DEEP LEARNING FROM PHYLOGENIES TO UNCOVER THE TRANSMISSION DYNAMICS OF EPIDEMICS

AUTHORS

Voznica J^{1,2,3,4*}, Zhukova A^{1,2,5*}, Boskova V⁶, Saulnier E^{1,2}, Lemoine F^{1,2,5}, Moslonka-Lefebvre M^{1,2}, Gascuel O^{1,2¶}

AFFILIATIONS

¹Unité de Bioinformatique Évolutive - Département Biologie computationnelle, Institut Pasteur, Paris, FRANCE

²Unité de Biologie Computationnelle, USR 3756 CNRS, Paris, FRANCE

³Université de Paris, Paris, FRANCE

⁴Institut de Biologie de l'École Normale Supérieure, Ecole Normale Supérieure, CNRS, INSERM, Université Paris Sciences et Lettres, 75005 Paris, FRANCE

⁵Hub de Bioinformatique et Biostatistique - Département Biologie computationnelle, Institut Pasteur, Paris, FRANCE

⁶Center for Integrative Bioinformatics Vienna, Max Perutz Labs, University of Vienna and Medical University of Vienna, Vienna, AUSTRIA

¶ OG: current address

* CO-CORRESPONDING AUTHORS

jakub.voznica@pasteur.fr (JV), anna.zhukova@pasteur.fr (AZ), gascuelolivier@gmail.com (OG)

ABSTRACT

Widely applicable, accurate and fast inference methods in phylodynamics are needed to fully profit from the richness of genetic data in uncovering the dynamics of epidemics. Standard methods, including maximum-likelihood and Bayesian approaches, which are both model specific, often rely on complex mathematical formulae and approximations, and do not scale well with dataset size. We develop a likelihood-free, simulation-based approach, which combines deep learning with (1) a large set of summary statistics measured on phylogenies or (2) a complete and compact vectorial representation of trees, which avoids potential limitations of summary statistics and applies to any phylodynamic model. Our method enables both model selection and estimation of epidemiological parameters. We demonstrate its speed and accuracy on simulated data, where it performs better than the state-of-the-art methods. To illustrate its applicability, we assess the dynamics induced by superspreading individuals in an HIV dataset of men-having-sex-with-men in Zurich.

KEYWORDS

deep learning, phylodynamics, molecular epidemiology, phylogenetic tree representation, summary statistics, HIV, computer simulations, birth-death models, superspreading

INTRODUCTION

Pathogen phylodynamics is a field combining phylogenetics and epidemiology^[1]. Viral or bacterial samples from patients are sequenced and used to infer a phylogeny, which describes the pathogen's spread among patients. The tips of such phylogenies represent sampled pathogens, and the internal nodes transmission events. Moreover, transmission events can be dated and thereby provide hints on transmission patterns. Such information is extracted by phylodynamic methods to estimate epidemiological and population dynamic parameters^[2-4], assess the impact of population structure^[2,5], and reveal the origins of epidemics^[6].

Birth-death models^[7] incorporate easily interpretable parameters common to standard infectious-disease epidemiology, such as basic reproduction number R_0 , infectious period, *etc.* In contrast to the standard epidemiological models, the birth-death models can be applied to estimate parameters from phylogenetic trees^[8]. In these models, births represent transmission events, while deaths represent removal events for example due to treatment or recovery. Upon a patient's removal their pathogen can be sampled, producing tips in the tree.

Here we focus on three specific, well-established birth-death models (**Fig. 1**): birth-death model (BD)^[8,9], birth-death model with exposed and infectious classes (BDEI)^[5,10,11], and birth-death model with superspreading (BDSS)^[5,12]. These models were deployed using BEAST2^[12,13] to study the phylodynamics of such diverse pathogens as Ebola virus^[10], Influenza virus^[12], Human Immunodeficiency Virus (HIV)^[5], Zika^[14] or Coronavirus^[15]. Using these models, we will demonstrate the reliability of our deep learning-based approach.

While a great effort has been invested in the development of new epidemiological models in phylodynamics, the field has been slowed down by the mathematical complexity inherent to these models. BD, the simplest model, has a closed form solution for the likelihood formula of a tree for a given set of parameters^[8,10], but more complex models, *e.g.* BDEI and BDSS, rely on a set of ordinary differential equations (ODEs) that cannot be solved analytically. To estimate parameter values through maximum-likelihood and Bayesian approaches, these ODEs must be approximated numerically for each tree node^[5,10-12]. These calculations become difficult as the tree size increases, resulting in numerical instability and inaccuracy^[12], as we will see below.

Inference issues with complex models are typically overcome by approximate Bayesian computation (ABC)^[16,17]. ABC is a simulation-based technique relying on a rejection algorithm^[18], where from a set of simulated phylogenies

within a given prior (values assumed for parameter values), those closest to the analyzed phylogeny are retained and give the posterior distribution of the parameters. This scheme relies on the definition of a set of summary statistics aimed at representing a phylogeny and on a distance measure between trees. This approach is thus sensitive to the choice of the summary statistics and distance metric, *e.g.* Euclidean distance. To address this issue *Saulnier et al.*^[19] developed a large set of summary statistics. In addition, they used a regression step to select the most relevant statistics and to correct for the discrepancy between the simulations and the analyzed phylogenies (see also *Blum et al.*^[20]).

We propose a likelihood-free, rejection-free, simulation-based method relying on deep learning from millions of trees of varying size simulated within a broad range of parameter values. To describe these trees and use them as input for the deep learner, we develop two tree representations: (1) a large set of summary statistics mostly based on *Saulnier et al.*^[19], and (2) a complete and compact vectorial representation of phylogenies, including both the tree topology and branch lengths. The summary statistics are derived from our understanding and knowledge of the epidemiological processes. However, they can be incomplete and thus miss some important aspects of the studied phylogenies, which can potentially result in low accuracy during inference. Moreover, it is expected that new phylodynamic models will require design of new summary statistics, as confirmed by our results with BDSS. In contrast, our vectorial representation is a raw data representation that preserves all information contained in the phylogeny and thus should be accurate and deployable on any new model, provided the model parameters are identifiable. Our vectorial representation naturally fits with deep learning methods, especially the convolutional architectures, which have already proven their ability to extract relevant features from raw representations, for example in image analysis^[21,22] or weather prediction^[23].

In the following, we introduce our vectorial tree representation and the new summary statistics designed for BDSS. We then present the deep learning architectures trained on these representations and evaluate their accuracy on simulated datasets in terms of both parameter estimation and model selection. The results are compared to those of the gold standard method, BEAST2^[12,13]. Lastly, we showcase our methods on an HIV dataset^[24,25] from men-having-sex-with-men (MSM) community from Zurich. All technical details are provided in Methods.

RESULTS

Neural networks are trained on numerical vectors from which they can learn regression and classification tasks. We trained such networks on phylogenetic trees to estimate epidemiological parameters (regression) and select phylodynamic models (classification). We undertook two strategies for representing phylogenetic trees as numerical vectors, which we describe first, before showing the results with simulated and real data.

