

# GHIST 2024: The 1st Genomic History Inference Strategies Tournament

Travis J. Struck<sup>1,†</sup>, Andrew H. Vaughn<sup>2,3,†</sup>, Austin Daigle<sup>4,5</sup>, Dylan D. Ray<sup>5</sup>, Ekaterina Noskova<sup>6,7</sup>, Jaison J. Sequeira<sup>8</sup>, Svetlana Antonets<sup>9</sup>, Elizaveta Alekseevskaya<sup>10</sup>, Elizaveta Grigoreva<sup>11</sup>, Evgenii Raines<sup>12</sup>, Eilish S. McMaster<sup>13</sup>, Toby G. L. Kovacs<sup>13</sup>, Aaron P. Ragsdale<sup>14</sup>, Andrés Moreno-Estrada<sup>15</sup>, Katie E. Lotterhos<sup>16</sup>, Adam Siepel<sup>17</sup>, Ryan N. Gutenkunst<sup>1,\*</sup>

<sup>†</sup>These authors contributed equally to this work.

\*To whom correspondence should be addressed: [rgutenk@arizona.edu](mailto:rgutenk@arizona.edu).

<sup>1</sup>Department of Molecular and Cellular Biology, University of Arizona, Tucson, AZ, USA;

<sup>2</sup>Center for Computational Biology, University of California, Berkeley, CA, USA;

<sup>3</sup>Bakar Computational Health Sciences Institute, University of California, San Francisco, CA, USA;

<sup>4</sup>Curriculum in Bioinformatics and Computational Biology, University of North Carolina, Chapel Hill, NC, USA;

<sup>5</sup>Department of Genetics, University of North Carolina, Chapel Hill, NC, USA;

<sup>6</sup>Department of Biology, University of Fribourg, Fribourg, Switzerland;

<sup>7</sup>Institute of Integrative Biology, ETH Zürich, Zürich, Switzerland;

<sup>8</sup>Department of Applied Zoology, Mangalore University, Karnataka, India;

<sup>9</sup>Genotek Ltd., Moscow, Russia;

<sup>10</sup>Institute of Cytology of the Russian Academy of Sciences, St. Petersburg, Russia;

<sup>11</sup>Gregor Mendel Institute of Molecular Plant Biology, Vienna, Austria;

<sup>12</sup>ITMO University, St. Petersburg, Russia;

<sup>13</sup>School of Life and Environmental Sciences, University of Sydney, Camperdown, Australia;

<sup>14</sup>Department of Integrative Biology, University of Wisconsin-Madison, Madison, WI, USA;

<sup>15</sup>Unidad de Genómica Avanzada (UGA-LANGEBIO), Centro de Investigación y Estudios Avanzados del IPN (Cinvestav), Guanajuato, México;

<sup>16</sup>Department of Marine and Environmental Sciences, Northeastern University Marine Science Center, Nahant, Massachusetts, USA;

<sup>17</sup>Simons Center for Quantitative Biology, Cold Spring Harbor Laboratory, Cold Spring Harbor, NY, USA

Key words: population genomics, demographic history, competition.

Evaluating population genetic inference methods is challenging due to the complexity of evolutionary histories, potential model misspecification, and unconscious biases in self-assessment. The Genomic History Inference Strategies Tournament (GHIST) is a community-driven competition designed to evaluate methods for inferring evolutionary history from population genomic data. The inaugural GHIST competition ran from July to November 2024 and featured four demographic history inference challenges of varying complexity: a bottleneck model, a split with isolation model, a secondary contact model with demographic complexity, and an archaic admixture model. Data were provided as error-free VCF files, and participants submitted numerical parameter estimates that were scored by relative root mean squared error. Approximately 60 participants competed, using diverse approaches. Results revealed the current dominance of methods based on site frequency spectra, while highlighting the advantages of flexible model-building approaches for complex demographic histories. We discuss insights regarding the competition and outline the next iteration, which is ongoing with expanded challenge diversity. By providing standardized benchmarks and highlighting areas for improvement, GHIST represents a substantial step toward more reliable inference of evolutionary history from genomic data.

Population genetic inference aims to reconstruct the recent evolutionary history of populations from genomic variation data. This field has seen explosive growth, driven by the increasing availability of whole-genome sequencing data from diverse groups of humans and other species (Pool et al. 2010). But population genetic inference is inherently challenging. First, the stochasticity of the evolutionary process means that the same history can produce different genetic patterns. Second, different histories can produce similar patterns of genetic variation, creating an identifiability problem (Myers et al. 2008; Lapierre et al. 2017; Lawson et al. 2018; Rosen et al. 2018). Third, real populations rarely conform to the simplified models typically used for inference, leading to potential biases when models are misspecified (Loog 2021; Momigliano et al. 2021). Finally, computational constraints often necessitate approximations that may impact accuracy.

Many methods for population genetic inference exist. For example, site frequency spectrum (SFS) methods examine the distribution of allele frequencies within and among populations (Marth et al. 2004; Gutenkunst et al. 2009; Excoffier et al. 2013). Linkage-based approaches analyze patterns of linkage disequilibrium or identity-by-descent (IBD) segments (Harris & Nielsen 2013; Browning & Browning 2015). Markovian coalescent methods reconstruct recent genealogical relationships among samples (Li & Durbin 2011; Schiffels & Durbin 2014), while ancestral recombination graph (ARG) methods explicitly reconstruct the genealogical history including recombination events (Rasmussen et al. 2014; Kelleher et al. 2019; Speidel et al. 2019). More recently, machine learning approaches apply supervised learning to haplotype matrices or summary statistics (Schrider & Kern 2018; Flagel et al. 2019; Sanchez et al. 2021; Tran et al. 2024). Each approach captures only a portion of the information contained in genomic data, and different methods excel in different scenarios.

Papers describing new inference methods typically benchmark against existing approaches, but these self-assessments are often biased (Norel et al. 2011; Boulesteix 2015), if unconsciously. First, method developers naturally focus on scenarios where their approaches excel, potentially masking weaknesses. Second, developers have intimate knowledge of optimal parameter settings for their own methods but may use default parameters for competing methods, leading to unfair comparisons. Lastly, developers benchmarking their own tools know the ground truth

they simulated, enabling unconscious bias toward that truth. Best-practice guidelines for benchmarking studies (Boulesteix 2015; Lotterhos et al. 2022) can reduce, but not eliminate, these biases.

