

Numerical Uncertainty in Analytical Pipelines Lead to Impactful Variability in Brain Networks

Gregory Kiar¹, Yohan Chatelain², Pablo de Oliveira Castro³, Eric Petit⁴, Ariel Rokem⁵, Gaël Varoquaux⁶, Bratislav Misic¹, Alan C. Evans^{1†}, Tristan Glatard^{2‡}

Abstract

The analysis of brain-imaging data requires complex processing pipelines to support findings on brain function or pathologies. Recent work has shown that variability in analytical decisions, small amounts of noise, or computational environments can lead to substantial differences in the results, endangering the trust in conclusions^{1–7}. We explored the instability of results by instrumenting a connectome estimation pipeline with Monte Carlo Arithmetic^{8,9} to introduce random noise throughout. We evaluated the reliability of the connectomes, their features^{10,11}, and the impact on analysis^{12,13}. The stability of results was found to range from perfectly stable to highly unstable. This paper highlights the potential of leveraging induced variance in estimates of brain connectivity to reduce the bias in networks alongside increasing the robustness of their applications in the classification of individual differences. We demonstrate that stability evaluations are necessary for understanding error inherent to brain imaging experiments, and how numerical analysis can be applied to typical analytical workflows both in brain imaging and other domains of computational science. Overall, while the extreme variability in results due to analytical instabilities could severely hamper our understanding of brain organization, it also leads to an increase in the reliability of datasets.

Keywords

Stability — Reproducibility — Network Neuroscience — Neuroimaging

¹Montréal Neurological Institute, McGill University, Montréal, QC, Canada; ²Department of Computer Science and Software Engineering, Concordia University, Montréal, QC, Canada; ³Department of Computer Science, Université de Versailles, Versailles, France; ⁴Exascale Computing Lab, Intel, Paris, France; ⁵Department of Psychology and eScience Institute, University of Washington, Seattle, WA, USA; ⁶Parietal project-team, INRIA Saclay-Île de France, France; [†]Authors contributed equally.

¹ The modelling of brain networks, called connectomics, has shaped our understanding of the structure and function of the brain across a variety of organisms and scales over the last decade^{11,14–18}. In humans, these wiring diagrams are obtained *in vivo* through Magnetic Resonance Imaging (MRI), and show promise towards identifying biomarkers of disease. ⁷ This can not only improve understanding of so-called “connectopathies”, such as Alzheimer’s Disease and Schizophrenia, ⁹ but potentially pave the way for therapeutics^{19–23}. ¹⁰ However, the analysis of brain imaging data relies on complex computational methods and software. Tools are trusted to ¹¹ perform everything from pre-processing tasks to downstream ¹² statistical evaluation. While these tools undoubtedly undergo ¹³ rigorous evaluation on bespoke datasets, in the absence of ¹⁴ ground-truth this is often evaluated through measures of reliability^{24–27}, proxy outcome statistics, or agreement with ¹⁵ ¹⁶

existing theory. Importantly, this means that tools are not necessarily of known or consistent quality, and it is not uncommon that equivalent experiments may lead to diverging conclusions^{1,5–7}. While many scientific disciplines suffer from a lack of reproducibility²⁸, this was recently explored in brain imaging by a 70 team consortium which performed equivalent analyses and found widely inconsistent results¹, and it is likely that software instabilities played a role.

The present study approached evaluating reproducibility from a computational perspective in which a series of brain imaging studies were numerically perturbed in such a way that the plausibility of results was not affected, and the implications of the observed instabilities on downstream analyses were quantified. We accomplished this through the use of Monte Carlo Arithmetic (MCA)⁸, a technique which enables characterization of the sensitivity of a system to small numerical perturbations. This is importantly distinct from data perturbation experiments where the underlying datasets are manipulated or pathologies may be simulated, and allows for the evaluation of experimental uncertainty in real-world settings. We explored the impact of numerical perturbations through the direct comparison of structural connectomes, the consistency of their features, and their eventual application in a neuroscience study. We also characterized the consequences of instability in these pipelines on the reliability of derived datasets, and discuss how the induced variability may be harnessed to increase the discriminability of datasets, in an approach akin to ensemble learning. Finally, we make recommendations for the roles perturbation analyses may play in brain imaging research and beyond.

Graphs Vary Widely With Perturbations

Prior to exploring the analytic impact of instabilities, a direct understanding of the induced variability was required. A subset of the Nathan Kline Institute Rockland Sample (NKIRS) dataset²⁹ was randomly selected to contain 25 individuals with two sessions of imaging data, each of which was sub-

sampled into two components, resulting in four samples per individual and 100 samples total ($25 \times 2 \times 2$ samples). Structural connectomes were generated with canonical deterministic and probabilistic pipelines^{30,31} which were instrumented with MCA, replicating computational noise either sparsely or densely throughout the pipelines^{4,9}. In the sparse case, a small subset of the libraries were instrumented with MCA, allowing for the evaluation of the cascading effects of numerical instabilities that may arise. In the dense case, operations are more uniformly perturbed and thus the law of large numbers suggests that perturbations will quickly offset one-another and only dramatic local instabilities will have propagating effects. Importantly, the perturbations resulting from the sparse setting represent a strict subset of the possible outcomes of the dense implementation. The random perturbations are statistically independent from one another across both settings and simulations. Instrumenting pipelines with MCA increases their computation time, in this case by multiplication factors of $1.2 \times$ and $7 \times$ for the sparse and dense settings, respectively⁴. The results obtained were compared to unperturbed (e.g. reference) connectomes in both cases. The connectomes were sampled 20 times per sample and once without perturbations, resulting in a total of 8,400 connectomes. Two versions of the unperturbed connectomes were generated and compared such that the absence of variability aside from that induced via MCA could be confirmed.

The stability of connectomes was evaluated through the normalized percent deviation from reference⁴ and the number of significant digits (Figure 1). The comparisons were grouped according to differences across simulations, subsampling of data, sessions of acquisition, or subjects, and accordingly sorted from most to least similar. While the similarity of connectomes decreases as the collections become more distinct, connectomes generated with sparse perturbations show considerable variability, often reaching deviations equal to or greater than those observed across individuals or sessions (Figure 1A; right). Interpreting these results with respect to

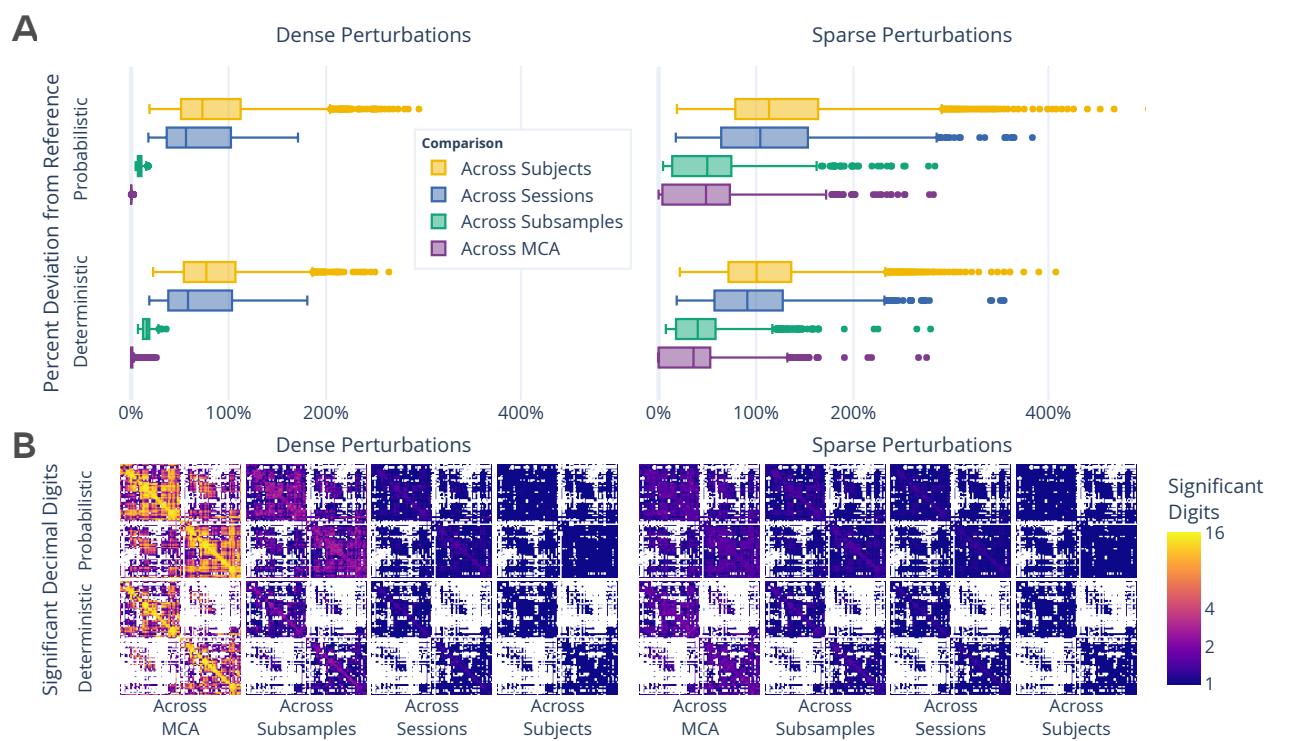


Figure 1. Exploration of perturbation-induced deviations from reference connectomes. **(A)** The absolute deviations between connectomes, in the form of normalized percent deviation from reference. The difference in MCA-perturbed connectomes is shown as the across MCA series, and is presented relative to the variability observed across subsamples, sessions, and subjects. **(B)** The number of significant decimal digits in each set of connectomes as obtained by evaluating the complete distribution of networks. In the case of 16, values can be fully relied upon, whereas in the case of 1 only the first digit of a value can be trusted. Dense and sparse perturbations are shown on the left and right, respectively.

