

WAYS OF KNOWING

Youth Edition

Twenty-four questions for learning to reason with evidence

Juan Segurola

Version v1.0 · 22 September 2026

© 2026 Juan Segurola. All Rights Reserved.

How to read this volume

Each chapter asks what we can really conclude, checks the evidence, marks where caution is needed and ends with a small research-style question. The language is shorter than in the scientific edition, but the evidential boundary is preserved.

Contents

KNO	Chapter	Page
KNO-001	A Mark on a Stone	4
KNO-002	Counting Before Numbers	5
KNO-003	From Marks to Symbols	6
KNO-004	From Symbols to Writing	7
KNO-005	Oral Language and Transmission	8
KNO-006	Maps and Spatial Representation	9
KNO-007	Calendars and Time Measurement	10
KNO-008	Instruments That Extend the Senses	11
KNO-009	Archives, Libraries and Collective Memory	12
KNO-010	Classification as a Knowledge Tool	13
KNO-011	Measuring Before Theorising	14
KNO-012	Experiments and Variable Control	15
KNO-013	Models, Maps and Reality	16
KNO-014	Causality and Correlation	17
KNO-015	Uncertainty and Probability	18
KNO-016	How Evidence Is Built: Putting the Tools Together	19
KNO-017	Is Mathematics a Language?	20
KNO-018	Images, Diagrams and Visualisation	21
KNO-019	Data, Tables and Records	22
KNO-020	Review, Consensus and Disagreement	23

KNO	Chapter	Page
KNO-021	Open Science and Reproducibility	24
KNO-022	Computation as External Memory	25
KNO-023	How Groups Know More Than Individuals	26
KNO-024	Lost Knowledge and Preservation Bias	27

KNO-001 · A Mark on a Stone

The question

What can we really conclude?

Short answer

A groove is first a physical trace. Microscopy, fracture and experiments may show that it was deliberately made, but intention does not by itself establish a symbol. A row of notches may be regular without counting anything; a tally may record units without representing language; writing is a broader conventional system. Ancient shells and ochres show deliberate patterning deep in the Pleistocene, while objects such as the Border Cave and Ishango bones invite—not compel—notational interpretations.

Evidence check

Microscopy, chemistry, experiments and context can establish that many ancient marks were deliberately produced and can often reconstruct tools, gestures and sequence. Cognitive experiments show that external states can redistribute memory and comparison. Documented administrative and writing systems demonstrate later conventional uses.

Where to be careful

For most early objects, we do not know what a mark meant, whether anyone revisited it, whether its rule was shared, or whether its units were numerical. We also cannot measure the vanished record on wood, fibre, skin or temporary surfaces. The earliest surviving example is therefore not necessarily the origin of the practice.

Think like a researcher

Look at three marks around you. Which facts can you observe directly, and which meanings would you only be guessing?

Remember

A deliberate mark can preserve a state beyond one brain, but symbol, tally, number and writing each require additional evidence.

KNO-002 · Counting Before Numbers

The question

What can we really conclude?

Short answer

Quantity can be handled before a language has words for every exact number. An approximate number system distinguishes roughly more from less, while a separate capacity tracks only a few individual items. One-to-one correspondence answers a different question: does each object in one set have a partner in another? A tally can preserve repeated events without naming a cardinal total, and a finger or body sequence can provide an ordered scaffold before a portable numeral exists.

Evidence check

Controlled experiments distinguish approximate magnitude, small-set tracking and learned exact-number practices. Developmental studies show that reciting count words precedes full understanding of cardinality. Cross-cultural studies demonstrate exact nonverbal performance under some conditions and the importance of words or external devices when quantities must be retained and transformed.

Where to be careful

Think like a researcher

Put 7 small objects on a table. Find two different ways to keep track of the exact quantity without saying the word seven.

Remember

Quantity can be managed without numerals, but exact counting requires a learned mapping among units, order and cardinal value; a repeated ancient mark proves none of those mappings by itself.

KNO-003 · From Marks to Symbols

The question

What can we really conclude?

Short answer

A scratch is not a symbol merely because it is visible or deliberate. Smoke can indicate fire through a causal relation; a portrait can resemble a person; a pointing gesture can direct attention; a conventional sign can refer even when it neither resembles nor is caused by its referent. These relations can overlap. A map icon may resemble an object and also acquire a conventional meaning.

