Locating a perspective in conceptual space

Orientation and justification in an interpretive atlas

Abstract

Conceptual maps represent interpretations as positions in a structured space. Two or more interpretations can nevertheless occupy the same position while differing in their grounds, content, and vulnerability to correction. This paper examines what follows for maps of interpersonal understanding. It distinguishes orientation from epistemic justification and states an elementary adequacy condition: a property is recoverable from a map only if it is constant across the interpretations that the map places together. A worked example separates introspection from extrospection and awareness of self from awareness of otherness. It shows how relative coordinates can conceal simultaneous engagement, how evidential improvement can occur without movement, and how movement can occur without improvement. I propose an atlas that retains reasons and revision conditions alongside coordinates. The resulting framework supports a limited form of conceptual exploration; it does not establish a universal taxonomy, a psychological measurement scale, or a neurological theory of consciousness.

Keywords: conceptual spaces; interpersonal understanding; representation; justification; conceptual engineering

Word count including captions and references, excluding alt text: approximately 5613

1 What a position can tell us

A person makes a joke. Another person becomes quiet. The speaker notices the silence and asks what it means. One interpretation concerns the speaker’s intentions: perhaps the joke was harmless and the response unexpected. Another concerns the listener’s experience: perhaps something painful has been recalled. A third concerns the encounter itself: perhaps the speaker has mistaken distraction for offence. The description can move between these orientations without any new observation being made. Conversely, the speaker may acquire decisive contextual information while remaining directed towards the same person and the same question.

A map could make these shifts visible. One axis might contrast introspection and extrospection; another might contrast consciousness of self and consciousness of otherness. A point would indicate how an interpretation is organised relative to the four terms. Moving it would express a revision in emphasis. This is an attractive device because it allows several descriptions to remain available together, without requiring an immediate verdict about a person’s character. It also makes a demand on its reader: explain why this interpretation belongs here.

The device becomes misleading if its position is allowed to answer questions for which the representation supplies no grounds. A movement towards otherness does not establish that the other person has been understood. A central point does not reveal whether both activities are vigorous, both are weak, or one category has been interpreted differently. A smooth path does not establish that every intermediate description is coherent. The appearance of spatial order must therefore be distinguished from the claims that the order warrants.

My organising observation is that two or more interpretations can occupy the same position on a conceptual map while having very different justifications. The plurality matters. Co-location is a structural possibility of a selective representation, rather than a peculiar accident involving one pair. A point can stand for an entire family of interpretations. Members of that family may agree in orientation while disagreeing in content, evidence, confidence, moral purpose, or conditions of revision.

This paper develops an interpretive atlas for examining that possibility. Its contribution is a proposed discipline for the use of conceptual maps: specify the question a chart answers, identify the distinctions its projection suppresses, and preserve separately the reasons required for the inferences users actually want to make. The formal condition introduced below is elementary. The philosophical work lies in deciding which distinctions matter for a task and in exposing cases where an attractive diagram encourages an inference that its construction cannot support.

The example examines a proposed atlas of mental faculties, treated here as a representational proposal. No inference is made from the diagram to cerebral anatomy, interhemispheric function, personality measurement, or a causal account of consciousness. All plotted placements in this paper are stipulated illustrations. They are neither observations of participants nor estimated psychological quantities.

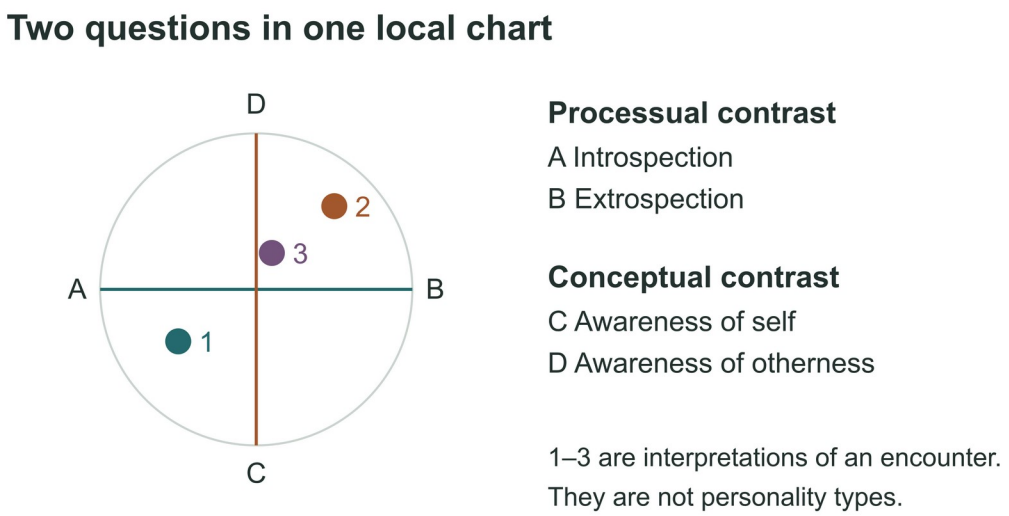


Figure 1. A proposed local chart of interpretive orientation. A and B mark introspection and extrospection; C and D mark awareness of self and awareness of otherness. The points represent descriptions of an encounter, rather than fixed identities of people. Their positions are schematic. The two axes distinguish questions that can interact: how an interpretation is approached and towards whom its attention is directed.

