

A continuous transcriptomic entropy score is associated with overall survival in WHO grade 3 IDH-mutant glioma: development and external validation across two sequencing platforms

Ivy Datta

dattaivyz@gmail.com

South Piedmont Community College <https://orcid.org/0009-0004-7634-3170>

Research Article

Keywords: IDH-mutant glioma, prognosis, transcriptomics, Shannon entropy, external validation

Posted Date: October 5th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11226034/v1>

License:  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Abstract

Purpose

Decisions in IDH-mutant lower-grade glioma rely on grade and MGMT promoter methylation, yet patients concordant for both show survival from under two years to over a decade. Existing transcriptomic methods assign discrete subtypes and give no per-sample score. We developed and validated a continuous one.

Methods

The Transcriptomic Ambiguity Index (TAMI) is the Shannon entropy of the posterior from a classifier separating glioblastoma from lower-grade glioma. A 200-gene panel was selected in The Cancer Genome Atlas (TCGA); values were rank-normalized within batch, scored, converted to entropy and standardized. IDH status came from PanCanAtlas molecular subtype, not histology. Validation used two Chinese Glioma Genome Atlas (CGGA) cohorts on separate platforms. Overall survival was analyzed by Cox regression with Benjamini-Hochberg false discovery rate (FDR) control at 0.05.

Results

In WHO grade 3 IDH-mutant glioma the hazard ratio per standard deviation (SD) of TAMI was 1.61 (1.25–2.06) in TCGA, 1.48 (1.23–1.79) in CGGA-693 and 1.93 (1.30–2.87) in CGGA-325, all surviving FDR correction. The association was stronger in molecularly confirmed than in histologically defined cases, so IDH-wildtype contamination does not explain it. In grade 2 disease it was null in two of three cohorts. Estimates were essentially unchanged after adjustment for 1p/19q codeletion, and TAMI remained independent in models containing codeletion, MGMT status, grade, age and treatment.

Conclusion

Classifier entropy carries prognostic information in grade 3 IDH-mutant glioma beyond grade, codeletion, MGMT status and treatment, and replicates across two platforms. It is not specific to grade 3 and requires standardization against the target cohort.

Introduction

The 2021 World Health Organization classification organizes adult-type diffuse gliomas carrying an isocitrate dehydrogenase (IDH) mutation into astrocytoma, IDH-mutant, and oligodendroglioma, IDH-mutant and 1p/19q-codeleted, each assigned a grade of 2, 3, or 4 [1]. Within that framework the decisions that most shape a patient's course after resection, namely how often to image and whether to give radiotherapy and alkylating chemotherapy immediately or defer them [2, 3], rest primarily on

histologic grade and, where measured, on O6-methylguanine-DNA methyltransferase (MGMT) promoter methylation status [1, 4]. Both are coarse. Grade is an ordinal summary of morphology, and MGMT methylation was established as predictive in glioblastoma rather than in lower-grade IDH-mutant disease [4].

The consequence is a wide spread in outcome among patients who look identical on paper: in TCGA, patients concordant for IDH mutation, grade and MGMT status show overall survival ranging from under two years to more than a decade [2]. A patient at the favorable end may be exposed to the late neurocognitive and vascular toxicity of radiotherapy for no gain; one at the unfavorable end may be observed too long to catch transformation while it remains tractable. Neither error is visible when the decision is made.

Existing transcriptomic frameworks for glioma are taxonomic by design: they place a tumor into one of a small number of molecular subtypes, and those subtypes differ in aggregate survival [5]. Subtype assignment describes a population well but is a poor instrument for an individual decision: it returns a label rather than a magnitude, and gives no ordering within a subtype. What is missing is a per-sample continuous score.

Such a score can be extracted from a quantity most classification pipelines discard: the confidence of the classifier itself. A classifier trained to separate two biologically distant reference states, glioblastoma and lower-grade glioma, assigns most lower-grade tumors to that class with a posterior close to one. A minority are not: their profiles sit between the reference states and the posterior falls near the decision boundary. The design premise is that such ambiguity is signal rather than noise, and that a profile resisting confident assignment to the lower-grade state has begun to acquire high-grade expression features histology has not yet registered. We formalize this as the Transcriptomic Ambiguity Index (TAMI), the Shannon entropy of the classifier posterior, which is continuous, bounded, defined per sample, monotone in classification uncertainty, and independent of which class is favored. This paper specifies the score, evaluates it in a development cohort, and tests it in two external cohorts sequenced on different platforms.

