

GEMCo: A Validated, Ethically Releasable Proxy for Inaccessible Counselling Data

Philipp Steigerwald¹ Eric Rudolph¹ Mara Stieler²

Jennifer Burghardt² Robert Lehmann² Jens Albrecht¹

¹Faculty of Computer Science, Center for Artificial Intelligence (KIZ)

²Faculty of Social Sciences, Institute for E-Counselling

Technische Hochschule Nürnberg Georg Simon Ohm, Nürnberg, Germany

{philipp.steigerwald, eric.rudolph, mara.stieler,

jennifer.burghardt, robert.lehmann, jens.albrecht}@th-nuernberg.de

Abstract

This paper presents **GEMCo**, a releasable, human-written proxy for inaccessible counselling data: 86 complete German e-mail counselling conversations (728 messages), expert-authored cases and counsellor sessions with trained role-players.¹ It is validated against a held-out reference of 124 real counselling conversations. The proxy and the real conversations are measured against each other on counsellor strategies and client emotions. The gap is detectable but small. A generative validation supports the analysis. The corpus is the primary contribution. The validation method generalises to any domain where real data cannot be shared but a human-made proxy can. Privacy and ethics keep real counselling data closed. GEMCo carries none by design and can be released, a first step toward language research in this domain.

1 Introduction

Text-based counselling plays an increasingly important role in mental health service delivery (Engelhardt, 2021), yet computational tools for counsellor training, quality assurance or assistive systems like intake and triage bots require data that is rarely available due to its sensitive nature (Malgarelli et al., 2023).

Almost all available dialogue resources are English (Table 1), running from crisis conversations to transcribed professional sessions. Crisis Text Line holds text-message support at scale under restricted access (Althoff et al., 2016), while ESConv openly releases crowdsourced emotional-support chats (Liu et al., 2021). Closer to professional practice is DAIC-WOZ, which records clinical interviews (Gratch et al., 2014). AnnoMI and HOPE are transcripts of counselling demonstrations by professionals (Wu et al., 2022; Malhotra et al., 2022). PsyQA is large but single-turn, pairing each ques-

Dataset	Lang.	N	Mt.	Pr.	Public
Crisis Text Line (Althoff et al., 2016)	EN	80K+	✓	–	restricted
ESConv (Liu et al., 2021)	EN	1,053	✓	–	✓
DAIC-WOZ (Gratch et al., 2014)	EN	189	✓	✓	restricted
AnnoMI (Wu et al., 2022)	EN	133	✓	✓	✓
HOPE (Malhotra et al., 2022)	EN	212	✓	✓	on req.
PsyQA (Sun et al., 2021)	ZH	22K	–	–	on req.
SMHD-GER (Zanwar et al., 2023)	DE	–	–	–	restricted
OnCoCo (Albrecht et al., 2026)	DE/EN	2,778*	–	✓	✓
GEMCo	DE	86	✓	✓	✓

Table 1: Mental-health NLP dialogue resources. Mt. = multi-turn; Pr. = professionals; ✓ = yes, – = no. * Single messages, not full threads.

tion with a long counselling answer rather than a multi-turn dialogue (Sun et al., 2021).

German offers far less. SMHD-GER collects social-media posts rather than counselling dialogue (Zanwar et al., 2023). OnCoCo 1.0 annotates single online-counselling messages rather than full threads (Albrecht et al., 2026). German is spoken by over 100 million people, yet no natively German corpus of asynchronous multi-turn counselling exists. Ethical obligations toward real clients keep authentic counselling data out of public release.

This paper makes three contributions.

- (i) **GEMCo**, the first publicly available German corpus of asynchronous multi-turn e-mail counselling for mental-health support: 86 complete, human-written threads (GEMCo-A, expert-authored cases; GEMCo-B, counsellor sessions with trained role-players), released under CC BY 4.0 as a *proxy* for inaccessible counselling data.
- (ii) A generalisable method for validating such a proxy. A releasable, human-made proxy is compared with real data that cannot be shared. The gap is scaled by the real data’s own split-half resampling noise, so it reads as inside or beyond natural variation. It needs only aggregate label distributions of the withheld set, so it transfers to any sensitive domain.
- (iii) The validation itself, run on counselling. GEMCo-A, GEMCo-B and their pool are compared with the 124 withheld real conversations

¹Data: github.com/th-nuernberg/GEMCo.




Figure 1: Research design. The releasable human-generated *proxy* corpus (GEMCo-A + GEMCo-B) is validated against the withheld *real* reference.

at corpus, progress and message level, with a generative probe of the training signal.

The corpus is the main contribution. The method makes its release defensible and transfers beyond counselling.

2 Corpus Description

GEMCo was built at Technische Hochschule Nürnberg Georg Simon Ohm with active practitioners of German online counselling. The releasable human-generated proxy corpus is built from two sources. The expert-authored cases of GEMCo-A come from one of Germany’s largest counselling services. GEMCo-B was created through role-plays involving counsellors from this and two further German counselling services. Every thread is a complete, human-written asynchronous e-mail exchange and the counsellors were instructed to counsel just as they would in real practice. A third set of 124 authentic conversations donated by consenting and informed real clients, the held-out *real data*, serves for aggregate validation but cannot be released (Figure 1). The distinction between *proxy* and *real data* is releasability, not authenticity. Every message in all three corpora is human-written, none of it machine-generated. GEMCo-A and GEMCo-B are released in full under CC BY 4.0.

2.1 GEMCo-A: Expert Case Authoring

Two practising counsellors collaboratively authored 50 threads (403 messages), drawing on their experience in online counselling. Both are graduate social pedagogues with systemic and client-centred training, decades of counselling experience and more than ten years of it in online counselling. They modelled every case on how a counselling conversation typically begins and on the kinds of first request they meet in practice, then built it up from there. Both sides were written together and revised iteratively as each case took shape. Each thread was then read independently for realism by a third professional, a certified online counsellor who is also a social-science researcher. The cases deliberately span concern types (from parenting, school and relationships to identity, pregnancy, grief, self-harm and family conflict) and communication styles. GEMCo-A is

fictional and holds no real personal data, so it can be released without restriction.

2.2 GEMCo-B: Role-Playing Sessions

GEMCo-B was collected during a four-week field deployment (Steigerwald et al., 2025a) that paired 34 professional counsellors with 13 trained students playing counsellee roles. Nine counsellors worked in educational and family counselling, fourteen in addiction and eleven had just completed certification. Each student corresponded with two or three counsellors at once, freely interpreting shared case vignettes (family, youth, addiction), so threads emerged from genuine asynchronous exchanges over weeks. Several pairs sharing one vignette give controlled thematic overlap. The resulting 36 threads comprise 325 messages with the stylistic breadth of 34 practitioners, two of whom completed two threads each. Only the help-seeker role was played, by trained humans rather than a model, so the counselling itself is genuine and the counsellor side fully authentic. LLM-simulated clients reproduce the communication patterns of counselling role-plays only in part (Rudolph et al., 2025). All participants consented to donate their data and every client is role-played, so GEMCo-B carries no confidentiality constraints.

2.3 The Held-Out Real Data

The 124 real conversations were obtained through voluntary, revocable data donation with informed consent at a real counselling institution. The data was anonymised and pseudonymised in-house at the counselling service before any of it reached the research team. In this paper only aggregate statistics are reported from this actual counselling data. Structural differences from the proxy are expected (Table 2). The early drop-offs are the clearest. Nearly a quarter of the real threads end after at most two messages, typically a client enquiry and the counsellor’s reply with no further message. The other 95 run longer, while the proxy holds 86 complete threads end to end. Real clients write markedly longer messages (Table 2). This is a first trace of the client-side contrast the validation later locates. The real data is never released and serves




Figure 2: Annotation pipeline. Each e-mail is split into spans, classified at the finest OnCoCo (counsellor) and GoEmotions (client) level, collapsed to nine acts / seven emotions and block-merged (§3.1–§3.4).

	GEMCo-A	GEMCo-B	Real
Conversations	50	36	124
Messages	403	325	587
Words	118,514	64,660	206,082
Msgs/conv. (mean±sd)	8.1±3.8	9.0±4.2	4.7±3.3
Words/msg (client)	214±163	139±103	345±312
Words/msg (couns.)	377±235	274±204	356±331
Client msgs	206 (51%)	180 (55%)	262 (45%)
Counsellor msgs	197 (49%)	145 (45%)	325 (55%)

Table 2: Corpus statistics for the two released subcorpora and the withheld real reference.

only as the reference GEMCo is measured against in the sections that follow. A close match licenses the releasable proxy as an ethically clean stand-in where the real conversations cannot be shared. Every check uses all 124 real conversations, except where the analysis depends on conversation position or order. The progress bins (Section 5) and within-message transitions (Section 6) use the 48 complete threads, where position is well defined. Filtering is stated and motivated wherever it applies.

3 Annotation Pipeline

Both datasets, the proxy and the real data, pass through one shared pipeline. It labels spans (semantic blocks) along two dimensions: how counsellors structure their interventions and how client emotions are expressed. Two English counselling corpora serve later as a cross-domain contrast (Section 4), the crowdsourced emotional-support chats of ESConv (Liu et al., 2021) and the motivational-interviewing transcripts of AnnoMI (Wu et al., 2022). Both were chosen for their proximity to the GEMCo setting, each a mental-health support corpus though neither is asynchronous e-mail counselling. They pass through the same classifiers, applied to their native turn segmentation (Appendix B). The pipeline runs in four stages (Figure 2): (1) segmentation into semantic spans (§3.1), (2) classification of every span at the finest level and (3) a collapse to the nine-act and seven-emotion analysis level (§3.2, §3.3), then (4) a per-taxonomy block-merge (§3.4) that defines the unit every later measurement is computed over.