Independent benchmarking studies can provide more reliable conclusions than developer-driven benchmarking (Boulesteix et al. 2013), and they have been conducted in population genomics, but limitations remain. While developing a data simulation framework for the community, the *stdpopsim* project compared methods for inferring demographic history, distributions of fitness effects, and selective sweeps, although not systematically (Adrien et al. 2020; Gower et al. 2025). For demographic history inference, parametric SFS-based methods have been compared with non-parametric SFS-based (Lapierre et al. 2017) and Markovian coalescent methods (Beichman et al. 2017). The confounding effects of background selection on such inference have been studied for SFS-based and Markovian coalescent methods (Johri et al. 2021) and ARG-based methods (Marsh & Johri 2024). Brandt et al. (2022) evaluated the accuracy of ARG inference methods in estimating coalescence times, Peng et al. (2025) evaluated ARG-based methods for predicting historical polygenic scores, and Patton et al. (2019) evaluated non-parametric methods for demographic history inference under varying genome assembly quality. Although these studies have investigated many different tools, each has been carried out by a small group of authors, and their expertise in the tools tested can strongly influence benchmark results (Lotterhos et al. 2016; Weber et al. 2019). And because each of these studies is singular, it is difficult to assess progress in the field from them.

Community-based competitions have proven effective at driving innovation across multiple domains of computational biology (Meyer et al. 2011). The Critical Assessment of Protein Structure Prediction (CASP), running since 1994, is perhaps the most successful (Moult et al. 1995). By providing semi-annual blind tests of protein structure prediction methods, CASP has catalyzed remarkable improvements, culminating in the 14th competition with AlphaFold 2's breakthrough performance that revolutionized structural biology (Jumper et al. 2021; Kryshchuk et al. 2021). Similarly, challenges from DREAM (Dialogue for Reverse Engineering Assessment and Methods) have addressed diverse problems in systems biology and genomics, from gene regulatory network inference to disease prediction (Stolovitzky et al. 2007;

Marbach et al. 2012; Saez-Rodriguez et al. 2016). More recently, the Critical Assessment of Genome Interpretation focuses on predicting phenotypic consequences of genetic variants, driving improvements in variant effect prediction (Critical Assessment of Genome Interpretation Consortium 2024). PrecisionFDA challenges evaluate methods for variant calling, genome assembly, and other genomics tasks, setting standards for precision medicine applications (Olson et al. 2022). These examples illustrate the power of competition-based assessments of computational biology methods. In evolutionary inference, real-world data for which the truth is known is typically lacking (see Randall et al. (2016) for an exception), but modern simulators capture enough features of real data to provide valuable insights (Baumdicker et al. 2022; Haller & Messer 2023).

The Genomic History Inference Strategies Tournament (GHIST, pronounced /gIst) adapts the successful competition model to address the specific challenges of population genetic inference. Here, we report results from the first competition, which consisted of four challenges focused on inferring demographic history. The competition attracted many participants, demonstrated the feasibility of the model, and revealed current community practices.

## Methods

The organizing team (authors Struck and Gutenkunst) and the design committee (authors Lotterhos, Moreno-Estrada, Ralph, and Siepel) collaborated closely to develop the structure of the first GHIST competition. Although creating highly complex challenges was tempting, we prioritized accessibility to ensure early community engagement and success. We chose demographic history inference as the competition’s focus, because it is foundational to many other population genetic analyses, it allows comparison across a variety of established methods, and aligns with the organizers’ expertise. To further encourage participation, we outlined a proactive communication strategy and offered authorship on the resulting paper as recognition for the top-performing competitors.

The structure of the first GHIST competition was developed in collaboration between the organizing team (authors Struck and Gutenkunst) and design committee (authors Lotterhos, Moreno-Estrada, Ralph, and Siepel). While it was tempting to develop extremely complex challenges, accessibility was

deemed important for early success of the competition. It was decided to focus on demographic history inference, because it is foundational for other population genetics inference tasks, there are many methods to compare, and because the organizers have specific expertise. To engage community participation, a preliminary plan for communications was also developed. Lastly, to incentivize participation, it was decided that top competitors would be offered authorship on the resulting paper.

The inaugural GHIST competition consisted of four demographic history inference challenges. These were a simple bottleneck, a simple split with migration, a complex split with secondary contact, and a complex archaic admixture scenario. Competitors could submit to any challenge(s) they chose, in any order. The scenarios were parameterized such that existing methods were expected to have good statistical power and sample sizes were set to be similar to contemporary non-human data sets. For all four challenges, the data were simulated using the Wright-Fisher coalescent method msprime (Baumdicker et al. 2022) and distributed as error-free Variant Call Format (VCF) files (Danecek et al. 2011), with only biallelic sites and correct ancestral states provided. To minimize complexity, mutation and recombination rates were uniform across the simulated regions, and selection was absent.

For each challenge, competitors reported estimates for a small number of key population genetic parameters, such as population sizes, divergence times, or admixture proportions. They were told the total size of the simulated region and the true simulated mutation and recombination rates. Entries were scored based on the relative root-mean-squared error between submitted  $\hat{\theta}$  and true parameter values  $\theta$ :

$$RRMSE = \sqrt{\sum_i \left( \frac{\hat{\theta}_i - \theta_i}{\theta_i} \right)^2}. \quad (1)$$

This interpretable metric allowed comparison across parameters of different scales and penalized both over- and under-estimation equally. For each challenge, the leaderboard was ranked based on RRMSE scores, with lower scores indicating better performance. To allow methodological refinement, competitors were allowed five submissions for each challenge. In addition to their inferences, competitors were asked to submit a brief write-up of their approach, including software tools used and the logical flow of their analyses. The scripts for generating the




Figure 1: Timeline of the first GHIST competition, including notable promotion events.

data and scoring submissions are available at <https://github.com/tjstruck/GHIST-2024-paper>.

