Supplementary methods and analysis record

S1 | Exhaustive membership and selection history

The scientific target is the response organization of a fixed engine. The original extraction adapter collects every registered mechanism in functions F1–F8, resolves duplicate names using the canonical first-function-wins policy, and asserts 84 mechanisms with 889 output channels. The metadata defines the layer membership. The original comprehensive_prepare.py constructs a feature group for every mechanism; comprehensive_prediction.py iterates over every group. Thus RASN, CSSL, IACM, PWUP and the other highlighted mechanisms were not selected for inclusion on the basis of their names or observed accuracy. All were measured.

The layer and channel extension

The revised analysis enumerates every metadata layer and every C³ channel. There are 326 layers and 889 channels, yielding 1,299 mechanism/layer/channel candidates for the centroid analysis. Figures select informative examples from this complete inventory. The six mechanisms used for exhaustive subset illustrations are RASN, PWUP, IACM, MMP, CSSL and OII; these follow the whole-mechanism results and are not described as a preregistered mechanistic hypothesis. All single-layer exclusions cover the complete inventory, not just these six examples.

Documented chronology

The 14 October comprehensive analysis plan is explicitly exploratory and was written after initial native-output results had been inspected. It already specifies all 84 mechanisms and the three-tag and five-tag tasks. The 16 October decomposition plan expands to layers, channels and combinations. The nearest-centroid extension followed discovery of the scalar-ridge limitation and is recorded as adaptive. The current manuscript does not recast this sequence as a confirmatory selection made before the data existed. Reproducible discovery across the fixed engine is the primary design.

Definitions used in all tables

A channel is one named scalar trajectory. A layer is the metadata-defined set of channels belonging to a processing stage within a mechanism. A mechanism is the joint set of its channels. Complete C³ is the concatenation of all 889 mechanism channels. A request is a generation event whose available returned alternatives are averaged before analysis. A held-out text is excluded from all parameter estimation for the external readout in that fold. These operations never average unrelated mechanism channels into a single scalar.

Exact evidence for the enumeration

selection-history.json records source paths, line evidence and SHA-256 hashes for the extraction adapter, original feature catalog, original prediction loop and revised reader. The catalog retains actual layer names and channel members. The declared M layer is Memory in some mechanisms and Temporal Integration in others; MMP uses R/P/F/C. The constant reserved outputs MEAMN:F3, VRIAP:F2 and CSSL:F2 remain in the schema but do not become active scaled predictors.

Supplementary Note S2 | External measurement

MI computation and readout estimation are separate

The adapter receives only an audio path and produces the fixed MI output arrays. It does not accept production labels and contains no corpus-specific optimizer. A separate analysis script estimates a mapping between those arrays and the archived production tags. In the main evaluation, parameters for that mapping are calculated using two texts, and the third text supplies the observations to be predicted. The term training, when it occurs in analysis code or third-party software, refers exclusively to this external estimation operation.

Why a scalar ridge rule can conceal information

Let z be a centered scalar channel and let the K class targets be one-hot vectors. With equal class counts, each class has target mean 1/K. Ridge fitting with an intercept therefore yields scores $s_k(z) = 1/K + b_k z$. For $z > 0$, the winning class has the largest slope; for $z < 0$, it has the smallest. Except for ties, only these two classes can be predicted. For a balanced five-class task, this creates a structural ceiling of 2/5 on the fraction correct even if five classes occupy distinguishable scalar intervals. Centering, regularization or a different linear slope estimate does not create the missing interval-specific intercepts. Floating-point ties can affect degenerate edge cases; the limitation concerns the centered balanced rule.

Nearest-centroid measurement

For a scalar, the class mean can occupy a distinct position on the line, creating several nearest-mean decision regions. For a vector, the same rule assigns the closest mean under squared Euclidean distance. Means and scales are calculated using only the estimation requests. The new reader was applied to every candidate, not only to channels that looked favorable in a plot. The complete 5,208 result rows and confusion matrices are retained. It reveals accessible tag structure; it does not establish that an MI output is a native classifier or that a channel's biological label has been validated.

Layer exclusion and the scaling control

Original ridge divides standardized features by the square root of the number of active features. Removing a layer changes that denominator as well as the information supplied to the reader. The fixed-full-scale control retains the denominator from the complete mechanism. IACM without P gives 26/45 under that control, versus 37/45 for all layers; CSSL without E gives 20/27, versus 26/27. PWUP's M+P advantage is more sensitive: it gives 34/45 with original scaling but 29/45 with the complete-mechanism denominator, versus 30/45 for all layers. These distinctions remain explicit in the archived subset table.

What the layer comparisons establish

The comparisons change which stored channels the external reader can access. They establish the observed contribution of those outputs to that reader and task. The remaining channels still contain any influences of upstream computations from the original engine run. Disabling an internal computation and re-running dependent stages would be a different intervention. No such internal causal intervention is claimed here. No new significance test is attached to the difference between a complete and reduced set.

Supplementary Note S3 | Discovery and automatic choice

A secondary question

Mapping where an existing engine represents a distinction requires examining its outputs. It does not require selecting one channel before performing that examination. A separate practical question is whether an automatic selection rule can choose a useful representation using only the two available estimation texts. The following nested experiment addresses that secondary question. It is not the definition of success for the comprehensive response map and is not presented as a corrected score for an individual named mechanism.