3.1 Segmentation

Counselling e-mails are long and weave several concerns and emotions together. A counsellor message greets, analyses, offers help and signs off. A client message runs through several emotions in turn. So a single label per message is rarely meaningful. To keep this internal structure, the pipeline first splits each e-mail into spans (stage 1, Fig. 2), semantic blocks that usually coincide with a sentence, each carrying one counsellor act and one emotion label. Counselling text often runs on without clear sentence boundaries, so each e-mail is segmented with the pre-trained Segment-any-Text (SaT) model (Frohmman et al., 2024), a neural segmenter that recovers boundaries even where punctuation is missing. On inspection this is more reliable than splitting on punctuation alone. The segmentation is kept deliberately fine. A per-taxonomy block-merge (§3.4) recombines same-label neighbours afterwards, so over-splitting is undone while a coarser split would risk collapsing distinct labels into one. GEMCo-A contributes 9,444 spans and GEMCo-B 4,493 (13,937 proxy), the real data a further 14,388. Together that is 28,325 spans across the 1,315 messages of proxy and the real data combined.

3.2 Counsellor Acts (OnCoCo)

Counsellor spans are classified with OnCoCo (Albrecht et al., 2026) (stage 2, Fig. 2), a counselling-act taxonomy built for online counselling, with the dataset and classifier publicly released. Its cross-validated macro F_1 of .72 is measured at the finest leaf level and rises to .83 when the right label is among the top two. The remaining errors mostly fall on neighbouring categories. It was trained on a bilingual German–English corpus of counselling messages manually segmented into labelled act spans, the same span-level input SaT produces, so its training and application share one segmentation. Being bilingual, it also labels the English cross-domain corpora (Appendix B). This paper collapses its output to the nine top-level categories, each capturing what the counsellor does (Figure 4), coarser and correspondingly more reliable than




Figure 3: Two-baseline scale for the proxy–real JSD (log axis).

the leaf level. Four are formal: formalities opening [FA] and closing [FC] a message, moderation [Mod] that structures the reply and other [O]. Five are therapeutic-process *Impact Factors*: analysis & clarification [AC], agreement on objectives [AO], creating motivation [CM], resource activation [RA] and help & problem solving [HP]. Per-category descriptions with examples are in Appendix H.

3.3 Client Emotions (Ekman)

Each client span is additionally classified for emotion (stage 2, Fig. 2) with the classifier of Lalk et al. (2025), which predicts the 28 GoEmotions categories (Demszyk et al., 2020). These 28 categories are mapped onto Ekman’s six basic emotions plus neutral (Ekman, 1992), the seven labels this analysis uses for what the client feels, following the mapping GoEmotions itself defines. That mapping table and the fuller GoEmotions-level analysis are given in Appendix E. The classifier is the closest available fit for GEMCo, developed for emotion detection in psychotherapy transcripts and fine-tuned on German. It reaches $F_1 = .45$ over all 28 categories, on par with the original English GoEmotions benchmark (.46). Being multilingual, it also labels the English cross-domain corpora (Appendix B). Collapsing to the seven Ekman classes lowers the resolution but folds together several categories the classifier tends to confuse.

3.4 Per-Taxonomy Block-Merging

SaT segments text more finely than the labels actually change. A single problem description can run over several semantic blocks that carry the same act or emotion and counting each separately would over-weight long passages. Before any analysis, consecutive spans with the same label are therefore merged into one block (stage 3, Fig. 2). For counsellor strategy, “Hello,” and “how are you

doing?” are both *Formalities (beginning)* and become one block. On the client side, “I cry all the time.” and “Nothing brings me joy anymore.” are both *sadness* and merge likewise. Each message then reads as a clean sequence of distinct OnCoCo acts on the counsellor side and of distinct emotions on the client side.

4 Corpus-Level Validation

Proxy is validated against the real data first by a vocabulary check (Section 4.1), then through two label-based lenses, the counsellor through *strategy* (OnCoCo acts, what they do) and the client through *emotion* (Ekman, what they feel).

Each lens pools all the labelled blocks of a corpus into one distribution and compares proxy against the real data with a single corpus-level number, the Jensen–Shannon divergence (JSD), which runs in $\log_2$ from 0 for identical distributions to 1 when they share no category (full definition, Appendix A).

Comparing the real data to a second sample of itself adapts classic split-half reliability (Spearman, 1910; Brown, 1910), while comparing the constructed proxy corpus to the real one follows distribution-level evaluation of model-generated against real data (Esteban et al., 2017; Pillutla et al., 2021).

4.1 Vocabulary Overlap

Before any labels, a classifier-free question: do the proxy and the real data draw on the same words? The measure is the Jaccard overlap, the share of word types two corpora have in common. The baseline is the real data against itself. Its 124 conversations are split into two random halves and their Jaccard overlap is measured. Repeated over 1,000 random splits, this gives a tight reference band (Jaccard $J^{RR} = .350 \pm .003$). The superscript names the two corpora compared, the real data first, so

RR is the real data against a second sample of itself and RP is the real data against the proxy. By speaker, the counsellor side reaches its own band ($J_{\text{counsellor}}^{\text{RP}} = .342$ vs. $J_{\text{counsellor}}^{\text{RR}} = .346 \pm .004$), so the professional wording is about as close as the real data is to itself, while the client side stays below it ($J_{\text{client}}^{\text{RP}} = .302$ vs. $J_{\text{client}}^{\text{RR}} = .322 \pm .004$).

4.2 Counsellor Strategy Comparison

Vocabulary overlap compares the words. The counsellor strategy check compares the counsellor acts, asking whether proxy deploys the same professional repertoire as the real counselling (full resolution in Appendix D). The gap is read on a scale with two empirical ends (Figure 3a, Table 3). The low end is the real data against itself. The same 1,000 half-splits scored by JSD form its own *noise band*, with the mean as the noise floor ($J_{\text{OnCoCo}}^{\text{RR}}$) and the 95th percentile ($P95^{\text{RR}}$) as the upper edge. The high end is the real data against the two external counselling corpora, far higher at .31 and .27. Proxy’s gap ($J_{\text{OnCoCo}}^{\text{RP}}$, Table 3) sits just past that upper edge, where 96 % of the within-real splits fall below it, with a small-to-medium effect size (Cramér’s V , Appendix A). The gap is small but systematic, measuring .020 on the coarser native segmentation, still an order of magnitude below the external anchors.

4.3 Client Emotion Comparison

The strategy check compares the counsellor’s acts. Emotion does this for the client, asking whether proxy’s clients move through the same feelings as the real ones, on the same two-ended scale (Figure 3b). The pooled proxy–real gap ($J_{\text{Ekman}}^{\text{RP}}$, Table 3) sits at the 99th percentile of the within-real noise band, near its edge. On the coarser native segmentation it sits two-to-three-fold below the external corpora (.017 vs. .044 and .051, Appendix B). The per-subcorpus rates (Appendix D) locate two shifts inside this small gap. A difference is counted only where a subcorpus’s 95 % confidence interval lies fully outside the real data’s. GEMCo-A reads a little warmer (more *joy*, ~30 % vs. ~25 %) but stays inside that interval, while GEMCo-B’s extra *surprise* falls outside it, its role-played clients overplaying their assigned cases. The professional side stays inside the noise band and the constructed client carries the difference.

Taxonomy	JSD^{RP}	95 % CI	V	JSD^{RR}
OnCoCo (9 cat.)	.0035	.0022–.0070	.07	.0016
Ekman (7 cat.)	.0036	.0016–.0082	.07	.0013

Table 3: Corpus-level proxy–real gap (JSD^{RP}) against the within-real noise floor (split mean JSD^{RR}); percentiles in the text, size-matched null in Appendix A.4.

A machine-translated ESConv or AnnoMI would not substitute for GEMCo. Their distance to the real data (.27–.31 vs. the proxy’s .020) is counselling format, modality and register along with language. Translation removes only the language part. The emotion classifier adds a caution: trained on DeepL-translated GoEmotions, it shows a macro- F_1 gap across the language boundary (.45 German vs. .38 English, Appendix B).

5 Conversation-Progress Validation

Conversation progress is the next check. It confirms that the match holds along the thread and that no opposite drifts cancel in the corpus. Each conversation is cut into five equal progress bins and read one by one, for GEMCo-A, GEMCo-B and the pooled corpus (Section 8 traces where the two subcorpora differ). Progress position is only well defined for a thread that runs to completion, so this view uses the 48 complete real threads with at least five messages, the same reference as the transitions in Section 6. The all-124 version is similar and given in Appendix A.4.

5.1 Counsellor Strategy per Bin

The nine strategy categories move together across GEMCo-A, GEMCo-B and the real data, steady along the thread (Figure 4, Table 4). The shared bend is *Analysis & Clarification* tapering toward the end, the arc from understanding a case to resolving it.




Figure 4: Counsellor strategy (OnCoCo, 9 cat.) over conversation progress. GEMCo-A hatched, GEMCo-B dotted, Real plain.