Methods

Cohorts

This was a retrospective analysis of three public glioma cohorts with paired transcriptomic and clinical data. Expression, phenotype and survival data for the TCGA GBM and lower-grade glioma series came from the UCSC Xena GBMLGG cohort [6], yielding 693 samples. Of these, 528 carried a lower-grade histology (astrocytoma, oligoastrocytoma or oligodendroglioma) and the rest glioblastoma.

IDH status was not inferred from histology. Molecular subtype calls came from the TCGA PanCanAtlas classification via the cBioPortal API [7], matched by patient identifier. Among the 528 histologic lower-grade gliomas, 256 were classified IDH-mutant non-codeleted, 169 IDH-mutant 1p/19q-codeleted, 94

IDH-wildtype, and 9 were unannotated. The development cohort was therefore the 425 confirmed IDH-mutant cases; the 94 IDH-wildtype cases were excluded from primary analyses and retained only for the sensitivity comparison. This matters because IDH-wildtype lower-grade glioma behaves like glioblastoma [2], and a histologically defined cohort would have carried roughly 18% such cases.

External validation used the Chinese Glioma Genome Atlas (CGGA) mRNAseq_693 and mRNAseq_325 cohorts [8], restricted to cases recorded as IDH-mutant with WHO grade II or III (n = 306 and n = 133). Neither contributed to panel selection, classifier training or standardization. Grades are reported as WHO grade 2 and 3; both sources were graded under pre-2021 criteria and not regraded. Baseline characteristics are in Table 1. The cohorts differ materially in event rate: 24% of grade 3 cases had died in the development cohort against 56% in CGGA-693 and 60% in CGGA-325, which affects comparability across cohorts but not the validity of any within-cohort estimate.

Table 1
Characteristics of the three analysis cohorts.

IDH-mutant cases, n	Development (TCGA)	CGGA-693	CGGA-325
	425	306	133
WHO grade 2, n	232	122	85
WHO grade 3, n	193	166	42
Deaths, grade 2, n (%)	33 (14.2)	33 (27.0)	33 (38.8)
Deaths, grade 3, n (%)	46 (23.8)	93 (56.0)	25 (59.5)
Age at diagnosis, median (IQR)	39 (32–49)	40 (34–46)	38 (34–43)
Female, n (%)	190 (44.7)	128 (41.8)	52 (39.1)
1p/19q codeleted, n (%)	169 (39.8)	109 (35.6)	57 (42.9)
MGMT promoter methylated, n (%) of tested	not available	154/229 (67.2)	73/120 (60.8)
Radiotherapy received, n (%) of recorded	not available	224/289 (77.5)	105/128 (82.0)
Temozolomide received, n (%) of recorded	not available	195/289 (67.5)	59/121 (48.8)
Median follow-up, months	28.4	75.6	130.9

Grade 2 and grade 3 counts are the cases entering the survival analyses and do not sum to the cohort total, because 18 CGGA-693 and 6 CGGA-325 cases lacked grade or survival data. Development-cohort IDH status is the TCGA PanCanAtlas molecular subtype call; CGGA IDH status is as recorded in the source clinical file. MGMT promoter methylation, 1p/19q codeletion and treatment percentages are of

cases with that field recorded, with the denominator shown; MGMT methylation and treatment are not recorded in the TCGA data release used here. Median follow-up is by reverse Kaplan-Meier. IQR, interquartile range.

All data are public and de-identified at source.

Score construction

Analysis used the 15,724 genes measured in all three cohorts. Within every batch submitted for scoring, values were converted to within-sample ranks across those genes and scaled to the unit interval. Rank normalization was chosen over cross-cohort quantile normalization or batch correction because a score intended for a single specimen cannot assume access to a contemporaneous cohort; every preprocessing step is computable from the submitted batch.

A single pipeline was used for all three cohorts. Within the development cohort, a 200-gene panel was selected by univariate F-test for differential expression between glioblastoma and lower-grade glioma and passed to a regularized logistic regression classifier with the penalty chosen by internal cross-validation. Feature selection and classifier fitting were both performed inside each fold of a five-fold stratified cross-validation, so that TCGA samples were scored only by models that had not seen them; the panel and weights were then refit on the full development cohort and applied unchanged to both CGGA cohorts. Out-of-fold discrimination between glioblastoma and lower-grade glioma was 0.9999 by area under the receiver operating characteristic curve. This near-perfect separation reflects the transcriptional distance between the two reference states rather than information leakage, since selection was nested within folds. Panel membership is listed in Online Resource 1.