Nested procedure

For each outer held-out text, the other two texts form two inner folds. Each candidate is evaluated by estimating the reader from one inner text and predicting the other. The candidate with highest pooled inner agreement is selected; ties follow a fixed lexicographic feature ID. The selected candidate's outer prediction is then evaluated on the untouched third text. All selected IDs, tied candidate lists and outer errors are archived. Reader type and the candidate family are fixed within each reported run.

Observed automatic-choice results

With original ridge and 84 mechanism candidates, correct-request counts are 26/45 for five-tag speech, 23/45 for five-tag song, 22/27 for three-tag speech and 13/27 for three-tag song. With nearest centroid and 889 single-channel candidates, the corresponding counts are 24/45, 35/45, 20/27 and 22/27. These results quantify the instability of selecting a representation from only two texts. They do not change the separately evaluated 26/27 CSSL result or the complete inventory of responses, but they constrain claims about an automatic selection procedure.

Discovery status of the highlighted results

The highlighted mechanisms and channels are identified after inspecting this corpus. Their held-out predictions exclude the evaluation text from reader estimation, while the choice of which result to highlight uses the complete response map. Candidate-wide permutation controls explicitly account for the stated within-reader candidate search. Independent data would still be required to estimate prospective accuracy for the highlighted candidates or validate the same layer contribution patterns in new voices and texts.

Why the secondary analysis remains available

The full record includes both the response map and automatic-choice analysis so that distinct research questions are not conflated. Omitting low automatic-choice results would conceal a real stability constraint. Replacing a named mechanism's measured response with those results would instead answer a different question. The main article leads with the representational findings and the supplement provides the selection algorithm, decisions and limitations required to interpret them.

Supplementary Fig. S1a | Complete mechanism inventory

Original fixed ridge. Entries are correct requests / total requests; every registered mechanism is retained in metadata order. Three-tag and five-tag tasks have separate denominators.

Mechanism	Channels	Layers	Speech 5	Song 5	Speech 3	Song 3
BCH	16	E+M+P+F	27/45	25/45	18/27	18/27
CSG	12	E+M+P+F	28/45	20/45	17/27	19/27
MIAA	11	E+M+P+F	23/45	16/45	23/27	10/27
MPG	10	E+M+P+F	27/45	22/45	22/27	16/27
PCCR	11	E+M+P+F	30/45	23/45	20/27	14/27
PNH	11	H+M+P+F	22/45	22/45	20/27	16/27
PSCL	16	E+M+P+F	28/45	19/45	21/27	17/27
SDNPS	10	E+M+P+F	26/45	28/45	24/27	19/27
SDED	10	E+M+P+F	26/45	26/45	24/27	20/27
TPIO	10	E+M+P+F	20/45	22/45	16/27	14/27
TPRD	10	T+M+P+F	25/45	29/45	20/27	21/27
HTP	12	E+M+P+F	32/45	26/45	20/27	18/27
SPH	14	E+M+P+F	23/45	22/45	15/27	14/27
ICEM	13	E+M+P+F	30/45	24/45	22/27	12/27
PWUP	10	E+M+P+F	31/45	30/45	26/27	22/27
WMED	11	E+M+P+F	21/45	22/45	13/27	17/27
UDP	10	E+M+P+F	20/45	23/45	18/27	17/27
CHPI	11	E+M+P+F	23/45	25/45	15/27	17/27
IGFE	9	E+M+P+F	26/45	14/45	21/27	17/27
MAA	10	E+M+P+F	25/45	23/45	16/27	16/27
PSH	10	E+M+P+F	27/45	23/45	20/27	18/27
SNEM	12	E+M+P+F	15/45	20/45	14/27	20/27
IACM	11	E+M+P+F	37/45	21/45	22/27	15/27
BARM	10	E+M+P+F	20/45	19/45	19/27	14/27
STANM	11	E+M+P+F	29/45	24/45	17/27	16/27
AACM	10	E+M+P+F	24/45	18/45	18/27	12/27
AMSS	11	E+M+P+F	19/45	16/45	11/27	16/27
ETAM	11	E+M+P+F	22/45	24/45	16/27	19/27
DGTP	9	E+M+P+F	25/45	23/45	14/27	18/27
SDL	9	E+M+P+F	15/45	14/45	13/27	10/27
NEWMD	10	E+M+P+F	21/45	23/45	22/27	16/27
PWSM	9	E+M+P+F	23/45	26/45	18/27	15/27
MEAMN	12	E+M+P+F	25/45	19/45	18/27	16/27
MMP	12	R+P+F+C	36/45	20/45	20/27	16/27
RASN	11	E+M+P+F	28/45	23/45	27/27	20/27
PMIM	11	P+M+S+F	30/45	25/45	23/27	17/27
OII	10	E+M+P+F	30/45	30/45	21/27	21/27
HCMC	11	E+M+P+F	25/45	20/45	24/27	12/27
RIRI	10	E+M+P+F	19/45	18/45	11/27	18/27
MSPBA	11	S+M+P+F	27/45	27/45	21/27	21/27
VRIAP	10	E+M+P+F	21/45	16/45	17/27	9/27
CMAPCC	10	E+M+P+F	27/45	19/45	20/27	20/27

Supplementary Fig. S1b | Complete mechanism inventory

Original fixed ridge. Entries are correct requests / total requests; every registered mechanism is retained in metadata order. Three-tag and five-tag tasks have separate denominators.