Summary statistics (SS) representation. We used a set of 83 SS developed by *Saulnier et al.*^[19]: 26 measures of branch lengths, such as median of both internal and tip branch lengths; 8 measures of tree topology, such as tree imbalance; 9 measures on the number of lineages through time, such as time and height of its maximum; and 40 coordinates representing the lineage-through-time plot. To capture more information on the phylogenies generated by the BDSS model, we further enriched these SS with 14 new statistics on transmission chains describing the distribution of the duration between consecutive transmissions (internal tree nodes). Our SS are diverse, complementary and somewhat redundant. We used feed-forward neural networks (FFNN) with several hidden layers (**Fig. 2 b (i)**) that select and combine relevant information from the input features. In addition to SS, we provide both the tree size, *i.e.* the number of tips, and the sampling probability used to generate the tree, as input to our FFNN (**Fig. 2 a (vi)**). We will refer to this method as FFNN-SS.

Compact vectorial tree representation. While converting raw information in a form of a phylogenetic tree into a set of SS, information loss is unavoidable. This means not only that the tree cannot be fully reconstructed from its SS, but also that depending on how much useful and relevant information is contained in the SS, the neural network may fail to solve the problem at hand. As an alternative strategy to SS, and to prevent information loss in the tree representation, we developed a representation called ‘Compact Bijective Ladderized Vector’ (CBLV).

Several vectorial representations of trees based either on polynomial^[26,27], Laplacian spectrum^[28] or F matrices^[29] have been developed previously. However, they represent the tree shape but not the branch lengths^[26,27], may lose information on trees^[28]. In addition, some of these representations^[27,29] require vectors or matrices of quadratic size with respect to the number of tips.

Inspired by these approaches, we designed our concise, easily computable, compact, and bijective (*i.e.* 1-to-1) tree representation that applies to trees of variable size and is appropriate as machine learning input (see **Methods** for

details). To obtain this representation, we first ladderize the tree, that is, for each internal node, the descending subtree containing the most recently sampled tip is rotated to the left, **Fig. 2 a (ii)**. This ladderization step does not change the tree but facilitates learning by standardizing the input data (**Supplementary Fig. 5**). Moreover, it is consistent with trees observed in real epidemiological datasets, for example Influenza, where ladder-like trees reflect selection and are observed for several pathogens^[1]. Then, we perform an inorder traversal^[30] of the ladderized tree, during which we collect in a vector for each visited internal node its distance to the root and for each tip its distance to the previously visited internal node. In particular, the first vector entry corresponds to the tree height. This transformation of a tree into a vector is bijective, as one can unambiguously reconstruct any given tree from its vectorial representation (**Supplementary Fig. 1**). The vector is as compact as possible, and its size grows linearly with the number of tips. We complete this vector with zeros to reach the representation length of the largest tree contained in our simulation set, and we add the known sampling probability (**Fig. 2 a (v), b (i)**).

Bijectivity combined with ladderization facilitates the training of neural networks, which do not need to learn that different representations correspond to the same tree. However, unlike our SS, this full representation does not have any high-level features. In CBLV identical subtrees will have the same representation in the vector whenever the roots of these subtrees have the same height, while the vector representation of the tips in such subtrees will be the same no matter the height of the subtree's root. Similar subtrees will thus result in repeated patterns along the representation vector. We opted for Convolutional Neural Networks (CNN), which are designed to extract information on patterns in raw data. Our CNN architecture (**Fig. 2 b (ii)**) includes several convolutional layers that perform feature extraction, and maximum and average pooling that select relevant features and keep reasonable dimensions of feature maps. The output of the CNN is then fed into an FFNN that combines the patterns found in the input to perform predictions. In the rest of the manuscript, we refer to this method as CNN-CBLV.

Simulated datasets

For each phylodynamic model (BD, BDEI, BDSS), we simulated 4 million trees, covering a large range of values for each parameter of epidemiological interest (R_0 , infectious period: $1/\gamma$, incubation period: $1/\epsilon$, the fraction at equilibrium of superspreading individuals: f_{SS} , and the superspreading transmission ratio: X_{SS}). Of the 4 million trees, 3.99 million were used as a training set, and 10,000 as a validation set for early stopping in the training phase^[31]. Additionally, we simulated another 10,000 trees, which we used as a testing set, out of which 100 were also evaluated

with the gold standard methods, BEAST2 and TreePar, which are more time consuming. Another 1 million trees were used to define confidence intervals for estimated parameters. For BD and BDEI we considered two settings: one with small trees (50 to 199 tips, in **Supplementary Fig. 2**) and a second one with large trees (200 to 500 tips, **Fig. 3**). For BDSS, we considered only the setting with large trees, as the superspreading individuals are at low fraction and cannot be detected in small trees (results not shown).

To increase the generality of our approach and avoid the arbitrary choice of the time scale (one unit can be a day, a week, or a year), we rescaled all trees and corresponding epidemiological parameters, such that the average branch length in a tree was equal to 1. After inference, we rescaled the estimated parameter values back to the original time scale.

Neural networks yield more accurate parameter estimates than gold standard methods

We compared accuracy of parameter estimates yielded by our deep learning methods and those yielded by two state-of-the-art phylodynamics inference tools, BEAST2^[12,13] and TreePar^[5]. The comparison shows that our deep learning methods trained with SS and CBLV are either comparable (BD) or more accurate (BDEI and BDSS) than the state-of-the-art inference methods (**Fig. 3**). The simple BD model has a closed form solution for the likelihood function, and thus BEAST2 results are optimal in theory^[8,9]. Our results with BD are similar to those obtained with BEAST2, and thus nearly optimal as well. For BDEI and BDSS our results are more accurate than BEAST2, which is likely explained by numerical approximations of likelihood calculations in BEAST2^[5,10,11] for these models. These approximations may lead to a lack of convergence (2% cases for BDEI and 15% cases for BDSS) or a convergence to local minima. We suspect BEAST2 of converging to local optima when it converged to values with high relative error (*i.e.* >1.0 ; 8% cases for BDEI and 11% cases for BDSS, **Fig. 3 b-c**). Furthermore, our deep learning approaches showed a lower bias in parameter estimation than BEAST2 (**Supplementary Table 3**). As expected, both approaches, FFNN-SS and CNN-CBLV, get more accurate with larger trees (**Supplementary Fig. 7**).

We tried to perform maximum likelihood estimation (MLE) implemented in the TreePar package^[5] on the same trees as well. While MLE under BD model on simulations yielded as accurate results as BEAST2, for more complex models it showed overflow and underflow issues (*i.e.* reaching infinite values of likelihood) and yielded inaccurate results, such as more complex models (BDEI, BDSS) having lower likelihood than a simpler, nested one (BD) for a part of

simulations (results not shown). These issues were more prominent for larger trees. TreePar developers confirmed these limitations and suggested using the last version of BEAST2 instead.

