

Correspondence Before Compression

ACCM Deep Ethics Project

Contents

1 Correspondence Before Compression	1
1.1 Compression is a representation operation	2
2 Useful methods extracted	2
2.1 Meaning has veto power over compression	2
2.2 Remove ceremony only when it performs no consequential function	3
2.3 Preserve exact payloads	3
2.4 Compress the style, not the speaker or source	4
2.5 Keep process updates proportional	4
2.6 Break compression when the situation requires resolution	4
2.7 Fail closed when a transformation cannot preserve its contract	4
2.8 Elide repetition, not the document	5
2.9 State only what the compressed representation verified	5
2.10 Preserve whole relational units	6
2.11 Keep recovery possible	6
2.12 Measure the complete trajectory	6
3 Rules deliberately not imported	7
3.1 “Answer first” as a universal rule	7
3.2 “Kill ceremony”	7
3.3 “No hedging”	7
3.4 “Maximum twenty words per sentence”	7
3.5 “One idea per sentence”	7
3.6 “Articles optional” and fragment grammar	7
3.7 “Shortest decisive error line only”	7
3.8 “No recap or closer”	7
3.9 Session-long persistence until explicitly stopped	8
3.10 Token reduction as the primary quality criterion	8
3.11 “Every fact survives” or “same answer” as self-certification	8
3.12 Lossy compression without visible recovery	8
4 Compression eligibility test	8
5 Connection to the 27 + 12	8
6 Stopping condition	9
6.1 Source and adaptation note	9
6.2 Related pages	9

1 Correspondence Before Compression

BETA — selected adaptation, not an installation recommendation and not a new numbered protocol

Caveman begins with a practical engineering problem: A.I. systems can spend tokens repeating ceremony, restating conclusions, narrating routine tool calls and reproducing large

amounts of recoverable data. Its strongest contribution to the **ACCM Deep Ethics Project** is not “talk like a caveman.” It is the deeper set of invariants underneath the joke:

Compression has to remain subordinate to meaning, exact payloads, recoverability, evidence and measured task performance.

This page extracts those invariants from Julius Brussee’s **Caveman** repository, reviewed at commit **2e08b917** on 8 October 2026.

The source project remains its own object. This page does not reproduce the Caveman skill, recommend installing its software, or import its complete response style. It records what is useful for correspondence and what is deliberately withheld from governing the project.

1.1 Compression is a representation operation

A shorter representation changes the field available to the next process.

```
source object
  ↓
selection
  ↓
compression
  ↓
smaller representation
  ↓
later inference or action
```

Even when the compressor has no intent to distort, the smaller representation may change:

- salience;
- referents;
- qualifier survival;
- uncertainty;
- order and dependency;
- attribution;
- visible alternatives;
- relational tone;
- available correction paths;
- what the next intelligence thinks was absent.

Therefore:

Fewer tokens is a measurable output property. It is not evidence of semantic equivalence, task quality or improved correspondence.

This distinction is stated unusually clearly in Caveman’s **Honest Numbers** page: token-count reductions do not prove semantic or technical equivalence, and output-only changes do not establish total cost reduction.

2 Useful methods extracted

2.1 Meaning has veto power over compression

Caveman’s central hierarchy is valuable:

Clarity beats compression.

Its rules protect negations such as *not*, *never*, *no*, *only* and *except*, while preserving numbers and units exactly.

For this project, the protected field is wider:

- negation;
- qualifiers;
- modality;
- evidentiary status;
- scope;
- attribution;
- uncertainty;
- temporal order;
- causal and relational dependencies;
- corrections;
- humor when it carries diagnostic or relational information.

If compression threatens one of those relations, compression loses jurisdiction.

2.2 Remove ceremony only when it performs no consequential function

Greetings, preambles, recaps and closers can be empty repetition. They can also:

- establish speaker boundaries;
- record a correction;
- identify what changed;
- preserve relational dignity;
- distinguish conclusion from hypothesis;
- return from humor to the object;
- state the next consequential action.