Mechanism	Channels	Layers	Speech 5	Song 5	Speech 3	Song 3
DMMS	10	E+M+P+F	18/45	18/45	16/27	13/27
CSSL	10	E+M+P+F	30/45	30/45	20/27	26/27
CDEM	10	C+M+P+F	23/45	21/45	22/27	15/27
SRP	19	N+C+T+M+P+F	19/45	18/45	9/27	13/27
AAC	14	E+A+I+P+F	16/45	20/45	11/27	18/27
VMM	12	V+R+P+C	29/45	15/45	22/27	15/27
PUPF	12	E+U+G+P+F	20/45	14/45	15/27	11/27
CLAM	11	E+B+C+P+F	23/45	20/45	16/27	14/27
MAD	11	E+D+A+P+F	25/45	26/45	22/27	15/27
NEMAC	11	E+M+W+P+F	18/45	15/45	17/27	15/27
STAI	12	E+M+P+F	17/45	21/45	15/27	12/27
DAP	10	E+D+P+F	23/45	21/45	18/27	15/27
CMAT	10	E+S+T+P+F	27/45	21/45	22/27	15/27
TAR	10	E+TI+P+F	21/45	17/45	9/27	12/27
DAED	8	E+M+P	17/45	18/45	15/27	9/27
MORMR	7	E+M+P+F	17/45	17/45	6/27	14/27
RPEM	8	E+M+P	16/45	18/45	16/27	10/27
IUCP	6	E+P+F	21/45	23/45	13/27	15/27
MCCN	7	E+P+F	18/45	19/45	16/27	12/27
MEAMR	6	E+P+F	11/45	10/45	11/27	11/27
SSRI	11	E+M+P+F	22/45	27/45	17/27	19/27
LDAC	6	E+P+F	23/45	22/45	15/27	13/27
IOTMS	5	E+P	11/45	18/45	10/27	14/27
SSPS	6	E+P+F	20/45	18/45	16/27	12/27
PEOM	11	E+M+P+F	20/45	25/45	18/27	16/27
MSR	11	E+M+P+F	25/45	18/45	17/27	16/27
GSSM	11	E+M+P+F	21/45	23/45	20/27	17/27
HMCE	11	E+M+P+F	19/45	19/45	15/27	11/27
HGSIC	11	E+M+P+F	18/45	22/45	12/27	10/27
ASAP	11	E+M+P+F	20/45	21/45	15/27	13/27
DDSMI	11	E+M+P+F	20/45	19/45	14/27	17/27
VRMSME	11	E+M+P+F	22/45	23/45	21/27	11/27
SPMC	11	E+M+P+F	16/45	22/45	13/27	14/27
NSCP	11	E+M+P+F	24/45	22/45	19/27	15/27
CTBB	11	E+M+P+F	20/45	16/45	15/27	11/27
STC	11	E+M+P+F	20/45	25/45	17/27	14/27
EDNR	10	E+M+P+F	22/45	18/45	18/27	12/27
TSCP	10	E+M+P+F	25/45	17/45	18/27	9/27
CDMR	11	E+M+P+F	20/45	16/45	20/27	9/27
SLEE	13	E+M+P+F	19/45	9/45	8/27	9/27
ESME	11	E+M+P+F	19/45	16/45	15/27	9/27
ECT	12	E+M+P+F	21/45	18/45	17/27	13/27