Evidence check

Developmental studies distinguish an object's physical identity from its representational role and document early joint attention and pointing. Interaction experiments show that feedback and common ground can stabilise novel signs. Transmission experiments show that learnability and communicative usefulness reshape systems across learners.

Where to be careful

Think like a researcher

Invent a symbol with a friend. What rule must both of you share for it to work?

Remember

A mark becomes a symbol not through its shape, but because a community can repeatedly use, recognise, transmit and correct a relation between that form and what it stands for.

KNO-004 · From Symbols to Writing

The question

What can we really conclude?

Short answer

A mark can be deliberate without being a symbol, and a symbol system can be powerful without encoding language. What converts a symbol system into writing, under the definition used here, is phonetic transfer: the discovery that a sign can stand for the sound of its name rather than for its meaning. Before that step a system grows only by adding signs; after it, the inventory stabilises and then contracts while the range of expressible content stops having a limit. Southwest Asia preserves the only stratified sequence in which tokens, sealed envelopes, impressed numerical tablets and proto-cuneiform can all be examined together, although that order records first attestation and not descent.

Evidence check

Stratified fourth-millennium material from southwest Asia preserves the only sequence in which tokens, sealed envelopes, impressed numerical tablets and proto-cuneiform can be examined together, though that order is one of first attestation rather than descent. Documented modern inventions establish the pace: the Cherokee syllabary was in use by the early 1820s and received council recognition in 1821, and the Vai syllabary followed within about a decade. Numerical, mathematical, chemical and musical notations show that a closed system can be unlimited within its domain and still encode no language at all.

Where to be careful

Think like a researcher

Try to record a short spoken sentence using only pictures. What parts of the language are hard to represent?

Remember

Writing begins when a sign can carry the sound of its name rather than its meaning; every other recording system is an alternative rather than a failed attempt.

KNO-005 · Oral Language and Transmission

The question

What can we really conclude?

Short answer

Oral transmission is not one practice but a family of techniques, and they are not comparable technologies. Casual report degrades quickly, because nothing in the practice resists that. Institutionally checked recitation is the strongest oral fidelity regime known. What produces the difference is constraint, not effort.

Evidence check

Cognitive work on strong traditions establishes that stability comes from many individually weak constraints acting at once — metre, melody, formula, imagery and narrative sequence — which together often leave one option where each alone leaves many. Parry and Lord's fieldwork shows a singer producing the same plot at similar length in substantially different words without treating it as failure. Vedic permutation recitation shows a community making error detectable rather than merely improbable.

Where to be careful

Think like a researcher

Tell a short story to someone, ask them to retell it, and compare what changed.

Remember

Strong oral fidelity is produced by constraining what can be said rather than by trying harder to remember it, which is why it favours verse and fails on arbitrary detail.

KNO-006 · Maps and Spatial Representation

The question

What can we really conclude?

Short answer

Spatial knowledge, external spatial representation and cartography are three separable things, and an expert navigator may hold the first at a high order while using no graphic artefact during the voyage. Place cells and grid cells are robust findings about neural activity, but they do not establish that human spatial thought takes the form of a graphic map; route knowledge can support successful travel without being a degraded survey map. Marshallese wave charts and the Caroline Islands technique known as etak show that spatial reasoning can be organised around models of ocean behaviour and around reference frames other than the fixed frame of a scaled plan. Geometry then imposes a hard limit: because a sphere and a plane differ in Gaussian curvature, no flat map of a sphere can preserve everything, so a projection encodes a decision about which distortions to accept.

Evidence check

- Grid cells are an established finding from single-unit recording in rats (Hafting et al. 2005). O'Keefe and Nadel proposed that the hippocampus supports an allocentric representation (O'Keefe and Nadel 1978). - Local isometries preserve Gaussian curvature, the Theorema Egregium (do Carmo 1976, p. 234).

Where to be careful

Think like a researcher

Draw a map of your room. Decide what you want to keep accurate: shape, distance, direction or useful landmarks.

Remember

No flat map of a sphere preserves every metric property, so the question to ask of one is not whether it distorts but which distortions it accepts.

KNO-007 · Calendars and Time Measurement

The question

What can we really conclude?