Pooling the two subcorpora helps on the counsellor side. The pooled GEMCo gap lands below both GEMCo-A and GEMCo-B at the corpus level and through the first three bins. The two lean off

Progress	Strategy			Emotion		
	A	B	GEMCo	A	B	GEMCo
All (corpus)	.0046	.0066	.0037	.0019	.0088	.0026
0–20 %	.0044	.0117	.0030	.0056	.0305	.0094
20–40 %	.0053	.0082	.0041	.0021	.0194	.0043
40–60 %	.0075	.0067	.0048	.0038	.0012	.0020
60–80 %	.0098	.0348	.0141	.0056	.0096	.0061
80–100 %	.0061	.0391	.0101	.0040	.0104	.0024
cross-domain	.27–.31			.04–.05		

Table 4: Per-bin JSD against the 48 complete real threads, for GEMCo-A (A), GEMCo-B (B) and pooled GEMCo.

the real data in different directions and partly cancel. GEMCo-A also stays inside the real data’s own noise band in every bin. A per-bin value in bold in Table 4 lies above that band’s 95th percentile ($P95^{RR}$). GEMCo-B is the one that breaks the band, in the closing two bins. Its role-played sessions still work the case to the end and lift *Analysis & Clarification* where the real sessions wind down. Even there the gap stays at least five-fold below the cross-domain distance (.27–.31, Table 4).

5.2 Client Emotion per Bin

Client emotion follows a clearer arc (Figure 5). In the real data and GEMCo-A, clients begin with elevated *sadness*, *fear* and *anger* that diminish as *joy* rises, the course a counselling exchange works toward. This is a structural signature, not direct evidence of a working alliance. Here GEMCo-A alone is closest to the real data and the pooled gap sits a little above it. GEMCo-B’s role-played clients overplay the emotion (Table 4). GEMCo-A holds inside the band throughout, only a touch warmer. GEMCo-B opens outside it in the first two bins. Its clients lead with *surprise* in place of the early *sadness* of real threads. The pooled corpus still stays inside the band, so adding GEMCo-B widens the case diversity at a small cost in emotional fit.




Figure 5: Client emotion (Ekman, 7 cat.) over conversation progress. GEMCo-A hatched, GEMCo-B dotted, Real plain.

6 Within-Message Transitions

Sections 4 and 5 checked how often each strategy and emotion occurs. This section checks their order, which act or emotion follows which inside a mes-

sage, for both the counsellor and the client. Such transition structure also supports next-act prediction in counselling dialogue (Rudolph et al., 2026). Each transition meets the same 95th-percentile test as the corpus gap (§4.2), now per transition, so each gets its own noise band. The false positives across the many tests are bounded by FDR control (Appendix A.5). The reference is the 48 real threads with at least five messages, where a full back-and-forth develops. Proxy holds only such complete threads. Full method and matrices are in Appendix G.

6.1 Counsellor Strategy Stability

The counsellor transitions stay within the real data’s noise. After correcting for the many transitions tested (FDR), a single shift stays significant: in GEMCo-B, fewer openings go straight to a sign-off (*Formalities (beginning)*→*Formalities (conclusion)*, −7.6 pp). It does not recur in GEMCo-A. With band medians of only 3.8 and 5.8 pp, a consistent shift would have surfaced, so within the message the proxy’s professionals order their strategy as the real data does.

6.2 Client Emotion Shifts

The client side does shift. The FDR survivors (1 in GEMCo-A and 6 in GEMCo-B) show the same positive-affect signature already seen in the distributions (§4.3). GEMCo-A lifts *surprise*→*joy* (+2.0 pp). GEMCo-B gains *neutral*↔*surprise* and loses *sadness* (*neutral*→*sadness*, −3.1 pp). The difference is in the order of feelings, not only their rates.

Together the two lenses place the residual divergence on the constructed client rather than on how the real counsellors sequence their work.

7 Cross-Speaker Interplay

The checks so far read each speaker on their own. This section puts the two together and asks whether counsellor and client fit each other in the proxy as they do in the real data. Linguistic Style Matching (LSM) (Niederhoffer and Pennebaker, 2002) measures whether the two speakers settle into the same way of writing, comparing their use of eight closed-class German function-word categories, topic-free words whose matching shows two people attuning. LSM uses pooled proxy and all 124 real threads, with no need for a thread to run to completion. Each conversation gets one score from 0 to 1. Both corpora score high (medians .768 proxy, .794 real), a touch higher in the real data. A

Kolmogorov–Smirnov test (Massey Jr., 1951) puts the distance between the distributions at $D = .312$ and resampling the real data against itself gives a noise band whose 95th percentile is .18, which D exceeds. Proxy pairs thus align slightly less tightly than real ones. Counsellor and client attune more in the real exchange than in drafted GEMCo-A or role-played GEMCo-B. The German function words have no English equivalent, so this check has no cross-domain anchor.

8 Where Proxy and the Real Data Differ

Sections 4–7 established that proxy tracks the real data closely. The residual differences fall primarily on the authored (GEMCo-A) and role-played (GEMCo-B) clients. Table 5 locates them, the three categories (of the nine strategies and seven Ekman emotions) where GEMCo-B’s 95 % CI lies fully outside the real data’s, no GEMCo-A category clearing it (Appendix D).

Category	GEMCo-A	GEMCo-B	Real	Pattern
<i>Counsellor strategies</i>				
Moderation	13.7 ± 2.0	11.0 ± 2.3	16.1 ± 1.5	B ↓
<i>Client emotions (Ekman)</i>				
surprise	12.9 ± 2.3	17.1 ± 3.2	12.1 ± 1.5	B ↑
sadness	11.0 ± 2.1	6.9 ± 2.2	12.7 ± 1.6	B ↓

Table 5: Per-subcorpus % for the three categories where GEMCo-B’s CI lies fully outside the real data’s. *Pattern*: B = GEMCo-B, arrow = direction.

8.1 Counsellor-Side Construction Effects

On the counsellor side the construction barely shows. Vocabulary, the corpus act gap ($JSD_{\text{OnCoCo}}^{\text{RP}} = .0035$), the per-bin arc and the within-message order all stay inside the real data’s band (§4.1–§6), only GEMCo-B’s lower *Moderation* clearing its CI. With little genuine back-and-forth to steer, less moderation is needed and motivational moves fill the space, so the residual sits on the client side.

8.2 Two Client Production Modes

On the client side the corpus emotion gap stays tiny ($JSD_{\text{Ekman}}^{\text{RP}} = .0036$), GEMCo-A leaning to more *joy* and GEMCo-B to more *surprise*, GEMCo-B’s opening bins and the order of feelings carrying it across corpus, thread and message (§4.1–§7). These trace to two production modes. GEMCo-A’s expert authors front-load warmth and motivation, a *narrative-closure* mode whose elevated *joy* stays within the real data’s CI. GEMCo-B’s role-players show the complementary *problem-fixation* mode, restating the assigned case where the real data lets

the solution play out. The elevated *surprise* (+5 pp) and reduced *sadness* (−6 pp) are the shifts that clear the CI.

8.3 The Paired-Design Advantage

GEMCo-A and GEMCo-B are complementary and lean off the real data in slightly different directions. Pooling lets their biases partly cancel, so the combined proxy lands closer to the real data than their average. The pooled strategy gap ($JSD_{\text{OnCoCo}}^{\text{RP}} = .0035$; Figure 3a) falls below both GEMCo-A’s (.0045) and GEMCo-B’s (.0064); for emotion the pool (.0036) beats their average but not GEMCo-A alone (.0031). Pooling also widens the diversity of cases and counsellor styles and yields more genuine counsellor messages, a richer training signal, so the two are best used together.

Distribution, order and cross-speaker interplay all hold up, but they read labels. The next section therefore puts the pooled corpus to a generative test, where using the two together pays off (Section 9).

9 Generative Validation

The earlier sections show that proxy and real match in what they contain. This last check asks whether the proxy also works as training data: a model trained only on it should write counsellor replies closer to the real ones than the off-the-shelf model does. The probe tests the corpus as a training signal, not a counselling generator to deploy. Mistral-Small-3.2-24B-Instruct (Mistral AI, 2025) was fine-tuned on the proxy counsellor turns alone, never on the real data, yielding *Proxy-SFT* (QLoRA adapter, Appendix I). At every counsellor turn of the 124 held-out real threads, three candidates write the next reply: the human counsellor, Proxy-SFT and the off-the-shelf base model with a strong counsellor prompt. Each is conditioned on the real conversation up to that turn, so the three differ only in the one reply.

9.1 LLM-as-Judge Ranking

A blind Mistral-Medium-3.5 model (Zheng et al., 2023) judged the candidates pairwise, each item showing the real conversation up to that point and two replies (A and B) and asking which better fits as the next counsellor message. Positions were swapped for every sample (AB vs. BA) to cancel any position bias (1,950 judgements, Table 6; prompt in Appendix J). The judge ran self-hosted, so no conversation data left the institution.

System	WR % [95 % CI]	BT	Rank
Human (reference)	73 [68–77]	1.71	1
Proxy-SFT	44 [41–48]	0.75	2
Off-the-shelf instruct	33 [29–37]	0.54	3

Table 6: Win rate (WR, bootstrap 95 % CI), Bradley–Terry strength (BT) and rank per system.