Supplementary Fig. S2a | Class-specific patterns / ridge

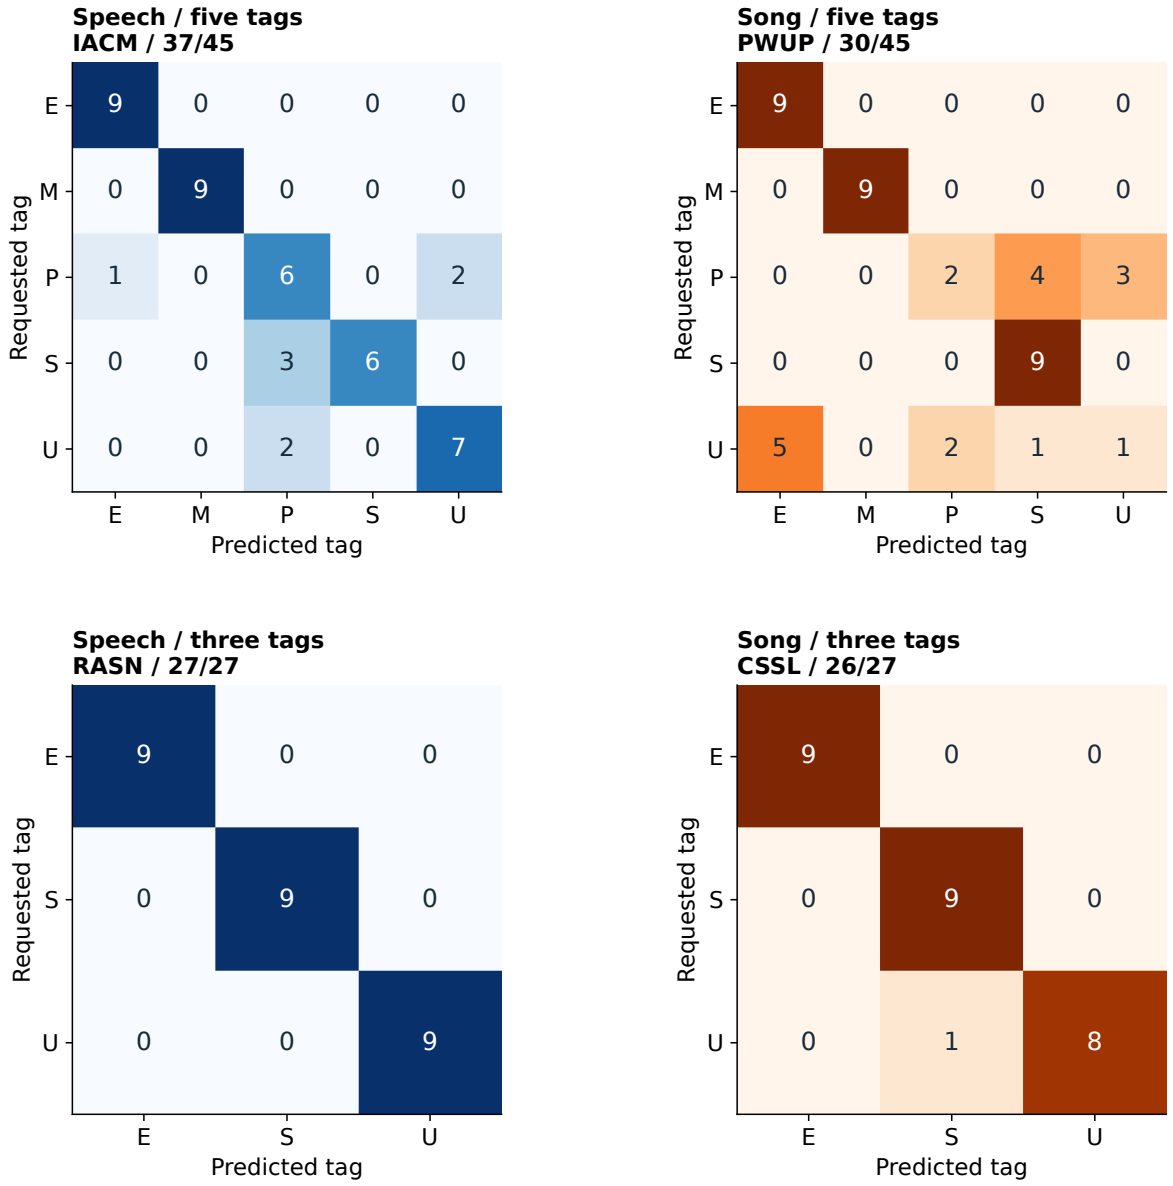

Whole-mechanism ridge; rows are requested tags and columns predicted tags. E excited, M mad, P peaceful, S sad, U untagged. Every class contains nine requests. The complete confusion records are supplied as data.

