

Readability index

Tom Rochette <tom.rochette@coreteks.org>

December 21, 2025 — f7070728

The readability index is a metadata field added to articles on this blog to help readers quickly assess whether an article is appropriate for their level of expertise and available attention.

1 The scale

Articles are rated on a 0-5 scale based on their readability:

- **0 - Personal notes:** Content that is only meaningful to me. These are often shorthand notes, context-dependent references, or incomplete thoughts that lack the necessary background for others to understand.
- **1 - Cryptic:** Readable but highly condensed or assumes significant context. These articles may use jargon without explanation, reference obscure concepts, or present ideas in a very terse manner. You might be able to extract value, but it requires effort and potentially external research.
- **3 - Specialized/Expert audience:** Articles written for readers with domain expertise. These assume familiarity with technical terminology, concepts, and background knowledge in a specific field (e.g., machine learning, software architecture, AGI research). The writing is clear if you have the prerequisite knowledge.
- **5 - General audience:** Articles written to be accessible to anyone with general reading ability. These explain concepts from first principles, define technical terms when used, and don't assume specialized background knowledge. They're structured for easy consumption.

2 Purpose

The readability index serves several purposes:

1. **Reader efficiency:** Helps readers quickly determine if an article matches their current context and expertise level before investing time in reading it.
2. **Content discovery:** Makes it easier to filter or search for articles at the appropriate level - whether you want deep technical content or accessible introductions.
3. **Author awareness:** Forces me to be conscious about my target audience when writing, which can improve clarity and focus.
4. **Archive navigation:** As this blog contains a mix of polished articles and personal research notes, the index helps distinguish between content types.

3 How it's used

This index is subjective and represents my assessment at the time of writing. Your mileage may vary - what I consider a "3" might be a "5" for experts in that field or a "1" for beginners.

4 Why not a continuous 0-5 scale?

You might notice the scale uses discrete values (0, 1, 3, 5) rather than every integer from 0 to 5. This is intentional:

- **2** would fall between “cryptic” and “specialized” - a fuzzy middle ground that’s hard to define
- **4** would be between “specialized” and “general” - again, unclear distinction

The four-point scale provides enough granularity to be useful without creating artificial precision. If an article truly feels intermediate, I’d likely rate it at the lower level (1 or 3) since it’s better to underpromise and overdeliver on accessibility.

5 Relationship to article status

The readability index is independent of the article’s `status` field:

- `status: draft` - Article is incomplete or actively being written
- `status: in progress` - Article is being updated and refined
- `status: finished` - Article is complete and unlikely to be revised

An article can be `status: finished` with `readability: 0` (polished personal notes) or `status: draft` with `readability: 5` (an in-progress accessible introduction). They measure different dimensions of the content.

6 Evolution

Like any other metadata on this blog, readability ratings may change over time as I revisit articles or as my sense of what constitutes each level evolves. The index is a tool for navigation, not a rigid categorization system.