For a tumor with classifier posterior p of belonging to the lower-grade class, TAMI is the binary Shannon entropy of that posterior,

$$TAMI = -p \log_2 p - (1 - p) \log_2 (1 - p),$$

which is 0 when assignment is certain and 1 bit when the posterior is indeterminate. Raw values were standardized to zero mean and unit variance against the distribution of the cohort being scored, and hazard ratios are reported per standard deviation of the standardized score. The pipeline is shown in Fig. 1. The classifier itself is frozen; only the standardization is recomputed, because entropy scale proved to be platform-dependent even after rank normalization of the inputs.

Statistical analysis

The primary endpoint was overall survival, from the start of follow-up as recorded in each source repository to death from any cause, censored at last follow-up. Association with TAMI was estimated by Cox proportional-hazards regression adjusted for age, fit separately within each grade since grade is constant within stratum. Proportional hazards were assessed by scaled Schoenfeld residuals; age

violated proportionality in CGGA-325, which was also fit with age-stratified baselines. Analyses were prespecified to be stratified by grade because the clinical question, detection of malignant progression not yet apparent histologically, applies to grade 3 in a way it does not to grade 2. Multiplicity was controlled by the Benjamini-Hochberg procedure at a false discovery rate of 0.05 over a prespecified family of six tests: three cohorts by two grades on the same continuous specification. Prespecified sensitivity analyses added 1p/19q codeletion to each grade 3 model. Incremental value was assessed in the CGGA cohorts, which annotate MGMT status, codeletion and treatment per case, by comparing Harrell's concordance index and by likelihood-ratio test between a fully adjusted model containing age, grade, codeletion, MGMT status, radiotherapy and temozolomide, and that model with TAMI added. Missing covariates were handled by complete-case analysis. P values are two-sided. Analyses used Python with lifelines, scikit-learn, statsmodels and pandas; versions are pinned in the archived code. A large language model assisted with writing part of the analysis code; details are in the Declarations. Reporting follows the TRIPOD + AI statement [9].

Results

In the development cohort restricted to molecularly confirmed IDH-mutant WHO grade 3 glioma (n = 193, 46 deaths), TAMI was associated with overall survival, with a hazard ratio of 1.61 per standard deviation (SD) of the score (95% confidence interval 1.25 to 2.06, P = 0.0002).

Predicted survival is shown in Fig. 2.

Applied without modification to the two CGGA cohorts, the score replicated: hazard ratio 1.48 per 1-SD (1.23 to 1.79, P < 0.001) in CGGA-693 grade 3 (n = 166, 93 deaths) and 1.93 per 1-SD (1.30 to 2.87, P = 0.001) in CGGA-325 grade 3 (n = 42, 25 deaths). All three associations survived Benjamini-Hochberg correction (q = 0.0006, 0.0002 and 0.002). In CGGA-325, where age violated proportional hazards, the age-stratified model gave an essentially unchanged estimate of 1.92 per 1-SD (1.27 to 2.90, P = 0.002). Estimates across cohorts are shown in Fig. 3, and all models are tabulated in Table 2.

Table 2
Cox proportional-hazards estimates for TAMI. Panel A: association with overall survival by grade and cohort, per 1 standard deviation of the standardized score, adjusted for age.

Model	n	Deaths	HR (95% CI)	P	FDR q
WHO grade 3					
TCGA, confirmed IDH-mutant	193	46	1.61 (1.25–2.06)	0.0002	0.0006
CGGA-693	166	93	1.48 (1.23–1.79)	<0.001	0.0002
CGGA-325	42	25	1.93 (1.30–2.87)	0.001	0.002
CGGA-325, age-stratified	42	25	1.92 (1.27–2.90)	0.002	n/a
TCGA, histologic cohort	270	93	1.49 (1.26–1.76)	<0.001	n/a
WHO grade 2					
TCGA, confirmed IDH-mutant	232	33	1.12 (0.96–1.31)	0.16	0.16
CGGA-693	122	33	1.51 (1.14–2.01)	0.004	0.007
CGGA-325	85	33	1.23 (0.92–1.63)	0.16	0.16

The false discovery rate family was prespecified as the six primary tests, three cohorts by two grades, all on the same continuous specification; the age-stratified and histologic-cohort rows are secondary and were not included in it. The histologic-cohort row retains IDH-wildtype cases and is shown for comparison only. CI, confidence interval; HR, hazard ratio.