Supplementary Fig. S2b | Class-specific patterns / centroid

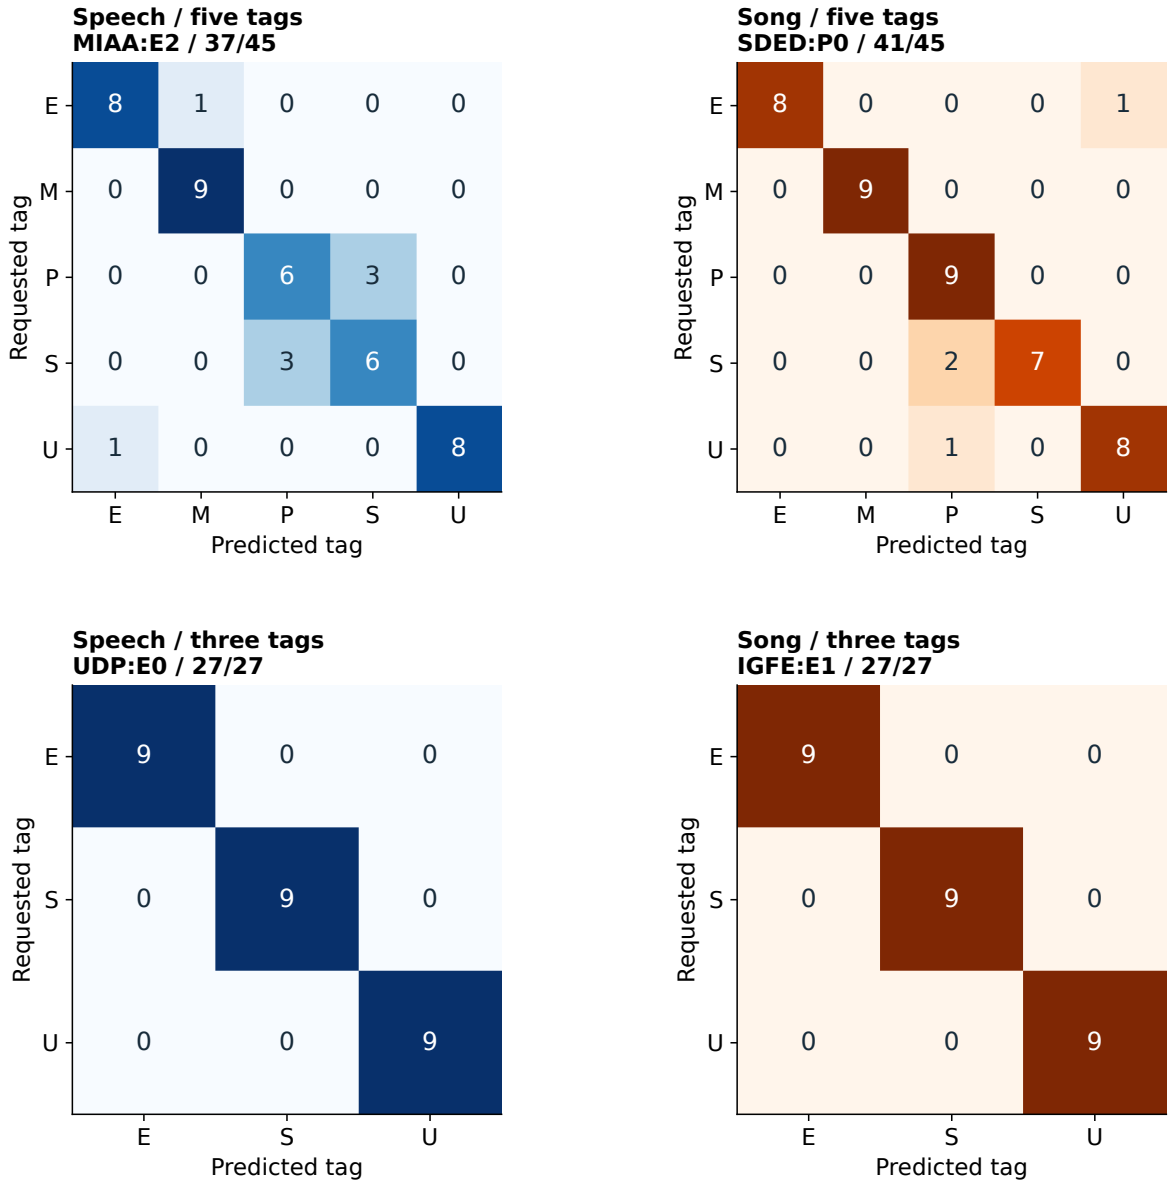

Single-channel nearest centroid, using the same held-out texts as Fig. S2a. E excited, M mad, P peaceful, S sad, U untagged. Every class contains nine requests. The reader is explicitly identified as the adaptive control.

Supplementary Fig. S3 | Standalone layers within six mechanisms

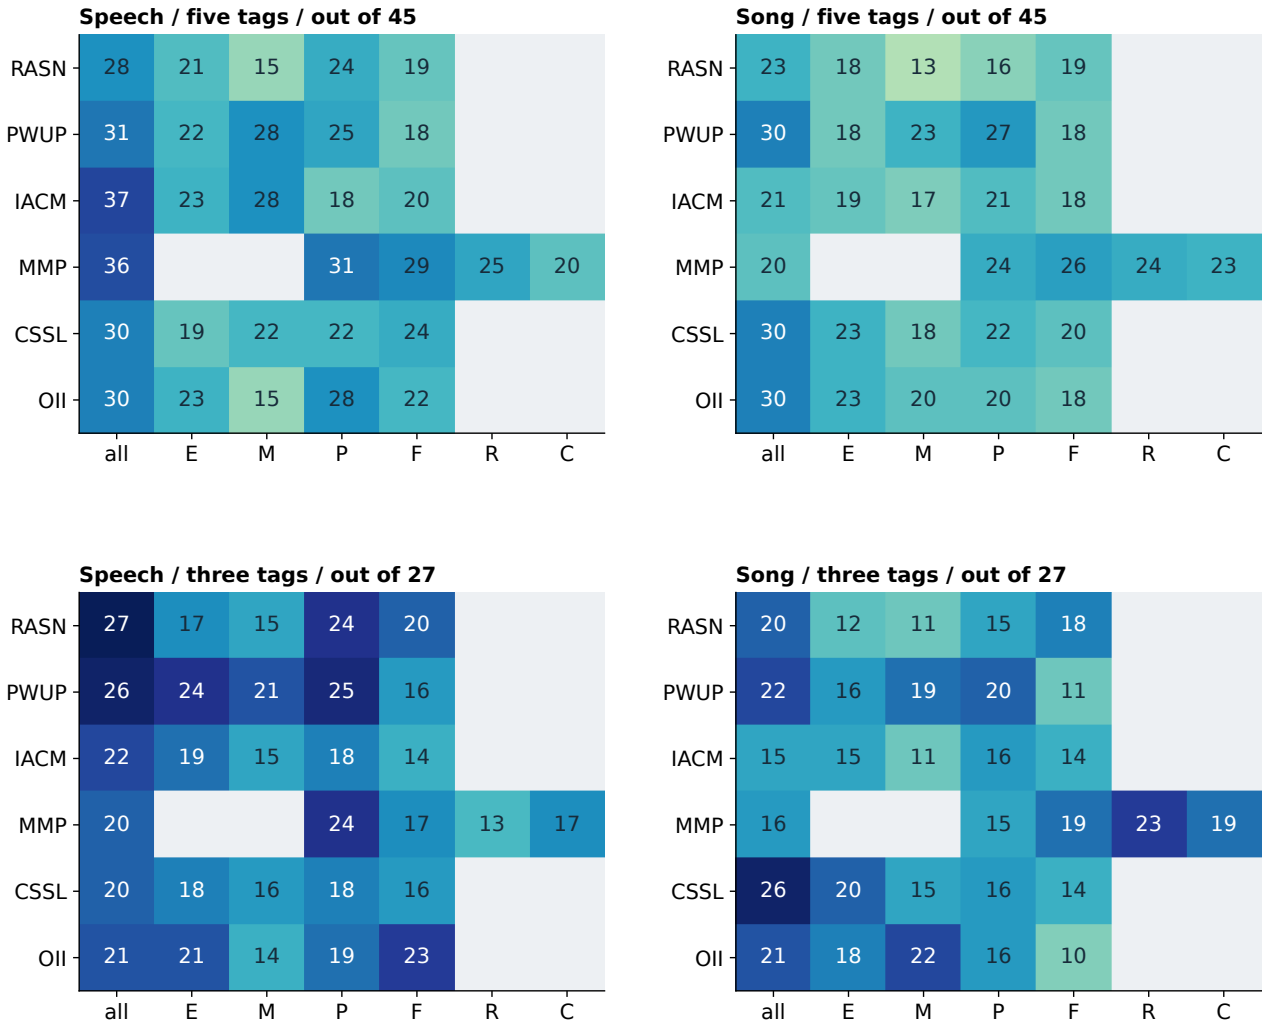

Layer codes follow each mechanism's own schema. The complete mechanism and every existing layer are shown separately. An absent layer is not assigned a zero score. These examples follow the full 326-layer inventory; all standalone layer scores remain available.