90 the distinct MCA environments used suggests that the tested 104 connectomes (Figure 1B) similarly decreases alongside the de-
 91 pipelines may not suffer from single dominant sources of 105 creasing similarity between comparison groups. While the
 92 instability, but that nevertheless there exist minor local in- 106 cross-MCA comparison of connectomes generated with dense
 93 stabilities which may propagate throughout the pipeline. 107 perturbations show nearly perfect precision for many edges
 94 Furthermore, this finding suggests that instabilities inherent 108 (approaching the maximum of 15.7 digits for 64-bit data),
 95 to these pipelines may mask session or individual differences, 109 this evaluation uniquely shows considerable drop off in per-
 96 limiting the trustworthiness of derived connectomes. While 110 formance when comparing networks across subsamplings
 97 both pipelines show similar performance, the probabilistic 111 (average of < 4 digits). In addition, sparsely perturbed con-
 98 pipeline was more stable in the face of dense perturbations 112 nectomes show no more than an average of 3 significant digits
 99 whereas the deterministic was more stable to sparse perturba- 113 across all comparison groups, demonstrating a significant lim-
 100 tions ($p < 0.0001$ for all; exploratory). As an alternative to 114 itation in the reliability of independent edge weights. The
 101 the normalized percent deviation, the stability of correlations 115 number of significant digits across individuals did not exceed
 102 between networks can be found in Supplemental Section S1. 116 a single digit per edge in any case, indicating that only the
 103 The number of significant digits per edge across con- 117 order of magnitude of edges in naively computed groupwise

118 average connectomes can be trusted. The combination of 154 indicating a dominant session-dependent signal for all individuals
119 these results with those presented in Figure 1A suggests that 155 despite no intended biological differences. However,
120 while specific edge weights are largely affected by instabilities, 156 while still significant relative to chance (score: 0.85 and 0.88;
121 macro-scale network structure is stable. 157 $p < 0.005$ for both), sparse perturbations lead to significantly lower
158 discriminability of the dataset ($p < 0.005$ for all). This reduction of the difference between sessions suggests that
159 the added variance due to perturbations reduces the relative impact of non-biological acquisition-dependent bias inherent
160 in the networks.

122 Sparse Perturbations Reduce Off-Target Signal

123 We assessed the reproducibility of the dataset through mimicking 163 Though the previous sets of experiments inextricably evaluate the interaction between data acquisition and tool, the
124 and extending a typical test-retest experiment²⁶ in which 164 use of subsampling allowed for characterizing the discriminability of networks sampled from within a single acquisition
125 the similarity of samples across sessions were compared to 165 (Hypothesis 3). While this experiment could not be evaluated using reference executions, the networks generated with
126 distinct samples in the dataset (Table 1, with additional experiments and explanation of the measure and its scaling 166 dense perturbations showed near perfect discrimination between subsamples, with scores of 0.99 and 1.0 ($p < 0.005$;
127 in Supplemental Section S2). The ability to discriminate 167 optimal: 0.5; chance: 0.5). Given that there was no variability
128 connectomes across subjects (Hypothesis 1) is an essential 168 in data acquisition, due to undesired effects such as participant
129 prerequisite for the application of brain imaging towards identifying individual differences¹⁸. In testing hypothesis 1, we 169 motion, or preprocessing, the ability to discriminate between
130 observe that the dataset is discriminable with a scaled score of 170 equivalent subsamples in this experiment may only be due
131 0.82 ($p < 0.001$; optimal score: 1.0; chance: 0.04) for both 171 to instability or bias inherent to the pipelines. The high variance
132 pipelines in the absence of MCA. We can see that inducing instabilities through MCA preserves the discriminability in the 172 ability introduced through sparse perturbations considerably
133 stabilities through MCA preserves the discriminability in the 173 lowered the discriminability towards chance (score: 0.71 and
134 dense perturbation setting, and discriminability decreased 174 0.61; $p < 0.005$ for all), further supporting this as an effective
135 slightly but remained above the unscaled reference value of 175 method for obtaining lower-bias estimates of individual
136 0.65 in the sparse case. This lack of significant decrease in 176 connectivity.

144 While the discriminability of individuals is essential for 181 Across all cases, the induced perturbations maintained the
145 the identification of individual brain networks, it is similarly 182 ability to discriminate networks on the basis of meaningful biological signal alongside a reduction in discriminability due to
146 reliant on network similarity – or lack of discriminability – 183 across equivalent acquisitions (Hypothesis 2). In this case, 184 of off-target signal in the sparse perturbation setting. This result
147 connectomes were grouped based upon session, rather than 185 appears strikingly like a manifestation of the well-known subject, and the ability to distinguish one session from another 186 bias-variance tradeoff³² in machine learning, a concept which
148 other based on subsamples was computed within-individual 187 observes a decrease in bias as variance is favoured by a model.
149 and aggregated. Both the unperturbed and dense perturbation 188 In particular, this highlights that numerical perturbations can
150 settings perfectly preserved differences between sessions with 189 be used to not only evaluate the stability of pipelines, but that
151 a score of 1.0 ($p < 0.005$; optimal score: 0.5; chance: 0.5), 190 the induced variance may be leveraged for the interpretation

Table 1. The impact of instabilities as evaluated through the discriminability of the dataset based on individual (or subject) differences, session, and subsample. The performance is reported as mean discriminability. While a perfectly discriminable dataset would be represented by a score of 1.0, the chance performance, indicating minimal discriminability, is 1/the number of classes. H_3 could not be tested using the reference executions due to too few possible comparisons. The alternative hypothesis, indicating significant discrimination, was accepted for all experiments, with $p < 0.005$.

Comparison			Unscaled Ref.		Scaled Ref.		Dense MCA		Sparse MCA	
	Chance	Target	Det.	Prob.	Det.	Prob.	Det.	Prob.	Det.	Prob.
H_1 : Across Subjects	0.04	1.0	0.64	0.65	0.82	0.82	0.82	0.82	0.77	0.75
H_2 : Across Sessions	0.5	0.5	1.00	1.00	1.00	1.00	1.00	1.00	0.88	0.85
H_3 : Across Subsamples	0.5	0.5					0.99	1.00	0.71	0.61

191 as a robust distribution of possible results.

192 **Distributions of Graph Statistics Are Reliable, But
193 Individual Statistics Are Not**

194 Exploring the stability of topological features of connectomes 220 ministic pipeline, though the probabilistic pipeline was more
195 is relevant for typical analyses, as low dimensional features are 221 stable for all comparisons ($p < 0.0001$; exploratory). In stark
196 often more suitable than full connectomes for many analytical 222 contrast, sparse perturbations led to highly unstable feature-
197 methods in practice¹¹. A separate subset of the NKIRS dataset 223 moments (Figure 2D), such that none contained more than
198 was randomly selected to contain a single non-subsampled ses- 224 5 significant digits of information and several contained less
199 sion for 100 individuals ($100 \times 1 \times 1$) using the pipelines and 225 than a single significant digit, indicating a complete lack of re-
200 instrumentation methods to generate connectomes as above. 226 liability. This dramatic degradation in stability for individual
201 Connectomes were generated 20 times each, resulting in a 227 measures strongly suggests that these features may be unre-
202 dataset which also contained 8,400 connectomes with the 228 liable as individual biomarkers when derived from a single
203 MCA simulations serving as the only source of repeated mea- 229 pipeline evaluation, though their reliability may be increased
204 surements. 230 when studying their distributions across perturbations. A sim-

205 The stability of several commonly-used multivariate graph 231 ilar analysis was performed for univariate statistics which
206 features¹⁰ were explored and are presented in Figure 2. The 232 obtained similar findings and can be found in Supplemental
207 cumulative density of the features was computed within in- 233 Section S3.

208 individuals and the mean cumulative density and associated

209 standard error were computed for across individuals (Fig- 234 **Uncertainty in Brain-Phenotype Relationships**

210 ures 2A and 2B). There was no significant difference between 235 While the variability of connectomes and their features was
211 the distributions for each feature across the two perturbation 236 summarized above, networks are commonly used as inputs to
212 settings, suggesting that the topological features summarized 237 machine learning models tasked with learning brain-phenotype
213 by these multivariate features are robust across both perturba- 238 relationships¹⁸. To explore the stability of these analyses, we
214 tion modes. 239 modelled the relationship between high- or low- Body Mass

215 In addition to the comparison of distributions, the stabil- 240 Index (BMI) groups and brain connectivity using standard di-