Alt text: A circular chart has horizontal introspection and extrospection poles and vertical self and otherness poles. Three example points occupy different positions.

2 Connections that constrain the construction

Geometrical representation of concepts is an established research programme. Gärdenfors (2024, §2) presents conceptual spaces as a research programme in semantics that represents concepts and their interrelations geometrically. It is therefore insufficient to present the use of coordinates, regions, or

similarity as an original discovery. The question here concerns a particular application and the conditions under which its simplifications are legitimate.

Gärdenfors's later account also discusses dimensional reduction in which regions of a higher-dimensional space are mapped onto points of a lower-dimensional representation (2024, §5). The existence of such compression is therefore part of the relevant literature. The present proposal asks a further, task-specific question: which epistemic differences must remain available when a chart is used to assess an interpretation? His distinction between integral and separable dimensions (§2) also cautions against equating the visual perpendicularity of our axes with psychological independence.

A recent application makes the methodological difference especially clear. Bendifallah, Abbou, Douven, and Burnett (2025, §3 and §4) construct conceptual spaces for feminist political identity using empirical categorisation material and dimensional modelling. They also discuss the difficulties of applying such methods to complex social concepts. Our axes are proposed in advance and our cases are constructed. We have no comparable empirical warrant for the dimensionality, metric, or placement of the present atlas.

Tversky (1977, opening sections) questions both dimensional and metric assumptions in representations of similarity and develops an alternative based on feature matching. His analysis includes the directionality of similarity judgements. It supplies a reason to examine what a distance claim commits us to. The fact that two marks on our page have a Euclidean distance does not establish that the corresponding interpretations have a psychologically meaningful, symmetric distance. This paper therefore uses coordinates to organise comparisons without treating plotted distance as an empirical similarity measure.

The purposes of the map also require scrutiny. Carnap's discussion of explication evaluates a proposed replacement through similarity to an earlier concept, exactness, fruitfulness, and simplicity (1950, §§2–3). These requirements provide questions to ask of an atlas: what use is retained, which ambiguity is removed, and what becomes possible through the new representation? They do not select our axes. Köhler and Veluwenkamp's account of normative functions in conceptual engineering (2024, §§2–4) adds a further question: what good reason have we to adopt the proposed representational device? A convenient display does not by itself settle that issue.

Historical connections sharpen the limits. Kant distinguishes the contributions of sensibility and understanding while insisting on their cooperation in cognition (1781/1787, Transcendental Logic, Introduction I). This offers an instructive case of complementary contributions, but supplies no derivation of introspection and extrospection as numerical poles. Kant's discussion of the inadequacy of merely formal criteria for material truth, in Introduction III, is a further warning against expecting an ordering scheme alone to establish an attribution. The present argument does not depend on accepting Kant's transcendental framework.

Bergson's objection is more directly directed at spatialisation. In Time and Free Will, chapter II, he argues that the representation of conscious states as mutually external items in a homogeneous medium can distort duration, whose phases interpenetrate. The consequence for our construction is a demand to state what a chart of successive descriptions omits. A path on a page records an ordering of representations; whether it captures the lived organisation of their transition remains a separate question.

These connections constrain rather than authorise the atlas. They establish neither its completeness nor the reality of its faculty divisions. They help identify the work that remains: defend the coding rule, specify a purpose, preserve relevant distinctions, and avoid treating a representation's formal properties as discoveries about its target.

3 The unit of representation

The first decision concerns what a point represents. A concept, an interpretation of an event, a person, and an episode of consciousness are different units. An atlas that moves silently among them can make an ordinary change of description appear to be a transformation of identity. I use an interpretation of a specified encounter as the unit. It includes a claim about what is happening, the interpreter's grounds, the relevant context, and an account of what would lead to revision.

Let D be a set of such interpretations under a stated coding procedure. Let M be the space used by a chart, and let f assign each interpretation to a point of M. The domain includes the contextual information needed by the procedure. If the context changes, the input may change; if the coding rule changes, we are dealing with a different function. Holding these fixed is essential when comparing positions.

$$f\colon D \to M, \qquad d \mapsto f(d)$$

For the example, the first coordinate expresses a stipulated relative emphasis between introspection and extrospection. Introspection concerns the interpreter's examination of their own experience or contribution. Extrospection concerns directed inquiry into outward expression, action, circumstances, or another person's account. These descriptions are deliberately modest: they neither resolve the epistemology of introspection nor reduce interpersonal understanding to observation of external movements.

The second coordinate expresses a relative emphasis between awareness of self and awareness of otherness. The first pair concerns an approach to inquiry; the second concerns its orientation towards subjects. They are not statistically independent variables established by research. Orthogonality on the page is a display convention that allows combinations to be inspected. Any claim that the categories carve psychological processes at their natural joints would require a further argument and different evidence.

