

NORMATIVE DETACHMENT THEORY

Contribution, Originality, and Lineage Atlas

Version 3.1

10 August 2026

Reader-facing positioning companion · generated from NDT in Context dataset 1.5

Formal Specification 1.4.4 controls the theory. Proposition Registry 1.6 controls claim status.

Executive answer

NDT does not claim to have invented practical inference, value-laden representation, relational morality, source-sensitive assessment, defeasible reasoning, framing effects, or computational modeling. Its strongest proposal advanced here is narrower: the descriptive moral type of a normative judgment depends on whether a qualifying represented moral relationship actually organized the evaluation used by that token. This criterion sits within a broader account of normative detachment, token identity, connected support, defeat, frame alignment, and polyphony.

The Atlas separates three questions: what NDT now commits to; what was already present in Beal's corpus through 2025; and what the 2026 formal and executable work added. A commitment can be continuous with the corpus, distinctive in relation to an external concept, and newly formalized in 2026 at the same time. The Atlas therefore runs in one direction: from an NDT component to corpus continuity, the 2026 change, and its relation to a source-established external concept.

The comparisons establish structured relations within the surveyed sources. They do not establish external priority, historical significance, empirical prevalence, or the truth of NDT or any external position.

1. Comparison discipline

Each row compares one specific NDT commitment with one source-established concept or construct. A theory, school, tradition, or author may identify the concept's provenance but is not itself the comparison object. 'NDT and Aristotle' is not a comparison; 'explicit-goal detachment is an exact embedding of the practical syllogism over a restricted domain' is. The comparison neither adjudicates the concept's parent theory nor attributes a scope claim to its originator.

A row is retained only if it passes two tests. The negation test asks whether an independently intelligible necessity, sufficiency, identity, or scope position about the concept would reject the stated relation. That position need not be held by the concept's originator, but it cannot be invented merely by negating the row. The consequence test asks whether the relation changes what NDT predicts, permits, excludes, or requires elsewhere. If verification removes either the source-established concept or the discriminating consequence, the row is removed.

Relation labels set the evidential ceiling. Exact relations hold only over their stated domains; restricted families leave NDT cases outside the family; compatible modules may supply content or computation without determining output or moral type; output approximations may recover a label without its derivation. Orthogonal relations correct category errors, and 'too vague to verify' records the absence of a stable typed translation. No relation label establishes external priority, empirical confirmation, or theory truth.

INSIDE-OUT ONLY



Exhibit 1. Inside-out comparison order. Each dataset row begins with an NDT component and its operative commitment, separates continuity through 2025 from the specific 2026 change, then names the external concept before stating its relation to NDT. This is the only direction printed in the Atlas.

2. Ordered relation structure

The controlled terms run from bidirectional output-preserving translation over a stated domain toward weaker, orthogonal, or indeterminate relations. This is an order of relation strength, not of importance, originality, confidence, or source quality. Counts show how many approved comparison rows use each term; zero-count types are absent.

MORE SPECIFIED



exact embedding	n = 2
exact submodule	n = 3
conditional exact case	n = 1
restricted decision-model family	n = 6
compatible module	n = 5
empirical signature	n = 1
output approximation	n = 1
tokenwise-only correspondence	n = 1
orthogonal	n = 1
too vague to verify	n = 1

WEAKER OR INDETERMINATE

Exhibit 2. Controlled relation order represented in dataset 1.5. Relation types represented in dataset 1.5, in the controlled taxonomy's authorized order. Numbers are row counts, not weights or distances; zero-count types are omitted.

3. Relation-ordered Atlas

Each subsection reproduces the controlled taxonomy definition and prints the corresponding dataset rows in stable row-ID order. The tables retain the single inside-out direction. Confidence concerns the present mapping, not the truth of NDT or the external source.

3.1 Exact embedding

Bidirectional output-preserving translation over a stated restricted domain.

Exhibit 3. Exact embedding comparisons. 2 rows generated from CTX001, CTX002 in dataset 1.5.