CNN-CBLV has high potential for application to new models

FFNN-SS and CNN-CBLV show similar accuracy across various settings (**Fig. 3, Supplementary Fig. 2, Supplementary Table 2**). The advantage of the CBLV is its generality, meaning there is no loss of information between the tree and its representation in CBLV regardless of which model the tree was generated under. This contrasts with the pre-designed SS, which might need additional specific SS depending on the model at hand. This was confirmed in our analyses of BDSS simulations. To estimate the parameters of this model, we added summary statistics on transmission chains on top of the SS taken from *Saulnier et al.*^[19]. This improved the accuracy of superspreading fraction estimates of the FFNN-SS (**Supplementary Fig. 6**), so that it was comparable to the CNN-CBLV.

Neural networks are fast inference methods

We compared the CPU time required by each of our inference approaches. For deep learning methods, while they require longer computing times in the learning phase (*e.g.* in the order of a couple of BEAST2 inferences on large trees under BDEI that we performed), inference is almost instantaneous and most of the time is spent on tree encoding into SS or CBLV. Averaged over 10,000 trees generated by the three birth-death models, the parameter inference with deep learning takes 6×10^{-5} CPU hours or ~ 0.2 seconds per tree.

For comparison, BEAST2 inference under the BD model with 5 million MCMC steps takes on average 0.18 CPU hours per tree, and 0.009 CPU hours if only the time to convergence is considered. Inference under BDEI and BDSS with 10 million MCMC steps takes 57 CPU hours (6 CPU hours to convergence) and 79 CPU hours (6 CPU hours to convergence), respectively.

Neural networks are accurate methods for model selection

We trained CNN-CBLV and FFNN-SS on simulated trees to predict the birth-death model under which they were simulated (BD or BDEI for small trees; BD, BDEI or BDSS for large trees). Note that for parameters shared between multiple models, we used identical parameter value ranges across all these models (**Supplementary Table 1**). Then, we assessed the accuracy of both of our approaches on 100 simulations obtained with each model and compared it

with the model selection under BEAST2 based on Akaike information criterion through Markov Chain Monte Carlo (AICM)^[32,33]. The AICM, similar to deviance information criterion (DIC) by *Gelman et al.*^[32], does not add computational load and is based on the average and variance of posterior log-likelihoods along the Markov Chain Monte Carlo (MCMC).

FFNN-SS and CNN-CBLV have similar accuracy (**Supplementary Fig. 3**), namely 92% for large trees (BD vs BDEI vs BDSS), and accuracy of 91% and 90%, respectively, for small trees (BD vs BDEI). BEAST2 yielded an accuracy of 91% for large trees and 88% for small trees. The non-converging simulations were not considered for any of these methods, *i.e.* 5% simulations for small trees and 24% for large trees.

The process of model selection with a neural network is as fast as the parameter inference, *i.e.* 6×10^{-5} CPU hours (~ 0.2 seconds) per tree. This represents a practical, fast and accurate way to perform model selection in phylodynamics.

Neural networks are well suited to learn complex models

To assess the complexity of learned models, we explored other inference methods, namely: 1] linear regression as a baseline model trained on summary statistics (LR-SS); 2] FFNN trained directly on CBLV (FFNN-CBLV); 3] CNN trained on Compact Random Vector (CNN-CRV), for which the trees were randomly rotated, instead of being ladderized as in **Fig. 2 (ii)**; and 4] two “null models”.

LR-SS yielded inaccurate results even for the BD model (**Supplementary Table 2**), which seems to contrast with previous findings^[19], where LR approach combined with ABC performed only slightly worse than BEAST2. This can be explained by the lack of rejection step in LR-SS, which enables to locally reduce the complexity of the relation between the representation and the inferred values to a linear one^[18]. However, the rejection step requires a metric, *e.g.* the Euclidean distance, which may or may not be appropriate depending on the model and the summary statistics. Moreover, rejection has a computational cost with large simulation sets.

Neural networks circumvent these problems with rejection and allow for more complex than linear relation between the tree representation and the inferred values to be captured. This also reflected in our results with FFNN-CBLV and CNN-CRV, which both proved to be generally more accurate than LR-SS. However, FFNN-CBLV was substantially less accurate than CNN-CBLV (**Supplementary Table 2, Supplementary Fig. 5**). This indicates the presence of repeated patterns that may appear all along the vectorial representation of trees, such as subtrees of any size, which

are better extracted by CNN than by FFNN. In its turn, CNN-CRV required larger training sets to reach an accuracy comparable to CNN-CBLV (**Supplementary Fig. 5**), showing that the ladderization and bijectivity of the CBLV helped the training.

To assess how much information is actually learned, we also measured the accuracy of two “null models”: FFNN trained to predict randomly permuted target values; and a random predictor, where parameter values were sampled from prior distributions. Results show that the neural networks extract a considerable amount of information for most of the estimated parameters (**Supplementary Table 2**). The most difficult parameter to estimate was the fraction of superspreading individuals in BDSS model, with accuracy close to random predictions with small trees, but better performance as the tree size increases (**Supplementary Fig. 7**).

Showcase study of HIV in MSM subpopulation in Zurich

The Swiss HIV Cohort is densely sampled, including more than 16,000 infected individuals^[24]. Datasets extracted from this cohort have often been studied in phylodynamics^[8,25]. We studied a dataset of MSM subpopulation from Zurich, which corresponds to a cluster of 200 sequences studied previously by *Rasmussen et al.*^[25], who focused on the degree of connectivity and its impact on transmission between infected individuals. Using coalescent approaches, they detected the presence of highly connected individuals at the beginning of the epidemic and estimated R_0 to be between 1.0 and 2.5. We used their tree as input for neural networks and BEAST2.

To perform analyses, one needs an estimate of the sampling probability. We considered that: 1] the cohort is expected to include around 45% of Swiss individuals infected with HIV^[24]; and 2] the sequences were collected from around 56% of individuals enrolled in this cohort^[34]. We used these percentages to obtain an approximation of sampling probability of $0.45 \times 0.56 \sim 0.25$ and used this value to analyze the MSM cluster. To check the robustness of our estimates, we also used sampling probabilities of 0.2 and 0.3 in our estimation procedures.

First, we performed a quick sanity check considering the resemblance of HIV phylogeny with simulations obtained with each model. All three considered birth-death models passed this check (**Supplementary Fig. 8**). Then, we performed model selection (BD vs BDEI vs BDSS) and parameter estimation using our two methods and BEAST2 (**Fig. 4 a-b**). Finally, we checked the model adequacy with a second, more specific, sanity check, derived from the inferred values (**Fig. 4 c**).

Model selection with CNN-CBLV and FFNN-SS resulted in the acceptance of BDSS (probability of 1.00 versus 0.00 for BD and BDEI), and the same result was obtained with BEAST2 and AICM. These results are consistent with what is known about HIV epidemiology, namely, the presence of superspreading individuals in the infected subpopulation^[35] and the absence of incubation period without infectiousness such as is emulated in BDEI^[36].