Both contrasts can mislead if understood as mutually exclusive capacities. Attending to my embarrassment may help me recognise that I have interpreted another person defensively. Listening to another person may alter my understanding of my own intentions. An interpretation can be strongly engaged with both self and other. For this reason, the local chart initially records relative emphasis rather than total capacity or an amount of empathy. There is no privileged moral location on it.

A chart must also permit an interpretation to remain unplaced. A reader may reject the categories, find a case underdescribed, or think that several positions fit equally well. Forcing every description into a point would convert the interface's demand for an answer into a purported fact about conceptual structure. In practical use, a region, several candidate positions, or an explicit refusal to code can be more informative than a precise mark.

4 Two or more interpretations at one position

Suppose a coding procedure places four interpretations at the same point because they have the same stipulated relative orientation towards outward inquiry and otherness. The first attributes distress after checking the listener's account and relevant circumstances. The second makes the same attribution by projecting a personally familiar experience onto the listener. The third relies on an uncorroborated report about the listener. The fourth attributes irritation rather than distress, while remaining equally oriented towards understanding the listener. Nothing in these two coordinates requires their separation.

The first interpretation has relevant evidence unavailable to the second and third. Whether it is ultimately justified depends on further details, including reliability and defeaters; the listener's testimony is not declared infallible. The fourth differs in content as well as grounds. The example makes no empirical claim that human interpreters actually receive these numerical positions. It shows what follows if an orientation chart gives these interpretations the same position while users want it to distinguish their justification.

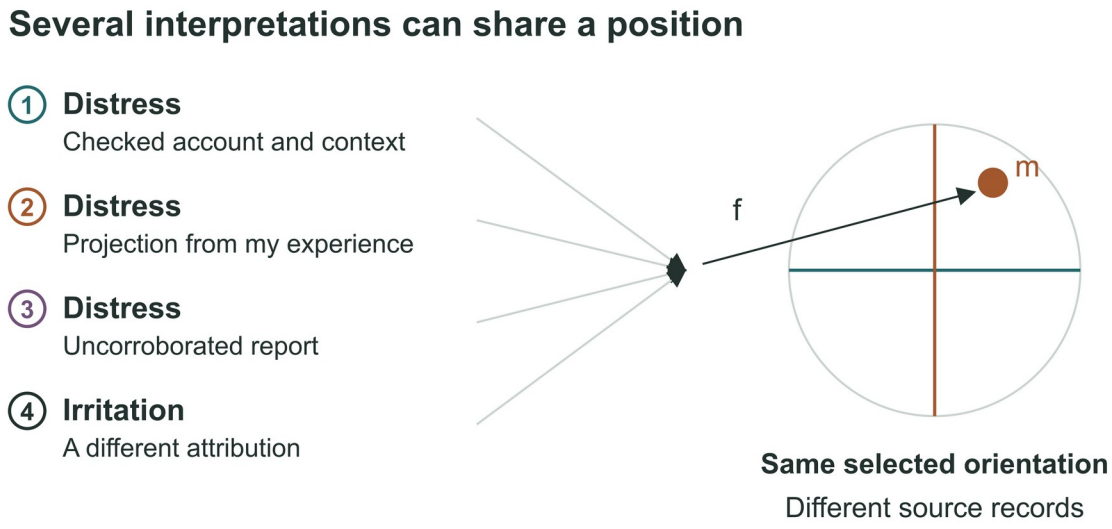


Figure 2. Four interpretations assigned to one position by an orientation-only coding rule. Their source records remain different. The labels do not rank every possible case: a report can be reliable, testimony can be defeated, and projection can accidentally reach a true conclusion. Co-location records agreement in the chart's selected respects, not identity of meaning or equality of justification.

Alt text: Four interpretation records with different claims or grounds converge through one coding rule to a shared point on an orientation chart.

The family of all interpretations mapped to a point m is its preimage, also called a fibre. The terminology introduces no claim about neural fibres or anatomical correspondence. It names a familiar mathematical relation between a representation and the different inputs that it treats alike.

$$F_m = \{d \in D : f(d) = m\}$$

Let P be a property that a user wants to recover from the chart, with values in a set V. It could distinguish assertible from insufficiently supported attributions in a specified context. It could instead

record whether the interpretation relies on another person’s testimony. The argument does not require every epistemic status to have a numerical value, nor does it provide a general theory of justification.

Adequacy condition. A rule h on the represented image can recover P from position alone exactly when P is constant on every fibre of f.

The right-hand condition below ranges over all pairs d_1 and d_2 in D; the recovery rule h has domain f(D).

$$\exists h \ (P = h \circ f) \quad \Leftrightarrow \quad [f(d_1) = f(d_2) \implies P(d_1) = P(d_2)]$$

The proof is immediate. If P is obtained by applying h to a position, identical positions produce identical values. Conversely, if P is constant on each fibre, assign to a represented position the common P-value of its interpretations. That rule is well-defined. The same reasoning applies to a fibre with two, four, or indefinitely many members. It is an elementary condition on information preservation, not a new mathematical theorem.