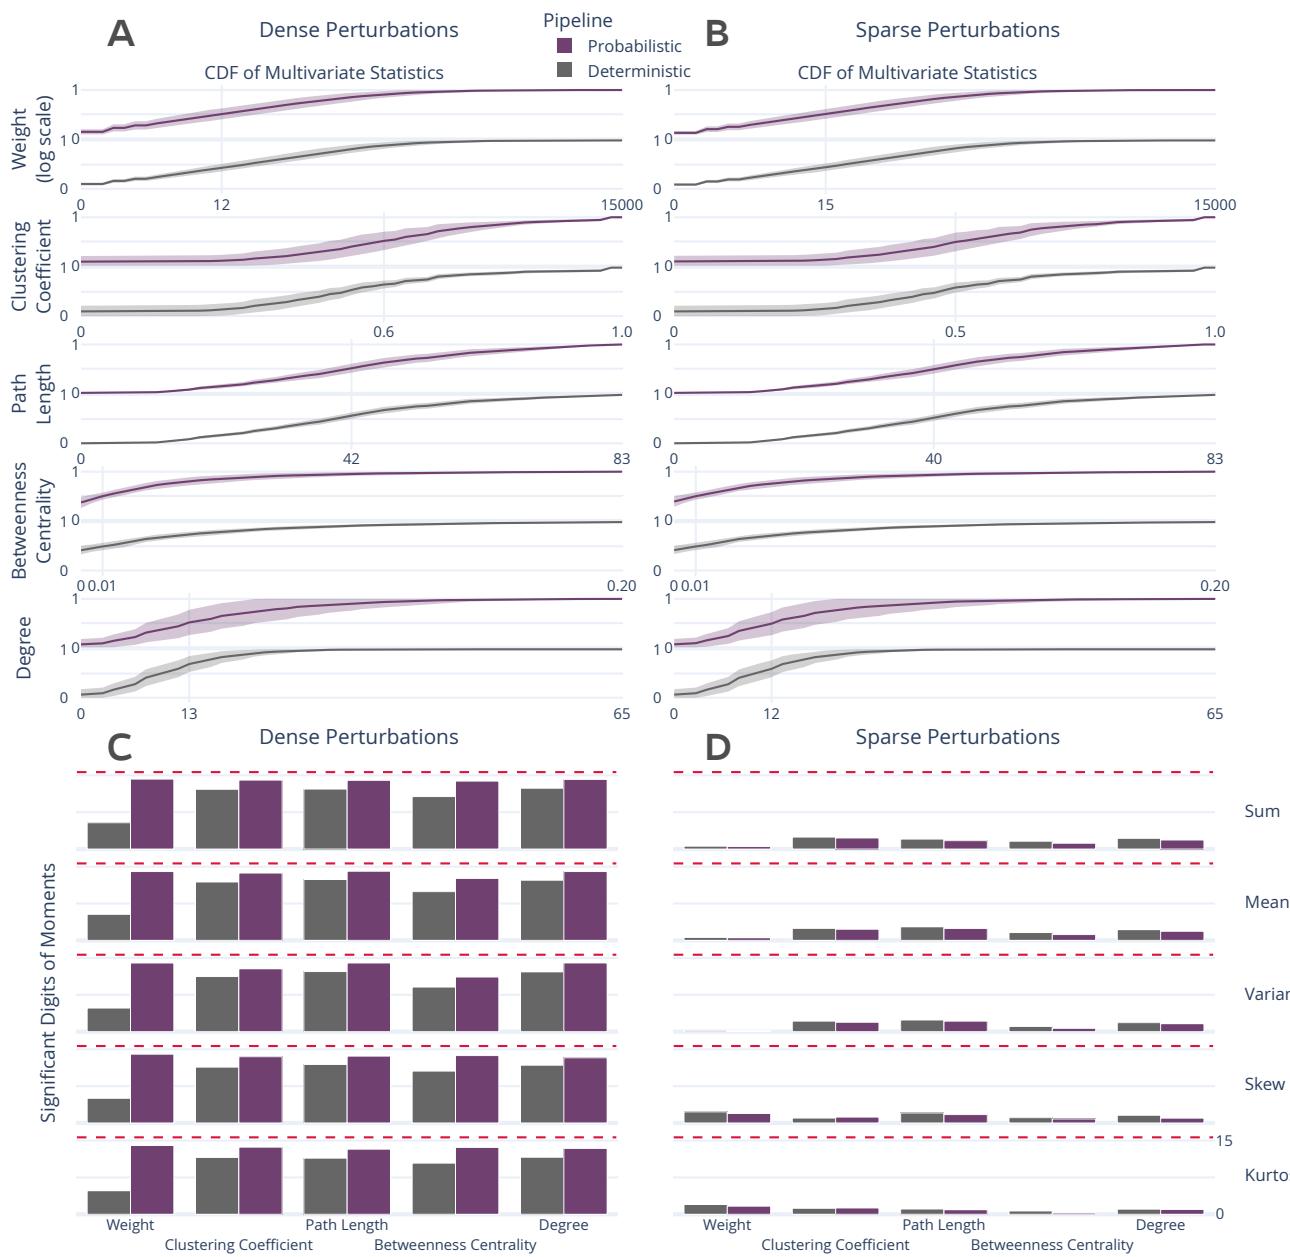


Figure 2. Distribution and stability assessment of multivariate graph statistics. (A, B) The cumulative distribution functions of multivariate statistics across all subjects and perturbation settings. There was no significant difference between the distributions in A and B. (C, D) The number of significant digits in the first 5 moments of each statistic across perturbations. The dashed red line refers to the maximum possible number of significant digits.

241 dimensionality reduction and classification tools^{12,13}, and com- 246 from 0.520 – 0.716 and 0.510 – 0.725, respectively, rang-
242 pared this to reference and random performance (Figure 3). 247 ing from at or below random performance to outperforming

243 The analysis was perturbed through distinct samplings of 248 performance on the reference dataset. This large variability
244 the dataset across both pipelines and perturbation methods. 249 illustrates a previously uncharacterized margin of uncertainty

250 in the modelling of this relationship, and limits confidence in

245 The accuracy and F1 score for the perturbed models varied

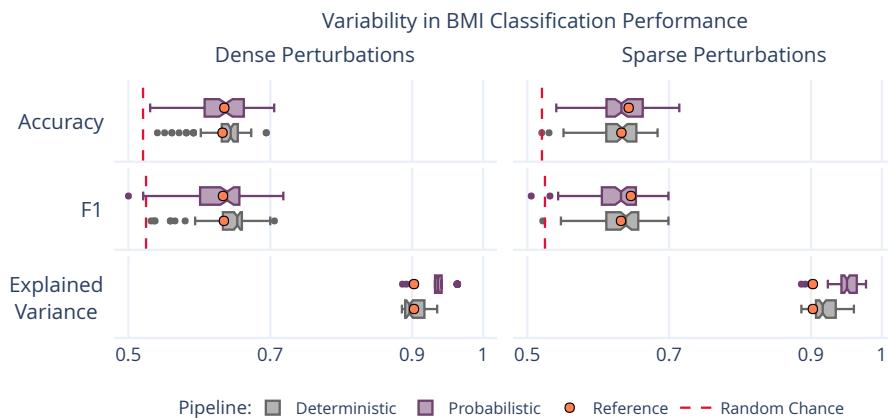


Figure 3. Variability in BMI classification across the sampling of an MCA-perturbed dataset. The dashed red lines indicate random-chance performance, and the orange dots show the performance using the reference executions.

251 reported accuracy scores on singly processed datasets. The 274 **Discussion**

252 portion of explained variance in these samples ranged from 275 The perturbation of structural connectome estimation pipelines
253 88.6% -- 97.8%, similar to the reference of 90.3%, suggest- 276 with small amounts of noise, on the order of machine error,
254 ing that the range in performance was not due to a gain or 277 led to considerable variability in derived brain graphs. Across
255 loss of meaningful signal, but rather the reduction of bias 278 all analyses the stability of results ranged from nearly per-
256 towards specific outcome. Importantly, this finding does not 279 fectly trustworthy (i.e. no variation) to completely unreliable
257 suggest that modelling brain-phenotype relationships is not 280 (i.e. containing no trustworthy information). Given that the
258 possible, but rather it sheds light on impactful uncertainty that 281 magnitude of introduced numerical noise is to be expected
259 must be accounted for in this process, and supports the use of 282 in computational workflows, this finding has potentially sig-
260 ensemble modeling techniques.

274 The perturbation of structural connectome estimation pipelines
275 with small amounts of noise, on the order of machine error,
276 led to considerable variability in derived brain graphs. Across
277 all analyses the stability of results ranged from nearly per-
278 fectly trustworthy (i.e. no variation) to completely unreliable
279 (i.e. containing no trustworthy information). Given that the
280 magnitude of introduced numerical noise is to be expected
281 in computational workflows, this finding has potentially sig-
282 nificant implications for inferences in brain imaging as it is
283 currently performed. In particular, this bounds the success of

261 One distinction between the results presented here and 284 currently performed. In particular, this bounds the success of
262 the previous is that while networks derived from dense pertur- 285 studying individual differences, a central objective in brain
263 bations had been shown to exhibit less dramatic instabilities 286 imaging¹⁸, given that the quality of relationships between
264 in general, the results here show similar variability in clas- 287 phenotypic data and brain networks will be limited by the
265 sification performance across the two methods. This consis- 288 stability of the connectomes themselves. This issue is accen-
266 tency suggests that the desired method of instrumentation may 289 tuated through the crucial finding that individually derived
267 vary across experiments. While sparse perturbations result 290 network features were unreliable despite there being no signif-
268 in considerably more variability in networks directly, the two 291 icant difference in their aggregated distributions. This finding
269 techniques capture similar variability when relating networks 292 is not damning for the study of brain networks as a whole, but
270 to this phenotypic variable. Given the dramatic reduction 293 rather is strong support for the aggregation of networks, either
271 in computational overhead, a sparse instrumentation may be 294 across perturbations for an individual or across groups, over
272 preferred when processing datasets for eventual application in 295 the use of individual estimates.

273 modelling brain-phenotype relationships.

296 **Underestimated False Positive Rates** While the instabil-
297 ity of brain networks was used here to demonstrate the lim-

298 iterations of modelling brain-phenotype relationships in the 335 paradigm shift. Given that MCA is data-agnostic, this tech-
299 context of machine learning, this limitation extends to classi- 336 nique could be used effectively in conjunction with, or in
300 cal hypothesis testing, as well. Though performing individual 337 lieu of, realistic noise models to augment existing datasets.
301 comparisons in a hypothesis testing framework will be accom- 338 While this of course would not replace the need for repeated
302 panied by reported false positive rates, the accuracy of these 339 measurements when exploring the effect of data collection
303 rates is critically dependent upon the reliability of the samples 340 paradigm or study longitudinal progressions of development
304 used. In reality, the true false positive rate for a test would be 341 or disease, it could be used in conjunction with these efforts to
305 a combination of the reported confidence and the underlying 342 decrease the bias of each distinct sample within a dataset. In
306 variability in the results, a typically unknown quantity. 343 contexts where repeated measurements are typically collected
307 When performing these experiments outside of a repeated- 344 to increase the fidelity of the dataset, MCA could potentially
308 measure context, such as that afforded here through MCA, it 345 serve as an alternative solution to capture more biological vari-
309 is impossible to empirically estimate the reliability of samples. 346 ability, with the added benefit being the savings of millions of
310 This means that the reliability of accepted hypotheses is also 347 dollars on data collection.