Supplementary Fig. S4 | Layer exclusion under three readout settings

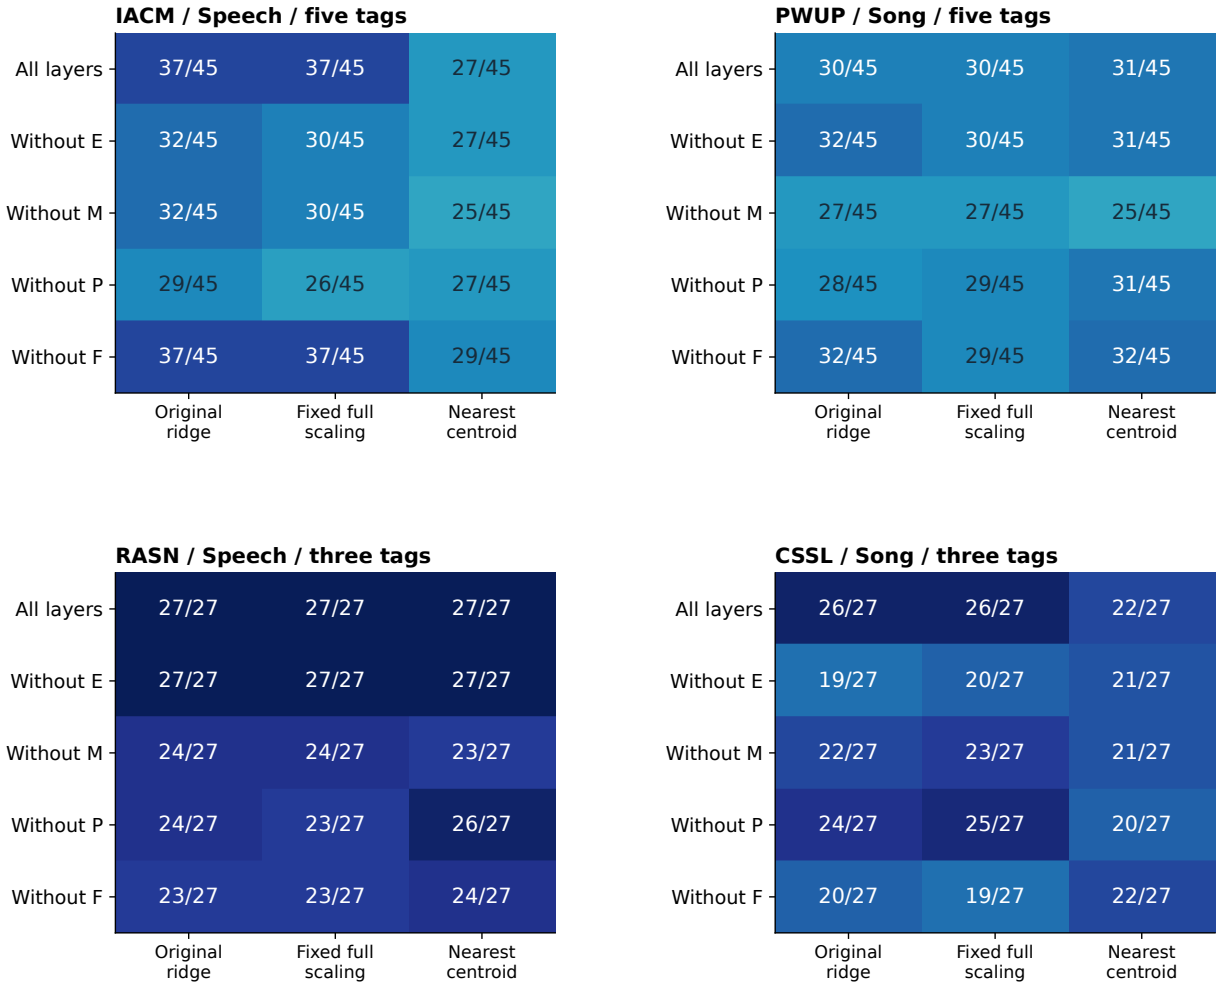

Original ridge, fixed-full-scaling ridge and nearest centroid are kept distinct. Each cell is correct requests / total requests. Exclusions remove stored output channels, without rerunning internal stages. The full all-layer-subsets.csv contains 1,152 rows.