The condition diagnoses a particular excess of inference. If relevantly different epistemic statuses occur within one fibre, a reader cannot recover that status from position alone. A chart can still be useful for a different question. Failure to recover justification does not imply failure to represent orientation; representing orientation successfully does not imply that justification has thereby been represented.

Nor does the condition show that a map could never encode epistemic status. A different function might include information about evidence, a categorical annotation might record the required distinction, or a sufficiently rich representational system might support the inference. Counting axes alone establishes no general impossibility result. The claim concerns the explicit mapping being used and the property demanded of it.

5 What the middle can conceal

A second limitation arises even before we ask about evidence. Suppose I and E are nonnegative illustrative weights for introspection and extrospection, and S and O are corresponding weights for self-directed and other-directed awareness. The weights are not measured psychological magnitudes. They make explicit one possible rule by which relative emphases could be converted into coordinates.

$$x = \frac{E - I}{E + I}, \qquad y = \frac{O - S}{O + S}$$

The ratios require nonzero denominators. If both relevant weights are zero, the corresponding coordinate is undefined; placing it at the centre would add a convention that must be stated. Where the denominators are positive, the centre of the chart can result from equal weights at many levels. Equal weak engagement and equal intense engagement have the same ratio. Their difference has been removed by normalisation.

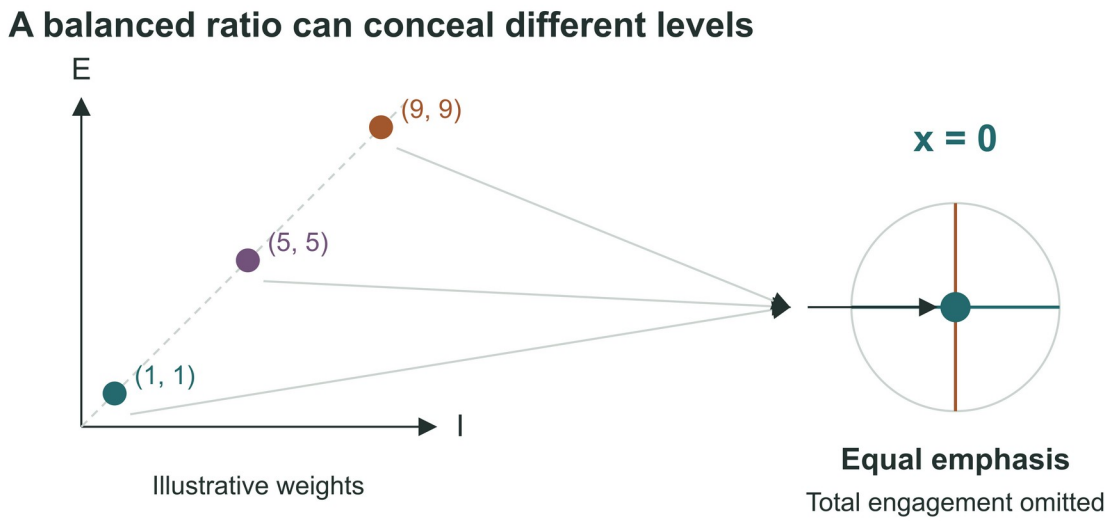


Figure 3. Relative emphasis suppresses common intensity. The illustrative pairs (I,E)=(1,1), (5,5), and (9,9) all produce x=0. The numbers are algebraic examples, not a psychological scale. A balanced position therefore does not distinguish weak joint engagement from strong joint engagement. At (0,0), the ratio is undefined.

Alt text: Three different equal-weight profiles project to the same balanced coordinate. Common intensity is lost by normalisation.

Positive rescaling makes the loss precise. Multiplying both I and E by any positive factor leaves x unchanged. Multiplying S and O by another positive factor leaves y unchanged. Each admissible pair of scaling factors therefore generates another member of the same positional family. For this illustrative coding rule, a single point can have infinitely many distinct weight profiles behind it.

One response is to retain the sums E+I and O+S alongside the ratios. Under the algebraic assumptions, a positive sum and the corresponding ratio determine the original pair of weights. This is a recovery of the stipulated numbers, not a validation of psychological quantities. If the weights have no defensible interpretation or cannot be assessed reliably, adding them to the display only produces a more elaborate unsupported representation.

A second response is to abandon bipolar coordinates where the task requires independent engagement. Separate dimensions for self and other permit both to be high. A third is to use qualitative annotations such as “simultaneous sustained attention to both”, leaving the quantitative appearance behind. The choice depends on the question and the grounds for the coding. There is no requirement that a defensible philosophical atlas have the same number of axes in every chart.

The distinction also clarifies a potentially confusing use of self-consciousness. An interpretation of how another person sees me concerns a socially mediated description of myself. It is not thereby an account of every form of self-awareness. Moving a description towards otherness may represent an increasing dependence on another person’s response while leaving other questions about first-person awareness unanswered. The position locates a description under selected contrasts; it does not exhaust the subject described.

6 What rotation adds

The proposed atlas began with four poles and the thought that their rotation could generate a conceptual sphere. This is a productive visual invitation, but several constructions must be distinguished. A sphere is a surface. Its interior belongs to a ball. A planar disk, a spherical surface, and a solid ball have different dimensions and different sets of points. The presence of four labels alone does not determine which of them is appropriate.