Panel B: fully adjusted models in each validation cohort, containing TAMI, age, WHO grade, 1p/19q codeletion status, MGMT promoter methylation status, radiotherapy and temozolomide receipt.

Covariate	HR (95% CI)	P
CGGA-693 (n = 197, 93 deaths)		
TAMl, per 1-SD	1.65 (1.36–2.01)	<0.001
WHO grade 3	2.90 (1.72–4.90)	<0.001
1p/19q codeleted	0.46 (0.28–0.76)	0.002
MGMT methylated	0.88 (0.55–1.41)	0.59
Temozolomide	0.68 (0.42–1.10)	0.12
Radiotherapy	0.91 (0.52–1.58)	0.74
Age, per year	1.01 (0.98–1.03)	0.58
Concordance	0.756	
CGGA-325 (n = 106, 48 deaths)		
TAMl, per 1-SD	1.90 (1.32–2.74)	<0.001
WHO grade 3	1.44 (0.68–3.06)	0.34
1p/19q codeleted	0.18 (0.08–0.38)	<0.001
MGMT methylated	0.73 (0.38–1.39)	0.33
Temozolomide	1.05 (0.56–1.97)	0.88
Radiotherapy	1.19 (0.56–2.54)	0.65
Age, per year	1.02 (0.98–1.06)	0.36
Concordance	0.786	

Models pool WHO grade 2 and grade 3 cases with complete data on every covariate and include grade as a term. Treatment terms are included to control confounding and should not be read as estimates of treatment efficacy, since allocation was by indication. Concordance without TAMl was 0.700 in CGGA-693 and 0.770 in CGGA-325; adding TAMl raised it to 0.756 and 0.786 (likelihood-ratio chi-square 21.5, $P < 0.001$ and 11.0, $P < 0.001$). CI, confidence interval; HR, hazard ratio.

The CGGA-325 point estimate exceeds the others, but we do not read this as better performance in that cohort: it has 42 grade 3 cases, its interval is correspondingly wide, and effect estimates conditional on

significance in small samples are biased upward [10]. The defensible reading is that direction and approximate magnitude replicate across platforms.

The association is not explained by IDH-wildtype misclassification

Because IDH-wildtype lower-grade glioma carries a glioblastoma-like course [2], a score trained to separate glioblastoma from lower-grade glioma could appear prognostic simply by identifying IDH-wildtype tumors present in a histologically defined cohort. We tested this directly. In the 270 histologically defined grade 3 cases, which include IDH-wildtype tumors, the hazard ratio was 1.49 per 1-SD (1.26 to 1.76, $P < 0.001$). Restricting to molecularly confirmed IDH-mutant cases, which removes those tumors, raised rather than lowered the estimate, to 1.61. The association therefore holds within genuinely IDH-mutant disease and is not an artifact of cohort contamination.

Behavior in grade 2 disease

In grade 2 tumors the picture is mixed. TAMI showed no association with overall survival in the development cohort (1.12 per 1-SD, 0.96 to 1.31, $P = 0.16$) or in CGGA-325 (1.23, 0.92 to 1.63, $P = 0.16$), but was clearly associated in CGGA-693 (1.51, 1.14 to 2.01, $P = 0.004$, surviving FDR correction at $q = 0.007$). We therefore do not claim that the score is specific to grade 3 disease. The effect is strongest and most consistent in grade 3, is absent in grade 2 in two of three cohorts, and is present in grade 2 in one. Whether the CGGA-693 grade 2 result reflects genuine prognostic signal in a cohort with a higher event rate, or residual grade misassignment under pre-2021 criteria, cannot be resolved with these data.

Adjustment for 1p/19q codeletion

IDH-mutant astrocytoma and 1p/19q-codeleted oligodendroglioma differ in natural history, so a score that merely separated the two would appear prognostic in a mixed grade 3 group. Adding codeletion status to each grade 3 model left the estimates essentially unchanged: 1.59 per 1-SD (1.23 to 2.04) in the development cohort, 1.46 (1.20 to 1.78) in CGGA-693 and 1.72 (1.11 to 2.66) in CGGA-325, all $P < 0.05$. Codeletion itself carried a substantial survival advantage, as expected, supporting the model specification.