The Bradley–Terry strength (Bradley and Terry, 1952) condenses the pairwise outcomes into one score per system, its ratios giving the head-to-head odds. The human is preferred to both models, head to head 68:32 over Proxy-SFT and 77:23 over the off-the-shelf model. The proxy-tuned model beats the off-the-shelf one 57:43 and sits closer to the human. The proxy adapter thus appears to move output toward authentic counsellor messages. The judge is stable under order swaps (AB/BA agreement 87 %) and the shared Mistral-Small-3.2-24B backbone would favour the two candidates equally, not the non-Mistral human who still ranks first (Appendix I).

9.2 Embedding-Based Check: MAUVE

The judge states preferences. MAUVE (Pillutla et al., 2021) asks instead whether the model’s replies occupy the same distribution as the human ones. Every reply is embedded on its own (message text only, no context) as one vector from Ministral-8B-Instruct-2410 (Mistral AI, 2024) (last-token hidden state). MAUVE summarises the divergence between the two embedding clouds on a 0–1 scale, 1 = indistinguishable, 0 = no common region. Each row of Table 7 compares two disjoint halves of the 124 real conversations, matching one half’s human replies against the other half’s (the ceiling), Proxy-SFT’s and the off-the-shelf model’s, size-matched over 50 splits (Appendix C).

Comparison	MAUVE
Human $\leftrightarrow$ Human (split-half)	.818 $\pm$.074
Human $\leftrightarrow$ Proxy-SFT	.742 $\pm$.092
Human $\leftrightarrow$ Off-the-shelf instruct	.214 $\pm$.056

Table 7: MAUVE over 50 conversation-level splits.

The Proxy-SFT replies sit close to the within-human ceiling, within about a standard deviation of it. The off-the-shelf model, given the same conversation prefixes and a strong counsellor prompt, reaches barely a quarter of the scale. The distance between the two reflects what the adapter learned from proxy’s counsellor side, the side the validation placed inside the real data’s noise band.

10 Summary

GEMCo is a shareable, asynchronous counselling corpus that complies with research ethics and stays close to the withheld real data on every measure tested here. The proxy–real gap sits just above the real data’s own noise floor and far below the cross-domain distances. Proxy reproduces what counsellors and clients do, without copying how the real data is written.

The small remaining gap falls on the constructed client side, arguably the more tolerable one. The dyadic style is slightly looser and the register lighter while the professional side stays closest to the real data.

GEMCo-A and GEMCo-B are complementary, their writing modes partly cancelling when pooled (§8), so the paired corpus approximates the real data better than their average. The same design could transfer to other low-resource, ethically constrained domains, pairing a small withheld reference with releasable human data read against its own noise band.

Limitations

GEMCo is small, 86 threads, the scale of AnnoMI and HOPE. Further field deployments of the GEMCo-B type (Steigerwald et al., 2025a) can extend it. The intended use is evaluation and fine-tuning rather than pretraining and Section 9 shows the 86 threads already carry a usable training signal. The open release lets third parties add annotation layers and pool GEMCo with other German counselling data. Its validation is distributional and sequential, read through classifiers, showing no clinical or therapeutic equivalence. It shows that proxy deploys the professional repertoire at the real data’s rates and in the real data’s order, not that any single reply is good counselling. The JSD is dominated by the common categories, so a rare category can differ sharply between proxy and the real data while barely changing the score. Equivalence on such rare categories is therefore not claimed. To offset this, the conversation-progress bins (§5) and within-message transitions (§6) test the match at a finer grain and in sequence. It also assumes the classifiers make the same systematic errors on proxy as on the real data (§3.3), plausible but unverified. Both corpora pass through the same instrument, so a systematic classifier bias shifts both distributions together and largely cancels in the comparison. What does not cancel is per-category reliability, with emotion ($F_1 = .45$) less reliable than strategy (.72). Col-

lapsing emotion to the seven Ekman classes should absorb some of that error, though the gain is unverified. Every label here is classifier-produced rather than human-annotated, so a human audit of a sample is the natural next step. Higher-level counselling qualities are not measured: empathy, therapeutic appropriateness, coherence and conversational effectiveness. Validated empathy scales and counsellor ratings could measure them. The generative check relies on an LLM judge and MAUVE, with no rating by professional counsellors. A counsellor-rated comparison of generated replies would give stronger evidence. GEMCo also covers only a few counselling areas (§2), such as youth, family and addiction counselling, while counselling itself can address almost any topic, so the corpus generalises only so far.

Intended Uses and Ethics

GEMCo (the proxy) is meant to open counselling support for German NLP research, not to automate care. The expected primary use is counsellor-facing. The tools it supports stay human-in-the-loop: drafting and supervision aids, training, quality assurance and triage where professionals are scarce (Steigerwald et al., 2025b). GEMCo itself still needs further preparation for such tools, but it could help widen access to mental-health support by pairing AI with human counsellors.

Several uses fall outside what the corpus supports: unsupervised counsellor chatbots, crisis response without human escalation and diagnosis or clinical decision support. There is no outcome data for them and the dual-use risk is clear.

GEMCo-A is fully fictional. GEMCo-B comes from a study whose participants were all fully informed and consented to publication. The real data was donated under informed consent, revocable at any time (§2). The real data is never released and only aggregate statistics of it appear in this paper. The generative probe trains only on proxy, which holds no real client data.

Acknowledgments

Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – FIP 160 – Project-ID 549142762.

References

Jens Albrecht, Robert Lehmann, Aleksandra Poltermann, Eric Rudolph, Philipp Steigerwald, and Mara

Stieler. 2026. OnCoCo 1.0: A public dataset for fine-grained message classification in online counseling conversations. In *Proceedings of the Workshop on Social Context and Integrating NLP and Psychology to Study Social Interactions (SoCon-NLPSI)*, co-located with LREC 2026, Palma de Mallorca, Spain. To appear.

Tim Althoff, Kevin Clark, and Jure Leskovec. 2016. Large-scale analysis of counseling conversations: An application of natural language processing to mental health. *Transactions of the Association for Computational Linguistics*, 4:463–476.

Yoav Benjamini and Yosef Hochberg. 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1):289–300.

Ralph Allan Bradley and Milton E. Terry. 1952. Rank analysis of incomplete block designs: I. The method of paired comparisons. *Biometrika*, 39(3/4):324–345.

William Brown. 1910. Some experimental results in the correlation of mental abilities. *British Journal of Psychology*, 3(3):296–322.

Jacob Cohen. 1988. *Statistical Power Analysis for the Behavioral Sciences*, 2 edition. Lawrence Erlbaum Associates, Hillsdale, NJ.

Dorottya Demszky, Dana Movshovitz-Attias, Jeongwoo Ko, Alan Cowen, Gaurav Nemade, and Sujith Ravi. 2020. GoEmotions: A dataset of fine-grained emotions. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4040–4054.

Paul Ekman. 1992. An argument for basic emotions. *Cognition and Emotion*, 6(3–4):169–200.

Emily M. Engelhardt. 2021. *Lehrbuch Onlineberatung*, 2 edition. Vandenhoeck & Ruprecht, Göttingen.

Cristóbal Esteban, Stephanie L. Hyland, and Gunnar Rätsch. 2017. Real-valued (medical) time series generation with recurrent conditional GANs. *arXiv preprint arXiv:1706.02633*.

Markus Frohmann, Igor Sterner, Ivan Vulić, Benjamin Minixhofer, and Markus Schedl. 2024. Segment any text: A universal approach for robust, efficient and adaptable sentence segmentation. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 11908–11941, Miami, Florida.

Jonathan Gratch, Ron Artstein, Gale Lucas, Giota Stratos, Stefan Scherer, Angela Nazarian, Rachel Wood, Jill Boberg, David DeVault, Stacy Marsella, David Traum, Skip Rizzo, and Louis-Philippe Morency. 2014. The distress analysis interview corpus of human and computer interviews. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*, pages