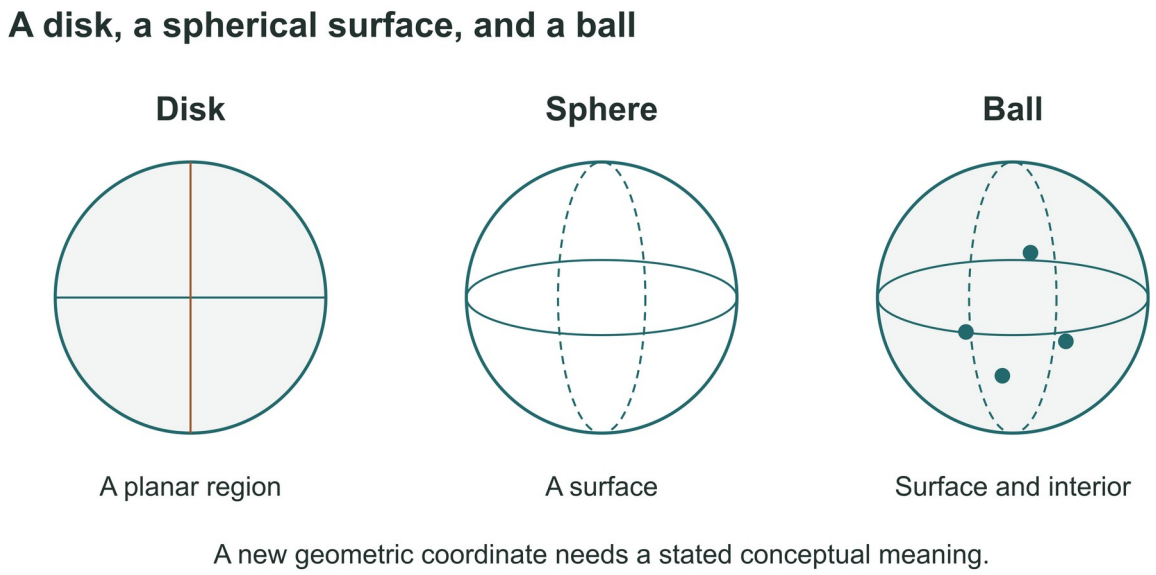


Figure 4. Three representational choices that must be specified. A disk is a two-dimensional region. A spherical surface requires an angular construction, such as rotating a semicircle about its diameter. A ball includes the interior, obtainable by rotating a filled half-disk about its diameter. These constructions define drawings and spaces; they do not identify an additional philosophical variable.

Alt text: A planar disk, a spherical surface, and a ball with interior points illustrate different geometric constructions.

Rotating one diameter around itself changes nothing. Rotating an orthogonal diameter around a fixed diameter sweeps a planar disk perpendicular to the fixed one; it does not by itself produce the entire spherical surface. If all orientations of fixed-radius endpoints are permitted, their sweep is a sphere. If every point of a diameter is allowed all orientations, their sweep is a ball. A construction must say which objects move and which transformations are permitted.

None of these operations supplies a new explanatory dimension merely by supplying a new screen coordinate. A third semantic dimension requires a specification of what it distinguishes. If depth is assigned to evidential source, for instance, the assignment must explain why that source has the chosen ordering and why interpolation is meaningful. Testimony, projection, and a behavioural observation need not occupy a single natural scale.

Rotation for inspection should also be separated from revision of an interpretation. A user may rotate a displayed sphere to see a label that was hidden. The record remains the same even though its screen position changes. A user who drags a point relative to the labelled axes is proposing a different coding. The two interactions require different feedback. Otherwise a change of viewpoint can be mistaken for a change in the philosophical description.

The safest use of the sphere is therefore conditional. It can communicate that interpretation has more structure than a single opposition and can make the interior available as an exploratory metaphor. Its geometry becomes explanatory only when the interpretation of its dimensions, boundaries, and transformations is independently defended. A circular boundary does not establish a legitimate faculty boundary; a point's being drawable does not establish its conceptual coherence.

7 Three changes that a chart should distinguish

Return to the joke. Suppose the speaker learns that the listener had just received distressing news. The relevant interpretation now has a different evidential basis. The speaker may continue to attend outwardly to the listener, so that an orientation chart registers no movement. Evidential improvement can therefore occur without positional change. Even the conclusion might remain the same, while the grounds change from an unsupported guess to a contextually supported attribution.

Now suppose the speaker becomes absorbed in the listener's imagined distress without obtaining any additional evidence. The point could move towards otherness, as the description becomes more other-directed. The change might be morally welcome or psychologically significant in some respect, but its direction alone does not establish better understanding. An increasingly elaborate projection could generate the same movement.

A third change is the rotation of the drawing just discussed. Here the same interpretation, coding, and evidence remain in place, while the user's view of the representation changes. It is a change in presentation. An interactive atlas should keep a record identifier fixed during rotation and should preserve the labelled coordinate values so that users can tell what happened.