We then inferred parameter values under the selected BDSS model (**Fig. 4 a-b**). The values obtained with FFNN-SS and CNN-CBLV are close to each other and the 95% CI are nearly identical. We inferred an R_0 of 1.6 and 1.7, and an infectious period of 10.2 and 9.8 years, with FFNN-SS and CNN-CBLV, respectively. Transmission by superspreading individuals was estimated to be around 9 times higher than by normal spreaders and superspreading individuals were estimated to account for around 7-8% of the population. Our R_0 estimates are consistent with the results of a previous study^[8] performed on data from the Swiss cohort, and the results of *Rasmussen et al.*^[25] with this MSM dataset^[25]. The infectious period we inferred is a bit longer than that reported by *Stadler et al.*, who estimated it to be 7.74 [95% CI 4.39-10.99] years^[8]. The infectious period is a multifactorial parameter depending on treatment efficacy and adherence, the times from infection to detection and to the start of treatment, *etc.* In contrast to the study by *Stadler et al.*, whose data were sampled in the period between 1998 and 2008, our dataset covers also the period between 2008 and 2014, during which life expectancy of patients with HIV was further extended^[37]. This may explain why we find a longer infectious period (with compatible CIs). Lastly, our findings regarding superspreading are in accordance with those of *Rasmussen et al.*^[25], and with a similar study in Latvia^[5] based on 40 MSM sequences analyzed using a likelihood approach. Although the results of the latter study may not be very accurate due to the small dataset size, they still agree with ours, giving an estimate of a superspreading transmission ratio of 9, and 5.6% of superspreading individuals. Our estimates were quite robust to the choice of sampling probability, *e.g.* $R_0 = 1.54$, 1.60 and 1.66, with a sampling probability of 0.20, 0.25 and 0.30, respectively (**Fig. 4b**).

Compared to BEAST2, the estimates of the infectious period and R_0 were similar for both approaches, but BEAST2 estimates were higher for the transmission ratio (14.5) and the superspreading fraction (10.6%). These values are in accordance with the bias of BEAST2 estimates that we observed in our simulation study, that is, positive bias for both the superspreading fraction and the transmission ratio, while our estimates were nearly unbiased (**Supplementary Table 3**).

Finally, we checked the adequacy of BDSS model by resemblance of HIV phylogeny to simulations. Using inferred 95% CI, we simulated 10,000 trees and performed Principal Component Analysis on SS, to which we projected the SS of our HIV phylogeny. This was close to simulations, specifically close to the densest swarm of simulations, confirming adequacy of both the inferred values and the selected model (**Fig. 4 c**).

PhyloDeep: Python package for parameter inference and model selection

FFNN-SS and CNN-CBLV parameter inference, model selection, 95% CI computation and *a priori* check are available via Python package PhyloDeep at <https://pypi.org>. They cover the parameter subspace as described in **Supplementary Table 1**. Version 0.2.5 is used for all results in this article. The input is a dated phylogenetic tree with 50-500 tips and presumed sampling probability. The output is a PCA plot for *a priori* check, and a csv file with probabilities for each model (for model selection) and point estimates and 95% CI values (for parameter inference).

PERSPECTIVES

In this manuscript, we presented new methods for parameter inference and model selection in phylodynamics based on deep learning from phylogenies. Using extensive simulations, we established that these methods are at least as accurate as the state-of-the-art methods. The main advantage of our approaches is the ease of deployment on new models and computational speed once the networks are trained. We also applied our deep learning methods to the Swiss HIV dataset from MSM and obtained results consistent with current knowledge of HIV epidemiology.

Using BEAST2, we obtained inaccurate results for some of the BDEI and BDSS simulations. While BEAST2 has been successfully deployed on many models and tasks, it clearly suffers from approximations in likelihood computation with these two models. However, these will likely improve in near future. In fact, we already witnessed substantial improvements done by BEAST2 developers to the BDSS model, while carrying out this research.

Both of our neural network approaches circumvent likelihood computation and thereby represent a new way of using molecular data in epidemiology, without the need to solve large systems of differential equations. This opens the door to novel phylodynamics models, which would make it possible to answer questions previously too complex to ask. This is especially true for CBLV representation, which does not require the design of new summary statistics, when applied to trees generated by a new mathematical model.

A direction of further research is to extend our approach to the family of phylodynamic models based on structured coalescent^[38,39], rather than birth-death models. Our methods could also be extended to the macroevolutionary birth-death models, which are used to study the diversification dynamics of species, and are closely related to epidemiological models^[40]. Other fields related to phylodynamics, such as population genetics, have been developing likelihood-free methods^[41], for which our approach might serve as a source of inspiration, too.

Other advantages of the deep learning approaches are that they yield close to immediate estimates and apply to trees of varying size. Collection of pathogen genetic data became standard in many countries, resulting in densely sampled infected populations. Examples of such datasets include HIV in Switzerland and UK^[24,42], 2013 Ebola epidemics^[6], several Influenza epidemics and the 2019 SARS-Cov-2 pandemic (www.gisaid.org)^[43]. For many such pathogens, trees can be efficiently and accurately inferred^[44-46] and dated^[47-49] using standard approaches. When applied to such dated trees, our methods can perform model selection and provide accurate phylodynamic parameter estimates within

304 a fraction of a second. Such properties are desirable for phylogeny-based real-time outbreak surveillance methods,
305 which must be able to cope with the daily influx of new samples, and thus increasing size of phylogenies, as the
306 epidemic unfolds, in order to study local outbreaks and clusters, and assess and compare the efficiency of healthcare
307 policies deployed in parallel.

ACKNOWLEDGEMENT: We would like to thank Dr Kary Ocaña and Tristan Dot for initiating experiments on machine learning and phylogenetic trees in our laboratory. We would like to thank Quang Tru Huynh for administrating GPU farm at Institut Pasteur and the INCEPTION program (Investissement d’Avenir grant ANR-16-CONV-0005) that financed the GPU farm. We would like to thank Dr Christophe Zimmer from Institut Pasteur, Sophia Lambert and Dr Hélène Morlon from Institut de Biologie de l’Ecole Normale Supérieure IBENS and Dr Guy Baele from Katholieke Universiteit KU Leuven for useful discussions and Dr Isaac Overcast from IBENS for critical reading of the manuscript. We would like to thank Dr Tanja Stadler and Jérémie Sciré for their help with BEAST2 and MLE approaches. JV is supported by Ecole Normale Supérieure Paris-Saclay and by ED Frontières de l’Innovation en Recherche et Education, Programme Bettencourt. VB would like to thank Swiss National Science Foundation for funding (Early PostDoc mobility grant P2EZP3_184543). OG is supported by PRAIRIE (ANR-19-P3IA-0001).

AUTHOR CONTRIBUTIONS: JV, AZ and OG conceived and set up the methods; JV and VB performed the experiments; JV and OG analyzed the results; JV and AZ wrote the Python package; JV and OG wrote the manuscript; JV, AZ VB and OG edited the manuscript; OG initiated and supervised the project; all authors helped in this research, discussed the results and read the final manuscript.

COMPETING INTERESTS: The authors have no competing interests to declare.