NDT component and commitment	Corpus continuity through 2025	Specific 2026 change	External concept and source-grounded statement	Relation rationale, opposed position, and downstream consequence
CTX001 · Explicit-goal detachment When an explicit goal is active and one feasible action is represented as a necessary means, NDT yields that means as guidance under the restricted conditions.	2018 dissertation and 2025 Is-to-Ought essay already use practical inference as a nonmoral normative case.	Formalized as a restricted exact case with feasibility, blockers, choice sets, and token anchoring.	Practical syllogism The practical syllogism derives an action from premises concerning the good and the possible. (Aristotle, n.d.-a)	Why this relation. The restricted action-guidance function is output-preserving in both directions; NDT contains cases outside the fragment. Opposed position. A noninferentialist model on which practical syllogisms only rationalize actions, without detaching an action-guiding conclusion from the premises, would reject the bidirectional embedding. What this changes downstream for NDT. The row predicts exact agreement only in explicit-goal, necessary-means cases; encoded evaluation, hard requirements, moral provenance, and polyphony lie outside. Role: Lineage · Boundary · Confidence: High

NDT component and commitment	Corpus continuity through 2025	Specific 2026 change	External concept and source-grounded statement	Relation rationale, opposed position, and downstream consequence
CTX002 • Explicit conative evaluation A separately tokened desire plus instrumental beliefs is an exact restricted NDT case, but NDT also allows encoded evaluative organization without a detachable desire token.	The 2025 essay distinguishes the motivational/evaluative problem and proposes multiple pathways from representation to ought.	Separated explicit, attributed, and encoded evaluation; added an independently identified control profile.	Reason–passion motivational structure In the reason–passion motivational structure, passion supplies the impulse to act while reasoning directs it through represented causal connections. (Hume, n.d.)	Why this relation. A separately tokened conative state plus instrumental causal beliefs realizes this structure exactly within NDT's explicit-evaluation subdomain; encoded evaluation lies outside that mapped subdomain. Opposed position. A separable-conation-necessity position predicts that encoded evaluative organization cannot motivate unless it is realized as a distinct passion or desire. What this changes downstream for NDT. The comparison predicts distinct signatures for guidance by an occurrent conative state and guidance by encoded evaluative control without a detachable desire token. Role: Boundary • Confidence: High

3.2 Exact submodule

One external formal operation is isomorphic to one NDT layer, not the whole theory.

Exhibit 4. Exact submodule comparisons. 3 rows generated from CTX003, CTX005, CTX020 in dataset 1.5.

NDT component and commitment	Corpus continuity through 2025	Specific 2026 change	External concept and source-grounded statement	Relation rationale, opposed position, and downstream consequence
CTX003 · Rule-conditioned detachment Explicit or attributed requirements can generate typed candidates without truth-functional implication.	Institutional and value-based rule reasoning is present throughout the corpus; the broad norm/nonmoral-norm distinction is explicit by 2025.	Formalized candidate formation, feasibility, provenance, defeat, and output survival around a rule-detachment submodule.	Input/output logic Input/output logic transforms represented conditions into normative outputs without treating the rule as ordinary material implication. (Makinson & van der Torre, 2007)	Why this relation. The rule-output operation is isomorphic to one NDT layer; I/O logic does not supply frames, actual support, moral relationships, or token type. Opposed position. A truth-functional reduction of norms would deny the need for the I/O operation. What this changes downstream for NDT. The comparison predicts agreement on explicit rule outputs but divergence on provenance, moral type, and challenge admission. Role: Lineage · Boundary · Confidence: High
CTX005 · Grounded challenge semantics Admitted challenges can block premises, links, or targets under a skeptical grounded semantics.	Earlier corpus recognizes polyphony and conflict but does not contain a full typed defeat algorithm.	Added independent challenge admission, typed attacks, grounded labeling, and explicit status as one implementation family.	Grounded argumentation semantics Grounded semantics computes a unique skeptical extension of an attack graph. (Dung, 1995)	Why this relation. The fixed-point calculation is exact for the admitted attack graph; NDT adds admission, target typing, provenance, and moral classification. Opposed position. A credulous or probabilistic account of challenge resolution denies that grounded semantics is the universal psychology. What this changes downstream for NDT. The relation predicts exact finite labels under the reference semantics while leaving alternative psychologies open. Role: Lineage · Confidence: High

NDT component and commitment	Corpus continuity through 2025	Specific 2026 change	External concept and source-grounded statement	Relation rationale, opposed position, and downstream consequence
CTX020 • Actual contribution under redundant support Moral type depends on whether a relationship route actually contributed to the token, including cases with redundant support.	The 2021 corpus makes the relationship constitutive and flags false positives; it does not yet provide a full redundant-support test.	Added actual support graphs, minimal sufficiency over organization-preserving profiles, profile-silence, and redundant-route witnesses.	Actual causation under redundant support NESS and structural-causal approaches distinguish actual causes from merely present conditions under sufficient-set or structural criteria. (Wright, 1985 ; Halpern & Pearl, 2005)	Why this relation. Contribution-sensitive causal structure supplies a formally similar subproblem; NDT applies it to judgment routes and constitutive profile change. Opposed position. Simple but-for accounts deny contribution when redundant routes are both active. What this changes downstream for NDT. The comparison predicts moralization by redundant relationship routes only when they alter the constitutive token profile under the fixed model instance. Role: Lineage • Confidence: High

3.3 Conditional exact case

Exact only when specified background institutions or rules are represented.