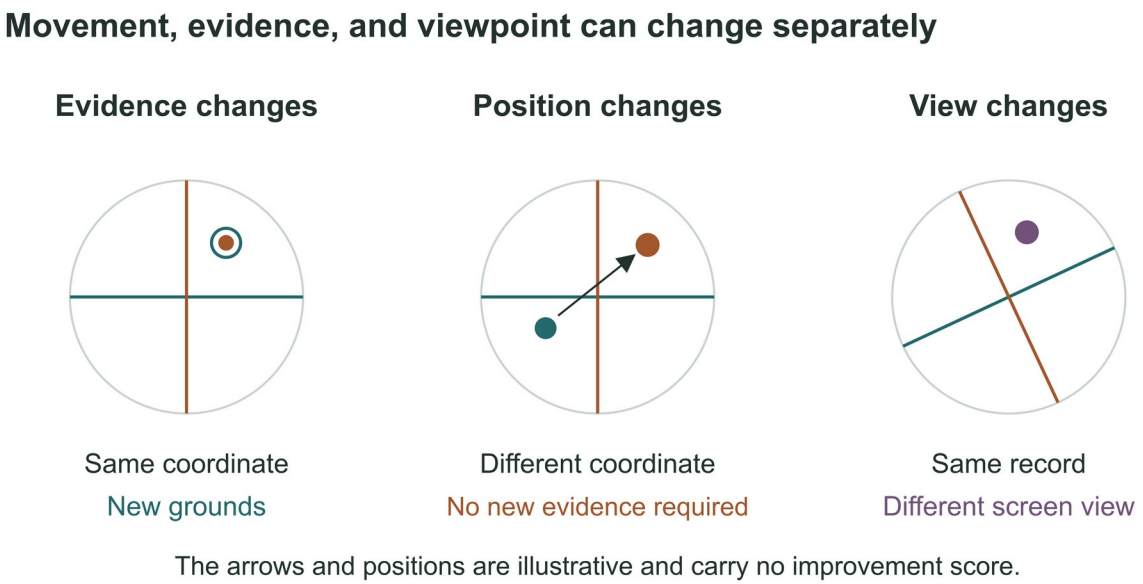


Figure 5. Three changes with different implications. Left: the coordinates stay fixed while evidence changes. Middle: the coordinates change without additional evidence. Right: the drawing rotates while the record and its coordinates relative to the labelled axes remain fixed. All three are schematic possibilities; an arrow is not an improvement score.

Alt text: Three charts distinguish new evidence at a fixed position, a changed position without new evidence, and a rotated display of a fixed record.

The temporal order of interpretations should also be recorded separately. A line connecting three points says which descriptions were considered in which order. It does not supply an explanation of the transition. The interpreter may have acquired information, changed aims, reinterpreted a term, or followed a prompt. Each would require a different account. Repetition of the same point can be meaningful if the grounds or context changed between visits.

An account of revision therefore needs at least a claim, a source record, relevant defeaters, and a reason for the change. For the present cases, one might record the listener's account, what was observed, what was inferred, and what remains uncertain. These fields should remain inspectable in ordinary language. Replacing them with a single confidence colour would conceal the distinctions the atlas is intended to expose.

8 Understanding an adversary

A rival is vulnerable to humiliation. An interpreter notices the vulnerability and predicts which remark would exploit it. Another interpreter notices the same vulnerability and chooses to avoid exploiting it. A third recognises the vulnerability but thinks intervention would worsen the situation. Their understanding of what the rival experiences could be similar while their purposes and actions differ. A chart of interpretive orientation cannot silently become a moral ranking.

Zahavi's analysis of empathy and interpersonal understanding provides a relevant constraint (2010, p. 295 and note 5 on p. 304). Access to another person's experience need not be a duplication of first-person access; understanding and concern must be distinguished. This supports keeping separate

questions about what is understood, how it is accessed, and what the interpreter intends to do. It does not vindicate the coordinates proposed here.

The adversary case also limits the term “empathy” as a label for the whole chart. Accounts of empathy disagree about the place of affect, perspective-taking, and concern. Instead of resolving that dispute by definition, the chart should state which sense is being examined in a particular exercise. A case may concern interpersonal prediction without satisfying a more demanding account of empathic understanding. Such a refusal is a useful boundary result.

The same caution applies to the socially mediated self. I can examine my embarrassment while trying to understand a rival’s fear; I can also use my embarrassment as the sole model for that fear. The first description permits self-awareness and other-directed inquiry to cooperate. The second risks projection. Their difference cannot be settled simply by how far a point lies from the self pole. It concerns the role my experience plays in the inference.

Nor does opposition itself remove the possibility of understanding. Agreement concerns endorsement of claims or aims. Concern concerns how another person’s good matters in deliberation. Interpersonal understanding concerns some grasp of their experience or reasons. These relations can interact, and particular theories may connect them more strongly, but a device that combines them must explain the connection. The present atlas lets those disagreements remain visible in the record rather than allocating moral authority to one corner.

9 Conditions for a useful atlas

An atlas is useful when it helps a reader make a distinction that the task requires and recognise when the representation has reached its limit. I propose six conditions for that use. They are methodological proposals to be tested against cases, rather than a completed theory of concepts.