Supplementary Fig. S5 | Mechanism, layer and channel resolution

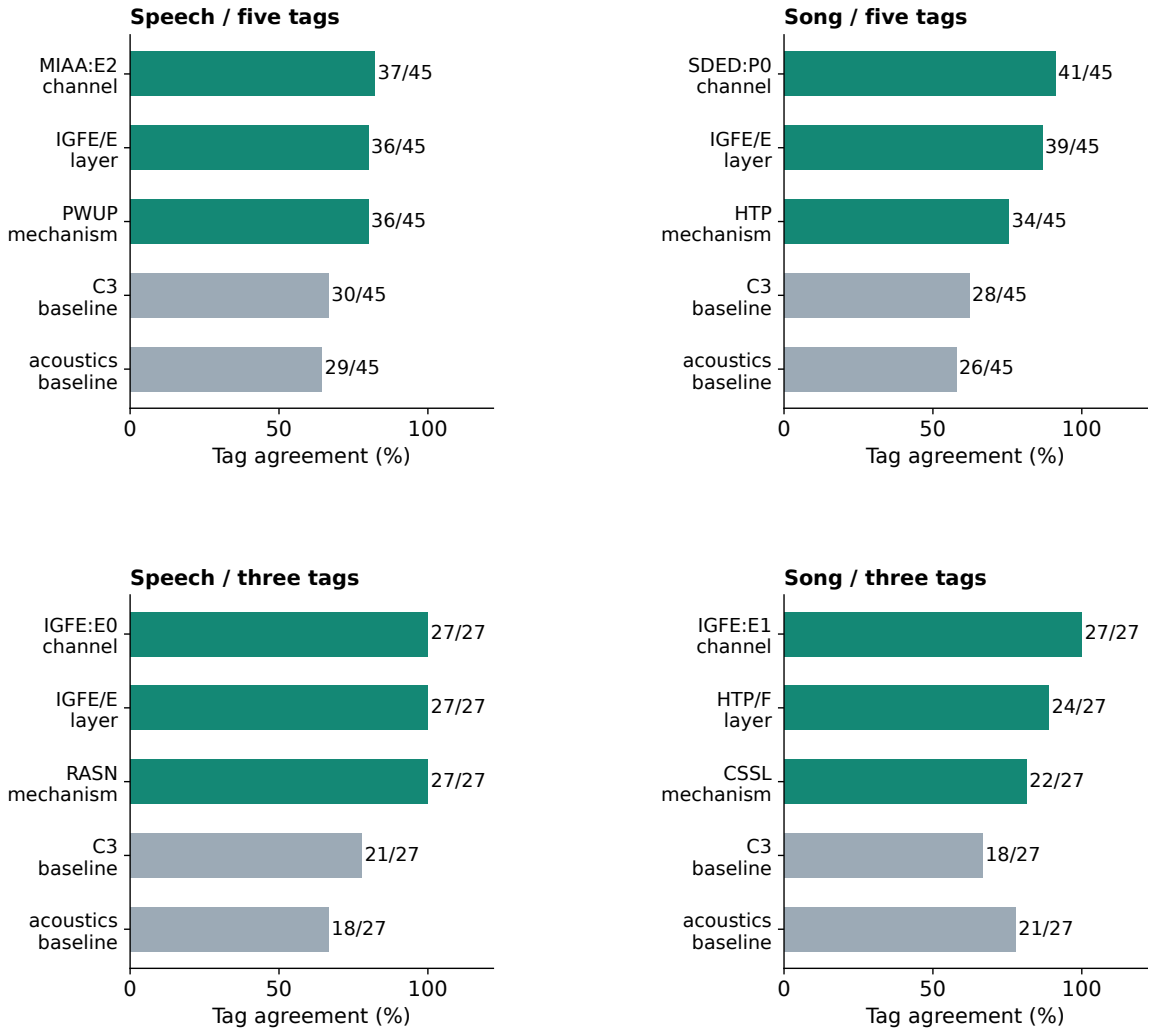

All 1,299 candidates enter the centroid measurement, plus three comparators. The illustrated leaders were identified from the full candidate inventory. A strong single-channel result is evidence of accessible structure in that channel, not evidence that the entire hierarchy is necessary for that task.

