

Supplementary Material

The Logical Structure of Clinical Trial Eligibility Criteria: A Corpus-Scale Measurement of Disjunction in Non-Small-Cell Lung Cancer Trials

This supplement collects three items referenced in the main text: a formal model of criterion structure that states precisely what a conjunctive reading assumes and where it fails (Supplementary Note S1), a profile of the corpus (Supplementary Table S1), and a set of worked examples of the disjunction label on real criteria (Supplementary Table S2). Section, definition, and table numbers carry an “S” prefix throughout; unprefixed cross-references point to the main manuscript.

S1 Supplementary Note S1: A Formal Model of Criterion Structure

To state precisely what a conjunctive reading assumes and where it fails, we give a lightweight formal model of the structure that the main paper measures. The model is deliberately minimal, sufficient only to define the disjunction rate and to make the mis-specification claim exact, and it is the point of contact between this descriptive pilot and the formal analysis of the eligibility decision that we defer to future work.

Let X denote the space of *patient profiles*, and let $\mathcal{A} = \{a_1, \dots, a_m\}$ be a finite set of *atomic predicates*, each a Boolean-valued map $a_i : X \rightarrow \{0, 1\}$ that evaluates one elementary condition, such as *ECOG performance status* ≤ 1 or the presence of an activating *EGFR* mutation. A *criterion* c is a Boolean formula built from atoms with the connectives $\wedge$, $\vee$, and $\neg$, and we write $c(x) \in \{0, 1\}$ for its value on a profile x ; without loss of generality we take c to be in negation normal form, so that negation applies only to atoms and c is monotone in its *literals*, a literal being an atom or a negated atom. We treat the atoms as logically independent, so that the logical equivalence and satisfiability used below are propositional; where two atoms are in fact semantically dependent, such as *ECOG 0* and *ECOG 0–1*, the statements below quantify over the consistent truth assignments rather than over all of X .

Definition S1 (Trial eligibility). A trial T is a pair (I_T, E_T) of finite sets of criteria, its *inclusion* and *exclusion* criteria. Its eligibility function $\Phi_T : X \rightarrow \{0, 1\}$ is

$$\Phi_T(x) = \left(\bigwedge_{c \in I_T} c(x) \right) \wedge \left(\bigwedge_{c \in E_T} \neg c(x) \right),$$

so a profile is eligible when it satisfies every inclusion criterion and no exclusion criterion.

The two conjunctions *across* criteria in Definition ?? are not the assumption at issue; they express the genuine semantics of a criterion *list*. The assumption at issue concerns the internal structure of each individual criterion.

Definition S2 (Disjunctiveness). A criterion c is *conjunctive* if it is logically equivalent to a single conjunction of literals $\ell_1 \wedge \dots \wedge \ell_k$, and *disjunctive* otherwise; equivalently, c is disjunctive if and only if its minimal disjunctive normal form contains more than one term. We write $\text{disj}(c) = 1$ when c is disjunctive and 0 otherwise.

Disjunctiveness in the sense of Definition ?? is the property our extractor labels, and it is invariant to surface phrasing: *brain or leptomeningeal metastases* and *measurable or evaluable disease* are disjunctive, whereas *ECOG 0 to 1 and adequate organ function* is conjunctive. To make the conjunctive reading explicit, define the *conjunctive projection* of a criterion c as $\kappa(c) = \bigwedge_{\ell \in \text{lit}(c)} \ell$, the conjunction of every literal occurring in c . A coverage or matching model that treats each mentioned condition as separately required computes $\kappa(c)$ in place of c ; for *measurable or evaluable disease*, κ demands both, where c demands either.

Proposition S1 (Criterion-level under-approximation). *For every criterion c in negation normal form, $\kappa(c) \models c$, that is, $\{x : \kappa(c)(x) = 1\} \subseteq \{x : c(x) = 1\}$. The inclusion is an equality if and only if c is conjunctive.*

Proof. Since c is in negation normal form it is monotone in its literals, so setting every literal true, as $\kappa(c)$ does, makes c true; hence $\kappa(c) \models c$. If c is conjunctive it is equivalent to $\kappa(c)$ and the two sets coincide. If c is disjunctive its minimal disjunctive normal form has at least two terms, so there is a profile satisfying one term while falsifying a literal that occurs only in another; that profile satisfies c but not $\kappa(c)$, and the inclusion is strict. $\square$

Corollary S1 (Trial-level consequence). *Reading every criterion conjunctively shrinks the satisfaction set of each inclusion criterion and enlarges the admitted set of each exclusion criterion. Relative to the true eligible set $\{x : \Phi_T(x) = 1\}$, the conjunctive reading therefore under-admits on account of disjunctive inclusion criteria and over-admits on account of disjunctive exclusion criteria. Writing Φ_T^κ for eligibility computed under the projection, if every exclusion criterion of T is conjunctive then $\Phi_T^\kappa \subseteq \Phi_T$ (the reading only under-admits); if every inclusion criterion is conjunctive then $\Phi_T^\kappa \supseteq \Phi_T$ (it only over-admits); in general the two sets are incomparable, and they coincide exactly when every criterion of T is conjunctive.*

Corollary ?? is the formal content of the paper’s central claim, and it carries a consequence that the empirical results sharpen. Because eligibility negates exclusion criteria (Definition ??), a disjunctive exclusion read conjunctively *over-admits*, that is, it passes patients the protocol meant to exclude, which is the more consequential, safety-relevant error; and exclusion criteria are, as the main text shows, the more disjunctive class. The quantity that governs the size of both errors across a corpus is the criterion-level disjunction rate, which is the estimand of the study.