Short answer

Twelve mean synodic months come to 354.36708 days, 10.87511 short of a tropical year, so a lunar calendar walks through the seasons and the Islamic calendar does that by design rather than by failure. Drop the lunar month and keep the year and the Egyptian civil year follows: twelve months of thirty days plus five epagomenal days, 365 in total, with no intercalated day and a slower drift of its own. Keep both and the price is intercalation, because nineteen tropical years and 235 synodic months differ by 0.08704 days, about 2.09 hours, which is close enough to build a nineteen year cycle on and not close enough to make it exact. The Julian year overshoots by 0.00781 days and the Gregorian by 0.00031, so 1582 reduced a residual rather than removing one, and it revised the lunar tables for Easter as well.

Evidence check

For epoch 2000 a standard astronomical reference gives a mean tropical year of 365.24219 days and a mean synodic month of 29.53059 days, and stresses that both are approximate means that change slowly. Everything this article derives from them is its own arithmetic, marked as such: twelve synodic months fall 10.87511 days short of a year, and nineteen tropical years differ from 235 synodic months by 0.08704 days. The Egyptian civil year ran 365 days with no intercalated day, and the Islamic calendar keeps months in relation to lunations without keeping its year with the Sun.

Where to be careful

Think like a researcher

Track the Moon for a week. Notice how a repeating cycle can be used to organise time.

Remember

A lunar calendar is not a failed solar one: astronomical periods do not fit exactly, and calendars that track them must choose how to handle the mismatch.

KNO-008 · Instruments That Extend the Senses

The question

What can we really conclude?

Short answer

Spectacles aid a sense, a telescope extends one, and a thermometer does neither: it converts a quantity no eye registers into a length the eye can read, and the observer infers. Believing such a reading is believing every link of that chain, and no point along it is settled by looking. Galileo's telescope was accepted through independent replication rather than on sight. Calibration is circular at the outset, because justifying an instrument needs the law that measuring with it is meant to establish; that circle was broken by iteration rather than by decree, and a published critique holds that reasoned argument can break it.

Evidence check

Experts at Bologna reported seeing little or nothing through Galileo's telescope in April 1610, the Collegio Romano withheld confirmation until October, and credibility was established by independent replication; the source for this was read only as far as its accessible pages go. The circularity of validating a measuring instrument without the law that justifies it is called the problem of nomic measurement and is attributed to a 2004 study of thermometry, which argues the circle was worn away by iterative self correcting refinement rather than broken by decree. Franklin argues from the gravitational wave episode that reasoned argument broke the regress.

Where to be careful

Think like a researcher

Choose a thermometer or ruler. Write down the chain from the thing being measured to the number you read.

Remember

An instrument rarely extends a sense; it ends a chain of conversions, and believing the reading means believing every link.

KNO-009 · Archives, Libraries and Collective Memory

The question

What can we really conclude?

Short answer

Records, archives and libraries are filed together and run on incompatible principles, so each is silent about something different: an archive about people who generated no records, a library about what never entered the circuits it collects from, a record about whatever its purpose did not require. What enters an archive is decided by appraisal, whose own professional framing is what to save and what to destroy, and the surviving series reads as complete because destruction leaves no gap in it. Trouillot separates four moments where silence enters, which matters because they are four failures with four remedies and better searching fixes only some. Description then governs what can be retrieved, which is a constraint and not a second appraisal, and it can be redone where appraisal cannot.

Evidence check

Appraisal is defined by the professional dictionary as identifying materials offered to an archives that have sufficient value to be accessioned, with the debate behind it framed as what to save and what to destroy. Archival theory has described the archivist moving from passive curator to active appraiser across shifting paradigms. Trouillot locates silence at four moments: fact creation, fact assembly, fact retrieval and retrospective significance.

Where to be careful

Think like a researcher

Imagine you can save only five objects from today for a museum. What future story would your choices create?

Remember

An archive's surviving holdings cannot reveal all their own omissions: what survives is evidence about the keeping as much as about the kept.

KNO-010 · Classification as a Knowledge Tool

The question

What can we really conclude?

Short answer

Classification looks like description and is not. A scheme is a set of boxes built so that some further operation can run, and its categories were shaped by that operation. Bowker and Star name three properties an ideal classification would have and report that none of the working systems they examined satisfies all three. A scheme that is not complete has to place what its categories did not anticipate, and where a residue sits tells a reader how that scheme handles it.