Incremental value over established markers and treatment

Incremental value was assessed in the CGGA cohorts, which annotate MGMT promoter methylation status, 1p/19q codeletion status, and radiotherapy and temozolomide receipt per case. In fully adjusted models containing age, grade, codeletion status, MGMT status, radiotherapy and temozolomide, TAMI retained a strong independent association in both cohorts: hazard ratio 1.65 per 1-SD (1.36 to 2.01, $P < 0.001$) in CGGA-693 ($n = 197$, 93 deaths) and 1.90 per 1-SD (1.32 to 2.74, $P < 0.001$) in CGGA-325 ($n = 106$, 48 deaths). Adding TAMI raised Harrell's concordance from 0.700 to 0.756 in CGGA-693 and from 0.770 to 0.786 in CGGA-325, and the likelihood-ratio test was significant in both (chi-square 21.5 and

11.0, both $P < 0.001$). The gain is substantial in one cohort and small in the other, where the covariate model was already strong; what replicates is independence rather than a large gain. Table 2 reports the full models.

TAMI was the strongest predictor after grade in CGGA-693 and the strongest in CGGA-325, while MGMT status, radiotherapy and temozolomide were not independently associated with survival in either. The treatment terms are confounding controls, not evidence about efficacy, since allocation was by indication; the point is that the TAMI estimate is not attributable to what patients received.

One qualification is that the CGGA cohorts are the same ones used to validate the age-adjusted association, so the incremental estimate is external to the development data but not independent of the validation data.

Standardization does not transfer across platforms

Applying the development cohort's standardization parameters directly to the CGGA cohorts placed their entire score distributions at the extreme of the development scale. The failure was not one of ranking, since within each cohort the ordering of samples by TAMI remained prognostic, but one of location and scale: entropy sits at a different level on a different sequencing platform even though expression had been rank-normalized within batch before scoring. Rank normalization removes the effect of platform on each gene's marginal distribution, but the posterior depends on the joint configuration of the panel genes, and how closely that configuration resembles the training data varies with platform; dissimilarity pushes posteriors toward the decision boundary and raises entropy across the board. Recomputing mean and standard deviation from the target cohort's own distribution, which is what all results above use, restores a usable score. The operational consequence is that TAMI transfers as a relative score requiring a reference set processed on the target platform, and cannot be issued with a universal absolute value.

Discussion

A continuous score derived from classifier entropy was associated with overall survival in WHO grade 3 IDH-mutant glioma in a development cohort and in two independent cohorts sequenced on separate platforms, survived correction for multiple testing, strengthened rather than weakened when the cohort was restricted to molecularly confirmed IDH-mutant disease, and added discrimination to a model containing codeletion status, MGMT status, grade, age and treatment received in both external cohorts.

Why classifier uncertainty should carry prognostic information has a parsimonious explanation: it is a proxy for transcriptional position along the trajectory toward high-grade disease. Malignant transformation in IDH-mutant glioma is progressive, and the expression program of a tumor on that trajectory should resemble high-grade disease before the morphologic features defining high grade are established and sampled. A classifier trained on the endpoints of that trajectory returns intermediate

posteriors for tumors between them, and entropy is the natural summary of intermediacy. Two observations fit this account. Preservation of ranking under platform change is the first: absolute entropy levels shifted, but within-cohort ordering stayed prognostic, as expected if the underlying quantity is a biological position that platform noise displaces rather than scrambles. The second is that the effect strengthened when IDH-wildtype tumors were removed, which is the opposite of what a contamination-driven artifact would do.

It is worth being explicit about what this score is not. It is not a subtype assignment and does not compete with the classifications that organize glioma taxonomically [5]; it sits downstream of them as a continuous within-class score. It is not a substitute for grade or MGMT status, since the fully adjusted models show it adds to those markers rather than replacing them. And it is not a thresholded assay. We examined tertile risk bands in the development cohort and found that the low and intermediate bands were indistinguishable, with the signal concentrated in the upper tail; no banding scheme was prespecified, and we report the continuous score rather than select a cutpoint after the fact.