Exhibit 5. Conditional exact-case comparisons. 1 row generated from CTX004 in dataset 1.5.

NDT component and commitment	Corpus continuity through 2025	Specific 2026 change	External concept and source-grounded statement	Relation rationale, opposed position, and downstream consequence
CTX004 · Institutional norm formation Represented institutions can construct obligations when the constitutive rule and uptake conditions are active.	The corpus treats institutions and roles as ontological frames that alter what is represented as required.	Added a conditional exact promise case, explicit exception handling, and separation between norm formation and moral type.	Institutional promise obligation A promise made under an active institution can generate an obligation from institutional facts. (Searle, 1964)	Why this relation. The promise case is exact only when the institution, uptake, feasibility, and absence of an undefeated exception are represented. Opposed position. A strict Humean is-ought-gap model on which no constitutive institutional fact can supply a normative bridge would reject Searle's local derivation. What this changes downstream for NDT. The row predicts loss of the obligation when the institutional bridge or uptake is absent, without deciding whether a moral relationship is present. Role: Lineage · Boundary · Confidence: High

3.4 Restricted decision-model family

A philosophical style occupies a defined region of NDT model space; NDT has cases outside it.

Exhibit 6. Restricted decision-model family comparisons. 6 rows generated from CTX006, CTX007, CTX008, CTX009, CTX010, CTX022 in dataset 1.5.

NDT component and commitment	Corpus continuity through 2025	Specific 2026 change	External concept and source-grounded statement	Relation rationale, opposed position, and downstream consequence
CTX006 · Outcome-maximizing guidance Some NDT model instances guide action by maximizing represented outcome orderings, but that guide form neither exhausts normative guidance nor determines moral type.	The 2025 essay distinguishes normative inference from moral relationship construction and treats abstract principles as downstream.	Formalized consequentialist-style cases as a restricted family and separated guide content from moral provenance.	Principle of utility The principle of utility assesses actions by their tendency to augment or diminish the happiness of the affected party. (Bentham, n.d.)	Why this relation. The principle of utility defines an outcome-ordered region of NDT model space; NDT also represents non-outcome guidance and classifies moral type independently of the guide's maximizing form. Opposed position. A utility-sufficiency position predicts that maximizing represented utility settles both normative rightness and moral type. What this changes downstream for NDT. The comparison predicts non-utilitarian normative tokens and utility-maximizing tokens that lack moral provenance. Role: Boundary • Confidence: High
CTX007 · Hard requirements An applicable non-overridden requirement can disqualify an outcome-better action.	The corpus treats norms and benders as capable of changing guidance independently of simple outcome preference.	Formalized hard requirements, applicability, source, scope, and violation.	Formula of Humanity The Formula of Humanity treats rational nature as an end in itself. (Kant, n.d.)	Why this relation. The Formula of Humanity occupies NDT's hard-requirement family when its demand is represented as an applicable requirement with no active override. Opposed position. A consequentialist-sufficiency position about the Formula of Humanity predicts that an outcome-better action cannot be excluded solely by violating the formula. What this changes downstream for NDT. The relation predicts divergence from outcome maximization while preserving a common detachment architecture. Role: Boundary • Confidence: High

NDT component and commitment	Corpus continuity through 2025	Specific 2026 change	External concept and source-grounded statement	Relation rationale, opposed position, and downstream consequence
CTX008 · Override and residue Priority can license one violation without erasing the fact that a lower-priority requirement was broken.	The earlier corpus recognizes conflicting moral demands and complex framed responsibility.	Added typed priority grounds, operative override, residue, and tragic fallback.	Prima-facie duty and moral residue A prima-facie duty can be overridden while retaining moral remainder. (Ross, 1930)	Why this relation. The family maps to priority-sensitive requirements; NDT additionally requires support, provenance, and well-formed priority grounds. Opposed position. A Kantian no-conflict theory of duty, on which only the stronger ground yields an actual obligation, denies that an overridden duty persists as moral residue. What this changes downstream for NDT. The comparison predicts distinct outputs and residual appraisal under otherwise identical outcome profiles. Role: Lineage · Boundary · Confidence: High
CTX009 · Relationship-grounded guidance A represented relationship can construct the operative requirement rather than merely supply a motive for an independently specified rule.	Moral relationships are central in the 2020 and 2021 corpus; the 2025 essay calls them builders.	Added frame-matched relationship tokens, same-chain organization, object/subject bindings, and actual contribution to the operative evaluation.	Care as a moral framework Care is treated as a basic moral value and as the wider moral framework within which justice is situated. (Held, 1995)	Why this relation. Care as a moral framework occupies NDT's relationship-grounded family when a represented caring relationship constructs or sustains the operative requirement. Opposed position. A rule-priority position about care predicts that caring relationships may motivate compliance but cannot construct or sustain the operative requirement. What this changes downstream for NDT. The row predicts same-content moral/nonmoral divergence when care provenance is present versus policy-only support. Role: Central Contribution · Confidence: High