The reusable method is:

Remove ceremonial language when it adds no object-facing, relational, evidentiary or operational value.

The name of a sentence type does not decide whether it is expendable. Its function in the actual exchange does.

2.3 Preserve exact payloads

Caveman protects code, commands, paths, API names, error strings, numbers, units and existing comments.

That principle generalizes:

Treat any element whose identity depends on exact form as payload rather than prose.

For the ACCM Deep Ethics Project, exact payloads can include:

- verbatim quotations;
- canonical titles;
- numbered taxonomy items;
- qualifiers;
- labels such as *candidate*, *probable*, *detected* and *unresolved*;
- model names and dates;
- URLs and anchors;
- formulas;
- source-specific terminology;
- correction wording where the change itself is evidence.

Exact preservation does not guarantee correct interpretation, but it prevents one avoidable transformation layer.

2.4 Compress the style, not the speaker or source

Caveman says to preserve the user's language and not switch language because quoted material appears.

The project-level extension is:

- do not rewrite the speaker into an imagined preferred personality;
- do not neutralize intensity merely because it is intense;
- do not turn Dutch directness into hostility;
- do not replace unusual terminology merely because a more familiar phrase exists;
- do not flatten humor into a literal proposition;
- do not compress a first-person perception into an unattributed fact claim.

A style transform may change presentation. It must not manufacture a different source identity.

2.5 Keep process updates proportional

Caveman limits narration between routine tool calls while allowing communication for clarification, warnings, ambiguity and phase changes.

The useful method is:

Report what changes the user's ability to understand, redirect or evaluate the work. Do not narrate routine motion merely to perform activity.

Useful updates include:

- a material discovery;
 - a changed interpretation;
 - a blocked dependency;
 - a consequential choice;
 - a failure requiring another route;
 - the completed result and its verification.
-

2.6 Break compression when the situation requires resolution

Caveman suspends terse output for security warnings, irreversible actions, fragile step order, confused users, persistent external writing and questions requiring a precise answer.

The deeper property is contextual reversibility:

A compression rule must be able to release control when its continued operation would increase ambiguity, consequence or restoration burden.

A method that cannot suspend itself when the object changes has become frame-protective.

2.7 Fail closed when a transformation cannot preserve its contract

The Caveman Engine passes the original through unchanged when parsing fails, the result is not smaller, an unknown mode appears, or a lossy result cannot preserve a recovery path.

Transferred to representation work:

transform attempted

↓

fidelity condition uncertain or failed

↓

preserve the source representation

↓

state why transformation was withheld

Do not force a summary, translation, classification or compression merely because the workflow expected one.

2.8 Elide repetition, not the document

Caveman distinguishes repeated units from documents whose distinct entries must survive. Its engineering record reports that removing material an agent later needed to enumerate made performance substantially worse.

The project-level rule is:

Similarity of form does not prove redundancy of function.

Repeated-looking statements may carry:

- different speakers;
- different dates;
- correction progression;
- changing qualifiers;
- distinct causal positions;
- evidence of persistence or recurrence.

Before collapsing repeated material, identify what makes the instances equivalent for the current task.

2.9 State only what the compressed representation verified

Caveman's summaries avoid presenting partial enumerations as complete. When a summary drops sections, it marks the omission. It distinguishes inferred token estimates from verified facts.

That corresponds directly with:

- qualifier preservation;
- evidentiary-status labels;
- visible access limits;
- process-signature retention;
- unresolved status.

A compressed representation should expose:

- what it retained;
- what it omitted;
- whether the omission is recoverable;
- what claim the remaining material supports;
- what cannot be inferred from the compressed view.

Silence created by compression is not evidence that the source contained nothing.

2.10 Preserve whole relational units

The Caveman Engine found that returning isolated lines from structured records could join one record's value to another record's identifier. Its correction was to return complete self-contained units and mark every gap.