The constraints on this work are specific. All data are retrospective and drawn from cohorts assembled for research rather than under a prespecified surveillance protocol, with the attendant risks of selection and informative censoring; treatment was allocated by indication rather than at random, so the fully adjusted models control for confounding by treatment without supporting any causal reading of the treatment terms. Median follow-up was 28 months in the development cohort against 76 and 131 months in the validation cohorts, which explains much of the difference in event rate and means the development estimate rests on shorter observation. Molecular subtype was unavailable for 9 development-cohort cases, and the subgroup with subtype annotation is not a random sample of the whole, so the confirmed cohort differs from the histologic one in more than IDH status. The two validation cohorts share a population and an assembling institution, so replication across them is strong evidence against overfitting and against platform-specific artifact but weak evidence about generalizability to other populations. CGGA-325 contributes only 42 grade 3 cases. Grade was taken as recorded under pre-2021 criteria without regrading, so grade misassignment relative to current criteria cannot be excluded, although adjustment for 1p/19q codeletion left the estimates unchanged. Deployment requires a reference set processed on the target platform. And there is no prospective validation and no evaluation against a decision a clinician actually makes; association with overall survival is necessary but not sufficient for clinical utility.

Three lines of work follow. The first is prospective evaluation in grade 3 IDH-mutant glioma under a prespecified surveillance protocol, with time to radiographic or histologic progression as the endpoint rather than overall survival alone. The second is a decision-analytic assessment of net benefit across the threshold probabilities a clinician would actually use, which is the appropriate way to ask whether a concordance gain of this size changes management. The third is methodological: the standardization requirement invites work on reference-free scaling, for instance by anchoring the entropy distribution to control profiles sequenced alongside each batch. Nothing in the construction of TAMI is glioma-specific,

requiring only two reference states separated by a progression trajectory, so whether the approach generalizes to other tumor types is open and testable.

Conclusion

A continuous transcriptomic entropy score is associated with overall survival in WHO grade 3 IDH-mutant glioma, replicates across two independent cohorts and sequencing platforms, holds within molecularly confirmed IDH-mutant disease, and adds discrimination beyond 1p/19q codeletion status, MGMT promoter methylation status, grade, age and treatment received. It is not specific to grade 3 disease, all evidence is retrospective, and it requires standardization against a reference cohort processed on the target platform. Prospective evaluation against a defined surveillance decision, and a formal assessment of net clinical benefit, are the appropriate next stages.

Declarations

Acknowledgements

The author thanks the investigators and patients who contributed to The Cancer Genome Atlas and the Chinese Glioma Genome Atlas, without whose public data release this analysis would not have been possible.

Funding

No funds, grants or other support were received for conducting this study or preparing this manuscript.

Competing interests

The author has no relevant financial or non-financial interests to disclose. No patent application has been filed and no commercial licence has been sought in relation to the score described here.

Author contributions

I.D. is the sole author and was responsible for conceptualization, methodology, software, formal analysis, investigation, data curation, visualization, and writing of the original draft and its revision. The author read and approved the final manuscript.

Ethics approval

This study analyzed only publicly available, de-identified data from The Cancer Genome Atlas and the Chinese Glioma Genome Atlas, in which informed consent was obtained by the original investigators. It involved no human participants, no human tissue and no identifiable data collected by the author. Ethics approval was therefore not required for this secondary analysis of previously collected, de-identified public data. The study was performed in line with the principles of the Declaration of Helsinki.

Consent to participate

Not applicable. All data were obtained from public repositories in which informed consent was obtained by the original investigators.

Consent to publish

Not applicable. The manuscript contains no individual person's data in any form.

Data availability

TCGA lower-grade glioma expression and clinical data are publicly available from UCSC Xena, and molecular subtype annotations from the cBioPortal PanCanAtlas study. The CGGA mRNAseq_693 and mRNAseq_325 cohorts are publicly available from the Chinese Glioma Genome Atlas portal. All analysis code is archived at Zenodo (<https://doi.org/10.5281/zenodo.22756344>)

Use of generative artificial intelligence

A large language model (Claude, Anthropic) assisted with writing part of the analysis and plotting code. No image-generating artificial intelligence was used; all figures were rendered from the study data by the archived code. The author designed the study, ran all analyses, verified every reported value against the analysis output, and takes responsibility for the integrity of the work. No artificial intelligence tool is listed as an author.