311 unknown, regardless of the reported false positive rate. In 348 **Shortcomings and Future Questions** Given the complex-
312 fact, it is a virtual certainty that the true false positive rate 349 ity of recompiling complex software libraries, pre-processing
313 for a given hypothesis exceeds the reported value simply as 350 was not perturbed in these experiments as the instrumentation
314 a result of numerical instabilities. This uncertainty inherent 351 of the canonical workflow used in diffusion image process-
315 to derived data is compounded with traditional arguments 352 ing would have added considerable technical complexity and
316 limiting the trustworthiness of claims³³, and hampers the 353 computational overhead to the large set of experiments per-
317 ability of researchers to evaluate the quality of results. The 354 formed here. Other work has shown that linear registration, a
318 accompaniment of brain imaging experiments with direct 355 core piece of many elements of pre-processing such as motion
319 evaluations of their stability, as was done here, would allow 356 correction and alignment, is sensitive to minor perturbations⁷.
320 researchers to simultaneously improve the numerical stability
321 of their analyses and accurately gauge confidence in them.
322 The induced variability in derived brain networks may be
323 leveraged to estimate aggregate connectomes with lower bias
324 than any single independent observation, leading to learned
325 relationships that are more generalizable and ultimately more
326 useful.

327 **Cost-Effective Data Augmentation** The evaluation of reli- 364 This paper does not explore methodological flexibility or
328 ability in brain imaging has historically relied upon the expen- 365 compare this to numerical instability. Recently, the nearly
329 sive collection of repeated measurements choreographed by 366 boundless space of analysis pipelines and their impact on out-
330 massive cross-institutional consortia^{34,35}. The finding that per- 367 comes in brain imaging has been clearly demonstrated¹. The
331 turbing experiments using MCA both preserved the discrim- 368 approach taken in these studies complement one another and
332 inability of the dataset due to biological signal and decreased 369 explore instability at the opposite ends of the spectrum, with
333 the discriminability due to off-target differences across ac- 370 human variability in the construction of an analysis workflow
334 quisitions and subsamples opens the door for a promising 371 on one end and the unavoidable error implicit in the digital

348 **Shortcomings and Future Questions** Given the complex-
349 ity of recompiling complex software libraries, pre-processing
350 was not perturbed in these experiments as the instrumentation
351 of the canonical workflow used in diffusion image process-
352 ing would have added considerable technical complexity and
353 computational overhead to the large set of experiments per-
354 formed here. Other work has shown that linear registration, a
355 core piece of many elements of pre-processing such as motion
356 correction and alignment, is sensitive to minor perturbations⁷.
357 It is likely that the instabilities across the entire processing
358 workflow would be compounded with one another, resulting
359 in even greater variability. While the analyses performed in
360 this paper evaluated a single dataset and set of pipelines, ex-
361 tending this work to other modalities and analyses, alongside
362 the detection of local sources of instability within pipelines,
363 is of interest for future projects.

364 This paper does not explore methodological flexibility or
365 compare this to numerical instability. Recently, the nearly
366 boundless space of analysis pipelines and their impact on out-
367 comes in brain imaging has been clearly demonstrated¹. The
368 approach taken in these studies complement one another and
369 explore instability at the opposite ends of the spectrum, with
370 human variability in the construction of an analysis workflow
371 on one end and the unavoidable error implicit in the digital

372 representation of data on the other. It is of extreme interest
373 to combine these approaches and explore the interaction of
374 these scientific degrees of freedom with effects from software
375 implementations, libraries, and parametric choices.

376 Finally, it is important to state explicitly that the work
377 presented here does not invalidate analytical pipelines used in
378 brain imaging, but merely sheds light on the fact that many
379 studies are accompanied by an unknown degree of uncertainty
380 due to machine-introduced errors. The presence of unknown
381 error-bars associated with experimental findings limits the
382 impact of results due to increased uncertainty. The desired
383 outcome of this paper is to motivate a shift in scientific com-
384 puting – both in neuroimaging and more broadly – towards a
385 paradigm that favours the explicit evaluation of the trustwor-
386 thiness of claims alongside the claims themselves.

Methods

388 Dataset

389 The Nathan Kline Institute Rockland Sample (NKI-RS)²⁹
390 dataset contains high-fidelity imaging and phenotypic data
391 from over 1,000 individuals spread across the lifespan. A
392 subset of this dataset was chosen for each experiment to both
393 match sample sizes presented in the original analyses and to
394 minimize the computational burden of performing MCA. The
395 selected subset comprises 100 individuals ranging in age from
396 6 – 79 with a mean of 36.8 (original: 6 – 81, mean 37.8),
397 60% female (original: 60%), with 52% having a BMI over 25
398 (original: 54%).

399 Each selected individual had at least a single session
400 of both structural T1-weighted (MPRAGE) and diffusion-
401 weighted (DWI) MR imaging data. DWI data was acquired
402 with 137 diffusion directions in a single shell; more informa-
403 tion regarding the acquisition of this dataset can be found in
404 the NKI-RS data release²⁹.

405 In addition to the 100 sessions mentioned above, 25 indi-
406 viduals had a second session to be used in a test-retest analysis.
407 Two additional copies of the data for these individuals were
408 generated, including only the odd or even diffusion direc-
409 tions (64 + 9 B0 volumes = 73 in either case) such that the
410 acquired data was evenly represented across both portions,
411 and each subsample represented a realistic complete acqui-
412 sition. This allowed for an extra level of stability evaluation to
413 be performed between the levels of MCA and session-level
414 variation.

415 In total, the dataset is composed of 100 subsampled ses-
416 sions of data originating from 50 acquisitions and 25 indi-
417 viduals for in depth stability analysis, and an additional 100
418 sessions of full-resolution data from 100 individuals for sub-
419 sequent analyses.

420 Processing

421 The dataset was preprocessed using a standard FSL³⁶ work-
422 flow consisting of eddy-current correction and alignment. The

423 MNI152 atlas³⁷ was aligned to each session of data via the 457
424 structural images, and the resulting transformation was ap- 458
425 plied to the DKT parcellation³⁸. Subsampling the diffusion 459
426 data took place after preprocessing was performed on full- 460
427 resolution sessions, ensuring that an additional confound was 461
428 not introduced in this process when comparing between down- 462
429 sampled sessions. The preprocessing described here was per- 463
430 formed once without MCA, and thus is not being evaluated. 464

431 Structural connectomes were generated from preprocessed 465
432 data using two canonical pipelines from Dipy³⁰: deterministic 466
433 and probabilistic. In the deterministic pipeline, a constant 467
434 solid angle model was used to estimate tensors at each voxel 468
435 and streamlines were then generated using the EuDX algo- 469
436 rithm³¹. In the probabilistic pipeline, a constrained spherical 470
437 deconvolution model was fit at each voxel and streamlines 471
438 were generated by iteratively sampling the resulting fiber ori-
439 entation distributions. In both cases tracking occurred with 8 472
440 seeds per 3D voxel and edges were added to the graph based 473
441 on the location of terminal nodes with weight determined by 474
442 fiber count.

443 The random state of both pipelines was fixed for all anal- 472
444 yses. Fixing this random state led to entirely deterministic 473
445 repeated-evaluations of the tools, and allowed for explicit 474
446 attribution of observed variability to limitations in tool preci- 475
447 sion as provoked by Monte Carlo simulations, rather than the 476
448 internal state of the algorithm.

449 Perturbations

450 All connectomes were generated with one reference execu-
451 tion where no perturbation was introduced in the processing.
452 For all other executions, all floating point operations were
453 instrumented with Monte Carlo Arithmetic (MCA)⁸ through
454 Verificarlo⁹. MCA simulates the distribution of errors im-
455 plicit to all instrumented floating point operations (flop). This
456 rounding is performed on a value x at precision t by:

$$457 \text{inexact}(x) = x + 2^{e_x - t} \xi \quad (1)$$

457 where e_x is the exponent value of x and ξ is a uniform ran-
458 dom variable in the range $(-\frac{1}{2}, \frac{1}{2})$. MCA can be introduced in
459 two places for each flop: before or after evaluation. Perform-
460 ing MCA on the inputs of an operation limits its precision,
461 while performing MCA on the output of an operation high-
462 lights round-off errors that may be introduced. The former is
463 referred to as Precision Bounding (PB) and the latter is called
464 Random Rounding (RR).

465 Using MCA, the execution of a pipeline may be performed
466 many times to produce a distribution of results. Studying the
467 distribution of these results can then lead to insights on the
468 stability of the instrumented tools or functions. To this end,
469 a complete software stack was instrumented with MCA and
470 is made available on GitHub at <https://github.com/verificarlo/fuzzy>.

472 The RR variant of MCA was used for all experiments.
473 As was presented in⁴, both the degree of instrumentation (i.e.
474 number of affected libraries) and the perturbation mode have
475 an effect on the distribution of observed results. For this work,

476 the RR-MCA was applied across the bulk of the relevant oper-
477 ations (those occurring in BLAS, LAPACK, Python, Cython,
478 and Numpy) and is referred to as dense perturbation. In this
479 case the bulk of numerical operations were affected by MCA.