Evidence check

Bowker and Star define a classification system as a set of boxes into which things can be put to then do some kind of work, set out three properties an ideal classification would have, and report that none of the working systems they examined satisfies all three. They observe that the residual phrase runs throughout the ICD but never as a top-level category, and they name torque for the mutual twisting of a biography and a process of metrication. Wells and colleagues report that species remain controversial in definition, delimitation and reality, and propose treating them as a heuristic.

Where to be careful

Think like a researcher

Sort ten objects in two different ways. How does the purpose change the categories?

Remember

A classification is built to make an operation possible, so its categories cannot be read as a neutral picture of the world.

KNO-011 · Measuring Before Theorising

The question

What can we really conclude?

Short answer

Two orders both exist and neither is the order this field follows: sometimes theory names a quantity and an experiment goes looking for it, sometimes the numbers accumulate until a theory has to reproduce them. Whether classifying already counts as measuring depends on which definition is in force, and the permissive and the restrictive answer are both named here and neither adopted. An instrument has to be trusted before anything exists that could check it, and in the episodes discussed the circle was not avoided rather than broken. A formula can be written to hit a target that measurement had already set: a communication about long wavelength behaviour came first and the equation followed, which is an order of events and not an account of what caused what.

Evidence check

The permissive definition takes measurement to be the assignment of numerals to objects or events according to rules, and separates four scale types by the empirical operation each supports. The restrictive objection is that under such a view any rule at all already defines a nominal scale, so whether the attribute is quantitative is never put to test. Competing radiation laws differed only at long wavelengths; a residual ray experiment across four cavities found none of them reproducing the data at all temperatures; a communication about the limit at long wavelengths preceded the equation that worked, and the two results were presented at the same meeting.

Where to be careful

Think like a researcher

Measure the same object several times. What stays stable and what changes?

Remember

Numbers often arrive before the theories that explain them, so ask what a measurement was answerable to before asking whether it is true.

KNO-012 · Experiments and Variable Control

The question

What can we really conclude?

Short answer

In an assigned comparative design, an experiment buys one thing: the investigator decides who receives what, so the reason one group got the treatment has a known answer instead of an unknown one. Experiments that assign nothing lie outside what follows. Everything else in the method protects that property. The decision itself is therefore the vulnerable point, and the word randomisation covers two different jobs at it: generating a sequence nobody can predict, and concealing that sequence from whoever enrolls participants.

Evidence check

Fisher titled section 9 of *The Design of Experiments* Randomisation; the Physical Basis of the Validity of the Test, and argued there that the simple precaution of randomisation suffices to guarantee the validity of the test of significance. That is a claim about the validity of a test and not about the balance of any particular experiment, and the surrounding context also makes clear that validity is not a claim that randomisation removes every other avoidable source of error. The CONSORT 2010 checklist carries sequence generation, allocation concealment mechanism and blinding as three separately numbered items.

Where to be careful

Think like a researcher

Design a fair test with two groups. Change one planned factor and keep the comparison as clear as possible.

Remember

In the assigned comparative designs treated here, assignment replaces the unknown reason why these units and not those with a procedure, and the rest defends that substitution.

KNO-013 · Models, Maps and Reality

The question

What can we really conclude?

Short answer

A model represents a chosen piece of the world, not the world, so asking whether it is good means asking what it was built for. Leaving something out is not the same as getting it wrong: an abstraction omits, an idealisation distorts on purpose, and only the second asserts a falsehood deliberately. The test is intent, because what a modeller would fix on noticing is an error and what they would defend has to earn its place. One developed position holds that the most unifying laws are the least descriptive, recorded here and not adopted.

Evidence check

The reference formulation states that many scientific models are representational and stand in for a selected part or aspect of the world, and that idealised models involve a deliberate simplification or distortion aimed at tractability. Cartwright is reported, from the publisher's description of the book rather than from its argument, as holding that fundamental explanatory laws do not describe regularities in nature while localised laws and theoretical entities may still be read realistically. Levins wrote in 1966 that generality, realism and precision cannot be maximised together in a manageable model, with realism there meaning retained detail.

Where to be careful

Think like a researcher

Make a simple model of something familiar. List one useful simplification and one thing the model cannot answer.