The competition was hosted on the Synapse platform developed by Sage Bionetworks, a not-for-profit organization that promotes open science and collaborative research. Synapse provided automated handling of competitor submissions, including timestamps, versioning, validation, and real-time leaderboards. The integrated wiki functionality was used for competition documentation and tutorials, and discussion boards enabled competitors to ask questions of the organizers. The Synapse site for the first GHIST competition is available at <https://synapse.org/Synapse:syn51614781>, and the main GHIST website is at <https://ghi.st>.

The inaugural GHIST competition ran from July to November 2024, to span the summer conference season and the beginning of the academic term (Fig. 1). It began with a kickoff workshop at the Society for Molecular Biology and Evolution (SMBE) conference in Puerto Vallarta, Mexico, where participants were introduced to the competition, analyzed data from the Bottleneck challenge using *dadi-cli* (Huang et al. 2023), and submitted their inferences. The competition extended into the academic term to enable new students to participate as a training opportunity. The competition was promoted in-person at the SMBE and Evolution conferences, through posts to the Evoldir, *dadi* user, and *fastsimcoal* user mailing lists, and through targeted emails to specific investigators known to the organizers. It was also promoted through posts on X and Bluesky by the organizers and SMBE.

## Results

The inaugural GHIST competition attracted approximately 60 participants spanning career stages from graduate students to senior faculty. Participation varied across challenges, with more entries for the simpler challenges. Competitors employed a variety of approaches, with top competitors mostly relying on the site frequency spectrum (SFS). A variety of software was employed, including custom pipelines.

### Bottleneck challenge

The first challenge involved a simple bottleneck (Fig. 2A), with competitors inferring the timing and magnitude of the population decline. Competitors were given 100 megabases (Mb) of data from 20 diploid individuals, yielding 219 thousand biallelic variants.

Submissions for the bottleneck challenge showed a range of strategies and accuracy. The RRMSE values of submissions spanned orders of magnitude (Fig. 2B). Almost all submissions successfully identified the presence of a bottleneck (Fig. 2C), but only a few were highly accurate.

The most accurate submissions for the Bottleneck challenge used site frequency spectrum (SFS) approaches (Fig. 2D). Competitor Vaughn developed a custom approach using an analytic result for piecewise constant demographic histories to calculate the expected model SFS (DeWitt et al. 2021) and the Kullback-Leibler (KL) divergence to measure differences between model and data spectra. Competitors McMaster and Kovacs used the SFS-based methods *dadi-cli* (Huang et al. 2023) and *fastsimcoal2* (Excoffier et al. 2021) and the Markovian coalescent tool *SMC++* (Terhorst et al. 2017) for their submissions. Competitor Noskova led a team using her GADMA (Noskova et al. 2020, 2023) framework, using the *dadi* (Gutenkunst et al. 2009), moments (Jouganous et al. 2017), and *moments* (Kamm et al. 2020) engines for calculating model spectra.

A surprise was that top competitor Vaughn metagamed the challenge by using the leaderboard to optimize his submissions. He made an excellent first submission (Fig. 2D) based on the provided data, but this would not have been enough to win the challenge. To improve his result, he correctly deduced that the challenge simulation used round parameter values, and he used his remaining four submissions to search through the parameter space using the leaderboard RRMSE score to converge on nearly the exact

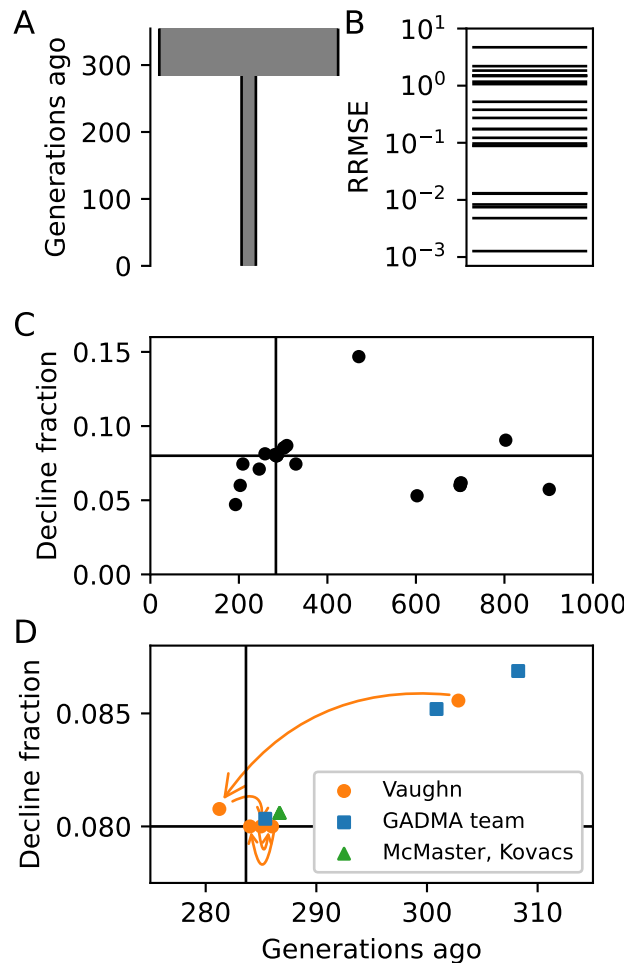


Figure 2: GHIST 2024 Bottleneck challenge. A) True simulated demographic history. B) Relative root mean square error scores of submissions. C) Parameter inferences of majority of submissions. True values are indicated by solid lines. D) Parameter inferences zoomed close to true values to indicate top competitors. Arrows indicate competitor Vaughn's leaderboard optimization procedure.

