AI News

Project Prompt

2026-05-02 - 13:17

Compiled with Claude AI

1 Project overview

AI News documents the evolution of artificial intelligence as it unfolds. The project is owned by Norbert Weber and compiled with assistance from *Claude* AI (**ANTHROPIC**). Outputs include monthly Word reports, companion PDFs, and a public website at **ai-news.sliver.lu**. This document records the conventions, methodology and workflow that produce them.

1.1 Purpose and scope

The publication produces one report per calendar month, retrospective by design: each edition is written after the month ends and covers events that took place within it. Reports cover large language models, generative media tools, coding assistants, and the broader industry context — funding, regulation, infrastructure, and strategic moves. The scope is the AI ecosystem in commercial operation, not academic research.

1.2 Intended audience

The reader is a working professional who uses or evaluates AI tools without specialising in machine learning. The register assumes familiarity with general technology vocabulary but not with academic literature. Where a term may be ambiguous, the glossary provides plain-language definitions. The publication does not target specialists who need depth from primary research papers, nor newcomers who need an introduction to the field.

1.3 What this project is not

Not a buying guide. Not an investment newsletter. Not a marketing channel. No product is endorsed and no entity sponsors the publication. Comparative claims are restricted to verifiable facts — release dates, prices, *benchmark* scores, regulatory rulings. The publication does not predict the future beyond surfacing announced or scheduled events.

2 Document family

The project produces four kinds of documents. Each has a defined role, a stable filename pattern, and a clear publication cadence.

2.1 Monthly Trends reports

The Trends report is the primary deliverable — one per calendar month covered. Filename pattern: AI - Trends - YYYY-MM (TS).docx, where TS is the timestamp of the most recent generation. The subject portion uses ISO month notation so that lexicographic sort matches chronological sort. Each report is self-contained: a reader picking up a single edition does not need to consult prior reports to understand it.

2.2 Glossary

A single, stable document covering technical vocabulary used across all Trends. Maintained as one file with one URL. New terms are added when they appear in monthly content; existing definitions are revised as the underlying

technology shifts. Entries are alphabetically sorted within thematic sections. The glossary has no monthly version — it is a living reference, not an archive.

2.3 Foreword

A standalone editorial document presenting the perspective from which this project is compiled. Where the Trends are factual, the Foreword acknowledges its own subjectivity. It is offered to readers as a counterpart, not as an introduction to be agreed with. Published once, revised rarely.

2.4 Project Prompt

This document. It records operational reference material for production sessions and provides methodological transparency for external readers. Updated whenever conventions, structure or workflow change.

3 Trends — structure and tone

Each Trends report follows a fixed seven-section structure. The structure is reused verbatim across editions to support comparison and pattern recognition over time. The tone is held constant for the same reason.

3.1 Section structure

The seven sections of every Trends report:

- **Executive Summary —** one page, three to five paragraphs covering the month's most consequential developments.
- **Large Language Models —** releases, benchmarks, pricing and access changes.
- **Generative Media —** image, video, music and audio. Subsections active only when activity warrants.
- **Coding and Developer Tools —** releases, pricing, adoption signals.
- **Industry Trends —** funding, regulation, mergers and acquisitions, infrastructure.
- **Monthly Recap —** chronological table of dated events with category and source. The only mandatory table in the report.
- **Outlook —** what is expected the following month: announced releases, regulatory deadlines, scheduled court dates, conferences.

3.2 Tone and register

The register approximates that of a sectoral briefing or a technical review. Adjective-heavy framing is avoided — words like revolutionary, groundbreaking, or stunning carry no information and are excluded by policy. Specificity does the work instead. A *benchmark* score with its baseline tells the reader more than the assertion that a model is much better; a price in dollars per million *tokens* tells the reader more than the assertion that a service is affordable. The reader is presumed competent enough to draw conclusions when given the facts.

3.3 Tables versus prose

Prose is the default. A table is used only when the tabular form is the only one that conveys the information without loss — typically a chronological event log or a strict numerical comparison across several dimensions. Releases, benchmarks, and pricing are integrated into prose paragraphs rather than tabulated, which keeps the reading rhythm continuous and avoids the visual clutter of repeated tables across editions.

4 Methodology and sources

This section is offered as transparency to external readers and as operational reference for production. It records how information enters the publication, where its limits lie, and how errors are corrected.

4.1 Information gathering

Information is gathered through web search using *Claude*'s tooling. Primary sources are preferred — vendor blog posts, regulatory filings, peer-reviewed papers, official press releases, exchange filings — over secondary aggregators. When a claim depends on a single source, that source is named explicitly. When sources disagree, the disagreement is reported rather than resolved by editorial fiat.

4.2 Knowledge limits and cutoff

Claude's reliable knowledge cutoff is the end of January 2026. For any month beyond that date, the report depends on web search rather than internal knowledge. This dependency is acknowledged once here rather than disclaimed inside every edition. The corollary is that a Trends report covering a month near the cutoff blends two information regimes; one covering a later month rests entirely on retrieved sources.