Definition S3 (Disjunction rate and per-trial share). For a finite corpus of criteria C , the criterion-level disjunction rate is

$$D(C) = \frac{1}{|C|} \sum_{c \in C} \text{disj}(c),$$

and for a single trial T the *per-trial disjunction share* is the same quantity restricted to that trial’s criteria, $D(T) = D(I_T \cup E_T)$.

The trial-level statement “ T contains at least one disjunctive criterion” is just $\mathbf{1}[D(T) > 0]$, a far coarser summary than $D(T)$ itself; for a trial with many criteria it is close to certain whenever $D(C)$ is bounded away from zero, so the distribution of $D(T)$, not the indicator, is the informative trial-level quantity. We do not observe $\text{disj}(c)$ directly; we estimate it with the validated extractor of the main text, and we estimate $D(C)$ from a random sample by the sample proportion $\hat{D}$ with a Wilson score interval, reporting it overall and within the strata of interest, namely inclusion versus exclusion, phase, sponsor class, and calendar era.

S2 Supplementary Table S1: Corpus profile

Two surface signals, visible before any extraction, motivated the analysis. The token “or” appears in 98.9 percent of trials, with a median of fifteen occurrences per trial, and parentheses, the orthographic marker of nested grouping, appear in 93.8 percent. These counts are only a loose upper signal, because many instances of “or” fall inside noun phrases that name a single concept rather than joining two eligibility alternatives, so the true logical disjunction rate cannot be read from string counting and must come from an extractor scored against human annotation.

Table S1: Profile of the NSCLC corpus (7,155 interventional trials; 7,154 with eligibility text; 188,706 criteria; median 23 criteria and 2,589 characters per trial). Counts are trials. Median criterion-section length is shown by era to support the discussion of the temporal trend in the main text.

<i>By phase</i>		<i>By sponsor class</i>	
Phase 2	2,874	Academic / other	3,633
Phase 1	1,435	Industry	3,035
Phase 3	914	NIH	205
Phase 1/2	843	Network	163
Not applicable	783	Other government	107
Phase 4	130	Federal	9
Phase 2/3	107		
Early Phase 1	68		

By five-year start block: trials (median criteria length, characters)

1990–94: 16 (1,649)	1995–99: 132 (1,683)	2000–04: 471 (1,838)	2005–09: 1,030 (1,980)
2010–14: 938 (2,069)	2015–19: 1,604 (3,189)	2020–24: 2,053 (3,143)	2025–: 871 (3,053)

S3 Supplementary Table S2: Worked examples of the disjunction label

Table S2: Worked examples of the disjunction label on real NSCLC criteria (text abbreviated; connectives italicized). *M* is the extractor’s label, *H* the human consensus. Rows 1–4 are true disjunctions across the surface forms the construct must span (noun-phrase “or”, a slash, “and/or”); row 5 is a numeric range, which is not a disjunction; row 6 is a model false negative, an “and/or” buried in a long sentence; rows 7–8 are two of the five criteria on which the three raters split 2–1, boundary cases where a surface connective and an eligibility-level alternative come apart. *H* is the majority label.

Criterion	Side	M	H
Active severe <i>or</i> uncontrolled chronic infection	exc	Y	Y
Histologically <i>or</i> cytologically documented NSCLC ... recurrent <i>or</i> progressive disease	inc	Y	Y
Arterial/venous thrombus events ... within 12 months	exc	Y	Y
Unstable angina <i>and/or</i> congestive heart failure requiring hospitalization	exc	Y	Y
ECOG performance status 0 to 2	inc	N	N
... agreement (by patient <i>and/or</i> partner) to use contraception ...	inc	N	Y
PD-L1 CPS ≥ 1 as determined by local <i>or</i> central IHC testing	inc	Y	Y*
Patients unable to swallow ... <i>and</i> patients with GI disorders	inc	Y	Y*

* 2–1 split among the three raters; the majority label is shown.

The written guidelines, used for both the human adjudication and the extractor prompt, mark as disjunctive an explicit “or” between eligibility options (*pregnant or breast feeding*), an “and/or”, an “either ... or” construction, a slash between genuine alternatives (*kidney injury/failure*), a list introduced as alternatives, and a parenthetical alternative. They mark as non-disjunctive a plain conjunction or single condition, a numeric range written with “to” or a hyphen, and an “or” that belongs to a fixed term or a definition rather than offering the patient a second way to qualify. The last case is the genuinely hard one, and it is the source of the residual disagreement reported in the main text (rows 7–8 above, resolved by the majority rule).