Specify the task and unit. State whether the chart concerns interpretations, concepts, uses of a word, persons, or experiences, and identify the question it is meant to answer. A placement of an episode should not become a verdict about a person’s enduring personality. The same chart can serve more than one task, but each inferential use requires its own adequacy assessment.

Disclose the projection. Explain how descriptions receive positions, how ambiguous cases are treated, and what the axes omit. The rule may be qualitative. If numerical values are used, their scale and operations require a defence. Coordinate precision should not exceed the precision of the coding. A claim of interpretability does not substitute for a reproducible procedure where reproducibility is part of the task.

Inspect the families at each position. Ask whether two or more co-located interpretations differ in a property that users are being invited to infer. An example of such a difference identifies a limit of that inference. The response can be to refine the mapping, retain a qualification, or restrict the question. There is no obligation to separate everything that differs in any respect.

Retain reasons and revision conditions. Keep evidence, defeaters, uncertainty, and the proposed reason for movement accessible beside the map. These records do not justify themselves by being displayed. They allow the reader to assess a claim rather than to infer its status from a location. An improved record may matter more than a longer trajectory.

Justify transformations and boundaries. State what remains invariant when the drawing rotates, what changes when a point moves, and what a border excludes. If a path is used, explain whether it records temporal order, a proposed conceptual transition, or a sequence of interface actions. A drawn connection establishes none of these interpretations automatically.

Admit failure at the right level. If a description is underdetermined, request or retain more context. If relevant information is systematically lost, revise the chart or narrow its use. If a proposed distinction does not fit the domain, permit refusal and an alternative representation. A system that counts every objection as another permissible point has removed a principal means of evaluating itself.

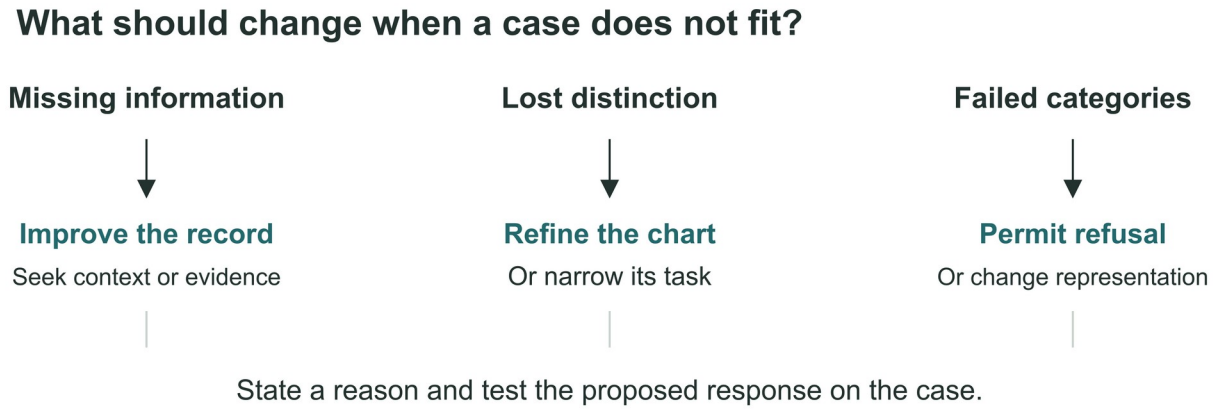


Figure 6. Responses to a case that the map does not resolve. Missing information calls for a better record; a distinction lost by the projection calls for a changed chart or a narrower task; a failed categorisation may call for refusal or a different representation. These are different interventions. Each requires a stated reason and a case on which its success can be assessed.

Alt text: Three columns connect missing information to improving the record, lost distinctions to refining the chart, and failed categories to refusal or a different representation.

These conditions also explain what the analogy with a periodic table can and cannot promise. An empty location becomes a prediction only if some independently supported constraint specifies what should occupy it and excludes alternatives. A visually available gap is otherwise an invitation to construct a description. An atlas does not acquire predictive force through its organisational ambition alone.

The ambition can nevertheless be productive. A proposed vacant region can prompt a search for an overlooked distinction. Its value can be assessed by whether the resulting description is coherent, useful, and different in a task-relevant way from existing descriptions. That is a programme of conceptual exploration. Stronger claims about discovery or exhaustiveness require a defined domain and independent success conditions.

10 Local charts and the limit of subjectivity

The wider proposed atlas proposes eight faculties: empathy, intuition, attention, creativity, memory, language, faith, and time. Nothing in the argument above derives this list or establishes that every philosophical concept belongs within it. Each faculty chart would need a separate account of its

distinctions, its domain, and its connection with neighbouring charts. A finite list of labels cannot establish exhaustive coverage of an unspecified domain.

There is a constructive way to connect charts. The same interpretive record can be examined with a chart of attention, a chart of memory, and a chart of interpersonal orientation. The record supplies the link; equal numerical coordinates in different charts do not. A reader might ask how selective attention affected which expression was noticed, how a remembered humiliation supplied a model, and how another person’s account changed that model. Those questions cross faculties without presupposing a shared unit of measurement.