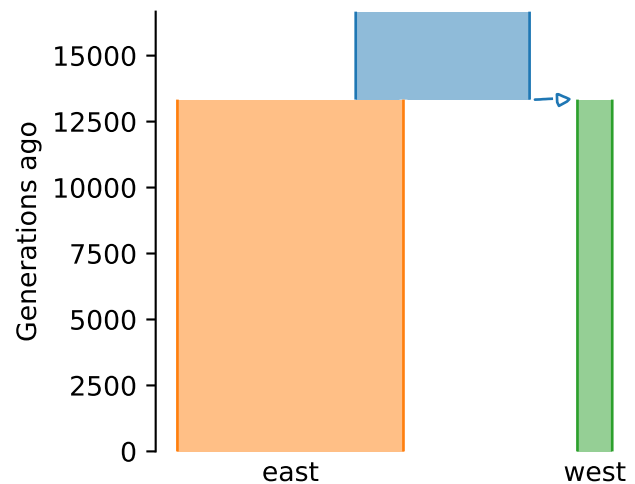


Figure 3: True demographic history for the Split with Isolation challenge, represented using demesdraw (Gower et al. 2022).

values.

### Split with Isolation challenge

The second challenge involved two populations that diverged without subsequent gene flow, representing geographic isolation (Fig. 3), with competitors inferring the contemporary population sizes and the timing of the split. They were given 100 Mb of data from 22 and 18 individuals from the two populations, yielding 1.2 million biallelic variants.

Performance on this challenge was generally strong, with several competitors achieving high accuracy for all three parameters (Table 1). SFS-based methods performed well in this challenge, with competitor Vaughn using *tskit* (Kelleher et al. 2016; Wong et al. 2024) to calculate expected spectra and KL divergence to fit the model and McMaster and Kovacs using *dadi-cli* for inference. The Team of Daigle and Ray used a machine learning approach. They first used *dadi* to identify the relevant ranges of parameter values, then simulated data over those ranges with *msprime*, and then used *scikit-allel* (Miles et al. 2024) and *pylibseq* (Thornton 2003) to calculate summary statistics, including statistics based on the SFS, haplotypes, and LD decay. These summary statistics were then passed to a multi-layer perceptron for inference. However, their best-scoring submission for this challenge simply employed *dadi*. As in the first Challenge, competitor Vaughn achieved the top score by strategically rounding his inferences.



| $RRMSE$ | $N_{east}$ | $N_{west}$ | $T$   | Competitor       | Approach   |
|---------|------------|------------|-------|------------------|------------|
| truth   | 130000     | 20000      | 13333 |                  |            |
| 0.005   | 130000     | 20100      | 13300 | Vaughn           | metagaming |
| 0.027   | 126686     | 19973      | 13233 | McMaster, Kovacs | dadi-cli   |
| 0.029   | 126941     | 19867      | 13119 | Daigle, Ray      | dadi       |
| 0.031   | 125992     | 20109      | 13301 | Vaughn           | tskit SFS  |

Table 1: Top submissions for the Split with Isolation challenge.

## Secondary Contact challenge

The third challenge involved secondary contact between isolated populations, with complexity in population size histories that no parametric model was expected to capture (Fig. 4A). Competitors were tasked with inferring the contemporary population sizes, timing of the split and recontact, and the rate of migration after recontact. Competitors were again given 100 Mb of data, from 22 diploid mainland individuals and 8 island individuals, for a total of 842 thousand biallelic sites.

As expected, this challenge was more difficult than the previous two, with no submission accurately estimating all parameters (Table 2). The team of Daigle and Ray did well with their machine learning approach based on summary statistics. Competitor Vaughn’s top submissions were all based on leaderboard optimization after his initial inference. All these submissions assumed simple constant population size histories, like the truth in the Split with Isolation challenge. The best performance came from the GADMA team, using the moments engine (Table 2). GADMA automatically builds and refines models of increasing complexity, and their best model allowed for growth in both populations (Fig. 4B), perhaps enabling their model to account for some of the effects of the true complex population size changes.

## Archaic Admixture challenge

To probe a distinct but related form of inference, the final challenge involved archaic admixture. Competitors were tasked with inferring the timing and magnitude of admixture into two modern populations (Fig. 5). They were given 250 Mb of data from 20 and 16 samples for the modern populations, along with 1 to 3 samples from each of the potential archaic contributors, sampled 17,500 to 100,000 simulated years ago, for a total of 1.7 million biallelic sites.

The top competitors accurately estimated admixture proportions but were less accurate when estimat-




Figure 4: GHIST 2024 Secondary Contact challenge. A: True demographic history, represented using demesdraw (Gower et al. 2022). B: Top-scoring model, from the GADMA team.

| $RRMSE$ | $N_{\text{main}}$ | $N_{\text{island}}$ | $T_{\text{split}}$ | $T_{\text{mig}}$ | $m$ | Competitor  | Approach                 |
|---------|-------------------|---------------------|--------------------|------------------|-----|-------------|--------------------------|
| truth   | 240000            | 36000               | 23000              | 1277             | 5.0 |             |                          |
| 0.92    | 284284            | 15567               | 24397              | 950              | 8.2 | GADMA team  | GADMA w/ moments         |
| 0.97    | 150000            | 35000               | 30000              | 200              | 5.0 | Vaughn      | metagaming               |
| 1.05    | 200000            | 12000               | 18000              | 310              | 5.0 | Vaughn      | metagaming               |
| 1.05    | 220000            | 60000               | 17000              | 300              | 5.0 | Vaughn      | metagaming               |
| 1.30    | 211999            | 10945               | 15818              | 215              | 1.8 | Daigle, Ray | summary stats perceptron |

Table 2: Top submissions for Secondary Contact challenge.

