AI in Research

# Part One: AI-Assisted Research

May 2026

Anton Holst,
Quantitative Researcher

## From Copy-Paste to Collaboration

When generative AI tools became widely accessible in late 2022, the initial reaction across the investment industry mirrored that of most other fields. Users primarily relied on copying and pasting text into the model to summarize documents, draft emails, or explain concepts. It was useful, but it was not research.

Over the following year, AI-enabled coding environments introduced automatic code suggestions, allowing researchers to write code faster with the model filling in gaps. This was a meaningful step forward, but the interaction remained reactive and limited to completing fragments. As models and their integration with development tools matured, the workflow evolved from autocompletion to generating complete functions and modules from scratch. Tasks could be described in natural language, with the AI producing working code in response. This noticeably changed the pace of research, though much of the orchestration remained manual.

The most recent shift has been toward more capable agents. With newer AI coding tools, much of the implementation is handled by the agent, allowing researchers to focus on structure, direction, and quality. Rather than writing code line by line, they set the objective, provide guidance, and review the output. The AI has evolved from a sophisticated autocomplete into a genuine collaborator.

Lynx began serious development of this kind of AI-assisted research workflow in early 2025, when the path forward was still uncertain. By the end of the year, its potential was evident. Today, AI is not a peripheral tool in our process. It is central to how we conduct research.

# *The Evolution of AI in Research*

Four stages, 2022–2026: how the researcher's relationship with the model deepened from one-way prompting to genuine collaboration.

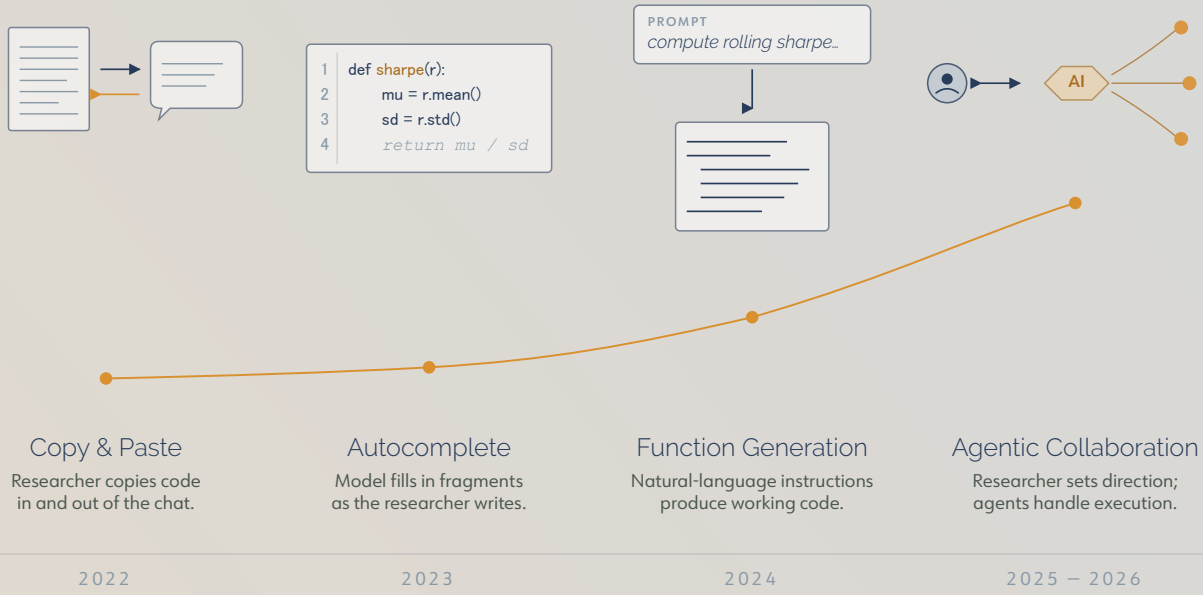


**Copy & Paste** (2022)
Researcher copies code in and out of the chat.

**Autocomplete** (2023)
Model fills in fragments as the researcher writes.

**Function Generation** (2024)
Natural-language instructions produce working code.

**Agentic Collaboration** (2025 – 2026)
Researcher sets direction; agents handle execution.

## A Shared Foundation of Knowledge

For an AI agent to be a useful research collaborator, it needs access to the same body of material a human researcher would draw on. A common misconception is that because information may be publicly available on the internet, or part of an AI model's training data, it can be reliably retrieved, referenced and reasoned about. This is not the case. Much of the world's knowledge sits behind paywalls, in subscription databases, or in formats that are difficult to parse, such as scanned PDFs, image-based documents, tables embedded in reports, or proprietary file types. Even openly available material is often buried deep in archives, poorly indexed, or scattered across sources. Large language models are trained on vast amounts of text, yet they remain limited in their ability to accurately recall or cite specific papers, books, or findings on demand. The training process produces general knowledge and reasoning ability, not a searchable index of sources. To produce reliable, grounded answers, the AI needs structured access to the underlying material.

To address this, the Lynx research department has built a centralized knowledge base that consolidates insights from thousands of sources - academic papers, reports, and other literature - along with non-text media such as podcasts. All of this material is structured, tagged, and enriched with extracted metadata to make it easily navigable by both human researchers and AI agents. An automated pipeline continuously reviews new documents, processes them, and ensures that the facts are searchable and ready for use. This knowledge base serves as a shared resource across the entire organization, accessible to all researchers and every agent.