REFERENCES

1. Grenfell, B.T. et al. Unifying the epidemiological and evolutionary dynamics of pathogens. *Science* **303**, 327-332 (2004).
2. Volz, E.M., Kosakovsky Pond, S.L., Ward, M.J., Leigh Brown, A.J., Frost, S.D. Phylodynamics of infectious disease epidemics. *Genetics* **183**, 1421-30 (2009).
3. Drummond, A.J., Rambaut, A., Shapiro, B., Pybus, O.G. Bayesian Coalescent Inference of Past Population Dynamics from Molecular Sequences. *Molecular Biology and Evolution*, **22**, 1185–1192 (2005).
4. Stadler, T. Birth–death skyline plot reveals temporal changes of epidemic spread in HIV and hepatitis C virus (HCV). *Proceedings of the National Academy of Sciences* **110**, 228-233 (2013)
5. Stadler, T., Bonhoeffer, S. Uncovering epidemiological dynamics in heterogeneous host populations using phylogenetic methods. *Philosophical Transactions of the Royal Society B: Biological Sciences* **368**, (2013)
6. Gire, S.K. et al. Genomic surveillance elucidates Ebola virus origin and transmission during the 2014 outbreak. *Science* **345**, 1369-72 (2014).
7. Boskova, V., Bonhoeffer, S., Stadler, T. Inference of Epidemiological Dynamics Based on Simulated Phylogenies Using Birth-Death and Coalescent Models. *PLOS Computational Biology* **10**, (2014).
8. Stadler, T. et al. Estimating the Basic Reproductive Number from Viral Sequence Data. *Mol. Biol. Evol.* **29**, 347–357 (2012).
9. Leventhal, G.E., Günthard, H.F., Bonhoeffer, S., Stadler, T. Using an Epidemiological Model for Phylogenetic Inference Reveals Density Dependence in HIV Transmission. *Mol. Biol. Evol.* **31**, 6–17 (2014).
10. Stadler, T., Kuhnert, D., Rasmussen, D.A., du Plessis, L. Insights into the early epidemic spread of Ebola in sierra leone provided by viral sequence data. *PLoS Curr.* **6**, (2014).
11. Kühnert, D., Stadler, T., Vaughan, T.G., Drummond, A.J. Phylodynamics with Migration: A Computational Framework to Quantify Population Structure from Genomic Data. *Mol. Biol. Evol.* **33**, 2102-16 (2016).

12. Sciré, J., Barido-Sottani, J., Kühnert, D., Vaughan, T.G., Stadler, T. Improved multi-type birth-death
 phylodynamic inference in BEAST 2. Preprint at
<https://www.biorxiv.org/content/10.1101/2020.01.06.895532v1.full.pdf> (2020).
13. Bouckaert, R. et al. BEAST 2: A Software Platform for Bayesian Evolutionary Analysis. *PLoS Computational Biology* **10**, (2014).
14. Boskova, V., Stadler, T., Magnus, C. The influence of phylodynamic model specifications on parameter
 estimates of the Zika virus epidemic. *Virus Evolution* **4**, January (2018).
15. Vaughan, T.G., Sciré, J., Nadeau, S.A., Stadler, T. Estimates of outbreak-specific SARS-CoV-2
 epidemiological parameters from genomic data. Preprint at
<https://www.medrxiv.org/content/10.1101/2020.09.12.20193284v1.full.pdf> (2020).
16. Rubin, D.B. Bayesianly Justifiable and Relevant Frequency Calculations for the Applies Statistician. *The Annals of Statistics* **12**, 1151-1172 (1984).
17. Beaumont, M.A., Zhang, W., Balding, D.J. Approximate Bayesian Computation in Population Genetics. *Genetics* **164**, 2025-2035 (2002).
18. Csilléry, K., Blum, M.G.B., Gaggiotti, O.E., François, O. Approximate Bayesian Computation (ABC) in
 practice. *Trends in Ecology & Evolution* **25**, 410-418 (2010).
19. Saulnier, E., Gascuel, O., Alizon, S. Inferring epidemiological parameters from phylogenies using
 regression-ABC: A comparative study. *PLoS Comp. Biol.* **13**, (2017).
20. Blum, M.G.B. *Handbook Of Approximate Bayesian Computation Ch. Regression approaches for ABC*. 71–
 85. (Chapman and Hall/CRC Press, Boca Raton, 2018).
21. LeCun, Y., Kavukcuoglu, K., Farabet, F. Convolutional networks and applications in vision. *Proc. IEEE Int. Symp. Circuits Syst.* 253-256 (2010).
22. Krizhevsky, K., Sutskever, I., Hinton, G.E. ImageNet Classification with Deep Convolutional Neural
 Networks. *Advances in neural information processing systems* 1097-1105 (2012).
23. Chattopadhyay, A., Hassanzadeh, P., Pasha, S. Predicting clustered weather patterns: A test case for
 applications of convolutional neural networks to spatio-temporal climate data. *Sci. Rep.* **10**, 1317 (2020)
24. The Swiss HIV Cohort Study et al. Cohort Profile: The Swiss HIV Cohort Study. *International Journal of Epidemiology* **39**, 1179–1189 (2010).

25. Rasmussen, D.A., Kouyos, R., Günthard, H.F., Stadler, T. Phylodynamics on local sexual contact networks. *PLOS Comp. Biol.* **13**, (2017).
26. Colijn, C. & Plazzotta, G. A metric on phylogenetic tree shapes. *Systematic Biology* **67**, 113–126 (2018).
27. Liu, P., Gould, M., Colijn, C. Polynomial Phylogenetic Analysis of Tree Shapes. Preprint at <https://doi.org/10.1101/2020.02.10.942367> (2020).
28. Lewitus, E. & Morlon, H. Characterizing and Comparing Phylogenies from their Laplacian Spectrum. *Systematic Biology* **65**, 495-507 (2016).
29. Kim, J., Rosenberg, N.A., Palacios, J.A. Distance metrics for ranked evolutionary trees. *Proceedings of the National Academy of Sciences* **117**, 28876-28886 (2020).
30. Cormen, T.H., Leiserson, C.E., Rivest, R.L., Stein, C. *Introduction To Algorithms*. 286-307 (The MIT Press, Cambridge, 2009).
31. Bengio, Y. *Neural Networks: Tricks Of The Trade, Ch. Practical Recommendations for Gradient-Based Training of Deep Architectures*. (Springer, Berlin, Heidelberg 2002).
32. Gelman, A., Carlin, J.B., Stern, H.S., Rubin, D.B. *Bayesian Data Analysis: Second Edition*. (Chapman and Hall/CRC Press, Boca Raton, 2004).
33. Baele, G. et al. Improving the accuracy of demographic and molecular clock model comparison while accommodating phylogenetic uncertainty. *Mol. Biol. Evol.* **29**, 2157-67 (2012).
34. Kouyos, R.D. et al. Molecular epidemiology reveals long-term changes in HIV type 1 subtype B transmission in Switzerland. *J. Infect. Dis.* **201**, 1488-97 (2010).
35. May, R.M. & Anderson, R.M. Transmission dynamics of HIV infection. *Nature* **326**, 137–142 (1987).
36. Brenner, B.G. et al. Quebec Primary HIV Infection Study Group. High rates of forward transmission events after acute/early HIV-1 infection. *J. Infect. Dis.* **195**, 951-9 (2007).
37. Gueller, A. et al. Swiss National Cohort Life expectancy in HIV-positive persons in Switzerland. *AIDS* **31**, 427-436 (2017).
38. Rasmussen, D.A., Volz, E.M., Koelle, K. Phylodynamic Inference for Structured Epidemiological Models. *PLoS Comput. Biol.* **10**, (2014).
39. Volz, E.M. & Siveroni, I. Bayesian phylodynamic inference with complex models. *PLoS Comput. Biol.* **14**, (2018).