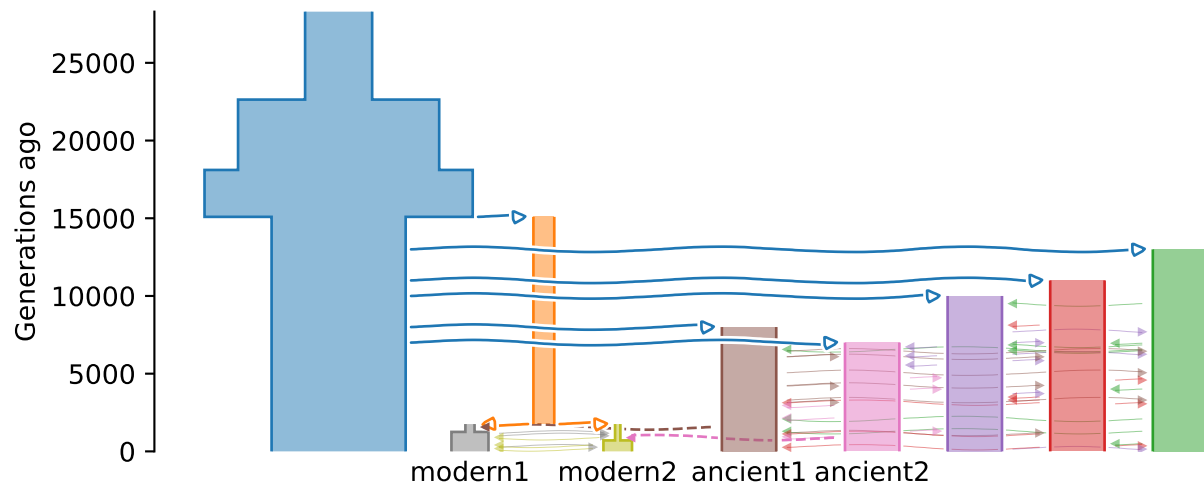


Figure 5: True demographic history for Archaic Admixture challenge, represented using demesdraw (Gower et al. 2022).

| $RRMSE$ | %admix <sub>1</sub> | $T_{\text{admix}_1}$ | %admix <sub>2</sub> | $T_{\text{admix}_2}$ | Competitor | Approach   |
|---------|---------------------|----------------------|---------------------|----------------------|------------|------------|
| truth   | 1.10                | 1566                 | 0.20                | 883                  |            |            |
| 0.10    | 1.00                | 1500                 | 0.20                | 900                  | Vaughn     | metagaming |
| 0.78    | 1.01                | 1096                 | 0.19                | 252                  | GADMA team | mom2       |
| 0.81    | 1.02                | 1066                 | 0.16                | 250                  | GADMA team | mom2       |
| 1.05    | 0.59                | 1000                 | 0.37                | 1000                 | GADMA team | mom2       |
| 1.93    | 1.06                | 2500                 | 0.22                | 2500                 | Vaughn     | tskit SFS  |

Table 3: Top submissions for Archaic Admixture challenge.

ing timings (Table 3). For each modern population, competitor Vaughn used *tskit* simulations to simulate a two-population model with archaic admixture from a ghost population and fit that to the SFS from the modern population. He then optimized the leaderboard to refine his estimates. The GADMA software does not support ancient samples, so it could not be applied to this challenge. But the GADMA team used *momi2* (Kamm et al. 2020) directly to fit models involving all five sample groups, achieving superior accuracy before metagaming.

## Discussion

The inaugural GHIST competition demonstrated the feasibility and value of a community-driven evaluation framework for population genetic inference methods. The Synapse platform proved robust and capable, and the range of challenges enabled accessibility while pushing the limits of existing inference methods. The conference-based launch and extended timeframe facilitated participation from diverse researchers, including students.

The GHIST competition provided several insights into the relative performance of inference approaches. Approaches based on the site frequency spectrum were most common and successful, because they are both accessible from established software tools and powerful for demographic inference. For the Bottleneck, Split with Isolation, and Secondary Contact challenges, the GADMA team directly compared SFS-based engines with the *moments.LD* engine that uses multi-population linkage disequilibrium statistics (Ragsdale & Gravel 2019, 2020), achieving better scores with SFS-based engines. Approaches based on machine learning showed promise but were not widely used by competitors. As those approaches become more accessible, we expect their representation and success to increase. Almost all approaches applied assumed prespecified parametric models, which may not capture the complexity of real demographic histories (Loog 2021). The exception was GADMA, and its success in the Secondary Contact challenge (Fig. 4A), which was designed to violate typical prespecification, highlights the importance of model flexibility when dealing with complex histories.

There were notable gaps in the methods employed by participants. Approaches based on ancestral recombination graphs show great promise for population genetics (Rasmussen et al. 2014; Kelleher et al. 2019; Speidel et al. 2019; Deng et al. 2024), but they

were not applied to this competition, perhaps because of their high computational cost or complexity. The Archaic Admixture challenge (Fig. 5) was designed to encourage the use of specialized methods based on lengths of admixture tracts (Pool & Nielsen 2009; Gravel 2012), but no competitors used them, perhaps due to insufficient outreach to the relevant subset of researchers. Methods for demographic history inference based on the Markovian coalescent (Li & Durbin 2011; Schiffels & Durbin 2014) that don't depend on a user-specified parametric model were also underrepresented relative to their popularity in the literature.

The second GHIST competition launched at the Evolution conference in June 2025 in Athens, Georgia and runs through November 2025, with expanded challenge types and refinements based on lessons from the inaugural tournament. To discourage metagaming while preserving the benefits of iterative submission, for each challenge multiple submissions are allowed on a testing data set, but only a single submission is allowed on the final data set. To increase the realism and difficulty of demographic history inference, two of the challenges include background selection, leveraging the *stdpopsim* framework for simulation (Gower et al. 2025). To expand the range of tasks, four challenges involve inferring single or multiple hard selective sweeps (Stephan 2019), under simple and complex demographic scenarios and with and without background selection. The Synapse site for this second competition is at <https://synapse.org/Synapse:syn65877330>.

Ultimately, the success of GHIST depends on community participation. The more methods developers, users, educators, and students engage with the competitions, the more the community will learn. The space of potential challenges is vast, including inferences such as distributions of fitness effects (Eyre-Walker & Keightley 2007), spatial models (Bradburd & Ralph 2019), and polygenic selection (Barghi et al. 2020), and including complications such as low-pass data (Crawford & Lazzaro 2012), polyploidy (Dufresne et al. 2014), and biased gene conversion (Pouyet et al. 2018). Within Synapse, submissions are scored using custom code executed on a cloud instance, so more complex evaluation metrics are possible. This initial competition has demonstrated the feasibility and utility of competitions for this community. Future competitions will enable deep insight into best practices for population genetics inference.