4.3 Error handling

Errors are inevitable. When a factual error is discovered after publication, the affected report is regenerated with the correction integrated in chronological context, and the new version replaces the previous one on the website. The replacement does not annotate the change within the document; the document presents the corrected record as if it had always been so. The previous version is not preserved publicly.

4.4 Retro-edition workflow

Reports may also be revised when a milestone that was missed at the original time of writing surfaces later. The process is the same as error correction: integration in chronological context, regeneration from scratch, replacement. The only marker that distinguishes a retro-edition is the file timestamp, which advances. The retro-edited content is presented as if written at full knowledge of the month covered. This convention is explicit so that the reader understands what a fresh timestamp on an old month means.

5 Glossary and brands

Two side files complement the Trends reports. The glossary defines technical vocabulary. The brand list registers structurally present companies. Both are loaded automatically by the document generator, so terms and brand names render consistently across every edition without manual intervention in the prose.

5.1 Glossary structure

Each glossary entry has a canonical term, a list of aliases, a stable HTML anchor, and a plain-language definition. Definitions avoid jargon by design: they target the reader who has heard the term but is not certain of its precise meaning. The anchor is generated once and never modified — older HTML pages with bookmarks to a term continue to resolve.

5.2 Brand list

A separate file lists the names of editors and platforms that recur across Trends. Names of products, of specific models, or of versions are not on the list — only the structurally present companies. Brand names are rendered in small caps bold throughout the publication for visual consistency. The list is conservative by intent; transient or geographically restricted entities are excluded until they prove durable.

5.3 Term and brand surveillance

At every production cycle, candidate terms and brands are surfaced for review before generation begins. A term qualifies for the glossary if it recurs across multiple editions and is not self-explanatory; a name qualifies for the brand list if the entity is structurally present in the ecosystem. Decisions are taken in advance so that a given term or brand renders consistently across all editions in which it appears, including those produced retroactively.

6 Site architecture

The public site at **ai-news.sliver.lu** is the primary point of access. Word and PDF documents are downloadable from each page. The site has no commenting, tracking, or advertising — only static pages.

6.1 Landing page rubriques

The landing page lists four sections. Section 01 — Foreword — leads to the editorial document. Section 02 — AI evolution by year — lists Trends reports in reverse chronological order, grouped by calendar year. Section 03 — Glossary — is a single permanent link. Section 04 — Project documents — currently contains only this Project *Prompt*; it may grow as additional methodological documents are added.

6.2 Per-edition page

Each Trends report has its own HTML page with the slug YYYY-MM. The slug is sortable lexicographically, which keeps the year groupings on the landing

page in correct order without external sorting logic. The page renders the report content with the same hierarchical structure as the Word document.

6.3 Companion PDF

Every Trends report, the Glossary, the Foreword, and this Project *Prompt* are accompanied by a PDF at /assets/pdf/. The PDF is a direct conversion of the Word document and preserves layout. It is regenerated and republished alongside the HTML page on every edition update.

6.4 Family domain routing

The publication is part of a small family of domains: **sliver.lu** and **hexi.lu**. Hyperlinks to these domains and their subdomains are routed to remain in the originating context — specifically, captured by the **Sliver.lu** Android WebView client. Hyperlinks to any other domain open in the system browser. This routing is invisible to the reader on a desktop browser; it matters only inside the Android client.

7 Naming and timestamps

All project files follow a single naming pattern, intended for sortability and visual consistency across years and document categories.

7.1 File naming pattern

The pattern: AI - Category - Subject (YYYY-MM-DD - HHhMM).ext. Categories in active use: Trends, Glossary, Brands, Project, Registry, Foreword. Spaces and hyphens only — never underscores. The Registry file is the only one without a timestamp in its name; it is the source of timestamps for all other files.

7.2 ISO date for sortability

Trends are subject-dated YYYY-MM rather than March 2026 or march-2026. Lexicographic sort then matches chronological sort with no special logic. The convention extends to all file paths and HTML slugs in the project. A directory listing reads in correct order; a list of links sorts itself.

7.3 Timestamp generation

The TS in filenames is the moment of generation in Luxembourg local time, derived programmatically. Any retroactive content edit produces a fresher TS — there is no concept of a frozen timestamp for a republished document. The generation timestamp is also embedded as document metadata, so the version of the stylesheet and the moment of compilation are recorded inside the file itself.

8 Production workflow

The workflow is monthly by default and event-driven on demand. Both flows converge on the same generation chain.

8.1 Monthly workflow

At the start of each month, the previous month is documented. Web research feeds a Trends report covering the month just ended. Glossary and brand candidates surfaced during research are reviewed before report generation. The Trends report is generated, validated, and deployed alongside any glossary or brand updates.

8.2 Retroactive edits

A retroactive edit follows the same chain but starts from a single trigger — a missed milestone or a factual correction — rather than from the calendar. The relevant Trends report is regenerated from scratch with the correction integrated in chronological context, the registry updated, and the site redeployed.

8.3 Validation

Each generated document passes structural validation before publication: filename consistency, heading-to-content level coherence, table column widths, document metadata. Validation is automated; a failure halts production rather than producing a degraded document. The intent is that every published file meets the same structural standard, regardless of when it was generated or by whom.