## How It Works in Practice

With this foundation in place, AI-assisted research at Lynx becomes a genuine collaboration between a human researcher and one or more AI agents. The researcher leads the process, generating hypotheses, assessing quality, and making strategic decisions, while the AI handles much of the implementation. A researcher can explore an idea through extended dialogue with an AI agent, shape it into a clear specification, and then work through each step together, from finding the right data to writing model code and producing backtests. The AI operates within our proprietary platform, ensuring that everything it builds meets the same standards and uses the same frameworks as work done entirely by humans.

To illustrate, consider a researcher who finds that a set of features is less robust than expected across market regimes. Instead of spending days manually reviewing the literature, the researcher can use multiple AI agents in parallel: one to summarize established methods from technical references, and another to review recent research and identify newly proposed techniques. Within a short time, the researcher gains a comprehensive view of both classical and emerging approaches and can make an informed decision on how to proceed.

By enabling parallel, targeted exploration, AI-assisted workflows significantly reduce implementation and research time. Tasks that previously required days can now often be completed in hours, allowing researchers to spend less time on mechanical work and more time on the judgment-intensive decisions where decades of experience in systematic investing matter most.

## Building Infrastructure That Moves with the Technology

The AI landscape evolves quickly. The leading model of today may be surpassed within months, and capabilities that recently seemed out of reach can become standard overnight. We have built our infrastructure to accommodate this pace. Our systems are designed to work across different AI models, allowing us to adopt new ones as they are released without rebuilding our platform. We can use different models for different tasks when one is better suited to a particular kind of work, and swap them as stronger options emerge. Our legal, compliance, and technology teams have pre-cleared the leading providers, so that when a more capable model arrives, it can be deployed immediately.

We focus our own development effort on the things that are specific to how we work. This includes how AI connects to our data pipelines, market databases, backtesting systems, and research workflows. These are domain-specific challenges not addressed by general model improvements, and where decades of hard-won knowledge about markets and systematic investing are embedded in everything we build.

## Compounding Knowledge Through Skills

As our researchers and agents work together, effective techniques, useful patterns, and hard-earned insights are captured as reusable "skills" that become available across the organization. When one project produces a better way to handle a particular data transformation or evaluation step, that knowledge is formalized and shared. Researcher and agents alike benefit from it going forward. This creates a compounding effect: the more the system is used, the more capable it becomes. Our institutional expertise does not just live in people's heads – it is embedded in our tools and workflows, growing more valuable with each completed project.

This is where we are today. AI-assisted research has fundamentally changed how quickly and thoroughly we explore ideas, while keeping human judgment at the center of every decision. In Part Two, we will look at what comes next: fully agentic research, where AI agents carry out entire research cycles with increasing autonomy, and the researcher's role shifts from hands-on collaborator to strategic director.

## *Skills Accumulate; The Network Densifies*

As researchers and agents work together, reusable techniques are captured as shared skills.
The more the system is used, the more capable it becomes.

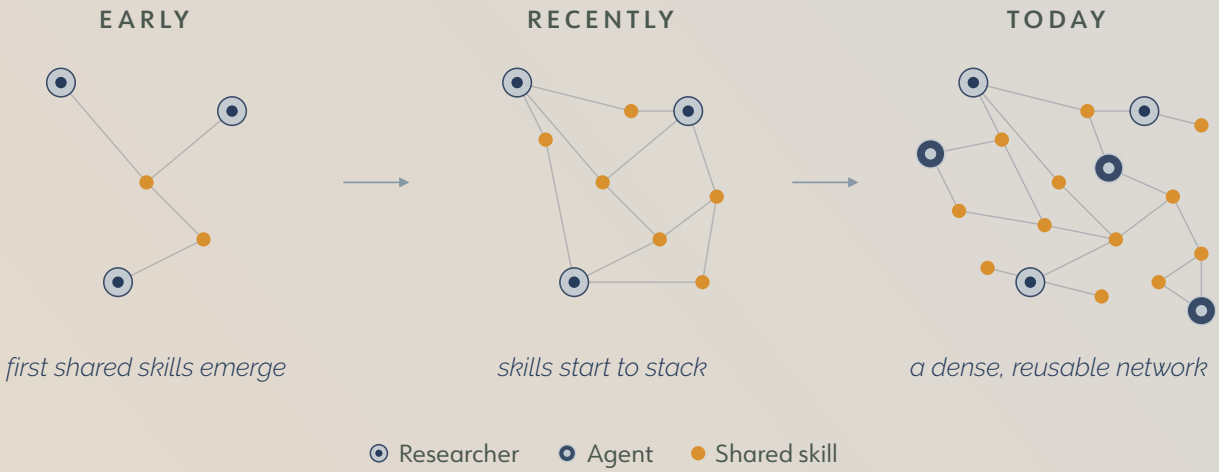

# Important Information

This document is provided for informational purposes only and describes certain internal technologies, workflows, and research practices related to the firm's use of artificial intelligence. Such practices may evolve over time, and no assurance is given as to the accuracy, completeness, or future applicability of the information. Nothing in this document should be construed as investment advice, an offer to sell, or a solicitation to purchase any security, financial instrument or service. The views expressed are those of the author and do not necessarily represent the views of Lynx Asset Management AB. There is no guarantee that any approach described will result in the achievement of any investment edge or alpha, and systematic investment strategies involve inherent risks.