40. MacPherson, A., Louca, S., McLaughlin, A., Joy, J.B., Pennell, M.W. A General Birth-Death-Sampling Model for Epidemiology and Macroevolution. Preprint at <https://www.biorxiv.org/content/10.1101/2020.10.10.334383v2> (2020).
41. Sanchez, T., Cury, J., Charpiat, G., Jay, F. Deep learning for population size history inference: Design, comparison and combination with approximate Bayesian computation. *Mol. Ecol. Resour.* **00**, 1-16. (2020).
42. Dunn, D. & Pillay, D. UK HIV drug resistance database: background and recent outputs. *J. HIV Ther.* **12**, 97–8 (2007).
43. Shu, Y. & McCauley, J. GISAID: Global initiative on sharing all influenza data - from vision to reality. *Euro Surveill.* **22**, 30494 (2017).
44. Minh, B.Q. et al. IQ-TREE 2: New models and efficient methods for phylogenetic inference in the genomic era. *Mol. Biol. Evol.* **37**, 1530–1534 (2020).
45. Kozlov, A.M., Darriba, D., Flouri, T., Morel, B., Stamatakis, A. RAxML-NG: a fast, scalable and user-friendly tool for maximum likelihood phylogenetic inference. *Bioinformatics* **35**, 4453–4455 (2019).
46. Guindon, S. et al. New Algorithms and Methods to Estimate Maximum-Likelihood Phylogenies: Assessing the Performance of PhyML 3.0. *Systematic Biology* **59**, 307-21 (2010).
47. Sagulenko, P., Puller, V., Neher, R.A. TreeTime: Maximum-likelihood phylodynamic analysis. *Virus Evol.* **4**, (2018).
48. To, T.H., Jung, M., Lycett, S., Gascuel, O. Fast Dating Using Least-Squares Criteria and Algorithms. *Syst Biol.* **65**, 82-97 (2016).
49. Volz, E.M. & Frost, S.D.W. Scalable relaxed clock phylogenetic dating. *Virus Evol.* **3**, (2017).

FIGURE LEGENDS

Fig. 1: Birth-death models

a, Birth-death model (BD)^[8,9], **b**, birth-death model with Exposed-Infectious individuals (BDEI)^[5,10,11] and **c**, birth-death model with SuperSpreading (BDSS)^[5,12]. BD is the simplest generative model, used to estimate R_0 and the infectious period ($1/\gamma$)^[8,9]. BDEI and BDSS are extended version of BD. BDEI enables to estimate latency period ($1/\epsilon$) during which individuals of exposed class E are infected, but not infectious^[5,10,11]. BDSS includes two populations with heterogeneous infectiousness: the so-called superspreading individuals (S) and normal spreaders (N). Superspreading individuals are present only at a low fraction in the population (f_{ss}) and may transmit the disease at a rate that is multiple times higher than that of normal spreaders (rate ratio = X_{ss})^[5,12]. Superspreading can have various complex causes, such as the heterogeneity of immune response, disease progression, co-infection with other diseases, social contact patterns or risk behavior, *etc.* Infectious individuals I (superspreading infectious individuals I_S and normal spreaders I_N for BDSS), transmit the disease at rate β ($\beta_{X,Y}$ for an individual of type X transmitting to an individual of type Y for BDSS), giving rise to a newly infected individual. The newly infected individual is either infectious right away in BD and BDSS or goes through an exposed state before becoming infectious at rate ϵ in BDEI. Infectious individuals are removed at rate γ . Upon removal, they can be sampled with probability s , becoming of removed sampled class R. If not sampled upon removal, they move to non-infectious unsampled class U.

Fig. 2: Pipeline for training neural networks on phylogenies

442

443 **a, Tree representations.** **a (i)**, simulated binary trees. Under each model from Fig. 1, we simulate many trees of
444 variable size (50 to 200 tips for ‘small trees’ and 200 to 500 tips for ‘large trees’). For illustration, we have here a tree
445 with 5 tips. We encode the simulations into two representations, either **a (ii-v)**, in a complete and compact tree
446 representation called ‘Compact Bijection Ladderized Vector’ abbreviated as CBLV or **a (vi)** with summary statistics
447 (SS). CBLV is obtained through **a (ii)** ladderization or sorting of internal nodes so that the branch supporting the most
448 recent leaf is always on the left and **a (iii)** an inorder tree traversal, during which we append to a real-valued vector
449 for each visited internal node its distance to the root and for each visited tip its distance to the previously visited
450 internal node. We reshape this representation into **a (iv)**, an input matrix in which the information on internal nodes
451 and leaves is separated into two rows. Finally, **a (v)**, we complete this matrix with zeros so that the matrices for all
452 simulations have the size of largest simulation matrices. For illustration purpose, we here consider that the maximal
453 tree size covered by simulations is 10, and the representation is thus completed with 0s accordingly. SS consists of **a**
454 **(vi)**, a set of 98 statistics: 83 published in *Saulnier et al.*^[19], 14 on transmission chains and 1 on tree size. The
455 information on sampling probability is added to both representations. **b**, Neural networks are trained on these
456 representations to estimate parameter values or to select the underlying model. For SS, we use, **b (i)**, a deep feed-
457 forward neural network (FFNN) of funnel shape (we show the number of neurons above each layer). For the CBLV
458 representation we train, **b (ii)**, Convolutional Neural Networks (CNN). The CNN is added on top of the FFNN. The
459 CNN combines convolutional, maximum pooling and global average pooling layers, as described in detail in **Methods**.

Fig. 3: Assessment of deep learning accuracy

Comparison of inference accuracy by BEAST2 (in blue), deep neural network trained on SS (in orange) and convolutional neural network trained on the CBLV representation (in green) on 100 test trees. The size of training and testing trees was uniformly sampled between 200 and 500 tips. We show the relative error for each test tree. The error is measured as the normalized distance between the median *a posteriori* estimate by BEAST2 or point estimates by neural networks and the target value for each parameter. We highlight simulations for which BEAST2 did not converge and whose values were thus set to median of the parameter subspace used for simulations by depicting them as red squares. We further highlight the analyses with a high relative error (>1.00) for one of the estimates as black diamonds. We compare the relative errors for **a**, BD-simulated, **b**, BDEI-simulated and **c**, BDSS-simulated trees. Average relative error is displayed for each parameter and method in corresponding color below each figure. The average error of a FFNN trained on summary statistics but with randomly permuted target is displayed as black dashed line and its value is shown in bold black below the x-axis. The accuracy of each method is compared by paired z-test; $P < 0.05$ is shown as thick full line; non-significant is not shown.