480 Conversely, the case in which RR-MCA was applied
481 across the operations in a small subset of operations (those
482 occurring in Python and Cython) is here referred to as sparse
483 perturbation. In this case, the inputs to operations within
484 the instrumented libraries were perturbed, resulting in less
485 frequent, data-centric perturbations. Alongside the stated the-
486 oretical differences, sparse perturbation is considerably less
487 computationally expensive than dense perturbation.

488 All perturbations targeted the least-significant-bit for all
489 data ($t = 24$ and $t = 53$ in float32 and float64, respectively⁹).
490 Perturbing the least significant bit importantly serves as a
491 perturbation of machine error, and thus is the appropriate
492 precision to be applied globally in complex pipelines. Simula-
493 tions were performed 20 times for each pipeline execution for

494 the 100 sample dataset and 10 times for the repeated measures 527 of differences in observed graphs relative to the original signal
 495 dataset. A detailed motivation for the number of simulations 528 intensity. A Pearson correlation coefficient⁴⁰ was computed
 496 can be found in³⁹. 529 in complement to normalized percent deviation to identify
 530 the consistency of structure and not just intensity between ob-
 531 served graphs, though the result of this experiment is shown
 532 only in Supplemental Section S1.

497 Evaluation

498 The magnitude and importance of instabilities in pipelines
 499 can be considered at a number of analytical levels, namely:
 500 the induced variability of derivatives directly, the resulting
 501 downstream impact on summary statistics or features, or the
 502 ultimate change in analyses or findings. We explore the na-
 503 ture and severity of instabilities through each of these lenses.
 504 Unless otherwise stated, all p-values were computed using
 505 Wilcoxon signed-rank tests. To avoid biasing these statistics in
 506 this unique repeated-measures context, tests were performed
 507 across sets of independent observations and then the results
 508 were aggregated in all cases.

509 Direct Evaluation of the Graphs

510 The differences between perturbation-generated graphs was
 511 measured directly through both a direct variance quantifica-
 512 tion and a comparison to other sources of variance such as
 513 individual- and session-level differences.

514 **Quantification of Variability** Graphs, in the form of adja-
 515 cency matrices, were compared to one another using three
 516 metrics: normalized percent deviation, Pearson correlation,
 517 and edgewise significant digits. The normalized percent devi-
 518 ation measure, defined in⁴, scales the norm of the difference
 519 between a simulated graph and the reference execution (that
 520 without intentional perturbation) with respect to the norm of
 521 the reference graph, and is defined as⁴:

$$\%Dev(A, B) = \sqrt{\sum_{i=1}^m \sum_{j=1}^n |a_{ij} - b_{ij}|^2} / \sqrt{\sum_{i=1}^m \sum_{j=1}^n |a_{ij}|^2}, \quad (2)$$

522 where A and B each represent a graph, and \square_{ij} are el-
 523 ements therein corresponding to row and column i and j , 554
 524 respectively. For these experiments, the A graph always refers 555 to the reference, where B represents a perturbed value. The 556
 526 purpose of this comparison is to provide insight on the scale 557 observation belonging to a given class will be more similar to

527 of differences in observed graphs relative to the original signal
 528 intensity. A Pearson correlation coefficient⁴⁰ was computed
 529 in complement to normalized percent deviation to identify
 530 the consistency of structure and not just intensity between ob-
 531 served graphs, though the result of this experiment is shown
 532 only in Supplemental Section S1.

533 Finally, the estimated number of significant digits, s' , for
 534 each edge in the graph is calculated as:

$$s' = -\log_{10} \frac{\sigma}{|\mu|} \quad (3)$$

535 where μ and σ are the mean and unbiased estimator of
 536 standard deviation across graphs, respectively. The upper
 537 bound on significant digits is 15.7 for 64-bit floating point
 538 data.

539 The percent deviation, correlation, and number of signifi-
 540 cant digits were each calculated within a single session of data,
 541 thereby removing any subject- and session-effects and provid-
 542 ing a direct measure of the tool-introduced variability across
 543 perturbations. A distribution was formed by aggregating these
 544 individual results.

545 **Class-based Variability Evaluation** To gain a concrete un-
 546 derstanding of the significance of observed variations we ex-
 547 plore the separability of our results with respect to understood
 548 sources of variability, such as subject-, session-, and pipeline-
 549 level effects. This can be probed through Discriminability²⁶,
 550 a technique similar to ICC²⁴ which relies on the mean of a
 551 ranked distribution of distances between observations belong-
 552 ing to a defined set of classes. The discriminability statistic is
 553 formalized as follows:

$$Disc. = Pr(\|g_{ij} - g_{i'j'}\| \leq \|g_{ij} - g_{i'j'}\|) \quad (4)$$

554 where g_{ij} is a graph belonging to class i that was measured
 555 at observation j , where $i \neq i'$ and $j \neq j'$.

556 Discriminability can then be read as the probability that an
 557 observation belonging to a given class will be more similar to

558 other observations within that class than observations of a different class. It is a measure of reproducibility, and is discussed in detail in²⁶. This definition allows for the exploration of deviations across arbitrarily defined classes that in practice can be any of those listed above. We combine this statistic with permutation testing to test hypotheses on whether differences between classes are statistically significant in each of these settings. This statistic is similar to *ICC*²⁴ in a two-measurement setting, however, given the dependence on a rank distribution from all measurements, discriminability scores do not become meaningless by the addition of more samples which are highly similar to the originals, whereas *ICC* scores would much more rapidly trend towards 1, making discriminability appropriate in this context. The scaling properties of discriminability are described more fully in Supplemental Section S2.

573 With this in mind, three hypotheses were defined. For each setting, we state the alternate hypotheses, the variable(s) which were used to determine class membership, and the remaining variables which may be sampled when obtaining multiple observations. Each hypothesis was tested independently for each pipeline and perturbation mode.

579 *H_{A1}*: Individuals are distinct from one another

580 Class definition: *Subject ID*

581 Comparisons: **Session (1 subsample)**, **Subsample (1 session)**, **MCA (1 subsample, 1 session)**

583 *H_{A2}*: Sessions within an individual are distinct

584 Class definition: *Session ID | Subject ID*

585 Comparisons: **Subsample**, **MCA (1 subsample)**

586 *H_{A3}*: Subsamples are distinct

587 Class definition: *Subsample | Subject ID, Session ID*

588 Comparisons: **MCA**

589 As a result, we tested 3 hypotheses across 6 MCA experiments and 3 reference experiments on 2 pipelines and 2 perturbation modes, resulting in a total of 30 distinct tests. While results from all tests can be found within Supplemental

593 Section S2, only the bolded comparisons in the list above have 594 been presented in the main body of this article. Correction for 595 repeated testing was performed.

596 Evaluating Graph-Theoretical Metrics

597 While connectomes may be used directly for some analyses, 598 it is common practice to summarize them with structural measures, that can then be used as lower-dimensional proxies 599 of connectivity in so-called graph-theoretical studies¹¹. We 600 explored the stability of several commonly-used univariate 601 (graphwise) and multivariate (nodewise or edgewise) features. 602 The features computed and subsequent methods for comparison 603 in this section were selected to closely match those computed in¹⁰.

606 Univariate Differences

For each univariate statistic (edge count, mean clustering coefficient, global efficiency, modularity of the largest connected component, assortativity, and mean path length) a distribution of values across all perturbations within subjects was observed. A Z-score was computed for each sample with respect to the distribution of feature values within an individual, and the proportion of "classically significant" Z-scores, i.e. corresponding to $p < 0.05$, was reported and aggregated across all subjects. There was no correction for multiple comparisons in these statistics, as they were not used to interpret a hypothesis but demonstrate the false-positive rate due to perturbations. The number of significant digits contained within an estimate derived from a single subject were calculated and aggregated. The results of this analysis can be found in Supplemental Section S3.

621 Multivariate Differences

In the case of both nodewise (degree distribution, clustering coefficient, betweenness centrality) and edgewise (weight distribution, connection length) features, the cumulative density functions of their distributions

were evaluated over a fixed range and subsequently aggregated across individuals. The number of significant digits for each moment of these distributions (sum, mean, variance, skew, and kurtosis) were calculated across observations within

629 a sample and aggregated.

630 Evaluating A Brain-Phenotype Analysis

631 Though each of the above approaches explores the instability of derived connectomes and their features, many modern 632 studies employ modeling or machine-learning approaches, for 633 instance to learn brain-phenotype relationships or identify differences 634 across groups. We carried out one such study and explored the instability of its results with respect to the upstream 635 variability of connectomes characterized in the previous sections. 636 We performed the modeling task with a single sampled 637 connectome per individual and repeated this sampling and 638 modelling 20 times. We report the model performance for 639 each sampling of the dataset and summarize its variance.

642 **BMI Classification** Structural changes have been linked to 643 obesity in adolescents and adults⁴¹. We classified normal- 644 weight and overweight individuals from their structural net- 645 works (using for overweight a cutoff of $BMI > 25$ ¹³). We 646 reduced the dimensionality of the connectomes through prin- 647 cipal component analysis (PCA), and provided the first N- 648 components to a logistic regression classifier for predicting 649 BMI class membership, similar to methods shown in^{12,13}. The 650 number of components was selected as the minimum set 651 which explained $> 90\%$ of the variance when averaged across 652 the training set for each fold within the cross validation of 653 the original graphs; this resulted in a feature of 20 compo- 654 nents. We trained the model using k -fold cross validation, 655 with $k = 2, 5, 10$, and N (equivalent to leave-one-out; LOO).