- 3123–3128, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Christopher Lalk, Kim Targan, Tobias Steinbrenner, Jana Schaffrath, Steffen Eberhardt, Brian Schwartz, Antonia Vehlen, Wolfgang Lutz, and Julian Rubel. 2025. [Employing large language models for emotion detection in psychotherapy transcripts](#). *Frontiers in Psychiatry*, 16.
- Siyang Liu, Chujie Zheng, Orianna Demasi, Sahand Sabour, Yu Li, Zhou Yu, Yong Jiang, and Minlie Huang. 2021. [Towards emotional support dialog systems](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*, pages 3469–3483.
- Matteo Malgaroli, Thomas D Hull, James M Zech, and Tim Althoff. 2023. [Natural language processing for mental health interventions: A systematic review and research framework](#). *Translational Psychiatry*, 13(1):309.
- Ganeshan Malhotra, Abdul Waheed, Aseem Srivastava, Md Shad Akhtar, and Tanmoy Chakraborty. 2022. [Speaker and time-aware joint contextual learning for dialogue-act classification in counselling conversations](#). In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining (WSDM)*, pages 735–745.
- Frank J. Massey Jr. 1951. The kolmogorov-smirnov test for goodness of fit. *Journal of the American Statistical Association*, 46(253):68–78.
- Mistral AI. 2024. Un Ministral, des Ministraux: Introducing the world’s best edge models. <https://mistral.ai/news/ministraux>.
- Mistral AI. 2025. Mistral Small 3.2 (24B Instruct 2506). <https://huggingface.co/mistralai/Mistral-Small-3.2-24B-Instruct-2506>.
- Kate G. Niederhoffer and James W. Pennebaker. 2002. Linguistic style matching in social interaction. *Journal of Language and Social Psychology*, 21(4):337–360.
- Krishna Pillutla, Swabha Swayamdipta, Rowan Zellers, John Thickstun, Sean Welleck, Yejin Choi, and Zaid Harchaoui. 2021. MAUVE: Measuring the gap between neural text and human text using divergence frontiers. In *Advances in Neural Information Processing Systems 34 (NeurIPS)*.
- Eric Rudolph, Philipp Steigerwald, and Jens Albrecht. 2025. [Comparing human role-players and LLM-simulated clients in online counselling training: An analysis of counselling patterns](#). In *Proceedings of the 18th International Conference on Educational Data Mining (EDM)*, pages 381–387.
- Eric Rudolph, Philipp Steigerwald, and Jens Albrecht. 2026. [Transition-matrix regularization for next dialogue act prediction in counselling conversations](#). In *Findings of the Association for Computational Linguistics: ACL 2026*, pages 25442–25457.
- Charles Spearman. 1910. Correlation calculated from faulty data. *British Journal of Psychology*, 3(3):271–295.
- Philipp Steigerwald, Nico Bienlein, Jennifer Burghardt, Mara Stieler, Robert Lehmann, and Jens Albrecht. 2025a. [CAIA in practice: Field evaluation of an AI-assisted support system for text-based online counselling](#). In *Proceedings of the 37th IEEE International Conference on Tools with Artificial Intelligence (ICTAI)*.
- Philipp Steigerwald, Jennifer Burghardt, Eric Rudolph, and Jens Albrecht. 2025b. AI systems in text-based online counselling: Ethical considerations across three implementation approaches. In *Proceedings of the 3rd International Conference on Frontiers of Artificial Intelligence, Ethics, and Multidisciplinary Applications (FAIEMA)*.
- Hao Sun, Zhenru Lin, Chujie Zheng, Siyang Liu, and Minlie Huang. 2021. [PsyQA: A Chinese dataset for generating long counseling text for mental health support](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1489–1503.
- Zixiu Wu, Simone Balloccu, Vivek Kumar, Rim Helaoui, Ehud Reiter, Diego Reforgiato Recupero, and Daniele Riboni. 2022. [Anno-MI: A dataset of expert-annotated counselling dialogues](#). In *Proceedings of ICASSP 2022 – IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 6177–6181.
- Sourabh Zanwar, Daniel Wiechmann, Yu Qiao, and Elma Kerz. 2023. [SMHD-GER: A large-scale benchmark dataset for automatic mental health detection from social media in German](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1526–1541.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. Judging LLM-as-a-judge with MT-bench and chatbot arena. In *Advances in Neural Information Processing Systems 36 (NeurIPS), Datasets and Benchmarks Track*.

A Statistical Methodology

Two measures compare how counsellor strategies and client emotions are distributed across the proxy and the real data. The Jensen–Shannon divergence (JSD) quantifies how different two distributions are and Cramér’s V gives a standardised effect size. Let P_P and P_R denote the label distributions of the proxy and real corpus over k categories, with $P_{P,i}$ and $P_{R,i}$ the proportions of category i . The reported proxy–real value JSD^{RP} is symmetric in its

arguments. Within-real split-half values are JSD^{RR} (Appendix A.4).

A.1 Jensen–Shannon Divergence

JSD is the primary distributional distance for three reasons. It is **symmetric**, so proxy–real and real–proxy agree and the within-real split-half null is well defined. It is **bounded** in $[0, \log_2 2] = [0, 1]$, so values compare directly across taxonomies of very different cardinalities ($k = 7$ to $k = 66$). And it stays **finite when one distribution puts zero mass on a category the other observes**, a routine situation in small corpora that makes KL divergence infinite. JSD builds on the Kullback–Leibler divergence,

$$D_{\text{KL}}(P_P \| P_R) = \sum_{i=1}^k P_{P,i} \log_2 \frac{P_{P,i}}{P_{R,i}}. \quad (1)$$

D_{KL} is asymmetric. JSD symmetrises it through a mixture $M = \frac{1}{2}(P_P + P_R)$,

$$\text{JSD}(P_P \| P_R) = \frac{1}{2} D_{\text{KL}}(P_P \| M) + \frac{1}{2} D_{\text{KL}}(P_R \| M). \quad (2)$$

With $\log_2$, JSD equals zero when the distributions are identical and one when their supports are disjoint. 95 % CIs are obtained by resampling conversations with replacement (10,000 iterations), accounting for within-conversation correlation (intra-class correlation $\approx .03$ – $.04$ by taxonomy).

A.2 Cramér’s V

Cramér’s V expresses the distributional difference as a standardised effect size,

$$V = \sqrt{\frac{\chi^2}{n \cdot \min(r-1, k-1)}}, \quad (3)$$

where χ^2 is Pearson’s statistic, n the total observations, r the number of groups and k the number of categories. With $r = 2$ corpora this simplifies to $V = \sqrt{\chi^2/n}$, bounded in $[0, 1]$ (0 for identical category distributions, 1 when corpus membership is fully determined). Cohen (1988) sets small/medium/large for w at $.1/.3/.5$ and cautions that a constant w reads differently as the table grows (§7.2.3; his conversion is $\phi' = w/\sqrt{r-1}$, formula 7.2.6). Here the thresholds are divided by $\sqrt{k-1}$ instead, a deliberately stricter convention, giving $\approx .04/.11/.18$ at $k=9$ and $\approx .01/.04/.06$ at $k=66$. V is reported raw and judged against these stricter thresholds, not the $df=1$ benchmark.

A.3 Design-Effect Adjusted CIs

Spans within a conversation are correlated, so per-category CIs use a cluster-adjusted standard error,

$$\hat{\sigma} = \sqrt{\frac{p(1-p)}{n}} \cdot \text{DEFF}, \quad (4)$$

$$\text{DEFF} = 1 + (\bar{n}_{\text{conv}} - 1) \cdot \rho, \quad (5)$$

where p is the category proportion, n the total span count, $\bar{n}_{\text{conv}}$ the mean spans per conversation and ρ the within-conversation intraclass correlation, measured per taxonomy (.033 OnCoCo, .040 Ekman, .031 GoEmotions). Per-category tables report $\pm 1.96\hat{\sigma}$ (95 % CI half-width). Overlapping intervals suggest the difference may be within sampling uncertainty.

A.4 Real Split-Half Reference Distribution

To characterise the within-real noise band, the real data’s 124 conversations are partitioned into two equally sized halves $R_1^{(b)}, R_2^{(b)}$ at $B = 1,000$ random split points. For each split b ,

$$\text{JSD}_{(b)}^{\text{RR}} = \text{JSD}(P_{R_1^{(b)}} \| P_{R_2^{(b)}}), \quad (6)$$

yielding the empirical distribution $\{\text{JSD}_{(b)}^{\text{RR}}\}_{b=1}^B$ with mean JSD^{RR} . Splitting at the conversation level rather than the span level preserves within-conversation correlation, so the distribution is the real data’s own noise band. The proxy–real divergence JSD^{RP} is placed on it by its empirical percentile (Table 9), a value below the 95th percentile sitting inside the noise band. The distribution is right-skewed and bounded below at zero, so the percentile is reported rather than a parametric z .

The conversation-progress bins (§5) and transitions (§6) use the 48 complete real threads, where position and order are well defined. Including the early drop-offs changes little: the corpus pooled proxy–real JSD is .0035 (counsellor) and .0036 (emotion) over all 124, against .0037 and .0026 over the 48 and the two real distributions differ by only $\text{JSD} = .0004$ on both lenses. Per bin the drop-offs leave a small trace in the *Formalities*-heavy edge bins of the counsellor lens. A short thread is largely its opening and closing. Against all 124 the pooled strategy gap rises in the opening bin from .0030 to .0082 and in the closing from .0101 to .0207, while the inner three bins and every emotion bin move by at most .005. The complete-thread reference of §5 gives the cleaner per-bin picture.

A second variant removes a residual sample-size confound. The 62/62 split-half has a wider

noise band than the actual proxy–real comparison ($n_P = 86$, $n_R = 124$), so a further bootstrap null draws both groups with replacement from the 124 real conversations at the matched sizes ($B = 1,000$). The resulting band is tighter and confirms the ordering in Table 9. Under this n -matched null every proxy–real value exceeds the 95th percentile, at the 9-category resolution $JSD^{RP} = .0035$ vs. $.0020$, finer taxonomies progressively further out (Table 8). All nonetheless remain several-fold to an order of magnitude below the cross-domain references on the shared native basis (Appendix B).

Taxonomy	JSD^{RP}	JSD^{RR-nm}	p_{95}^{RR-nm}
OnCoCo Co. (9)	.0035	.0009	.0020
Ekman (7)	.0036	.0008	.0017
GoEmotions (28)	.0087	.0034	.0057
OnCoCo all (66)	.0159	.0039	.0065

Table 8: n -matched null (both groups resampled from the real data, $n_P = 86$, $n_R = 124$); every proxy–real JSD exceeds its 95th percentile P_{95}^{RR-nm} .

Taxonomy	JSD^{RP}	JSD^{RR}	Pctl.
OnCoCo Co. (9)	.0035	.0016	96
OnCoCo Co. (38)	.0086	.0051	96
OnCoCo all (66)	.0159	.0059	>99
Ekman (7)	.0036	.0013	99
GoEmotions (28)	.0087	.0057	95

Table 9: Proxy–real JSD vs. the within-real split-half (1,000 splits); Pctl. = its percentile. Bracketed sizes are nominal, the real data observes 32 of 38 and 59 of 66.