# Acknowledgments

This work was supported by the National Institute of General Medical Sciences of the National Institutes of Health (R35GM149235 to R.N.G.). A.D. received support from NIGMS predoctoral training grant 5T32 GM067553. We thank Pablo Meyer Rojas for introducing us to the Synapse platform and the Society for Molecular Biology and Evolution for hosting the kickoff workshop and promoting the competition on their social media channels. We thank all competitors for their participation.

This paper was written with the assistance of generative artificial intelligence (AI). MacWhisper was used to transcribe the audio from author Gutenkunst's talk about GHIST at the 2024 Probabilistic Modeling in Genomics conference. Anthropic's Claude 4 Sonnet model was then given that transcript and asked to generate a detailed outline for the paper, including leveraging its existing knowledge base, within a Project containing previous papers by author Gutenkunst. Author Gutenkunst also used Claude's research mode to generate a report on published papers that independently compared population genetics inference approaches, which yielded a few studies he was previously unaware of. Author Gutenkunst edited or generated all text in the final manuscript and verified all references.

# References

Adrión JR, Cole CB, Dukler N, Galloway JG, Gladstein AL, Gower G, Kyriazis CC, Ragsdale AP, Tsambos G, Baumdicker F, Carlson J, Cartwright RA, Durvasula A, Gronau I, Kim BY, McKenzie P, Messer PW, Noskova E, Ortega-Del Vecchyo D, Racimo F, Struck TJ, Gravel S, Gutenkunst RN, Lohmueller KE, Ralph PL, Schrider DR, Siepel A, Kelleher J, Kern AD (2020) A community-maintained standard library of population genetic models. *eLife* 9:e54967.

Barghi N, Hermisson J, Schlötterer C (2020) Polygenic adaptation: a unifying framework to understand positive selection. *Nature Reviews Genetics* 21:769.

Baumdicker F, Bisschop G, Goldstein D, Gower G, Ragsdale AP, Tsambos G, Zhu S, Eldon B, Ellerman EC, Galloway JG, Gladstein AL, Gorjanc G, Guo B, Jeffery B, Kretzschmar WW, Lohse

K, Matschiner M, Nelson D, Pope NS, Quinto-Cortés CD, Rodrigues MF, Saunack K, Sellinger T, Thornton K, Van Kemenade H, Wohns AW, Wong Y, Gravel S, Kern AD, Koskela J, Ralph PL, Kelleher J (2022) Efficient ancestry and mutation simulation with msprime 1.0. *Genetics* 220:iyab229.

Beichman AC, Phung TN, Lohmueller KE (2017) Comparison of single genome and allele frequency data reveals discordant demographic histories. *G3* 7:3605.

Boulesteix AL (2015) Ten Simple Rules for Reducing Overoptimistic Reporting in Methodological Computational Research. *PLOS Computational Biology* 11:e1004191.

Boulesteix AL, Lauer S, Eugster MJA (2013) A Plea for Neutral Comparison Studies in Computational Sciences. *PLoS ONE* 8:e61562.

Bradburd GS, Ralph PL (2019) Spatial Population Genetics: It's About Time. *Annual Review of Ecology, Evolution, and Systematics* 50:427.

Brandt DY, Wei X, Deng Y, Vaughn AH, Nielsen R (2022) Evaluation of methods for estimating coalescence times using ancestral recombination graphs. *Genetics* 221:iyac044.

Browning SR, Browning BL (2015) Accurate non-parametric estimation of recent effective population size from segments of identity by descent. *American Journal of Human Genetics* 97:404.

Crawford JE, Lazzaro BP (2012) Assessing the Accuracy and Power of Population Genetic Inference from Low-Pass Next-Generation Sequencing Data. *Frontiers in Genetics* 3:66.

Critical Assessment of Genome Interpretation Consortium (2024) CAGI, the critical assessment of genome interpretation, establishes progress and prospects for computational genetic variant interpretation methods. *Genome Biology* 25:53.

Danecek P, Auton A, Abecasis G, Albers CA, Banks E, DePristo MA, Handsaker RE, Lunter G, Marth GT, Sherry ST, McVean G, Durbin R, 1000 Genomes Project Analysis Group (2011) The variant call format and VCFtools. *Bioinformatics* 27:2156.