Remember

A model is a deliberate misdescription with a job, so ask what job, what was left out, what was distorted on purpose, and what that cost.

KNO-014 · Causality and Correlation

The question

What can we really conclude?

Short answer

Association alone does not tell us which causal story generated a pattern. A cause may produce an outcome, a common cause may produce both variables, the outcome may influence the supposed exposure, or selection and measurement may create the dependence. Causal inference therefore begins by defining the intervention or contrast of interest and asking what design makes that contrast identifiable. Randomisation can protect against baseline confounding when properly implemented.

Evidence check

Association and causal effect are distinct targets. Randomisation, causal diagrams, potential outcomes, negative controls, instrumental variables and target-trial emulation each make different parts of the causal problem inspectable. Their evidential force comes from design plus assumptions, not from the statistical estimate alone.

Where to be careful

Observed data cannot by themselves establish causal direction, the absence of unmeasured confounding, a valid exclusion restriction or the correctness of an assumed causal graph. Those questions require external knowledge, design information and tests whose scope is itself limited.

Think like a researcher

Think of two things that often happen together. List at least two possible explanations besides direct causation.

Remember

Causal inference begins when an association is linked to a clearly defined intervention question through explicit, testable design and assumptions.

KNO-015 · Uncertainty and Probability

The question

What can we really conclude?

Short answer

Uncertainty is not one thing. A measurement can be uncertain because of calibration, repeatability or resolution; an estimate can vary across samples; model parameters can be imprecise; structural choices can change the result; future observations can add prediction uncertainty; and some unknowns lie outside the probability model entirely. Confidence intervals and Bayesian credible intervals answer different questions because they rely on different inferential frameworks. A p-value is not the probability that a hypothesis is true.

Evidence check

Uncertainty can arise from measurement, sampling, parameters, prediction, model structure and judgement, and these sources are not interchangeable. Statistical intervals and probabilities have meanings fixed by their inferential framework, while sensitivity analysis exposes dependence on assumptions.

Where to be careful

No finite interval automatically contains every relevant unknown. Structural omissions, unmodelled mechanisms and genuinely unforeseen possibilities can remain outside a probability model even when the arithmetic inside that model is correct.

Think like a researcher

Repeat a measurement and write down what could make the result uncertain.

Remember

Represent uncertainty in the form that matches the claim, and state explicitly which important unknowns the calculation does not contain.

KNO-016 · How Evidence Is Built: Putting the Tools Together

The question

What can we really conclude?

Short answer

Evidence becomes stronger when the links between claim, observation, design, model and inference are explicit enough to test. A precise claim must be vulnerable to failure. Measurements are produced by instruments and procedures, so their calibration and transformation matter. Study design determines which alternative explanations can be excluded; models connect observations to quantities or mechanisms; causal claims require assumptions beyond association; uncertainty belongs inside the result rather than after it.

Evidence check

Evidence is built by linking precise claims to observations through measurement, design, models and inferential assumptions. Replication, reproducibility, triangulation, directness and transparency strengthen different links, especially when their failure modes are not shared.

Where to be careful

No audit can guarantee that every relevant alternative has been imagined, that every source of bias is independent or that a well-supported model is the only possible explanation. Evidence can therefore become stronger without becoming logically final.

Think like a researcher

Take one claim from a news story and draw a chain: claim → evidence → method → assumptions.

Remember

Strong evidence is not a pile of results; it is a chain of support whose assumptions, error modes and tests can be inspected.

KNO-017 · Is Mathematics a Language?

The question

What can we really conclude?

Short answer

Mathematics contains formal languages, but mathematical practice is broader than symbol strings. Formal notation has syntax: rules determine which expressions are well formed and which transformations are permitted. Meaning enters through interpretation, definitions, conventions and connections to mathematical structures or empirical systems. Proof is not merely the presence of symbols; it is a rule-governed argument whose validity depends on the system in which it is read.

Evidence check

Mathematics uses formal languages with syntax and interpretation, and notation can make reasoning more inspectable and computationally tractable. Proof, meaning, modelling and explanation nevertheless depend on practices and interpretations that are not exhausted by symbol strings.

Where to be careful

The language analogy does not settle philosophical questions about what mathematics is, why mathematical structures apply so effectively in science or whether mathematical entities are discovered, invented or described in some other way.