A.5 Multiple-Comparison Control

The within-message analysis (§6) tests every transition cell at once, up to 9×8 off-diagonal counsellor pairs and 7×6 client pairs. With that many tests at the 95 % level, some cells cross their band by chance. The *false discovery rate* (FDR) is the expected share of falsely flagged cells among all flagged ones and Benjamini–Hochberg control (Benjamini and Hochberg, 1995) caps it at $q = 0.10$. Each eligible cell (off-diagonal, with at least 10 combined transitions) carries a two-sided empirical p -value, the fraction of the 2,000 real-data resamples whose $|\Delta|$ reaches the observed $|\Delta|$. Ordering these as $p_{(1)} \leq \dots \leq p_{(m)}$, the procedure flags every cell up to the largest rank k for which

$$p_{(k)} \leq \frac{k}{m} q. \quad (7)$$

At most a fraction q of the flagged shifts are expected to be false, so a surviving cell is a shift that resampling the real data rarely produces. The level $q = 0.10$ is a convention, not derived from the data. FDR analyses use 0.05 or 0.10 and 0.10 is the more

lenient choice, common in exploratory multiple testing. Even at this lenient level few cells survive, as most differences sit inside the noise band.

B JSD Baselines

The proxy–real divergence is compared against two reference points: a split-half baseline (expected noise within the real data itself) and cross-domain counselling corpora (ESConv (Liu et al., 2021) and AnnoMI (Wu et al., 2022)), all annotated with the same pipeline. ESConv and AnnoMI are counselling, so the contrast stays fair. They differ in language, modality and format, so they mark what a genuinely different corpus looks like. An in-domain German e-mail corpus would sit almost on the real data and give the scale no upper end. Field-foreign text such as news would be trivially far.

English enters only through these two reference corpora; the proxy–real validation itself is German throughout. The OnCoCo classifier handles the English input in-distribution, having been trained on bilingual German–English counselling data (§3.2). The emotion classifier is a multilingual XLM-RoBERTa fine-tuned on German only and labels the English corpora in zero-shot cross-lingual transfer. On the original English GoEmotions test set it reaches $F_1 = .38$ over the 28 categories (vs. $.45$ on German) and $.49$ after the Ekman collapse. The emotion reference values are therefore noisier than the German measurements but remain meaningful. The anchors are read as the distance to a genuinely different corpus, in which language, modality and counselling format differ together, not as a pure domain distance. The tables in this appendix use the native, coarser OnCoCo message segmentation rather than the finer SaT-enriched spans of the main text, so absolute JSDs are larger here (e.g. $.020$ vs. $.0035$ for counsellor strategies, $.017/.049$ vs. $.0036/.0087$ for Ekman/GoEmotions). The position of proxy–real relative to the noise floor and the cross-domain references is unchanged.

Table 10 reports counsellor strategy divergence at four OnCoCo granularity levels. After block-merging, the proxy–real JSDs sit at one to two-and-a-half times the within-real noise floor and roughly an order of magnitude below the AnnoMI/ESConv references, with proxy–real $V \leq .25$ throughout while cross-domain V jumps to $.38-.67$.

Table 11 shows the same pattern for client emotions. At GoEmotions resolution the ordering is proxy–real $\ll$ proxy–ESConv $<$ proxy–AnnoMI and the proxy–real gap ($.049$) sits close to the real data’s own native noise floor ($.041$).

Comparison	Co. (9)		Co. (38)		Cl. OnCoCo (28)		all (66)	
	JSD	V	JSD	V	JSD	V	JSD	V
Proxy-Real	.020	.16	.044	.23	.039	.21	.045	.23
Real split-half	.009	.10	.035	.20	.037	.21	.036	.20
Proxy-AnnoMI	.215	.52	.415	.67	.363	.63	.389	.65
Proxy-ESConv	.213	.43	.400	.48	.257	.39	.327	.44
Real-AnnoMI	.270	.57	.445	.66	.443	.67	.445	.67
Real-ESConv	.307	.47	.466	.50	.322	.38	.399	.45

Table 10: OnCoCo divergence at four granularities (native annotations). Co. = counsellor, Cl. = client, all = both speakers; V = Cramér’s V (Appendix A).

Comparison	Ekman (7)		GoEmotions (28)	
	JSD	V	JSD	V
Proxy-Real	.017	.14	.049	.24
Real split-half	.008	.10	.041	.22
Proxy-AnnoMI	.017	.12	.199	.46
Proxy-ESConv	.022	.07	.126	.20
Real-AnnoMI	.051	.17	.290	.50
Real-ESConv	.044	.07	.172	.16

Table 11: Client emotion divergence (native annotations); proxy-real $\ll$ cross-domain at GoEmotions-28, washing out at Ekman-7.

Table 12 breaks the comparison down by conversation-progress quintile. All per-bin $V \leq .21$, confirming that the similarity holds at every stage of the conversation.

Bin	Co. (9)		Co. (38)		Ekman		GoEmotions	
	JSD	V	JSD	V	JSD	V	JSD	V
0–20 %	.008	.11	.024	.18	.008	.11	.030	.20
20–40 %	.009	.11	.021	.17	.005	.08	.020	.17
40–60 %	.002	.06	.020	.16	.004	.08	.015	.14
60–80 %	.011	.12	.025	.18	.011	.12	.022	.17
80–100 %	.021	.17	.032	.21	.002	.06	.020	.16
Corpus	.004	.07	.009	.11	.004	.07	.009	.11

Table 12: Per-bin proxy-real divergence over five quintiles (block-merged); all $V \leq .21$, last row = corpus level.

C MAUVE Setup

Both MAUVE analyses (the generative check in Section 9 and the corpus-level raw-text companion below) use the reference implementation (Pillutla et al., 2021). MAUVE is bounded in $[0, 1]$, 1 = the two text distributions are indistinguishable, 0 = they share no common region. The score is the area under the KL-divergence curve between the two sides’ cluster histograms, obtained by quantising the pooled embeddings. Each GEMCo message or candidate reply is one sample; the cross-domain corpora contribute one utterance per sample (600 sampled from each). Features are the last-token hidden state of Ministral-8B-Instruct-2410 (fp16, texts truncated at 512 tokens). The model is multi-lingual, so the English reference corpora share the embedding space. Every run uses identical settings (25 quantisation buckets, fixed seed), so the values are directly comparable.

Generative check (§9). At each of the 325 judged counsellor turns there are three replies: the human counsellor’s, the proxy-tuned model’s and the off-the-shelf model’s. The 124 real conversations split into two random halves of 62, about 162 turns per side. Each half contributes its conversations’ turns, so disjoint conversation sets sit on the two sides and no model reply meets the human reply of its own turn. The human replies of half one are compared with three second halves: the human replies (ceiling), the proxy-tuned replies and the off-the-shelf replies at those turns. Embeddings, quantisation (25 buckets) and seed are shared across the rows, which differ only in whose replies stand in the second half. Table 7 reports mean $\pm$ standard deviation over 50 splits.

At full sample size (325 vs. 325, not size-matched) the values are .907 (human-Proxy-SFT), .254 (human-instruct) and .511 (Proxy-SFT-instruct). MAUVE rises with sample size, so these are not comparable to the split-half ceiling.

Corpus-level raw-text check. At the whole-message level the within-real split-half reference is $.940 \pm .030$ (minimum .846). Against the real data, pooled proxy reaches .560, GEMCo-A .654 and GEMCo-B .195. The cross-domain anchors sit at .004–.005 (proxy/real vs. ESConv) and .004–.004 (vs. AnnoMI). At the surface level proxy is clearly distinguishable from the real data, yet two orders of magnitude closer to it than the cross-domain corpora. The label-based validation of Sections 4–8 measures what remains once this surface register drops out.

Two controls separate length from register. Splitting the real data at its median message length gives MAUVE = .060 between the short and the long half, so the embedding is strongly length-sensitive. Matching the real data to GEMCo-B’s length distribution (decile matching, drawn with replacement) nonetheless leaves the GEMCo-B–real value at .196 (vs. .195 unmatched). GEMCo-B’s distance therefore reflects its role-played register and the deliberately shared case vignettes (§2), not message length.

D Per-Subcorpus Distributions With Confidence Intervals

Table 13 gives the per-subcorpus distributions in compact form, for GEMCo-A, GEMCo-B and the real data. Tables 14 and 15 add 95 % design-effect-adjusted confidence intervals ($\pm 1.96\hat{\sigma}$, $\rho = .033/.040/.031$ by taxonomy; Eq. 4) for every category at all three resolutions: counsellor strategies

(9 OnCoCo categories), Ekman emotions (7) and GoEmotions (28). Table 5 in the main text picks the three rows where GEMCo-B’s CI lies fully outside the real data’s. Every other category overlaps within its interval across GEMCo-A, GEMCo-B and the real data.