Deng Y, Nielsen R, Song YS (2024) Robust and Accurate Bayesian Inference of

- Genome-Wide Genealogies for Large Samples. <https://www.biorxiv.org/content/10.1101/2024.03.16.585311>.
- DeWitt WS, Harris KD, Ragsdale AP, Harris K (2021) Nonparametric coalescent inference of mutation spectrum history and demography. *Proceedings of the National Academy of Sciences* 118:e2013798118.
- Dufresne F, Stift M, Vergilino R, Mable BK (2014) Recent progress and challenges in population genetics of polyploid organisms: an overview of current state-of-the-art molecular and statistical tools. *Molecular Ecology* 23:40.
- Excoffier L, Dupanloup I, Huerta-Sánchez E, Sousa VC, Foll M (2013) Robust demographic inference from genomic and SNP data. *PLoS genetics* 9:e1003905.
- Excoffier L, Marchi N, Marques DA, Matthey-Doret R, Gouy A, Sousa VC (2021) *fastsimcoal2*: demographic inference under complex evolutionary scenarios. *Bioinformatics* 37:4882.
- Eyre-Walker A, Keightley PD (2007) The distribution of fitness effects of new mutations. *Nature Reviews Genetics* 8:610.
- Flagel L, Brandvain Y, Schrider DR (2019) The unreasonable effectiveness of convolutional neural networks in population genetic inference. *Molecular Biology and Evolution* 36:220.
- Gower G, Pope NS, Rodrigues MF, Tittes S, Tran LN, Alam O, Cavassim MIA, Fields PD, Haller BC, Huang X, Jeffrey B, Korfmann K, Kyriazis CC, Min J, Rebollo I, Rehmann CT, Small ST, Smith CCR, Tsambos G, Wong Y, Zhang Y, Huber CD, Gorjanc G, Ragsdale AP, Gronau I, Gutenkunst RN, Kelleher J, Lohmueller KE, Schrider DR, Ralph PL, Kern AD (2025) Accessible, realistic genome simulation with selection using stdpopsim. <https://www.biorxiv.org/content/10.1101/2025.03.23.644825>.
- Gower G, Ragsdale AP, Bisschop G, Gutenkunst RN, Hartfield M, Noskova E, Schiffels S, Struck TJ, Kelleher J, Thornton KR (2022) Demes: a standard format for demographic models. *Genetics* 222:iyac131.
- Gravel S (2012) Population Genetics Models of Local Ancestry. *Genetics* 191:607.
- Gutenkunst RN, Hernandez RD, Williamson SH, Sankaranarayanan CD (2009) Inferring the Joint Demographic History of Multiple Populations from Multidimensional SNP Frequency Data. *PLoS Genetics* 5:e1000695.
- Haller BC, Messer PW (2023) SLiM 4: Multispecies Eco-Evolutionary Modeling. *The American Naturalist* 201:E127.
- Harris K, Nielsen R (2013) Inferring Demographic History from a Spectrum of Shared Haplotype Lengths. *PLoS Genetics* 9:e1003521.
- Huang X, Struck TJ, Davey SW, Gutenkunst RN (2023) dadi-cli: Automated and distributed population genetic model inference from allele frequency spectra.
- Johri P, Riall K, Becher H, Excoffier L, Charlesworth B, Jensen JD (2021) The Impact of Purifying and Background Selection on the Inference of Population History: Problems and Prospects. *Molecular Biology and Evolution* 38:2986.
- Jouganous J, Long W, Ragsdale AP, Gravel S (2017) Inferring the joint demographic history of multiple populations: beyond the diffusion approximation. *Genetics* 206:1549.
- Jumper J, Evans R, Pritzel A, Green T, Figurnov M, Ronneberger O, Tunyasuvunakool K, Bates R, Židek A, Potapenko A, Bridgland A, Meyer C, Kohl SAA, Ballard AJ, Cowie A, Romera-Paredes B, Nikolov S, Jain R, Adler J, Back T, Petersen S, Reiman D, Clancy E, Zielinski M, Steinegger M, Pacholska M, Berghammer T, Bodenstein S, Silver D, Vinyals O, Senior AW, Kavukcuoglu K, Kohli P, Hassabis D (2021) Highly accurate protein structure prediction with AlphaFold. *Nature* 596:583.
- Kamm J, Terhorst J, Durbin R, Song YS (2020) Efficiently Inferring the Demographic History of Many Populations With Allele Count Data. *Journal of the American Statistical Association* 115:1472.
- Publisher: Informa UK Limited.
- Kelleher J, Etheridge AM, McVean G (2016) Efficient Coalescent Simulation and Genealogical Analysis for Large Sample Sizes. *PLOS Computational Biology* 12:e1004842.
- Kelleher J, Wong Y, Wohns AW, Fadil C, Albers PK, McVean G (2019) Inferring whole-genome histories in large population datasets. *Nature Genetics* 51:1330.

- Kryshtafovych A, Schwede T, Topf M, Fidelis K, Mout J (2021) Critical assessment of methods of protein structure prediction (CASP)—Round XIV. *Proteins: Structure, Function, and Bioinformatics* 89:1607.
- Lapierre M, Lambert A, Achaz G (2017) Accuracy of Demographic Inferences from the Site Frequency Spectrum: The Case of the Yoruba Population. *Genetics* 206:439.
- Lawson DJ, Van Dorp L, Falush D (2018) A tutorial on how not to over-interpret STRUCTURE and ADMIXTURE bar plots. *Nature Communications* 9:3258.
- Li H, Durbin R (2011) Inference of human population history from individual whole-genome sequences. *Nature* 475:493.
- Loog L (2021) Sometimes hidden but always there: the assumptions underlying genetic inference of demographic histories. *Philosophical Transactions of the Royal Society B* 376:20190719.
- Lotterhos KE, Fitzpatrick MC, Blackmon H (2022) Simulation Tests of Methods in Evolution, Ecology, and Systematics: Pitfalls, Progress, and Principles. *Annual Review of Ecology, Evolution, and Systematics* 53:113.
- Lotterhos KE, François O, Blum MG (2016) Not just methods: User expertise explains the variability of outcomes of genome-wide studies. *bioRxiv* : the preprint server for biology <https://www.biorxiv.org/content/10.1101/055046v1>.
- Marbach D, Costello JC, Küffner R, Vega NM, Prill RJ, Camacho DM, Allison KR, DREAM5 Consortium (2012) Wisdom of crowds for robust gene network inference. *Nature Methods* 9:796.
- Marsh JJ, Johri P (2024) Biases in ARG-based inference of historical population size in populations experiencing selection. *Molecular Biology and Evolution* 41:msae118.
- Marth GT, Czabarka E, Murvai J, Sherry ST (2004) The allele frequency spectrum in genome-wide human variation data reveals signals of differential demographic history in three large world populations. *Genetics* 166:351.
- Meyer P, Alexopoulos LG, Bonk T, Califano A, Cho CR, De La Fuente A, De Graaf D, Hartemink AJ, Hoeng J, Ivanov NV, Koepl H, Linding R, Marbach D, Norel R, Peitsch MC, Rice JJ, Royyuru A, Schacherer F, Sprengel J, Stolle K, Vitkup D, Stolovitzky G (2011) Verification of systems biology research in the age of collaborative competition. *Nature Biotechnology* 29:811.
- Miles A, bot p, Rodrigues MF, Ralph P, Kelleher J, Schelker M, Pisupati R, Rae S, Millar T (2024) scikit-allele: v1.3.13.
- Momigliano P, Florin AB, Merilä J (2021) Biases in demographic modelling affect our understanding of recent divergence. *Molecular Biology and Evolution* 38:msab047.
- Mout J, Pedersen JT, Judson R, Fidelis K (1995) A large-scale experiment to assess protein structure prediction methods. *Proteins: Structure, Function, and Bioinformatics* 23:ii.
- Myers S, Fefferman C, Patterson N (2008) Can one learn history from the allelic spectrum? *Theoretical Population Biology* 73:342.
- Norel R, Rice JJ, Stolovitzky G (2011) The self-assessment trap: can we all be better than average? *Molecular Systems Biology* 7:537.
- Noskova E, Abramov N, Iliutkin S, Sidorin A, Dobrynin P, Ulyantsev VI (2023) GADMA2: more efficient and flexible demographic inference from genetic data. *GigaScience* 12:giad059.
- Noskova E, Ulyantsev V, Koepfli KP, O'Brien SJ, Dobrynin P (2020) GADMA: Genetic algorithm for inferring demographic history of multiple populations from allele frequency spectrum data. *GigaScience* 9:giaa005.
- Olson ND, Wagner J, McDaniel J, Stephens SH, Westreich ST, Prasanna AG, Johanson E, Boja E, Maier EJ, Serang O, Jáspez D, Lorenzo-Salazar JM, Muñoz-Barrera A, Rubio-Rodríguez LA, Flores C, Kyriakidis K, Medina-Gómez C, Rivadeneira F, Uitterlinden AG, Ikram MA, van Rooij JGJ, Kanai M, Okada Y, Takahashi A, Momozawa Y, Hirata M, Matsuda K, Murakami Y, Kubo M, Kamatani Y, Li Y, Tam V, Shrine N, Guyatt AL, John C, Shrine N, Guyatt AL, John C, Packer R, Wain LV, Tobin MD, Sayers I, Hall IP, Brightling C, Singapuri A, Soler Artigas M, Obeidat M, Sin DD, Bossé Y, Postma DS, Vonk JM, Boezen HM, Shrine N, Guyatt AL, John C, Packer R, Brightling