Fig. 4: Parameter inference on HIV data sampled from MSM Zurich

475

476 **Inference on Zurich HIV MSM dataset.** Using BDSS model with BEAST2 (in blue), FFNN-SS (in orange), and
 477 CNN-CBLV (in green) we infer, **a (i)**, basic reproduction number, **a (ii)**, infectious period (in years), **a (iii)**,
 478 superspreading transmission ratio and, **a (iv)**, superspreading fraction. For FFNN-SS and CNN-CBLV, we show the
 479 distributions and the 95% CIs obtained with a fast approximation of the parametric bootstrap (**Methods**). For
 480 BEAST2, the distributions and 95% CI were obtained considering all reported steps (9,000 in total) excluding the 10%
 481 burn-in. Arrows show the position of the original point estimates obtained with FFNN-SS and CNN-CBLV and the
 482 median *a posteriori* estimate obtained with BEAST2. Circles show lower and upper boundaries of 95% CI. **b**, these
 483 values are reported in a table, together with point estimates obtained while considering lower and higher sampling
 484 probability (0.20 and 0.30). **c**, 95% CI boundaries obtained with FFNN-SS are used to perform an *a posteriori* model
 485 adequacy check. We simulated 10,000 trees with BDSS while resampling each parameter from a uniform distribution,
 486 whose upper and lower bounds were defined by the 95% CI. We then encoded these trees into SS, performed PCA
 487 and projected SS obtained from the HIV MSM phylogeny (red stars) on these PCA plots. We show here the projection
 488 into **c (i)**, first two components of PCA, **c (ii)**, the 3rd and 4th components, together with the associated percentage of
 489 variance displayed in parentheses. Warm colors correspond to high density of simulations.

METHODS

TABLE OF CONTENTS

	Page
TREE REPRESENTATION USING SUMMARY STATISTICS (SS)	26
<i>Saulnier et al. 's summary statistics</i>	26
<i>Additional summary statistics</i>	26
COMPLETE AND COMPACT TREE REPRESENTATION (CBLV)	27
<i>Tree ladderization</i>	27
<i>Tree traversal and encoding</i>	28
<i>Properties of CBLV</i>	28
<i>Alternative tree representations</i>	29
TREE RESCALING	29
REDUCTION AND CENTERING OF SUMMARY STATISTICS REPRESENTATION	30
PARAMETER INFERENCE USING NEURAL NETWORKS	30
<i>Deep feedforward neural network architecture for SS</i>	30
<i>Deep convolutional neural network for CBLV</i>	31
<i>Neural network setting and training</i>	31
<i>Preventing overfitting: Early stopping and Dropout</i>	31
<i>Neural networks for model selection</i>	31
CONFIDENCE INTERVALS (95% CI)	32
<i>Computation of 95% CI</i>	32
<i>Assessment of CI accuracy and width</i>	32
MODEL ADEQUACY	33
<i>A priori check</i>	33
<i>A posteriori check</i>	34

515	MODELS	34
516	<i>Constant rate birth-death model with incomplete sampling</i>	34
517	<i>Birth-death model with exposed-infectious classes</i>	34
518	<i>Birth-death model with superspreading</i>	35
519	SIMULATIONS	35
520	METHOD COMPARISON	36
521	<i>Parameter inference with BEAST2</i>	36
522	<i>Model selection with BEAST2</i>	38
523	<i>Linear regression</i>	38
524	<i>FFNN-CBLV</i>	39
525	<i>TreePar</i>	39
526	<i>Null models</i>	39
527	PERFORMANCE ASSESSMENT	40
528	<i>Mean relative error MRE</i>	40
529	<i>Mean relative bias MRB</i>	41
530	<i>Model selection accuracy</i>	41
531	<i>Comparison of time efficiency</i>	41
532	HIV DATASET	42
533	ADDITIONAL REFERENCES	42

TREE REPRESENTATION USING SUMMARY STATISTICS (SS)

We use a set of 98 summary statistics (SS), to which we add the sampling probability, summing to a vector of 99 values.

Saulnier et al. summary statistics

We use the 83 SS proposed by *Saulnier et al.*^[19]:

- 8 SS on tree topology
- 26 SS on branch lengths
- 9 SS on Lineage-Through-Time (LTT) plot
- 40 SS providing the coordinates of the LTT plot

The computing time of these statistics grows linearly with tree size. For details, see the original paper.

Additional summary statistics

In addition to *Saulnier et al.*^[19] statistics, we designed 14 SS on transmission chains. Moreover, we provide the number of tips in the tree as input resulting in $83+14+1 = 98$ SS in total.

The statistics on transmission chains are designed to capture information on the superspreading population. A superspreading individual transmits to more individuals within a given time period than a normal spreader. We thus expect that with superspreading individuals we would have shorter transmission chains. To have a proxy for the transmission chain length, we look at the sum of 4 subsequent shortest times of transmission for each internal node. This gives us a distribution of time-durations of 4-transmission chains. We assume that information on the transmission dynamics of superspreading individuals is retained in the lower, *i.e.* left, tail of 4-transmission-chain lengths distribution which contains relatively many transmissions with short time to next transmission), while the information on normal spreaders should be present in the rest of the distribution.

The implementation of this 4-transmission-chain SS is the following. For each internal node, we sum the distances from the internal node to its closest descendant nodes, descending exactly four times, that is, we take first the distance from the given internal node to its closest child node (of level 1), then from the (level 1) child node, we take its distance to its own closest child node (of level 2), *etc.* If one of the closest descendant nodes is a tip (except for the last one in

the chain), we do not retain any value for the given internal node. Other options, like the shortest 4-edge pathway, could have been used as well and would likely give comparable results.

On the obtained distribution of 4-transmission-chain lengths, we compute 14 statistics:

- number of 4-transmission chains in the tree
- 9 deciles of 4-transmission-chain lengths distribution
- minimum and maximum values of 4-transmission-chain lengths distribution
- mean value of 4-transmission-chain lengths
- variance of 4-transmission-chain lengths

Adding the same summary statistics but on chains comprising 2, 3 and 5 consecutive transmissions had a negligible impact on parameter inference accuracy (data not shown).

COMPLETE AND COMPACT TREE REPRESENTATION (CBLV)

Simulated dated trees are encoded in the form of real-valued vectors, which are then used as input for the neural networks. The representation of a tree with n tips is a vector of length $2n-1$, where one single real-valued scalar corresponds to one internal node or tip. This representation thus scales linearly with the tree size. The encoding is achieved in two steps: tree ladderization and tree traversal.