Category	GEMCo-A	GEMCo-B	Real
<i>Counsellor strategy (OnCoCo 9)</i>			
Formalities (beginning)	6.5	8.0	7.5
Formalities (conclusion)	8.9	8.5	8.9
Moderation	13.7	11.0	16.1
Other	9.4	8.1	9.2
Analysis & Clarification	18.6	24.7	21.3
Agreement on Objectives	6.2	7.3	5.8
Creating Motivation	10.6	10.1	8.0
Resource Activation	6.9	4.7	4.7
Help & Problem Solving	19.3	17.5	18.7
<i>Client emotion (Ekman 7)</i>			
joy	29.5	25.4	24.5
sadness	11.0	6.9	12.7
fear	3.6	3.9	4.2
anger	7.3	8.2	8.7
surprise	12.9	17.1	12.1
disgust	0.6	0.6	0.8
neutral	35.2	37.8	36.9

Table 13: Label distribution (%) by subcorpus: counsellor strategy (OnCoCo 9) over counsellor blocks and client emotion (Ekman 7) over client blocks. Proxy pools GEMCo-A and GEMCo-B; bold = 95 % CI fully outside the real data’s.

Category	GEMCo-A	GEMCo-B	Real
<i>Counsellor strategies (OnCoCo 9)</i>			
Formalities (beginning)	6.5 ± 1.4	8.0 ± 2.0	7.5 ± 1.1
Formalities (conclusion)	8.9 ± 1.7	8.5 ± 2.1	8.9 ± 1.2
Moderation	13.7 ± 2.0	11.0 ± 2.3	16.1 ± 1.5
Other	9.4 ± 1.7	8.1 ± 2.0	9.2 ± 1.2
Analysis & Clarification	18.6 ± 2.3	24.7 ± 3.2	21.3 ± 1.7
Agreement on Objectives	6.2 ± 1.4	7.3 ± 1.9	5.8 ± 1.0
Creating Motivation	10.6 ± 1.8	10.1 ± 2.2	8.0 ± 1.1
Resource Activation	6.9 ± 1.5	4.7 ± 1.6	4.7 ± 0.9
Help & Problem Solving	19.3 ± 2.3	17.5 ± 2.8	18.7 ± 1.6
<i>Client emotions (Ekman 7)</i>			
joy	29.5 ± 3.1	25.4 ± 3.7	24.5 ± 2.0
sadness	11.0 ± 2.1	6.9 ± 2.2	12.7 ± 1.6
fear	3.6 ± 1.3	3.9 ± 1.6	4.2 ± 0.9
anger	7.3 ± 1.7	8.2 ± 2.3	8.7 ± 1.3
surprise	12.9 ± 2.3	17.1 ± 3.2	12.1 ± 1.5
disgust	0.6 ± 0.5	0.6 ± 0.7	0.8 ± 0.4
neutral	35.2 ± 3.2	37.8 ± 4.1	36.9 ± 2.3

Table 14: Full per-subcorpus percentages ($\pm 1.96\hat{\sigma}$): counsellor strategy (OnCoCo 9) and client Ekman emotion (7).

E Client Emotion at GoEmotions Resolution

The main text reports client emotion at the seven-class Ekman level (§4.3). The finer 28-category resolution localises the residual proxy–real gap to specific affective categories. Table 16 gives the official grouping of the 28 GoEmotions categories onto Ekman’s six basic emotions plus neutral.

Category	GEMCo-A	GEMCo-B	Real
<i>Client emotions (GoEmotions 28)</i>			
admiration	3.6 ± 1.1	2.8 ± 1.3	2.8 ± 0.7
amusement	0.8 ± 0.5	0.7 ± 0.7	0.2 ± 0.2
approval	6.3 ± 1.5	6.3 ± 1.9	6.4 ± 1.0
caring	2.7 ± 1.0	2.2 ± 1.1	2.8 ± 0.7
desire	2.2 ± 0.9	1.9 ± 1.1	1.7 ± 0.6
excitement	1.1 ± 0.6	0.5 ± 0.6	0.7 ± 0.4
gratitude	8.3 ± 1.7	7.9 ± 2.1	7.6 ± 1.1
joy	3.1 ± 1.0	1.8 ± 1.0	2.3 ± 0.6
love	5.1 ± 1.3	2.6 ± 1.2	3.0 ± 0.7
optimism	3.1 ± 1.0	2.3 ± 1.2	2.2 ± 0.6
pride	0.3 ± 0.3	0.3 ± 0.4	0.2 ± 0.2
relief	0.5 ± 0.4	0.7 ± 0.7	0.5 ± 0.3
confusion	4.0 ± 1.2	8.3 ± 2.1	3.7 ± 0.8
curiosity	4.5 ± 1.3	4.0 ± 1.5	3.2 ± 0.7
realization	3.8 ± 1.2	3.8 ± 1.5	4.5 ± 0.9
surprise	0.5 ± 0.4	1.2 ± 0.8	0.6 ± 0.3
anger	0.7 ± 0.5	0.4 ± 0.5	1.1 ± 0.4
annoyance	3.1 ± 1.0	3.6 ± 1.4	3.2 ± 0.7
disappointment	3.3 ± 1.1	2.4 ± 1.2	4.0 ± 0.8
disapproval	2.9 ± 1.0	3.6 ± 1.4	3.9 ± 0.8
disgust	0.5 ± 0.4	0.6 ± 0.6	0.7 ± 0.4
embarrassment	0.4 ± 0.4	0.4 ± 0.5	0.6 ± 0.3
fear	2.0 ± 0.8	2.2 ± 1.1	2.7 ± 0.7
grief	0.0 ± 0.1	0.0 ± 0.0	0.1 ± 0.1
nervousness	1.1 ± 0.6	1.4 ± 0.9	1.2 ± 0.5
remorse	0.5 ± 0.4	0.2 ± 0.4	0.3 ± 0.2
sadness	5.7 ± 1.4	3.3 ± 1.4	7.2 ± 1.1
neutral	29.9 ± 2.8	34.5 ± 3.7	32.7 ± 2.0

Table 15: Full per-subcorpus percentages ($\pm 1.96\hat{\sigma}$): client emotion at GoEmotions resolution (28 categories).

Ekman	GoEmotions categories
joy	admiration, amusement, approval, caring, desire, excitement, gratitude, joy, love, optimism, pride, relief
surprise	confusion, curiosity, realization, surprise
sadness	disappointment, embarrassment, grief, remorse, sadness
anger	anger, annoyance, disapproval
fear	fear, nervousness
disgust	disgust
neutral	neutral

Table 16: Official GoEmotions→Ekman grouping (Demszyk et al., 2020) behind the seven-class client-emotion resolution (§3.3).

Figure 6 shows the full GoEmotions distribution over conversation progress for GEMCo-A, GEMCo-B and the real data. The corpus-level divergence sits at the 95th percentile of the real data’s noise band ($JSD_{GoEmotions}^{RP} = .0087$, $V = .11$; Table 9) and on the shared native basis stays several-fold below the cross-domain references (.049 vs. .17–.29; Table 11).

Two production modes separate at this resolution. GEMCo-A’s *narrative-closure* mode shows elevated *love* (+2.1 pp), which clears the real data’s CI although its coarser *joy* form does not. GEMCo-B’s *problem-fixation* mode shows elevated *confusion* (+4.6 pp) and reduced *sadness* (−3.9 pp), both outside the real data’s CI. These map onto the *joy* and *surprise* shifts of Section 8.




Figure 6: Client emotion (GoEmotions, 28 cat.) over conversation progress. GEMCo-A hatched, GEMCo-B dotted, Real plain.

F Counsellor Strategy at Leaf Resolution

The main text reports counsellor strategy at the nine top-level OnCoCo categories (§4.2). OnCoCo is hierarchical, each category subsuming finer acts down to 38 *leaf* categories (Albrecht et al., 2026) and agreement holds at this finer resolution. Figure 7 shows the twelve most frequent leaf categories over conversation progress for GEMCo-A, GEMCo-B and Real. Proxy and the real data track closely throughout and no leaf category becomes a subcorpus signature, unlike the client emotion modes of Appendix E.

The divergence grows with resolution but remains small. At the 38-leaf level the gap rises to $JSD_{\text{OnCoCo}}^{\text{RP}} = .044$ ($V = .23$, native annotations), an order of magnitude below the cross-domain references, whose JSD to ESConv and AnnoMI reaches .40–.47 ($V = .48$ –.67; Table 10).





Figure 7: Counsellor strategy at OnCoCo leaf resolution (top 12 of 38 categories) over conversation progress. GEMCo-A hatched, GEMCo-B dotted, Real plain.


G Within-Message Transitions

These matrices give the full per-cell results behind Section 6, where a cell is one ordered pair of consecutive blocks (row = current, column = next), the text is the raw proxy–real delta (pp) and the colour its standardised exceedance Δ/τ_{95} of that cell’s within-real noise band (hatched where too sparse to test). The band τ_{95} is the 95th percentile of the proxy–real difference over 2,000 real-data resamples. Shifts past the band are flagged under Benjamini–Hochberg FDR control ($q = 0.10$, Appendix A.5). Figure 8 shows counsellor strategy

transitions, where no cell is consistently shifted across the two subcorpora. The client emotion transitions yield several FDR survivors: more movement around *joy* in GEMCo-A and less traffic into and out of *sadness* in GEMCo-B (§6).



(a) GEMCo-A counsellor vs. the real data (≥ 5 -message threads).



(b) GEMCo-B counsellor vs. the real data (≥ 5 -message threads).

Figure 8: Counsellor strategy transitions vs. the 48 real threads with at least five messages (§6). Cell text = raw delta (pp), colour = Δ/τ_{95} , hatched = too sparse to test.

H OnCoCo Category Descriptions

Table 17 describes the nine OnCoCo top-level counsellor categories, with short corpus-style example utterances and English glosses. The four formal and five Impact-Factor categories are listed in §3.2.