References

1. Louis DN, Perry A, Wesseling P et al (2021) The 2021 WHO Classification of Tumors of the Central Nervous System: a summary. *Neuro Oncol* 23:1231–1251. <https://doi.org/10.1093/neuonc/noab106>
2. Cancer Genome Atlas Research Network, Brat DJ, Verhaak RGW et al (2015) Comprehensive, integrative genomic analysis of diffuse lower-grade gliomas. *N Engl J Med* 372:2481–2498. <https://doi.org/10.1056/NEJMoa1402121>
3. Buckner JC, Shaw EG, Pugh SL et al (2016) Radiation plus procarbazine, CCNU, and vincristine in low-grade glioma. *N Engl J Med* 374:1344–1355. <https://doi.org/10.1056/NEJMoa1500925>
4. Hegi ME, Diserens AC, Gorlia T et al (2005) MGMT gene silencing and benefit from temozolomide in glioblastoma. *N Engl J Med* 352:997–1003. <https://doi.org/10.1056/NEJMoa043331>
5. Ceccarelli M, Barthel FP, Malta TM et al (2016) Molecular profiling reveals biologically discrete subsets and pathways of progression in diffuse glioma. *Cell* 164:550–563. <https://doi.org/10.1016/j.cell.2015.12.028>
6. Goldman MJ, Craft B, Hastie M et al (2020) Visualizing and interpreting cancer genomics data via the Xena platform. *Nat Biotechnol* 38:675–678. <https://doi.org/10.1038/s41587-020-0546-8>

7. Cerami E, Gao J, Dogrusoz U et al (2012) The cBio Cancer Genomics Portal: an open platform for exploring multidimensional cancer genomics data. *Cancer Discov* 2:401–404. <https://doi.org/10.1158/2159-8290.CD-12-0095>
8. Zhao Z, Zhang KN, Wang Q et al (2021) Chinese Glioma Genome Atlas (CGGA): a comprehensive resource with functional genomic data from Chinese glioma patients. *Genomics Proteom Bioinf* 19:1–12. <https://doi.org/10.1016/j.gpb.2020.10.005>
9. Collins GS, Moons KGM, Dhiman P et al (2024) TRIPOD + AI statement: updated guidance for reporting clinical prediction models that use regression or machine learning methods. *BMJ* 385:e078378. <https://doi.org/10.1136/bmj-2023-078378>
10. Ioannidis JPA (2008) Why most discovered true associations are inflated. *Epidemiology* 19:640–648. <https://doi.org/10.1097/EDE.0b013e31818131e7>

Figures



Figure 1

TAMI pipeline. Expression of the 200-gene panel is rank-normalized within each batch and scored by a logistic regression classifier trained to separate glioblastoma from lower-grade glioma. TAMI is the Shannon entropy of the classifier output. The gene panel and classifier weights were fit once in TCGA and applied unchanged to the CGGA cohorts (solid boxes). Only the final standardization step is repeated in each cohort (dashed box).

Processing pipeline of the Transcriptomic Ambiguity Index (TAMI). Expression values for the 200-gene panel are rank-normalized within the submitted batch, evaluated by a regularized logistic regression classifier trained to discriminate glioblastoma from lower-grade glioma, and summarized as the Shannon entropy of the classifier posterior. Solid outlines mark the steps that are fixed: the gene panel and classifier weights are estimated once in the development cohort and applied unchanged to every external cohort. The dashed outline marks the only step recomputed for each cohort, because entropy scale is platform-dependent even after rank normalization of the inputs




Figure 2

Predicted overall survival in TCGA grade 3 IDH-mutant glioma at the 10th, 50th and 90th percentiles of TAMI, from the age-adjusted Cox model with age held at the cohort mean.

Model-predicted overall survival at representative values of TAMI in WHO grade 3 IDH-mutant glioma from the development cohort, shown at the 10th, 50th and 90th percentile of the standardized score with all other covariates held at their cohort means. Curves are derived from the fitted Cox model rather than from grouped observations, because TAMI is analyzed as a continuous variable and no risk-band cutpoints were prespecified




Figure 3

Hazard ratios per 1-SD increase in TAMI with 95% confidence intervals. Squares, age-adjusted models in grade 3 IDH-mutant glioma. Diamonds, models further adjusted for 1p/19q codeletion, MGMT methylation, grade, radiotherapy and temozolomide. Log scale; dashed line marks HR = 1.

Hazard ratios for overall survival per 1 standard deviation increase in TAMI. Squares denote the age-adjusted estimate in WHO grade 3 IDH-mutant glioma in the development cohort and in each external validation cohort; diamonds denote the estimate from fully adjusted models additionally containing 1p/19q codeletion status, MGMT promoter methylation status, histologic grade, age, and radiotherapy and temozolomide receipt. Horizontal lines are 95% confidence intervals and the dashed vertical line marks the null. The x-axis is on a logarithmic scale