656 Data & Code Provenance

657 The unprocessed dataset is available through The Consortium 658 of Reliability and Reproducibility (http://fcon_1000.projects.nitrc.org/indi/enhanced/), including 659 both the imaging data as well as phenotypic data which may 660 be obtained upon submission and compliance with a Data Us- 661 age Agreement. The connectomes generated through simula- 662 tions have been bundled and stored permanently (<https://doi.org/10.5281/zenodo.4041549>), and are made

665 available through The Canadian Open Neuroscience Platform 666 (<https://portal.conp.ca/search>, search term "Kiar").

667 All software developed for processing or evaluation is 668 publicly available on GitHub at <https://github.com/gkpapers/2020ImpactOfInstability.Experiments>. 669 Experiments 670 were launched using Boutiques⁴² and Cloudr⁴³ in Compute 671 Canada's HPC cluster environment. MCA instrumentation 672 was achieved through Verificarlo⁹ available on Github at 673 <https://github.com/verificarlo/verificarlo>.

674 A set of MCA instrumented software containers is available 675 on Github at <https://github.com/gkiar/fuzzy>.

676 Author Contributions

677 GK was responsible for the experimental design, data pro- 678 cessing, analysis, interpretation, and the majority of writing. 679 All authors contributed to the revision of the manuscript. YC, 680 POC, and EP were responsible for MCA tool development and 681 software testing. AR, GV, and BM contributed to experimen- 682 tal design and interpretation. TG contributed to experimental 683 design, analysis, and interpretation. TG and ACE were re- 684 sponsible for supervising and supporting all contributions 685 made by GK. The authors declare no competing interests for 686 this work. Correspondence and requests for materials should 687 be addressed to Tristan Glatard at tristan.glatard@concordia.ca.

689 Acknowledgments

690 This research was financially supported by the Natural Sci- 691 ences and Engineering Research Council of Canada (NSERC) 692 (award no. CGSD3-519497-2018). This work was also sup- 693 ported in part by funding provided by Brain Canada, in partner- 694 ship with Health Canada, for the Canadian Open Neuroscience 695 Platform initiative.

696 References

697 [1] R. Botvinik-Nezer, F. Holzmeister, C. F. Camerer, A. Dreber, J. Huber, 698 M. Johannesson, M. Kirchler, R. Iwanir, J. A. Mumford, R. A. Adcock 699 *et al.*, "Variability in the analysis of a single neuroimaging dataset by 700 many teams," *Nature*, pp. 1–7, 2020.

701 [2] C. M. Bennett, M. B. Miller, and G. L. Wolford, “Neural correlates of 746 [15] M. Xia, Q. Lin, Y. Bi, and Y. He, “Connectomic insights into topologi-
702 interspecies perspective taking in the post-mortem Atlantic salmon: An 747 cally centralized network edges and relevant motifs in the human brain,”
703 argument for multiple comparisons correction,” *Neuroimage*, vol. 47, no. 748 *Frontiers in human neuroscience*, vol. 10, p. 158, 2016.
704 Suppl 1, p. S125, 2009.

705 [3] A. Eklund, T. E. Nichols, and H. Knutsson, “Cluster failure: Why 750 [16] J. L. Morgan and J. W. Lichtman, “Why not connectomics?” *Nature
706 fMRI inferences for spatial extent have inflated false-positive rates,” 751 methods*, vol. 10, no. 6, p. 494, 2013.
707 *Proceedings of the national academy of sciences*, vol. 113, no. 28, pp. 752

708 7900–7905, 2016.

709 [4] G. Kiar, P. de Oliveira Castro, P. Rioux, E. Petit, S. T. Brown, A. C. 753 [17] M. P. Van den Heuvel, E. T. Bullmore, and O. Sporns, “Comparative
710 Evans, and T. Glatard, “Comparing perturbation models for evaluating 754 connectomics,” *Trends in cognitive sciences*, vol. 20, no. 5, pp. 345–361,
711 stability of neuroimaging pipelines,” *The International Journal of High 755 2016.
712 Performance Computing Applications*, 2020.

713 [5] A. Salari, G. Kiar, L. Lewis, A. C. Evans, and T. Glatard, “File-based 756 [18] J. Dubois and R. Adolphs, “Building a science of individual differences
714 localization of numerical perturbations in data analysis pipelines,” *arXiv 757 from fMRI,” *Trends Cogn. Sci.*, vol. 20, no. 6, pp. 425–443, Jun. 2016.
715 preprint arXiv:2006.04684*, 2020.

716 [6] L. B. Lewis, C. Y. Lepage, N. Khalili-Mahani, M. Omidyeganeh, S. Jeon, 758 [19] A. Fornito and E. T. Bullmore, “Connectomics: a new paradigm for
717 P. Bermudez, A. Zijdenbos, R. Vincent, R. Adalat, and A. C. Evans, 759 understanding brain disease,” *European Neuropsychopharmacology*,
718 “Robustness and reliability of cortical surface reconstruction in CIVET 760 vol. 25, no. 5, pp. 733–748, 2015.
719 and FreeSurfer,” *Annual Meeting of the Organization for Human Brain 761
720 Mapping*, 2017.

721 [7] T. Glatard, L. B. Lewis, R. Ferreira da Silva, R. Adalat, N. Beck, C. Lep- 762 [20] G. Deco and M. L. Kringelbach, “Great expectations: using whole-
722 age, P. Rioux, M.-E. Rousseau, T. Sherif, E. Deelman, N. Khalili-Mahani, 763 brain computational connectomics for understanding neuropsychiatric
723 and A. C. Evans, “Reproducibility of neuroimaging analyses across op- 764 disorders,” *Neuron*, vol. 84, no. 5, pp. 892–905, 2014.
724 erating systems,” *Front. Neuroinform.*, vol. 9, p. 12, Apr. 2015.

725 [8] D. S. Parker, *Monte Carlo Arithmetic: exploiting randomness in floating- 765 [21] T. Xie and Y. He, “Mapping the alzheimer’s brain with connectomics,”
726 point arithmetic*. University of California (Los Angeles). Computer 766 *Frontiers in psychiatry*, vol. 2, p. 77, 2012.
727 Science Department, 1997.

728 [9] C. Denis, P. de Oliveira Castro, and E. Petit, “Verificarlo: Checking 767 [22] M. Filippi, M. P. van den Heuvel, A. Fornito, Y. He, H. E. H. Pol,
729 floating point accuracy through monte carlo arithmetic,” *2016 IEEE 768 F. Agosta, G. Comi, and M. A. Rocca, “Assessment of system dys-
730 23rd Symposium on Computer Arithmetic (ARITH)*, 2016. 769 function in the brain through mri-based connectomics,” *The Lancet
731 [10] R. F. Betzel, A. Griffa, P. Hagmann, and B. Mišić, “Distance-dependent 770 Neurology*, vol. 12, no. 12, pp. 1189–1199, 2013.

732 consensus thresholds for generating group-representative structural brain 771 [23] M. P. Van Den Heuvel and A. Fornito, “Brain networks in schizophrenia,”
733 networks,” *Network neuroscience*, vol. 3, no. 2, pp. 475–496, 2019. 772 *Neuropsychology review*, vol. 24, no. 1, pp. 32–48, 2014.

734 [11] M. Rubinov and O. Sporns, “Complex network measures of brain con- 773 [24] J. J. Bartko, “The intraclass correlation coefficient as a measure of
735 nectivity: uses and interpretations,” *Neuroimage*, vol. 52, no. 3, pp. 774 reliability,” *Psychol. Rep.*, vol. 19, no. 1, pp. 3–11, Aug. 1966.

736 1059–1069, Sep. 2010.

737 [12] B.-Y. Park, J. Seo, J. Yi, and H. Park, “Structural and functional brain 775 [25] A. M. Brandmaier, E. Wenger, N. C. Bodammer, S. Kühn, N. Raz,
738 connectivity of people with obesity and prediction of body mass index 776 and U. Lindenberger, “Assessing reliability in neuroimaging research
739 using connectivity,” *PLoS One*, vol. 10, no. 11, p. e0141376, Nov. 2015. 777 through intra-class effect decomposition (ICED),” *Elife*, vol. 7, Jul. 2018.

740 [13] A. Gupta, E. A. Mayer, C. P. Sanmiguel, J. D. Van Horn, D. Woodworth, 778 [26] E. W. Bridgeford, S. Wang, Z. Yang, Z. Wang, T. Xu, C. Craddock,
741 B. M. Ellingson, C. Fling, A. Love, K. Tillisch, and J. S. Labus, “Pat- 779 J. Dey, G. Kiar, W. Gray-Roncal, C. Coulantoni *et al.*, “Eliminating
742 terns of brain structural connectivity differentiate normal weight from 780 accidental deviations to minimize generalization error: applications in
743 overweight subjects,” *Neuroimage Clin*, vol. 7, pp. 506–517, Jan. 2015. 781 connectomics and genomics,” *bioRxiv*, p. 802629, 2020.

744 [14] T. E. Behrens and O. Sporns, “Human connectomics,” *Current opinion 782 [27] G. Kiar, E. Bridgeford, W. G. Roncal, V. Chandrashekhar, and others, “A High-Throughput pipeline identifies robust connectomes but
745 in neurobiology*, vol. 22, no. 1, pp. 144–153, 2012. 783 troublesome variability,” *bioRxiv*, 2018.

746 [28] M. Baker, “1,500 scientists lift the lid on reproducibility,” *Nature*, 2016.

747 [29] K. B. Noonan, S. J. Colcombe, R. H. Tobe, M. Mennens *et al.*, “The
748 NKI-Rockland sample: A model for accelerating the pace of discovery
749 science in psychiatry,” *Front. Neurosci.*, vol. 6, p. 152, Oct. 2012.

750 [30] E. Garyfallidis, M. Brett, B. Amirbekian, A. Rokem, S. van der Walt,
751 M. Descoteaux, I. Nimmo-Smith, and Dipy Contributors, “Dipy, a library
752 for the analysis of diffusion MRI data,” *Front. Neuroinform.*, vol. 8, p. 8,
753 Feb. 2014.