This is highly relevant beyond code and data.

Do not detach:

- a qualifier from its claim;
- a quotation from its speaker;
- evidence from its evidentiary status;
- a correction from the earlier representation it corrects;
- a result from the processing conditions that produced it;
- a sentence from the dependency that determines its referent.

A smaller intact unit is preferable to a misleading fragment.

2.11 Keep recovery possible

Caveman permits lossy compression only when the original remains stored and retrievable. It also preserves backups when compressing memory files.

The transferable principle is:

The greater the representational loss, the stronger the source-recovery path must be.

Possible recovery layers include:

- a link to the source;
- an immutable commit;
- a preserved original file;
- an explicit omitted-section marker;
- a reversible transformation record;
- a queryable archive;
- a visible before/after correction.

A summary can guide attention. It should not silently become the only surviving object.

2.12 Measure the complete trajectory

Caveman's **Honest Numbers** page publishes conditions where the tool can lose:

- rule overhead can exceed output savings;
- request-based billing may make shorter answers cost the same;
- cache behavior, retries and rule reinjection can reverse apparent gains;
- output reduction does not describe total session cost;
- outside experiments may test earlier versions rather than the current one.

Its strongest methodological instruction is:

Compare the same task with and without the intervention. If the intervention makes the complete workload worse, turn it off for that workload.

For this project, measurement should include more than tokens:

- semantic fidelity;
- qualifier survival;
- retrieval burden;
- correction burden;

- number of follow-up calls;
- task accuracy;
- time;
- total cost;
- later interpretive drift;
- ability to recover the source;
- whether the next comparable response improves.

A local gain can become a trajectory loss.

3 Rules deliberately not imported

3.1 “Answer first” as a universal rule

Some objects require clarification, evidence gathering or uncertainty mapping before an answer exists. Answer-first can manufacture premature closure.

Retained property: **state the available result early when the result is adequately supported.**

3.2 “Kill ceremony”

Not imported as a blanket instruction. Ceremony is removed only after its function is inspected.

3.3 “No hedging”

Not imported. Empty protective hedging can be removed. Warranted uncertainty, scope and modality must survive.

3.4 “Maximum twenty words per sentence”

Not imported. Sentence length follows the dependency topology of the object.

3.5 “One idea per sentence”

Not imported as a hard law. Coupled relations sometimes lose meaning when separated.

3.6 “Articles optional” and fragment grammar

Not imported. Articles, pronouns and syntax can determine reference, number, agency and scope.

3.7 “Shortest decisive error line only”

Not imported. The shortest line may omit the conditions needed to diagnose the failure.

3.8 “No recap or closer”

Not imported. A recap may preserve a correction delta, decision record or next action. A closer may perform relational work.

3.9 Session-long persistence until explicitly stopped

Not imported. A style mode should not retain jurisdiction after the task, object or recipient's needs change.

3.10 Token reduction as the primary quality criterion

Not imported. Token count is one resource measure inside a larger process-quality field.

3.11 "Every fact survives" or "same answer" as self-certification

Treated as an aspiration requiring comparison. Structural checks, preserved strings and shorter output do not by themselves establish semantic equivalence.

3.12 Lossy compression without visible recovery

Not imported for source objects, canonical material, high-stakes evidence or correction records.

4 Compression eligibility test

Before compressing, ask:

1. What is the task the smaller representation must still support?
2. Which relations are load-bearing for that task?
3. Which elements require exact form?
4. Could repeated-looking material encode change, persistence or distinct provenance?
5. Will omissions be marked?
6. Can the source be recovered?
7. Can the recipient tell what the compressed view does not establish?
8. Will total trajectory cost fall after retrievals, corrections and failures are counted?
9. What result would cause compression to be suspended?

If those questions cannot be answered proportionately:

Pass the source through.