NDT component and commitment	Corpus continuity through 2025	Specific 2026 change	External concept and source-grounded statement	Relation rationale, opposed position, and downstream consequence
CTX010 · Complaint-sensitive guidance Operationalized complaint or rejectability structures can occupy a restricted NDT family without defining all normative cognition.	The earlier corpus does not make contractualism foundational; it permits abstract values and principles as downstream structures.	Added complaint-sensitive decision-model witnesses and explicit family boundaries.	Reasonable rejectability Reasonable rejectability assesses whether an act is justifiable to others on terms they could not reasonably reject. (Scanlon, 1998)	Why this relation. A complaint-sensitive operationalization of reasonable rejectability occupies a restricted NDT family; cardinal minimax is a model witness, not an analysis of reasonableness. Opposed position. A rejectability-necessity position predicts that normative detachment requires reasonable rejectability as the operative standard. What this changes downstream for NDT. The row predicts agreement only where rejectability is the represented operative standard. Role: Boundary · Confidence: Medium
CTX022 · Bargaining-derived decision models NDT can represent a restricted family in which bargaining-derived welfare standards populate outcome evaluation, bargaining-derived action standards populate operative requirements, and actual, simulated, or cached bargaining supplies the decision route; membership in that family does not settle moral type.	The corpus treats consequentialist, deontological, and contractualist styles as downstream reasoning structures rather than the essence of moral cognition.	Decision-model tokens now separate action-outcome maps, outcome orderings, requirements, priorities, model routes, and moral provenance, making the agreement-approximation region locatable without making it constitutive.	Resource-rational agreement approximation Many action-based and welfare-based moral judgments can be explained as resource-rational approximations of agreement-based judgments. (Levine et al., 2026)	Why this relation. Agreement-derived welfare standards map to NDT's outcome ordering, action standards to requirements, and explicit, simulated, or cached agreement processing to decision routes; NDT tests moral provenance separately. Opposed position. An agreement-approximation-sufficiency position predicts that instantiating agreement-based computation or one of its resource-rational approximations settles a token's moral type. What this changes downstream for NDT. Holding the agreement-derived decision route and output fixed while varying represented relationship provenance, NDT predicts a moral/nonmoral split within the same approximation family. Role: Implementation · Boundary · Confidence: Medium

3.5 Compatible module

Supplies content or computation used inside NDT but is insufficient for output or moral type.

Exhibit 7. Compatible-module comparisons. 5 rows generated from CTX011, CTX012, CTX014, CTX015, CTX016 in dataset 1.5.

NDT component and commitment	Corpus continuity through 2025	Specific 2026 change	External concept and source-grounded statement	Relation rationale, opposed position, and downstream consequence
CTX011 • Moral content modules Foundation-like content can populate requirements or appraisals but does not by itself determine token-level moral type.	The 2020 critique argues that abstract moral values are not the irreducible units; ontological frames are prior.	Formalized content modules separately from builders, provenance, and token type.	Moral foundations The Moral Foundations Questionnaire operationalizes five sets of moral intuitions: harm/care, fairness/reciprocity, ingroup/loyalty, authority/respect, and purity/sanctity. (Graham et al., 2011)	Why this relation. Foundation-specific intuitions can supply content to NDT requirements or appraisals, but NDT classifies token type by evaluative organization and provenance rather than content membership alone. Opposed position. A foundation-content-sufficiency position predicts that instantiating foundation-specific intuitive content settles token-level moral type. What this changes downstream for NDT. The comparison predicts that identical foundation-specific content can occur in tokens of different moral type when their evaluative organization or provenance differs. Role: Boundary • Confidence: High