Think like a researcher

Explain a mathematical idea once with symbols and once with words or a picture. What changes?

Remember

Mathematics contains languages and uses them powerfully, but mathematical knowledge is broader than the formal symbol systems in which it is written.

KNO-018 · Images, Diagrams and Visualisation

The question

What can we really conclude?

Short answer

A visual display can reduce cognitive work because spatial relations can make patterns explicit that prose leaves implicit. But visual channels differ in precision: position on a common scale is generally decoded more accurately than area or angle, and colour scales can create false boundaries or unequal emphasis. Similar summary statistics can conceal very different data structures, which is why plots and numerical summaries complement one another. A graph is still a transformation, not raw reality.

Evidence check

Visual representations transform data into spatial and perceptual encodings. Those choices affect what patterns are easy to detect, how accurately quantities are decoded and which aspects of acquisition or processing remain visible to the reader.

Where to be careful

No single visual encoding is optimal for every task, audience or dataset, and an attractive or reproducible rendering does not establish the correctness of the scientific interpretation. Important structure can remain hidden outside the chosen view.

Think like a researcher

Draw the same small dataset in two different graph types. Which pattern becomes easier or harder to see?

Remember

A scientific figure is evidence only when its transformation from source data to display is appropriate, inspectable and open to reconstruction.

KNO-019 · Data, Tables and Records

The question

What can we really conclude?

Short answer

Data become durable evidence only when their structure and history remain interpretable. Tables make some relationships easy to inspect, but spreadsheet convenience can also hide units, types, formulas or missing-value conventions. Metadata preserve the context needed to understand what values mean. Provenance records how data were produced or transformed.

Evidence check

We know that explicit tabular organization can make variables and observations easier to analyse , that spreadsheet conventions can reduce avoidable ambiguity and mixing of raw and derived content , that provenance can represent entities, activities and agents involved in producing data , and that data citation requires specificity sufficient to identify the object or version being cited . NIST RDaF Version 2.0 explicitly treats metadata, provenance, versioning, instruments, processing operations and reproducibility as parts of research-data management . The National Academies likewise connect reproducibility with access to the data and computational information needed to recompute results .

Where to be careful

Think like a researcher

Create a tiny table and add notes explaining units, date, source and any changes you make.

Remember

Data preserve knowledge only when the record preserves enough meaning, provenance and version context to make the values interpretable.

KNO-020 · Review, Consensus and Disagreement

The question

What can we really conclude?

Short answer

Peer review, publication, consensus and disagreement are stages or signals within scientific evaluation, not substitutes for the underlying evidence. Reviewers can expose weaknesses before publication, but different reviewers may reasonably emphasize different risks. Publication widens scrutiny rather than ending it. Consensus becomes more informative when it records convergence across evidence, methods and critical examination, while remaining open to revision.

Evidence check

- Structured peer review can expose problems before publication, but reviewer judgments vary and are not certificates of truth. - Editorial decisions can incorporate conflicting reviewer advice rather than simply count votes. - Reproducibility and replicability extend scrutiny beyond publication and can strengthen, qualify or challenge claims. - Formal scientific consensus processes can combine evidence review, expert deliberation and independent review. - Corrections can preserve an auditable record of errors found after publication.

Where to be careful

Think like a researcher

Take a disagreement and list what evidence could actually discriminate between the competing explanations.

Remember

Review, consensus and disagreement organize criticism; their epistemic value depends on the evidence, procedures and discriminating tests they make visible.

KNO-021 · Open Science and Reproducibility

The question

What can we really conclude?

Short answer

Open science is not one switch labelled “share everything.” It is a set of controls that make scientific work more inspectable. Reproducibility usually asks whether the same data and computational procedures can yield the same result; replicability asks whether new data collected for the same question produce results that are consistent within the relevant uncertainty. Preregistration can separate predictions from analyses chosen after seeing the data, but it cannot make a weak design strong. Registered reports move part of evaluation before results are known, while open data and code let others inspect transformations, exclusions and implementation.

Evidence check

- Reproducibility and replicability test different kinds of consistency and should not be treated as synonyms. - Preregistration can make the timing of hypotheses and analysis choices visible. - Open data, code and provenance can make more of an analysis inspectable, but access without context is insufficient. - Versioning and visible corrections preserve which scientific object a claim refers to.