790 [31] E. Garyfallidis, M. Brett, M. M. Correia, G. B. Williams, and I. Nimmer- 836 [44] H. Huang and M. Ding, “Linking functional connectivity and structural
791 Smith, “QuickBundles, a method for tractography simplification,” *Front. 837
792 Neurosci.*, vol. 6, p. 175, Dec. 2012. 838
793 [32] S. Geman, E. Bienenstock, and R. Doursat, “Neural networks and the
794 bias/variance dilemma,” *Neural computation*, vol. 4, no. 1, pp. 1–58,
795 1992.
796 [33] J. P. Ioannidis, “Why most published research findings are false,” *PLoS
797 medicine*, vol. 2, no. 8, p. e124, 2005.
798 [34] D. C. Van Essen, S. M. Smith, D. M. Barch, T. E. Behrens, E. Yacoub,
799 K. Ugurbil, W.-M. H. Consortium *et al.*, “The WU-Minn human connec-
800 tome project: an overview,” *Neuroimage*, vol. 80, pp. 62–79, 2013.
801 [35] X.-N. Zuo, J. S. Anderson, P. Bellec, R. M. Birn, B. B. Biswal,
802 J. Blautzik, J. C. Breitner, R. L. Buckner, V. D. Calhoun, F. X. Castel-
803 lanos *et al.*, “An open science resource for establishing reliability and
804 reproducibility in functional connectomics,” *Scientific data*, vol. 1, no. 1,
805 pp. 1–13, 2014.
806 [36] M. Jenkinson, C. F. Beckmann, T. E. J. Behrens, M. W. Woolrich, and
807 S. M. Smith, “FSL,” *Neuroimage*, vol. 62, no. 2, pp. 782–790, Aug.
808 2012.
809 [37] J. L. Lancaster, D. Tordesillas-Gutiérrez, M. Martinez, F. Salinas,
810 A. Evans, K. Zilles, J. C. Mazziotta, and P. T. Fox, “Bias between mni
811 and talairach coordinates analyzed using the icbm-152 brain template,”
812 *Human brain mapping*, vol. 28, no. 11, pp. 1194–1205, 2007.
813 [38] A. Klein and J. Tourville, “101 labeled brain images and a consistent
814 human cortical labeling protocol,” *Front. Neurosci.*, vol. 6, p. 171, Dec.
815 2012.
816 [39] D. Sohier, P. De Oliveira Castro, F. Févotte, B. Lathuilière, E. Petit, and
817 O. Jamond, “Confidence intervals for stochastic arithmetic,” Jul. 2018.
818 [40] J. Benesty, J. Chen, Y. Huang, and I. Cohen, “Pearson correlation coef-
819 ficient,” in *Noise Reduction in Speech Processing*, I. Cohen, Y. Huang,
820 J. Chen, and J. Benesty, Eds. Berlin, Heidelberg: Springer Berlin
821 Heidelberg, 2009, pp. 1–4.
822 [41] C. A. Raji, A. J. Ho, N. N. Parikhshak, J. T. Becker, O. L. Lopez, L. H.
823 Kuller, X. Hua, A. D. Leow, A. W. Toga, and P. M. Thompson, “Brain
824 structure and obesity,” *Hum. Brain Mapp.*, vol. 31, no. 3, pp. 353–364,
825 Mar. 2010.
826 [42] T. Glatard, G. Kiar, T. Aumentado-Armstrong, N. Beck, P. Bellec,
827 R. Bernard, A. Bonnet, S. T. Brown, S. Camarasu-Pop, F. Cervenansky,
828 S. Das, R. Ferreira da Silva, G. Flandin, P. Girard, K. J. Gorgolewski,
829 C. R. G. Guttmann, V. Hayot-Sasson, P.-O. Quirion, P. Rioux, M.-É.
830 Rousseau, and A. C. Evans, “Boutiques: a flexible framework to inte-
831 grate command-line applications in computing platforms,” *Gigascience*,
832 vol. 7, no. 5, May 2018.
833 [43] G. Kiar, S. T. Brown, T. Glatard, and A. C. Evans, “A serverless tool
834 for platform agnostic computational experiment management,” *Front.
835 Neuroinform.*, vol. 13, p. 12, Mar. 2019.

839

S1. Graph Correlation

840 The following presents a quantification of deviations of generated connectomes from the reference execution, similar to shown
841 in Figure 1. However, in this case, the “percent deviation” measure was replaced with the Pearson correlation coefficient.
842 The correlations between observed graphs (Figure S1) across each grouping follow the same trend to as percent deviation, as
843 shown in Figure 1. However, notably different from percent deviation, there is no significant difference in the correlations
844 between dense or sparse instrumentations. By this measure, the probabilistic pipeline is more stable in all cross-MCA and
845 cross-directions except for the combination of sparse perturbation and cross-MCA ($p < 0.0001$ for all; exploratory).

846 The marked lack in drop-off of performance across these settings, inconsistent with the measures show in Figure 1 is likely
847 due to the nature of the measure and the structure of graphs being compared. Given that structural graphs are sparse and contain
848 considerable numbers of zero-weighted edges, the presence or absense of edges dominated the correlation measure where it
849 was less impactful for the others. For this reason and others⁴⁴, correlation is not a commonly used measure in the context of
850 structural connectivity, and thus this analysis was demoted to the supplement material.

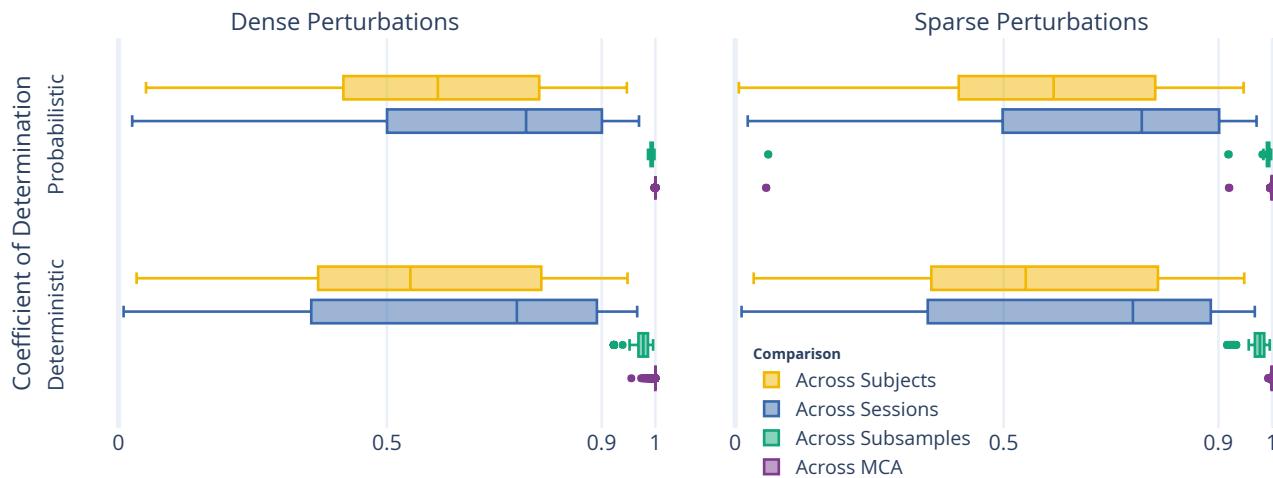


Figure S1. The correlation between perturbed connectomes and their reference.

851

S2. Complete Discriminability Analysis

Table S1. The complete results from the Discriminability analysis, with results reported as mean \pm standard deviation Discriminability. As was the case in the condensed table, the alternative hypothesis, indicating significant separation across groups, was accepted for all experiments, with $p < 0.005$.

Exp.	Subj.	Sess.	Samp.	Unscaled Reference		Dense Perturbations		Sparse Perturbations	
				Det.	Prob.	Det.	Prob.	Det.	Prob.
1.1	All	All	1	0.64 \pm 0.00	0.65 \pm 0.00	0.82 \pm 0.00	0.82 \pm 0.00	0.77 \pm 0.00	0.75 \pm 0.00
1.2	All	1	All	1.00 \pm 0.00	1.00 \pm 0.00	1.00 \pm 0.00	1.00 \pm 0.00	0.93 \pm 0.02	0.90 \pm 0.02
1.3	All	1	1			1.00 \pm 0.00	1.00 \pm 0.00	0.94 \pm 0.02	0.90 \pm 0.02
2.4	1	All	All	1.00 \pm 0.00	1.00 \pm 0.00	1.00 \pm 0.00	1.00 \pm 0.00	0.88 \pm 0.12	0.85 \pm 0.12
2.5	1	All	1			1.00 \pm 0.00	1.00 \pm 0.00	0.89 \pm 0.11	0.84 \pm 0.12
3.6	1	1	All			0.99 \pm 0.03	1.00 \pm 0.00	0.71 \pm 0.07	0.61 \pm 0.05

852 The complete discriminability analysis includes comparisons across more axes of variability than the condensed version.
853 The reduction in the main body was such that only axes which would be relevant for a typical analysis were presented. Here,
854 each of Hypothesis 1, testing the difference across subjects, and 2, testing the difference across sessions, were accompanied
855 with additional comparisons to those shown in the main body.

856 **Subject Variation** Alongside experiment 1.1, that which mimicked a typical test-retest scenario, experiments 1.2 and 1.3
857 could be considered a test-retest with a handicap, given a single acquisition per individual was compared either across
858 subsamples or simulations, respectively. For this reason, it is unsurprising that the dataset achieved considerably higher
859 discriminability scores.