NDT component and commitment	Corpus continuity through 2025	Specific 2026 change	External concept and source-grounded statement	Relation rationale, opposed position, and downstream consequence
CTX012 • Dyadic harm appraisal NDT can represent perceived dyadic harm as a third-personal appraisal, but the agent–patient template neither determines token-level moral type nor exhausts the moral structures NDT permits.	The 2020 and 2021 corpus engages dyadic and harm-centered accounts while arguing for broader object and relationship frames.	Formalized agent-implicit cases, subject gates, nonharm moral objects, and agentless negative appraisals.	Dyadic harm template The dyadic harm template represents perceived harm as an intentional agent causing damage to a vulnerable patient. (Schein & Gray, 2018)	Why this relation. The template can supply agent, patient, harm, and causation content inside NDT, but that content is insufficient for token type and does not exhaust the structures NDT can represent. Opposed position. A dyadic-harm-identity position predicts that matching the intentional-agent/damaged-patient template settles token-level moral type. What this changes downstream for NDT. The comparison predicts crossed classification: dyadic harm appraisals without NDT moral provenance and NDT moral tokens lacking the intentional-agent/damaged-patient structure. Role: Boundary • Confidence: High
CTX014 • Cooperation content and evolutionary function Cooperation can shape moral contents or learning while remaining neither necessary nor sufficient for token-level moral type.	The 2020 article explicitly criticizes cooperation as a criterion for the moral domain.	Formalized cooperation as a content/evolutionary module and supplied false-positive boundaries.	Cooperation-derived moral domains Recurrent cooperative problems and their solutions form the basis of moral domains. (Curry, Mullins & Whitehouse, 2019)	Why this relation. Modest evolutionary/content claims are compatible; an exhaustive constitutive claim contradicts NDT. Opposed position. A cooperation-sufficiency position predicts that membership in a cooperative solution class settles moral type. What this changes downstream for NDT. The comparison predicts moral tokens outside cooperation and cooperative judgments or behavior that are nonmoral in cognitive type. Role: Boundary • Confidence: High

NDT component and commitment	Corpus continuity through 2025	Specific 2026 change	External concept and source-grounded statement	Relation rationale, opposed position, and downstream consequence
CTX015 · Moral-relationship qualification NDT defines a qualifying represented moral relationship by its builder magnitudes and subject gate, not by its relational-model type.	The 2020–2021 publications decompose moral cognition into basic nonmoral elements, allow moral judgment to emerge from object perception within moral space, and treat relational structure as central.	Recomposed it into molecular builders, relation qualifiers, and actual-contribution conditions.	Relational-model individuation of moral motives Distinct moral motives correspond to four basic types of social relationships. (Rai & Fiske, 2011)	Why this relation. The constructs overlap but neither determines the other: NDT permits qualifying moral relationships instantiating none of the four models, and permits instances of the four models that fail its builder conditions. Opposed position. A relational-model-sufficiency position predicts that instantiating one of the four relational models settles qualification as a moral relationship. What this changes downstream for NDT. The comparison requires crossing relational-model type with NDT builder magnitudes and subject-gate status; either classification can vary while the other is held fixed. Role: Central Contribution · Confidence: High
CTX016 · Learned social-value weighting Learned welfare weights can populate outcome evaluation without settling moral type.	The corpus permits learned values as benders/attractors and distinguishes them from relationship builders.	Added outcome maps, control profiles, and moral provenance boundaries.	Hierarchical Bayesian moral-weight inference Hierarchical Bayesian inference uses observed behavior to infer the weights that other agents attach to abstract moral principles. (Kleiman-Weiner, Saxe & Tenenbaum, 2017)	Why this relation. The inferred weights can supply evaluative content to NDT's outcome ordering, but neither the weights nor the inference process determines moral type. Opposed position. A learned-value-sufficiency position predicts that the social-value weights used in a token's outcome evaluation settle its moral type. What this changes downstream for NDT. Holding learned social-value weights and output fixed while varying represented relationship provenance, NDT predicts different moral types. Role: Implementation · Boundary · Confidence: Medium

3.6 Empirical signature

Describes observable judgment profiles without constituting moral cognition.

Exhibit 8. Empirical-signature comparisons. 1 row generated from CTX013 in dataset 1.5.