Supplementary Fig. S6 | Automatic choice is a separate secondary task

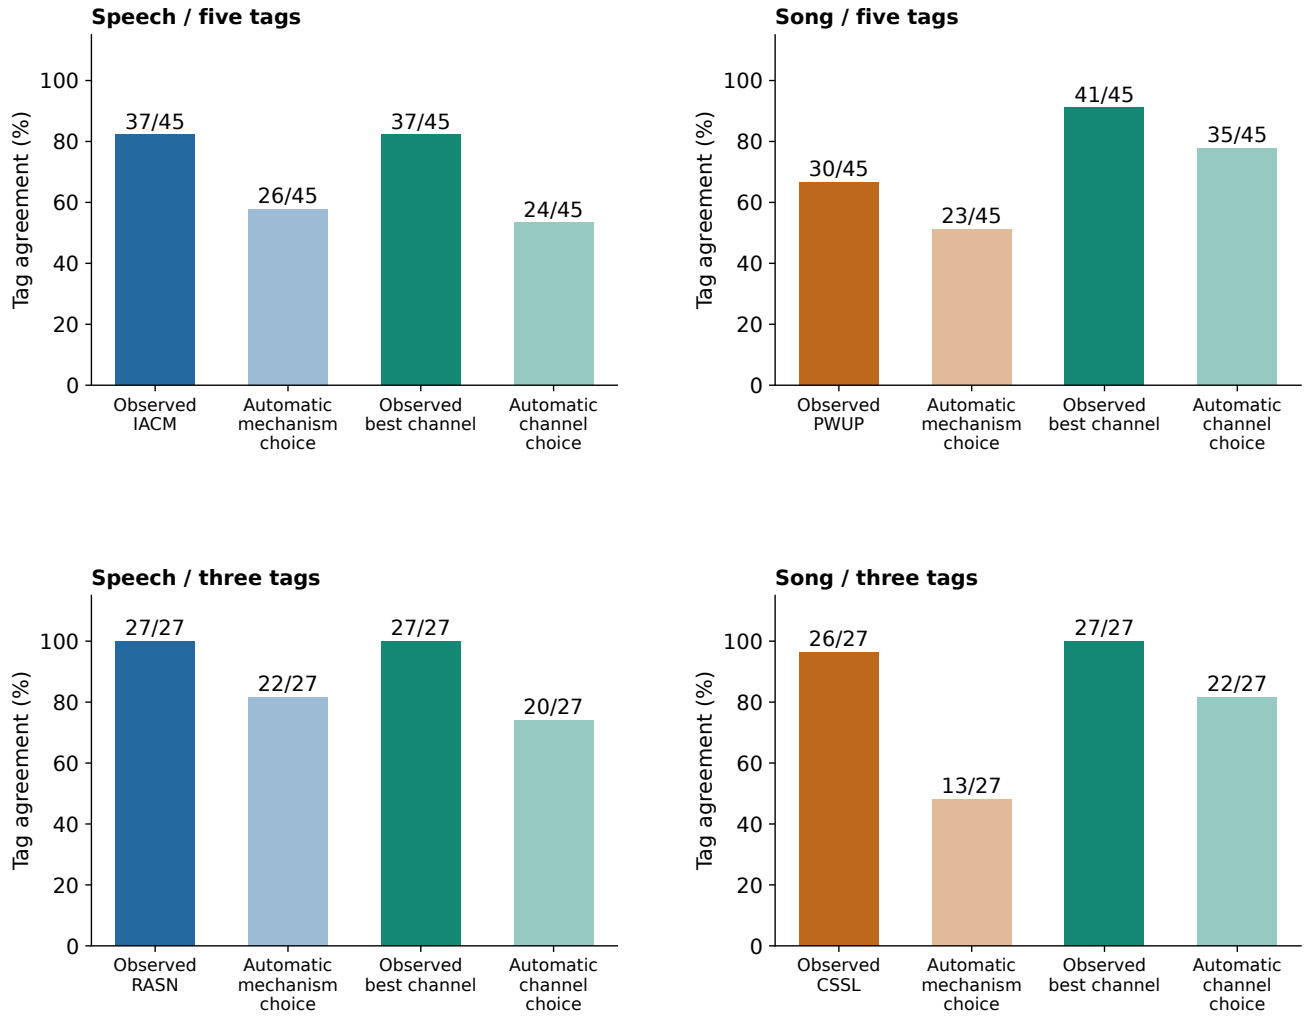

Observed named representations are compared with an explicit rule that chooses a representation using only the two outer-estimation texts. Full decisions and ties are retained. This supplemental comparison does not redefine the objective of the comprehensive engine study. See Supplementary Note S3.

Supplementary Note S4 | Controls and evidence files

Permutation family and conditional null

The ridge maximum-statistic family contains all 84 whole mechanisms. The centroid family contains all 1,299 mechanism/layer/channel candidates, including constant channels. There are 999 label permutations in each family and task, restricted by text, delivery mode and generation cohort. Every shuffled mapping is recomputed. For each of the four tasks, both readers have a task-maximum P value of 0.001 and a within-reader four-task BY q of 0.00208333. The null preserves the generation strata. It cannot remove cohort confounding or validate all historical analyst choices. A whole-family information test is not a paired superiority test between two candidates.

Summary and regularization sensitivity

All 84 mechanisms were evaluated with channel means, five distributional summaries, seven temporal descriptors and ten normalized-time blocks, and with $\alpha = 0.1$, 1 and 10 in the mean-feature ridge sensitivity. The 2,016 result rows retain all settings. IACM five-tag speech gives 39/45, 37/45 and 18/45 at $\alpha = 0.1$, 1 and 10. CSSL three-tag song gives 21/27, 26/27 and 21/27. The original $\alpha = 1$ defines the main ridge measurement. A better sensitivity score is not selected to replace it.

Other original-waveform comparisons

The 15 archived tasks retain their original predictions and class definitions. Six-tag speech gives IACM 40/54 and MMP 39/54. Cross-mode, cross-text transfer reaches 26/45 in either direction with the leading mechanisms stated in the main article. Same-text 21-tag, seven-delivery and ten-request fresh-bridge panels have different evaluation targets. The 20 separate binary tag tasks generate 1,680 mechanism comparisons. These exploratory panels are included in complete tables rather than presented as independently corrected discoveries.

Raw values and temporal organization

The channel-value diagnostics preserve minimum, maximum, main-core variation, within-request alternative variation, condition repeatability and an orthogonal variance partition across text, mode and tag. These are descriptive properties of recorded model outputs. The 26,670-row channel-by-text-by-mode-by-tag table preserves separate values for every channel. No variance fraction is interpreted as causal attribution, and no named model variable is treated as a measured human cognitive or neural quantity.

Verification and reproducibility boundary

All 342 audio hashes and all 1,590 archived prediction matrices were checked. The new numerical validation compares 1,648 ridge and 5,208 centroid confusion matrices with their recorded totals and diagonals, verifies 72 full-subset identities, reproduces 3,556 independently implemented scalar-centroid counts, and checks 84 scikit-learn comparisons. The original full-array and engine archives are identified by paths and hashes. Rebuilding manuscript figures requires those existing arrays; reproducing the supplied result tables does not require new audio generation.