Such connections require restraint. A chart transition is a proposed redescription of the case until a rule specifies how information is preserved between representations. Calling the collection an atlas does not make it a differential-geometric atlas, and no smooth transition functions are asserted here. The ordinary cartographic term denotes a collection of local representations whose purposes and losses must be declared.

Claims about personality and identity require further distinctions. A record can change substantially while the interpreter remains the same person. Two interpreters can occupy the same position while having different histories and dispositions. A trajectory is therefore insufficient to establish either identity or its fracture. Any numerical identity threshold in the literary construction remains part of its fictional convention unless an independent argument and appropriate evidence support a different use.

Finally, an improved description of relations among interpretations is not already an explanation of what experience is like for its subject. Nagel’s account of subjective character makes the issue explicit, while his closing proposal of a more objective phenomenology leaves room for partial structural description (1974, pp. 435–436 and 448–449). The relevant lesson is twofold: structural access can be valuable, and its limits must not be converted into a claim to reproduce another subject’s experience. The present projection argument proves no permanent metaphysical barrier and does not decide between theories of consciousness.

The atlas can therefore pursue a substantial but bounded aim. It can make competing descriptions inspectable, reveal the assumptions behind their placement, distinguish changes of orientation from changes of evidence, and identify cases that force revision of the representation. Its success lies in the quality of the comparisons it supports and the mistakes it helps expose.

11 Conclusion

Two or more interpretations can occupy the same position on a conceptual map while differing in content, grounds, and conditions of revision. This is not automatically a defect. A selective representation groups together what is alike in the respects it has chosen. The difficulty arises when the representation is used to support an inference about a property that varies within those groups.

The elementary adequacy condition makes that difficulty explicit. The worked cases give it philosophical use: relative coordinates can suppress joint intensity; evidence can improve without a point moving; a point can move without improvement; a rotated drawing can leave the entire

interpretive record unchanged. These distinctions motivate an atlas that keeps reasons and revision conditions available beside its coordinates.

The next question for each proposed chart is therefore concrete: which inference does this representation permit, and what would count as a case in which that permission fails? An atlas that can answer that question offers a disciplined way to explore conceptual possibilities. Its capacity to register its own losses is part of its philosophical value.

References

Bendifallah, Lina, Julie Abbou, Igor Douven, and Heather Burnett. 2025. “Conceptual Spaces for Conceptual Engineering? Feminism as a Case Study.” Review of Philosophy and Psychology 16: 199–229. Published online 19 December 2023. https://doi.org/10.1007/s13164-023-00708-7.

Bergson, Henri. 1910. Time and Free Will An Essay on the Immediate Data of Consciousness. Translated by F. L. Pogson. London: George Allen. First French edition 1889. Chapter II. https://www.gutenberg.org/ebooks/56852.

Carnap, Rudolf. 1950/1962. Logical Foundations of Probability. Second edition consulted, 1963 impression. Chicago: University of Chicago Press. §§2–3. https://www.fitelson.org/confirmation/carnap_logical_foundations_of_probability.pdf.

Gärdenfors, Peter. 2024. “Natural Concepts and the Economics of Cognition and Communication.” Philosophia 52: 865–882. https://doi.org/10.1007/s11406-024-00734-4.

Kant, Immanuel. 1781/1787. Critique of Pure Reason. Translated by J. M. D. Meiklejohn. Transcendental Logic, Introduction I and III. https://www.gutenberg.org/ebooks/4280.

Köhler, Sebastian, and Herman Veluwenkamp. 2024. “Conceptual Engineering: For What Matters.” Mind 133 (530): 400–427. https://doi.org/10.1093/mind/fzad064.

Nagel, Thomas. 1974. “What Is It Like to Be a Bat?” The Philosophical Review 83 (4): 435–450. https://doi.org/10.2307/2183914.

Tversky, Amos. 1977. “Features of Similarity.” Psychological Review 84 (4): 327–352. https://doi.org/10.1037/0033-295X.84.4.327. Reprint consulted at https://cseweb.ucsd.edu/~gary/PAPER-SUGGESTIONS/tversky-psych-rev-1977.pdf.

Zahavi, Dan. 2010. “Empathy, Embodiment and Interpersonal Understanding: From Lipps to Schutz.” Inquiry 53 (3): 285–306. https://doi.org/10.1080/00201741003784663. Author copy at https://cfs.ku.dk/staff/zahavi-publications/Empathy_-_Embodiment.PDF.

Method and AI disclosure

OpenAI Codex was used on 27 October 2026 for source discovery, substantial drafting and revision of the text, mathematical exposition, original explanatory diagram design, and document preparation. The exact underlying model build identifier was not available in the session. The author directed the conceptual project and reviewed the resulting manuscript. All cases and numerical examples are constructed illustrations. No participant study was conducted, and the figures depict no empirical observations. The proposed faculty taxonomy, coding rules, and anatomical interpretation are not validated by these illustrations.

The factorisation condition is elementary; no new mathematical theorem is claimed. The argument concerns its application to interpretive justification and revision. Neither taxonomic exhaustiveness nor a causal account of consciousness is asserted.