NDT component and commitment	Corpus continuity through 2025	Specific 2026 change	External concept and source-grounded statement	Relation rationale, opposed position, and downstream consequence
CTX013 • Domain signatures Observed differences among moral, conventional, and personal judgments may be empirical signatures without constituting the cognitive type.	The 2021 Nonmoral Conditions paper treats false positives and the moral/conventional distinction as a definitional problem.	Separated output signatures from constitutive support and moral provenance.	Domain-differentiated wrongness ranking Acts in the moral domain are ranked as more wrong than conventional acts, which are ranked as more wrong than personal acts. (Turiel, 1983)	Why this relation. The differentiated wrongness ranking is an observable judgment signature; NDT permits the same signature in tokens whose represented relationship provenance gives them different moral types. Opposed position. A wrongness-ranking-sufficiency position predicts that a canonical moral-domain wrongness response is sufficient for a judgment token to be moral. What this changes downstream for NDT. Holding the wrongness response fixed while varying represented relationship provenance, NDT predicts moral and nonmoral tokens with the same ranking signature. Role: Clarification • Confidence: High

3.7 Output approximation

Matches labels or contents without recovering derivation.

Exhibit 9. Output-approximation comparisons. 1 row generated from CTX023 in dataset 1.5.

NDT component and commitment	Corpus continuity through 2025	Specific 2026 change	External concept and source-grounded statement	Relation rationale, opposed position, and downstream consequence
CTX023 • Associative moralization Association graphs may model salience, attractors, or moral-label propagation but do not reconstruct a complete decision model or token provenance.	Attractors and ideological association are explicit in the 2025 essay.	Separated attractors from builders and actual support; registered output approximation as a boundary.	Moral Association Graph The Moral Association Graph uses a word-association network and psychological data to infer moral relevance. (Ramezani & Xu, 2025)	Why this relation. MAG can implement associative propagation or predict moral relevance; NDT requires framed agents, objects, alternatives, evaluation, and contribution. Opposed position. A moral-relevance-sufficiency position about MAG predicts that graph-based inference of moral relevance is sufficient to fix moral type without relationship provenance. What this changes downstream for NDT. The comparison predicts similar labels with different support routes and different moral type. Role: Boundary • Confidence: Medium

3.8 Tokenwise-only correspondence

Maps individual cases but does not define a stable model family.

Exhibit 10. Tokenwise-only correspondence. 1 row generated from CTX024 in dataset 1.5.

NDT component and commitment	Corpus continuity through 2025	Specific 2026 change	External concept and source-grounded statement	Relation rationale, opposed position, and downstream consequence
CTX024 • Tokenwise contextual variation NDT permits context-sensitive tokenwise guidance without treating particularism as one stable restricted family.	Ontological framing and context-sensitive moral relations are central through 2025.	Added typed model components that can vary case by case while still supporting general architectural claims.	Moral particularism Moral particularism denies that moral thought depends on exceptionless general principles. (Dancy. 2004)	Why this relation. Individual cases may map, but particularism alone does not specify a stable decision or provenance model. Opposed position. Strong generalism denies irreducible tokenwise context sensitivity. What this changes downstream for NDT. The row predicts that some particularist cases fit while the label itself does not determine NDT structure. Role: Boundary • Confidence: Medium

3.9 Orthogonal

Addresses truth/status rather than cognitive generation.

Exhibit 11. Orthogonal comparison. 1 row generated from CTX025 in dataset 1.5.

NDT component and commitment	Corpus continuity through 2025	Specific 2026 change	External concept and source-grounded statement	Relation rationale, opposed position, and downstream consequence
CTX025 • Metaethical neutrality NDT classifies cognitive generation without deciding whether a judgment is true, justified, or objectively valid.	The dissertation and 2025 relation essay explicitly separate moral-psychological type from normative correctness.	Formalized nested type, endorsement, and truth/status boundaries.	Objective-value denial Objective-value denial is the thesis that objective moral values or requirements do not exist. (Mackie, 1977)	Why this relation. Objective-value denial concerns whether objective moral values or requirements exist; NDT classifies how a judgment token is generated and leaves that status unsettled. Opposed position. A semantic-individuation position about objective-value denial predicts that purporting to represent objective value is sufficient to determine moral-judgment type. What this changes downstream for NDT. NDT can classify a moral judgment's generation without deciding whether objective moral values or requirements exist. Role: Clarification • Confidence: Medium

3.10 Too vague to verify

No stable typed translation and output criterion are specified; the formalization burden is recorded rather than simulated.

Exhibit 12. Too-vague-to-verify comparison. 1 row generated from CTX026 in dataset 1.5.