- C, Singapuri A, Soler Artigas M, Hall IP, Wain LV, Tobin MD, Sayers I, Obeidat M, Sin DD, Bossé Y, Postma DS, Vonk JM, Boezen HM, Zeberg H (2022) precisionFDA truth challenge V2: Calling variants from short- and long-reads in difficult-to-map regions. *Cell Genomics* 2:100129.
- Patton AH, Margres MJ, Stahlke AR, Hendricks S, Lewallen K, Hamede RK, Ruiz-Aravena M, Ryder O, McCallum HI, Jones ME, Hohenlohe PA, Storer A (2019) Contemporary Demographic Reconstruction Methods Are Robust to Genome Assembly Quality: A Case Study in Tasmanian Devils. *Molecular Biology and Evolution* 36:2906.
- Peng D, Mulder OJ, Edge MD (2025) Evaluating ARG-estimation methods in the context of estimating population-mean polygenic score histories. *Genetics* 229:iyaf033.
- Pool JE, Hellmann I, Jensen JD, Nielsen R (2010) Population genetic inference from genomic sequence variation. *Genome Research* 20:291.
- Pool JE, Nielsen R (2009) Inference of Historical Changes in Migration Rate From the Lengths of Migrant Tracts. *Genetics* 181:711.
- Pouyet F, Aeschbacher S, Thiéry A, Excoffier L (2018) Background selection and biased gene conversion affect more than 95% of the human genome and bias demographic inferences. *eLife* 7.
- Ragsdale AP, Gravel S (2019) Models of archaic admixture and recent history from two-locus statistics. *PLOS Genetics* 15:e1008204.
- Ragsdale AP, Gravel S (2020) Unbiased Estimation of Linkage Disequilibrium from Unphased Data. *Molecular Biology and Evolution* 37:923.
- Randall RN, Radford CE, Roof KA, Natarajan DK, Gaucher EA (2016) An experimental phylogeny to benchmark ancestral sequence reconstruction. *Nature Communications* 7:12847.
- Rasmussen MD, Hubisz MJ, Gronau I, Siepel A (2014) Genome-Wide Inference of Ancestral Recombination Graphs. *PLoS Genetics* 10:e1004342.
- Rosen Z, Bhaskar A, Roch S, Song YS (2018) Geometry of the sample frequency spectrum and the perils of demographic inference. *Genetics* 210:665.
- Saez-Rodriguez J, Costello JC, Friend SH, Kellen MR, Mangravite L, Meyer P, Norman T, Stolovitzky G (2016) Crowdsourcing biomedical research: leveraging communities as innovation engines. *Nature Reviews Genetics* 17:470.
- Sanchez T, Cury J, Charpiat G, Jay F (2021) Deep learning for population size history inference: design, comparison and combination with approximate Bayesian computation. *Molecular Ecology Resources* 21:2645.
- Schiffels S, Durbin R (2014) Inferring human population size and separation history from multiple genome sequences. *Nature Genetics* 46:919.
- Schrider DR, Kern AD (2018) Supervised machine learning for population genetics: a new paradigm. *Trends in Genetics* 34:301.
- Speidel L, Forest M, Shi S, Myers SR (2019) A method for genome-wide genealogy estimation for thousands of samples. *Nature Genetics* 51:1321.
- Stephan W (2019) Selective Sweeps. *Genetics* 211:5.
- Stolovitzky G, Monroe D, Califano A (2007) Dialogue on reverse-engineering assessment and methods: the DREAM of high-throughput pathway inference. *Annals of the New York Academy of Sciences* 1115:1.
- Terhorst J, Kamm JA, Song YS (2017) Robust and scalable inference of population history from hundreds of unphased whole genomes. *Nature Genetics* 49:303.
- Thornton K (2003) libsequence: a C++ class library for evolutionary genetic analysis. *Bioinformatics* 19:2325.
- Tran LN, Sun CK, Struck TJ, Sajan M, Gutenkunst RN (2024) Computationally efficient demographic history inference from allele frequencies with supervised machine learning. *Molecular Biology and Evolution* 41:msae077.
- Weber LM, Saelens W, Cannoodt R, Soneson C, Hapfelmeier A, Gardner PP, Boulesteix AL, Saeys Y, Robinson MD (2019) Essential guidelines for computational method benchmarking. *Genome Biology* 20:125.
- Wong Y, Ignatieva A, Koskela J, Gorjanc G, Wohns AW, Kelleher J (2024) A general and efficient representation of ancestral recombination graphs. *Genetics* 228:iyae100.