860 **Session Variation** Similar to subject variation, the session variation was also modelled across either both or a single
861 subsample in experiments 2.4 and 2.5. In both of these cases the performance was similar, and the finding that sparse
862 perturbations reduced the off-target signal was consistent.

863 S2.1 Scaling of discriminability with N

864 When samples were added to the dataset across perturbed executions, the discriminability statistic inflated to a plateau even
865 when no information was added (e.g. the dataset was replicated). This effect is demonstrated for the reference executions and is
866 shown in Figure S2. As we can see, the reference discriminability scores without data duplication (unscaled) were 0.64 and
867 0.65 for the deterministic and probabilistic pipelines, respectively. After duplicating the dataset 20 times, matching the size of
868 the 20-sample perturbed dataset, we can see that this (scaled) score plateaus at 0.82 for both pipelines. For consistency, in the
869 main body of the text the reference execution performance was communicated as the scaled quantity.

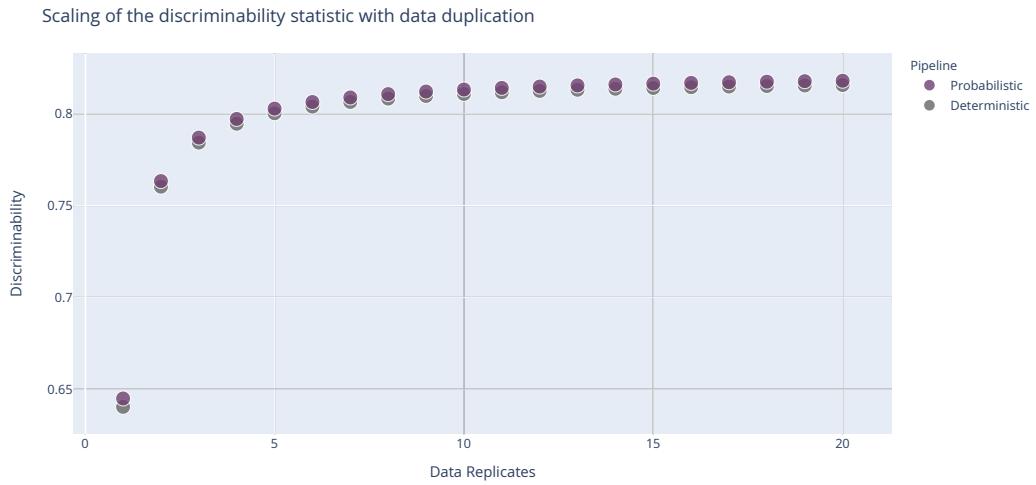


Figure S2. Scaling behaviour of the discriminability statistic with data duplication.

870

S3. Univariate Graph Statistics

871 Figure S3 explores the stability of univariate graph-theoretical metrics computed from the perturbed graphs, including modularity,
872 global efficiency, assortativity, average path length, and edge count. When aggregated across individuals and perturbations, the
873 distributions of these statistics (Figures S3A and S32B) showed no significant differences between perturbation methods for
874 either deterministic or probabilistic pipelines, consistent with the comparison of the cumulative density of the multivariate
875 statistics compared in 2.

876 However, when quantifying the stability of these measures across connectomes derived from a single session of data, the
877 two perturbation methods show considerable differences. The number of significant digits in univariate statistics for dense
878 perturbation instrumented connectome generation exceeded 11 digits for all measures except modularity, which contained more
879 than 4 significant digits of information (Figure S3C). When detecting false-positives from the distributions of observed statistics
880 for a given session, the rate (using a threshold of $p = 0.05$) was approximately 2% for all statistics with the exception of
881 modularity which again was less stable with an approximately 10% false positive rate. The probabilistic pipeline is significantly
882 more stable than the deterministic pipeline ($p < 0.0001$; exploratory) for all features except modularity. When similarly
883 evaluating these features from connectomes generated in the sparse perturbation setting, no statistic was stable with more than
884 3 significant digits or a false positive rate lower than nearly 6% (Figure S3D). The deterministic pipeline was more stable than
885 the probabilistic pipeline in this setting ($p < 0.0001$; exploratory).

886 Two notable differences between the two perturbation methods are, first, the uniformity in the stability of the statistics, and
887 second, the dramatic decline in stability of individual statistics in the sparse perturbation setting despite the consistency in the
888 overall distribution of values. This result is consistent with that obtained from the multivariate exploration performed in the
889 body of this article. It is unclear at present if the discrepancy between the stability of modularity in the pipeline perturbation
890 context versus the other statistics suggests the implementation of this measure is the source of instability or if it is implicit to
891 the measure itself. The dramatic decline in the stability of features derived from sparse perturbed graphs despite no difference
892 in their overall distribution both shows that while individual estimates may be unstable the comparison between aggregates or
893 groups may be considered much more reliable; this finding is consistent with that presented for multivariate statistics.

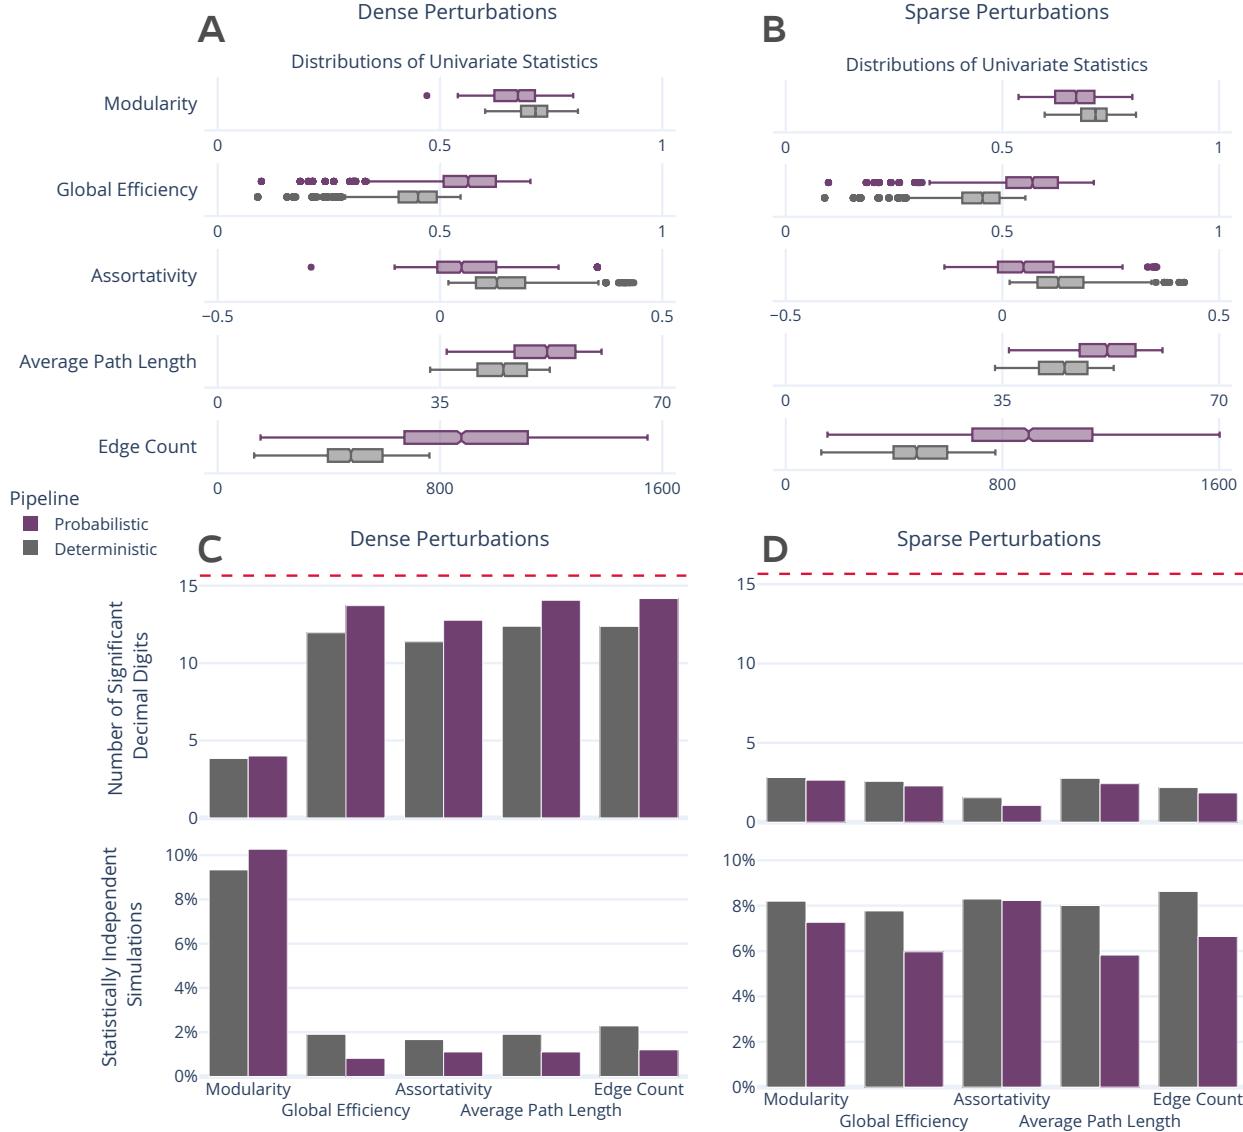


Figure S3. Distribution and stability assessment of univariate graph statistics. (A, B) The distributions of each computed univariate statistic across all subjects and perturbations for dense and sparse settings, respectively. There was no significant difference between the distributions in A and B. (C, D; top) The number of significant decimal digits in each statistic across perturbations, averaged across individuals. The dashed red line refers to the maximum possible number of significant digits. (C, D; bottom) The percentage of connectomes which were deemed significantly different ($p < 0.05$) from the others obtained for an individual.