NDT component and commitment	Corpus continuity through 2025	Specific 2026 change	External concept and source-grounded statement	Relation rationale, opposed position, and downstream consequence
CTX026 • Normative polyphony Distinct locally coherent frame activations can support outputs that conflict only when the system supplies the relevant act, target, and appraisal bridge.	The 2018 dissertation and 2022 Polyphony Principle develop many-voiced moral cognition.	Added frame-event and judgment-event identity, comparison-readiness, alignment coverage, and five pair outcomes.	Bakhtinian polyphony Bakhtinian polyphony consists of independent, unmerged voices and consciousnesses. (Bakhtin. 1984)	Why this relation. Bakhtin supplies the acknowledged term and structural intuition; NDT adds frame-event identity, comparison bridges, alignment coverage, and typed pair outcomes, with no output-preserving translation claimed. Opposed position. Bakhtin's monologic model, in which a single authorial consciousness subsumes the voices, would reject independent local outputs that remain unresolved. What this changes downstream for NDT. The comparison supports conceptual lineage only; NDT predicts direct conflict only after comparison bridges survive, while otherwise plurality remains unaligned. Role: Lineage • Confidence: High

4. What the mapping establishes

Six rows establish exact or conditionally exact relations over stated domains: 2 embeddings, 3 submodules, and 1 conditional case. 6 rows locate established decision concepts inside restricted NDT families. 6 treat external constructs as compatible content, computation, or empirical signatures. The remaining 4 record output approximation, tokenwise correspondence, orthogonality, or indeterminate conceptual lineage.

The resulting originality claim is bounded. NDT's contribution does not consist in inventing the compared concepts. It lies in the architecture that connects normative detachment to token-specific evaluative organization, represented moral-relationship provenance, actual contribution, frame alignment, and polyphony, while marking where established concepts embed exactly, occupy restricted regions, supply compatible modules, or remain external to the classificatory question.

The continuity and change columns prevent formal maturation from being mistaken for origin. They distinguish commitments documented through 2025 from the specific formal, executable, or methodological additions made in 2026.

5. Authority and evidential limits

Formal Specification 1.4.4 controls the theory; Proposition Registry 1.6 controls claim status. Data contract 1.2 controls the comparison unit and fields. Dataset 1.5 controls every comparative row, the relation taxonomy controls the terms and definitions, and source register 2.6 controls citations, links, and verifying evidence. The figures explain reading order and taxonomy order; they add no comparison claims.

The Atlas is a claim-level positioning document. A high-confidence mapping records confidence in the stated relation, not empirical confirmation or philosophical truth. The bibliography is comprehensive for register 2.6, including entries not cited by a surviving comparison row.

6. Bibliography

Entries follow the corpus convention: source-register order, S01 through S45. Each author-date citation in the Atlas links to its entry; each entry links to the registered source.

- S01.** Aristotle, *De Motu Animalium*, Part 7 (trans. Farquharson). [Source link](#)
- S02.** David Hume, *A Treatise of Human Nature*, 2.3.3.3, SBN 414. [Source link](#)
- S03.** David Makinson & Leendert van der Torre (2007), "What is Input/Output Logic?" [Source link](#)
- S04.** John R. Searle (1964), "How to Derive Ought from Is." [Source link](#)
- S05.** Phan Minh Dung (1995), "On the Acceptability of Arguments..." [Source link](#)
- S06.** Jeremy Bentham, *An Introduction to the Principles of Morals and Legislation*, ch. I, §II. [Source link](#)
- S07.** Immanuel Kant, *Groundwork of the Metaphysics of Morals*, Formula of Humanity. [Source link](#)