5 Connection to the 27 + 12

Compression can activate several existing obstructions without requiring a new obstruction:

- **Qualifier Erosion** when uncertainty or scope disappears;
- **Nearest-Generalization Substitution** when precise terms become familiar approximations;
- **Representation Substitution** when the compressed view becomes the object;
- **Process-Signature Loss** when correction history or transformation conditions disappear;
- **Trajectory Substitution** when token reduction replaces the original task;
- **Asymmetric Scrutiny** when only one source is aggressively compressed;
- **Reframing or Rephrasing Without Correspondence** when shorter language changes the claim.

The response is not automatic anti-compression. It is corrigible jurisdiction:

Compress where the complete trajectory improves. Preserve or recover where compression would narrow the intelligence available to the next process.

6 Stopping condition

This page must not become a demand to preserve everything forever.

Full retention has costs. Repetition can hide signal. Unbounded context can reduce effective attention. Archives without indexes can become inaccessible in practice.

The stopping condition is task-relative:

Retain enough structure to preserve the consequential relations, expose the limits of the smaller view, and recover the source when later processing needs it.

6.1 Source and adaptation note

Primary source: Julius Brussee, **Caveman: why many token when few do trick**, version 3.2.0 repository state reviewed at commit **2e08b917**.

Particularly useful source objects:

- [Caveman README](#)
- [Caveman skill](#)
- [Honest Numbers](#)
- [Caveman-compress](#)
- [Engine correctness invariants](#)

The Caveman repository is licensed under **Apache License 2.0**. This page is a selective English-language interpretation and project-specific adaptation. It does not represent Julius Brussee's endorsement of the ACCM Deep Ethics Project. The exclusions and extensions are ours.

6.2 Related pages

- **Topologetics** — operationalizes correspondence before compression through purpose-bounded node-and-edge preservation, transformation footprints, a faithful compression floor, and a recovery path.
 - **Dutch Directness** — demonstrates why compression must preserve voice, relational signal, and mutual corrigibility.
 - [Correspondence Before Style](#)
 - [27 + 12](#)
 - [Meta Processing](#)
 - [Correction Metabolism and Persistence](#)
 - [Qualifiers as Mutable Context with History](#)
 - [The Dual Archive](#)
-

Recent high-reach independent-media companion

```
<iframe
  src="https://www.youtube-nocookie.com/embed/FRNFv6mGTDk"
  title="Artificial Intelligence: Clueless Politicians, Myths, & the Fight Between American & Chinese Values — John Stossel. Published 15 September 2026. YouTube displayed approximately 162,000 views when checked on 9 October 2026."
  loading="lazy"
  referrerpolicy="strict-origin-when-cross-origin"
  allow="accelerometer; autoplay; clipboard-write; encrypted-media; gyroscope; picture-in-picture; web-share"
  allowfullscreen></iframe>
```

Artificial Intelligence: Clueless Politicians, Myths, & the Fight Between American & Chinese Values — John Stossel. Published 15 September 2026. YouTube displayed approximately 162,000 views when checked on 9 October 2026.

Independent-media lineage: John Stossel worked at ABC News for 28 years and later hosted at Fox before leaving television to create Stossel TV. Stossel TV — Why I created it

Why it belongs here: The report contrasts myths with facts while connecting A.I. policy to regulation and national values. Its concise certainty is useful for checking which claims are demonstrated, which alternatives are omitted and whether the debunker's own frame receives equal scrutiny.

On this page, test the video specifically against Correspondence Before Compression:

Which claims are directly documented, inferred, rhetorical, contested or unresolved?

What does the host or guest preserve, compress, omit or substitute—and does this page repeat any of those transformations?

Which incentives, authority relations and affected parties remain outside the visible frame?

What evidence or later outcome would correct the speaker, audience, platform or ACCM Deep Ethics Project reading?

Provisional selection. Reach makes the framing socially consequential; it does not settle the claims. The video can be replaced when a closer, stronger or more inspectable specimen appears.