Where to be careful

- How much each open-science intervention improves reliability across different disciplines and study designs. - Which combinations of openness and controlled access best preserve both auditability and legitimate privacy. - How long contemporary research objects remain executable as software and hardware ecosystems change. - How incentive systems alter the real use, rather than nominal adoption, of open practices.

Think like a researcher

Imagine another person must reproduce your school experiment. What files, notes or steps would they need?

Remember

Open science does not certify truth; it preserves more of the evidential path so that other people can inspect, reproduce, challenge and correct it.

KNO-022 · Computation as External Memory

The question

What can we really conclude?

Short answer

Humans routinely move cognitive work into the environment: we write lists, draw diagrams, set reminders and arrange objects so that the world carries part of the task. Cognitive-offloading research treats this as a general strategy rather than an anomaly. Computers extend the same logic. They can retain exact symbolic states, rerun procedures and transform representations at scales that unaided memory cannot manage.

Evidence check

- People use external representations and actions to reduce internal cognitive demand. - External representations can change the structure and cost of reasoning, not only storage requirements. - Expectations of external access can affect what people remember about information and its location. - Executable code can preserve procedures with greater operational specificity than prose alone.

Where to be careful

- How different forms of cognitive offloading reshape long-term learning rather than immediate task performance. - Which computational dependencies are necessary to preserve for century-scale executability. - How much automation changes users' ability to detect conceptual rather than implementation errors. - Where responsibility should sit in highly distributed human-computer workflows.

Think like a researcher

Write a short set of instructions a computer could repeat exactly. Where could an unnoticed mistake be repeated too?

Remember

Computation becomes epistemically powerful when it preserves not only results, but the states, procedures and provenance needed to understand how those results were produced.

KNO-023 · How Groups Know More Than Individuals

The question

What can we really conclude?

Short answer

A group can gain epistemic power in several ways. Different members can know different facts; independent estimates can cancel some individual errors; discussion can expose contradictions; specialized roles can preserve knowledge no one person holds completely. Experiments show that groups can outperform their best members on some structured problem-solving tasks. Models by Hong and Page show how diversity of problem-solving approaches can matter under specified assumptions.

Evidence check

- Groups can outperform their best members on some structured intellectual tasks. - Diverse problem-solving heuristics can improve group search under specified model assumptions. - Social influence can reduce diversity of estimates and undermine some wisdom-of-crowds benefits. - Interaction patterns such as conversational turn-taking are associated with collective performance in laboratory groups.

Where to be careful

- How well laboratory measures of collective intelligence generalize to long-lived institutions and high-stakes expert teams. - Which diversity dimensions matter for which kinds of scientific and engineering problems. - How to weight expertise without suppressing informative minority evidence. - How digital communication and AI mediation change independence, conformity and information pooling in human groups.

Think like a researcher

Ask several people to estimate something independently before they discuss it. Compare the individual and group answers.

Remember

Groups know more than individuals only when their structure preserves and combines distributed information better than conformity, shared bias or premature consensus destroy it.

KNO-024 · Lost Knowledge and Preservation Bias

The question

What can we really conclude?

Short answer

Records survive through filters. In archaeology, objects move from active use into archaeological context through cultural and natural formation processes, so what is excavated is already transformed. In paleontology, weathering, transport, burial and mineralization make some organisms and environments much more likely to enter the fossil record than others. Archives introduce another layer: people and institutions decide what to document, retain, classify and permit others to access.

Evidence check

- Archaeological records are shaped by formation processes between use and recovery. - Bone weathering and other taphonomic processes affect which biological traces survive and how they are interpreted. - Archives embody decisions about creation, appraisal, description and access rather than functioning as neutral containers. - Research datasets can become unavailable over time when preservation depends on individual custodians.

Where to be careful

- How much undocumented material has already been lost from contemporary digital research ecosystems. - How preservation probabilities vary across underrepresented communities and informal record-keeping systems. - Which digital preservation strategies will remain interpretable over century-scale technological change. - How best to quantify preservation bias when creation, loss and discovery probabilities are all uncertain.

Think like a researcher

List three reasons why an object or record from the past might disappear before anyone can study it.

Remember

A surviving record is evidence about the past and about the processes that selected what could survive; absence is interpretable only after both are considered.