- S08.** W. D. Ross (1930), *The Right and the Good*, ch. 2, p. 28. [Source link](#)
- S09.** Aristotle, *Nicomachean Ethics*, Books I–II (trans. W. D. Ross). [Source link](#)
- S10.** Virginia Held (1995), “The Meshing of Care and Justice.” [Source link](#)
- S11.** T. M. Scanlon (1998), *What We Owe to Each Other*, ch. 5, p. 189. [Source link](#)
- S12.** Robert Nozick, *Anarchy, State, and Utopia* (1974); SEP summary. [Source link](#)
- S13.** Thomas Aquinas, *Summa Theologiae* I–II, q.94, a.2. [Source link](#)
- S14.** Robert Merrihew Adams, modified divine-command theory; IEP summary. [Source link](#)
- S15.** Jonathan Dancy (2004), *Ethics Without Principles*, ch. 1, p. 7. [Source link](#)
- S16.** Graham et al. (2011), “Mapping the Moral Domain,” *Journal of Personality and Social Psychology* 101(2), 366–385, quote p. 366. [Source link](#)
- S17.** Schein & Gray (2018), “The Theory of Dyadic Morality,” *Personality and Social Psychology Review* 22(1), 32–70, quote p. 32. [Source link](#)
- S18.** Elliot Turiel (1983), *The Development of Social Knowledge: Morality and Convention*, ch. 4, p. 61. [Source link](#)
- S19.** Curry, Mullins & Whitehouse (2019), “Is It Good to Cooperate?” [Source link](#)
- S20.** Rai & Fiske (2011), “Moral Psychology Is Relationship Regulation,” *Psychological Review* 118(1), 57–75, p. 58. [Source link](#)
- S21.** Kleiman-Weiner, Saxe & Tenenbaum (2017), “Learning a Commonsense Moral Theory,” *Cognition* 167, 107–123, author manuscript p. 2. [Source link](#)
- S22.** Friston et al. (2017), “Active Inference: A Process Theory.” [Source link](#)
- S23.** Bai et al. (2022), “Constitutional AI: Harmlessness from AI Feedback.” [Source link](#)
- S24.** Hendrycks et al. (2021), “Aligning AI With Shared Human Values.” [Source link](#)
- S25.** Bree Beal (2020), “What Are the Irreducible Basic Elements of Morality?” [Source link](#)
- S26.** Bree Beal (2021), “The Nonmoral Conditions of Moral Cognition.” [Source link](#)
- S27.** Beal & Gogia (2021), “Cognition in Moral Space: A Minimal Model.” [Source link](#)
- S28.** Bree Beal (2022), “The Polyphony Principle.” [Source link](#)
- S29.** Beal, B. L. (2018). *Raskolnikov’s Confession: A Dostoevskian Model of Moral Psychology* [Doctoral dissertation, Emory University]. [Source link](#)
- S30.** Beal, B. (2025, January 20). *The Is-to-Ought Model of Moral Cognition*. Bree’s Substack. [Source link](#)
- S31.** Beal, B. (2025, January 31). *On the Relation Between Moral Philosophy and Moral Psychology*. Bree’s Substack. [Source link](#)
- S32.** Beal, B. (2026). *A Thousand Ramanujans*. Bree’s Substack. [Source link](#)
- S33.** Barbara Herman (1993), *The Practice of Moral Judgment*, ‘On the Value of Acting from the Motive of Duty,’ pp. 1–22, quote p. 21. [Source link](#)
- S34.** Nomy Arpaly (2003), *Unprincipled Virtue*, ‘Moral Worth,’ pp. 67–116, quote p. 72. [Source link](#)
- S35.** Julia Markovits (2010), ‘Acting for the Right Reasons,’ *The Philosophical Review* 119(2), 201–242, quote p. 202. [Source link](#)
- S36.** Paulina Sliwa (2016), ‘Moral Worth and Moral Knowledge,’ *Philosophy and Phenomenological Research* 93(2), 393–418, quote p. 393. [Source link](#)
- S37.** Hart, H. L. A., & Honore, T. (1985). *Causation in the Law* (2nd ed.). Oxford University Press. [Source link](#)
- S38.** Richard W. Wright (1985), ‘Causation in Tort Law,’ *California Law Review* 73(6), 1735–1828, quote p. 1774. [Source link](#)
- S39.** Joseph Y. Halpern & Judea Pearl (2005), ‘Causes and Explanations: A Structural-Model Approach. Part I: Causes,’ *British Journal for the Philosophy of Science* 56(4), 843–887. [Source link](#)

- S40.** Sydney Levine, Nick Chater, Joshua B. Tenenbaum & Fiery Cushman (2026), 'Resource-rational contractualism: A triple theory of moral cognition,' Behavioral and Brain Sciences 49, e133, manuscript p. 13. [Source link](#)
- S41.** Aida Ramezani & Yang Xu (2025), 'Moral Association Graph: A Cognitive Model for Automated Moral Inference,' Topics in Cognitive Science 17(1), 120–138, quote p. 120. [Source link](#)
- S42.** Pollock, J. L. (1995). Cognitive Carpentry: A Blueprint for How to Build a Person. MIT Press. [Source link](#)
- S43.** Horta, J. F. (2012). Reasons as Defaults. Oxford University Press. [Source link](#)
- S44.** Mikhail Bakhtin (1984), Problems of Dostoevsky's Poetics, ed. and trans. Caryl Emerson, ch. 1, p. 6. [Source link](#)
- S45.** J. L. Mackie (1977), Ethics: Inventing Right and Wrong, ch. 1, §2, p. 17. [Source link](#)