graph to grow – as the corpus expands to Lenin, Mao, and lesser-known texts that demand reasoning across authors. That expansion is the subject of the next section.

The intention is not to keep RedMarx as an internal tool. The conditions that produced it – the difficulty of the Marxist corpus, the inadequacy of generic AI, the centrality of dialectical thinking, and the need for source fidelity – are not unique to one institution. RedMarx will prove useful to other research collectives, party intellectuals, popular educators, and independent scholars who work within or alongside the Marxist tradition.

4. The Road Ahead

A research system is never finished. Several directions of work are already underway, in rough order of immediacy.

Deepening quality and the research loop: The nearest-term work is on rigour itself. The aim is to mature the research loop so that the system can review its own draft, identify what is missing, return to the sources, and try again for a limited number of rounds before a human reviews the result. In other words, the system should begin to ask of its own work what a careful researcher would ask: is the argument complete, are the sources sufficient, and has anything important been missed?

Broadening the corpus: In parallel, GSI is steadily integrating further primary sources and well-curated secondary literature. Marxism is not a closed canon but a living tradition, with major contributions across languages, regions, and historical moments. GSI will continue to add carefully selected works, with attention to translation quality, editorial provenance,

and licensing, scaling the corpus from the political economy core towards the wider canon of Lenin, Mao, and others, so that RedMarx's retrieval reflects the full breadth of the tradition rather than a narrow slice. Because so much of that tradition lives in translation, expanding across languages is a central part of this work, not an afterthought.

Richer analytical frameworks and information-gathering strategies: Beyond the current knowledge graph, GSI is developing additional analytical frameworks – conjunctural analysis, world-systems decomposition, value-chain analysis, and social-formation analysis – that a researcher can invoke when the question calls for them. Each brings its own way of reading the corpus.

Cross-traditional and cross-temporal research methods: Some of the most important contemporary questions cannot be answered from inside a single theoretical school or a single historical period. RedMarx is being extended with methods that allow comparison across traditions – reading dependency theory and world-systems theory against each other on the same question, for example – and across time. These methods help trace how a category evolves from its classical formulation through its mid-century reconstruction to its contemporary use.

Connecting to deeper research: Finally, RedMarx does not have to work alone. A faithfully sourced, dialectically structured analysis of a theoretical question is exactly the kind of input that a larger research process needs. Work is underway to let RedMarx feed its grounded output into more comprehensive research systems, so that rigorous theoretical reasoning becomes a building block of larger investigations rather than an endpoint.

The deeper reason for building the instrument in-house, rather than waiting for commercial platforms to serve the field, is that ownership shapes possibility. When the system belongs to the institution that uses it, the methodology embedded in it can be the one that institution actually believes in. The researcher can ask the system harder questions, configure it for finer-grained work, and, most importantly, direct human attention to the questions that only humans can answer. The mechanical work of locating passages, cross-checking citations, and mapping conceptual relations can be entrusted to the system. What remains for the researcher is the harder and more valuable work of judgement, synthesis, and political imagination.

5. The System Retrieves; the Human Thinks

A tradition is kept alive not by being quoted but by being thought with. For a century and a half, the Marxist canon has demanded more of its readers than any single lifetime can give: more volumes than one scholar can hold in view, more contexts than one mind can keep straight, more debates than one career can master. The danger is not that the texts disappear, but that the collective capacity to read them as a whole begins to thin. RedMarx is built against that attrition.

It also matters who builds the instrument. The capacity to work with a tradition is not a neutral convenience; it is infrastructure. When the tools that mediate access to the canon are built on someone else's terms and behind someone else's paywall, they shape the questions researchers can ask and the sources they can reach. Building RedMarx is a way of refusing that dependence; it is a small act of reclamation that returns part of the means of reading the tradition to those who work within it and gives a movement more control over the conditions in which it thinks.

RedMarx is also built by people, for people. Movement leaders and cadre posed the political questions the system exists to answer; scholars chose the most authoritative sources; retrieval and domain experts shaped the dialectical framework together; those who know the tradition drew the knowledge graph by hand; and AI specialists engineered the research loop that binds it together. No part of RedMarx thinks on its own; every part extends someone's thinking. The human is its origin and its end – the question begins in a real political need, as Mariana's mandate from the MST did, and the

answer is meant to be carried back into a real argument and put to work. What looks like a machine is, in truth, a collaboration among experts from different fields, each multiplying the others' reach.

RedMarx should not be mistaken for magic, or for a substitute for the disciplined reading that serious scholarship has always required. It is not a shortcut around that labour but a tool for absorbing the mechanical work of locating passages, checking citations, and mapping conceptual relations. Ultimately, judgement, synthesis, and political imagination remain the domain of the researcher. The system retrieves; the human thinks. That division of labour is not a weakness of the tool but the reason for building it.

Researchers and organisations interested in RedMarx, or in adapting the approach to another body of work, can write to GSI at contact_gsi@thetricontinental.org.