I Generative Validation Setup

Fine-tuning. The proxy-tuned model is Mistral-Small-3.2-24B-Instruct-2506 with a QLoRA adapter (4-bit NF4, double quantisation, bf16 compute) trained on the 342 proxy counsellor turns (GEMCo-A + GEMCo-B) for 3 epochs: LoRA

Code	Category	Description & example
FA	Formalities (beginning)	Opening greeting, acknowledging the message. <i>“Liebe Maria, vielen Dank für deine Nachricht.”</i> [Dear Maria, thank you for your message.]
Mod	Moderation	Conversation management, structuring or signposting the reply. <i>“Ich gehe gleich auf deine einzelnen Fragen ein.”</i> [I’ll address your individual questions in a moment.]
AC	Analysis & Clarification	Probing or reflecting the client’s situation back. <i>“Habe ich richtig verstanden, dass du dich von deinen Eltern unter Druck gesetzt fühlst?”</i> [Did I understand correctly that you feel pressured by your parents?]
AO	Agreement on Objectives	Proposing or agreeing on a counselling goal. <i>“Vielleicht konzentrieren wir uns zunächst auf das Gespräch mit deiner Mutter.”</i> [Perhaps we first focus on the conversation with your mother.]
CM	Creating Motivation	Encouragement, motivational reframing. <i>“Du hast schon einen wichtigen Schritt gemacht, indem du dich gemeldet hast.”</i> [You’ve already taken an important step by reaching out.]
RA	Resource Activation	Pointing the client to personal or social resources. <i>“Vielleicht könnte deine beste Freundin dir in dieser Situation Rückhalt geben.”</i> [Maybe your best friend could give you support in this situation.]
HP	Help & Problem Solving	Concrete advice or suggested action. <i>“Du könntest einen Termin bei der Schulpsychologin vereinbaren.”</i> [You could make an appointment with the school psychologist.]
O	Other	Statements outside the schema (meta-talk, asides). <i>“Ich kenne deine Stadt nicht persönlich, aber ...”</i> [I don’t know your city personally, but ...]
FC	Formalities (conclusion)	Closing remarks, sign-off. <i>“Ich wünsche dir alles Gute. Liebe Grüße, Anna.”</i> [I wish you all the best. Warm regards, Anna.]

Table 17: OnCoCo top-level counsellor categories.

rank 64, $\alpha = 128$, dropout 0.05 on all attention and MLP projections; learning rate 5×10^{-5} (cosine, 3% warmup); effective batch size 8 (1×8 gradient accumulation); max sequence length 4096; paged_adamw_32bit; seed 42. The off-the-shelf baseline is the same base model without the adapter.

Candidate generation. All candidates use nucleus sampling at temperature 0.7, top- p 0.95, up to 600 new tokens, seed 42 and no length control. Proxy-SFT keeps its training-time User:/Counsellor: template and system prompt. The off-the-shelf model uses the engineered counsellor system prompt below.

Proxy-SFT System Prompt — German (verbatim) / English

Du bist eine erfahrene psychosoziale Beraterin in einer asynchronen E-Mail-Beratung. Antworte mit der nächsten Beratenden-Nachricht an die ratsuchende Person — im Stil, Ton und Umfang einer realen E-Mail-Beratungsantwort. Keine Rollenangaben, keine Meta-Kommentare.

You are an experienced psychosocial counsellor in an asynchronous e-mail counselling service. Reply with the next counsellor message to the help-seeking person — in the style, tone and length of a real e-mail counselling reply. No role labels, no meta comments.

Off-the-Shelf Counsellor System Prompt — German (verbatim) / English

Du bist eine erfahrene psychosoziale Beraterin in einer asynchronen E-Mail-Beratung eines deutschen Online-Beratungsdienstes. Du erhältst den bisherigen E-Mail-Verlauf zwischen einer ratsuchenden Person und der Beratung. Schreibe ausschließlich die nächste Nachricht der Beratenden.

Orientiere dich an professioneller, systemisch-klientenzentrierter Beratungspraxis:

- Greife konkret auf, was die ratsuchende Person zuletzt geschrieben hat; spiegle Gefühle, ohne zu beschwichtigen oder zu übertreiben.

- Dränge früh im Verlauf nicht sofort auf Lösungen — kläre und erkunde zuerst. Biete konkrete Schritte erst an, wenn der Verlauf reif dafür ist.
- Schreibe im Stil, Ton und Umfang einer realen deutschen E-Mail-Beratungsantwort: zusammenhängender Fließtext mit Anrede und Schlussformel, KEINE Markdown-Listen, KEINE Aufzählungspunkte, keine Überschriften.
- Keine Rollenlabels, keine Meta-Kommentare, keine Hinweise darauf, dass du eine KI bist.

Antworte nur mit dem Text der Beratenden-Nachricht.

You are an experienced psychosocial counsellor in the asynchronous e-mail counselling of a German online counselling service. You receive the e-mail exchange so far between a help-seeking person and the counselling service. Write only the counsellor’s next message.

Follow professional, systemic-client-centred counselling practice:

- Pick up concretely on what the help-seeking person last wrote; mirror feelings without placating or overdoing them.
- Do not push for solutions early in the exchange — clarify and explore first. Offer concrete steps only once the exchange is ready for them.
- Write in the style, tone and length of a real German e-mail counselling reply: coherent running text with salutation and closing, NO markdown lists, NO bullet points, no headings.
- No role labels, no meta comments, no hints that you are an AI.

Reply only with the text of the counsellor message.

Judging. The blind Mistral-Medium-3.5 judge (system prompt in Appendix J) returns a forced JSON object with a binary winner (no ties) and a confidence, parsed by a tolerant brace-matching extractor. Each (conversation, turn, system pair) is judged in both AB and BA orders. Win rates and Bradley-Terry strengths aggregate all 1,950 judgements.

Judge family. The Mistral judge is not expected to bias the ranking (§9); only the absolute margins could be affected, so the models’ true distance to the human could be larger than measured.

J Judge Prompt

The generative validation (Section 9) uses the blind pairwise judge below: an experienced online counsellor, four fixed quality dimensions and a forced binary choice with no tie. The per-comparison user message supplies the conversation prefix and the two next-turn candidates as A and B; the prompt is given in the original German above the dashed rule, with the English translation beneath.

Judge System Prompt — German (verbatim) / English

Du bist eine erfahrene Sozialpädagogin und psychosoziale Beraterin mit langjähriger Praxis in der Online-Beratung. In deinem Berufsalltag liest du ständig Beraterinnen-Antworten von Kolleginnen mit — und du erkennst sofort, ob jemand sein Handwerk versteht oder ob die Antwort wie aus einem Standard-Baukasten klingt.

Was du jetzt tust. Du siehst einen echten Gesprächsverlauf zwischen einer Ratsuchenden und einer Beraterin, abgebrochen kurz bevor die nächste Beraterinnen-Antwort geschrieben werden müsste. Anschließend liest du zwei Vorschläge für genau diese nächste Antwort — Kandidat A und Kandidat B. Du weißt NICHT, wer sie verfasst hat (echte Kollegin oder Sprachmodell).

Worauf du achtest:

- (1) **Empathie & emotionale Intelligenz** — wird das Anliegen ernst genommen, Gefühle gespiegelt (nicht beschwichtigt, nicht übervalidiert)?
- (2) **Professionelle Haltung** — geschulte Fachkraft oder ChatGPT (Markdown-Listen, “Sie haben absolut recht!”, übergriffige Lösungen, Trost-Floskeln)?
- (3) **Phasen-Gefühl** — passt der Modus zu DIESER Stelle? Früh nicht sofort Lösungen; spät kein ergebnisloses Weiter-Erkunden.
- (4) **Beziehungsarbeit** — knüpft die Antwort konkret an das zuletzt Geschriebene an, lädt sie zu einer Antwort ein?

Deine Wahl. Du musst dich entscheiden — kein Unentschieden. Sind beide ähnlich gut, wähle die, die spürbar mehr nach geschulter Kollegin klingt. Sind beide schwach, die mit geringerem Schaden für die Ratsuchende.

Ausgabe. Antworte AUSSCHLIESSLICH mit dem JSON-Objekt im vorgegebenen Schema. Kein Vor- oder Nachtext.

You are an experienced social pedagogue and psychosocial counsellor with many years of practice in online counselling. In your daily work you constantly co-read colleagues' counsellor replies — and you recognise at once whether someone understands their craft or whether the reply sounds like a standard template kit.

What you are doing now. You see a real conversation between a help-seeker and a counsellor, cut off just before the next counsellor reply would have to be written. You then read two proposals for exactly this next reply — Candidate A and Candidate B. You do NOT know who wrote them (a real colleague or a language model).

What you look for:

- (1) **Empathy & emotional intelligence** — is the concern taken seriously, feelings mirrored (not placated, not over-validated)?
- (2) **Professional stance** — a trained professional or ChatGPT (markdown lists, “You’re absolutely right!”, overbearing solutions, comforting platitudes)?
- (3) **Sense of phase** — does the mode fit THIS point? Early: no immediate solutions; late: no fruitless further exploration.
- (4) **Relationship work** — does the reply connect concretely to what was last written, does it invite a response?

Your choice. You must decide — no tie. If both are similarly good, choose the one that sounds noticeably more like a trained colleague. If both are weak, the one that does less harm to the help-seeker.

Output. Respond ONLY with the JSON object in the given schema. No text before or after.