Tree ladderization

The tree ladderization consists of ordering each node's children. Child nodes are sorted based on the sampling time of the most recently sampled tip in their subtrees: for each node, the branch supporting the most recently sampled subtree is rotated to the left, as in **Fig. 2 a (i-ii)**.

We considered several alternatives with different criteria for child (subtree) sorting instead of ladderization: sampling time of the most anciently sampled tip, subtree length (*i.e.* sum of all branch lengths including the rooting branch), diversification (*i.e.* number of tips), normalized branch lengths (*i.e.* subtree length divided by the number of tips), *etc.* These did not yield better results than CBLV. We show in **Supplementary Fig. 5** the comparison of CBLV with Compact Random Vector (CRV), for which internal nodes were sorted randomly before the tree traversal, showing that CRV yields poorer results than CBLV, as expected.

Tree traversal and encoding

Once the tree is sorted, we perform an inorder tree traversal, using a standard recursive algorithm from the depth first family^[30]. When visiting a tip, we add its distance to the previously visited internal node or its distance to the root, for the tip that is visited first (*i.e.* the tree height due to ladderization). When visiting an internal node, we add its distance to the root. Examples of encoding are shown in **Fig. 2 a (ii-iii)**. This gives us the Compact Bijective Ladderized Vector (CBLV). We then separate information relative to tips and to internal nodes into two rows (**Fig. 2 a (iv)**) and complete the representation with zeros until reaching the size of the largest simulated tree for the given simulation set (**Fig. 2 a (v)**).

Properties of CBLV

CBLV has favorable features for deep learning. Ladderization does not actually change the input tree (phylogenies are unordered trees), but by ordering the subtrees it standardizes the input data and facilitates the learning phase, as observed with CRV (**Supplementary Fig. 5**). Then, the inorder tree traversal procedure is a bijective transformation, as it transforms a tree into a vector, from which the (ordered) tree can be reconstructed unambiguously, using a simple path-agglomeration algorithm shown in **Supplementary Fig. 1**. CBLV is “as concise as possible” being composed of $2n-1$ real values **Fig. 2 a (iii)**, where n is the number of tips. A rooted tree has $2n-2$ branches, and thus $2n-2$ entries are needed to represent the branch lengths. In our $2n-1$ vectorial encoding of trees, we not only represent the branch lengths, but also the tree topology using only 1 additional entry.

The compactness and bijectivity of tree representation reduce the number of simulations required for training the neural network (**Supplementary Fig. 5**). This is because the number of parameters to be trained remains reasonable with compact representation. Moreover, the networks do not need to learn that several different inputs correspond to the same tree.

Our neural networks are intended to apply to trees of variable sizes, *e.g.* trees of 200 to 500 tips in our experiments with large trees. Thus, they are trained on representations of different lengths (*e.g.* a vector of length 399 for a tree of 200 tips), that we complete with zeroes to reach the length of the largest trees (*i.e.* 999 for 500 tips). We add an additional zero to obtain a two-row matrix (500*2 for 500 tips).

Alternative tree representations

Our CBLV tree representation could likely be improved to ease the learning phase and obtain even better parameter estimates. We tested several alternative representations, some inspired by the polynomial representation of small subtrees^[26,27], the Laplacian spectrum^[28] and additive distance matrices that are equivalent to trees^[50]. None was by far as convincing as CBLV, which is likely due to their large size (n^2 for distance matrices and polynomials) or numerical instabilities and potential loss of information (for Laplacian spectrum).

Moreover, the margin for improvement of the accuracy of CNN-CBLV for the BD model, and likely for other models, is low. This is due to the observation that the accuracy of CNN-CBLV is the similar to that of likelihood-based approaches for the BD model, and the fact that we have an analytical likelihood formula for the BD model, making the likelihood-based approach itself optimal^[8,9].

TREE RESCALING

Before encoding, the trees are rescaled so that the average branch length is 1, that is, each branch length is divided by the average branch length of the given tree, called rescale factor. The values of the corresponding time-dependent parameters, *i.e.* infectious period and incubation period, are divided by the rescale factor too. The NN is then trained to predict these rescaled values. After parameter prediction, the predicted parameter values are multiplied by the rescale factor and thus rescaled back to the original time scale.

This step enables us to overcome problems of arbitrary time scales of input trees and makes a pre-trained NN more generally applicable. More specifically, an input tree with a time scale in days will be associated naturally with the same output as the same tree with a time scale in years, since both these trees will be rescaled to the same intermediate tree of average branch length of 1. Rescaling thus makes it possible to apply the same pre-trained NN to phylogenies reconstructed from sequences of a pathogen associated with an infectious period on the scale of days (*e.g.* EboV) or years (*e.g.* HIV), without the need to simulate new phylogenies and train a new NN.

REDUCTION AND CENTERING OF SUMMARY STATISTICS REPRESENTATION

Before training our NN and after having rescaled the trees to unit average branch length (see the sub-section above), we reduce and center every summary statistic by subtracting the mean and scaling to unit variance. To achieve this, we use the standard scaler from the scikit-learn package^[51], which is fitted to the training set.

PARAMETER AND MODEL INFERENCE USING NEURAL NETWORKS

We implemented deep learning methods in Python 3.6 using Tensorflow 1.5.0^[52], Keras 2.2.4^[53] and scikit-learn 0.19.1^[51] libraries. For each network, several variants in terms of number of layers and neurons, activation functions, regularization, loss functions and optimizer, were tested. In the end, we decided for two specific architectures that best fit our purpose: one deep FFNN trained on SS and one CNN trained on CBLV tree representation.

Deep feedforward neural network architecture for SS

The network consists of one input layer (of 99 input nodes both for trees with 50-199 and 200-500 tips), 4 sequential hidden layers organized in a funnel shape with 64-32-16-8 neurons and 1 output layer of size 2-4 depending on the number of parameters to be estimated. The neurons of the last hidden layer have linear activation, while others have exponential linear activation^[54].

Layer (type)	Output Shape	Param #
dense_1 (Dense)	(None, 64)	6400
dense_2 (Dense)	(None, 32)	2080
dense_3 (Dense)	(None, 16)	528
dense_4 (Dense)	(None, 8)	136
dense_5 (Dense)	(None, 2)	18
Total params: 9,162		
Trainable params: 9,162		
Non-trainable params: 0		

Architecture: Feedforward neural network architecture. Example of FFNN trained on large trees to estimate the parameters of the BD model (R_0 and infectious period $1/\gamma$). ‘Dense’ layer means that for each neuron, all the inputs are multiplied by learned weights, summed together with the bias term. The activation function is then applied to the weighted sum before being output to the next layer. Dense_1 to dense_4 are layers with neurons of exponential linear activation, while dense_5 is composed either of softmax (in case of model selection) or linear neurons (in case of parameter estimation). The number of trainable

parameters in each layer is displayed (Param #): for example in the first layer, we have 99 input values and 1 bias for each of the 64 neurons, giving us in total $(99+1)*64=6,400$ trainable parameters. Output by Keras^[53], the 'None' in the 'Output Shape' means the network can input more than one training example at the time and that there is no constraint on the batch size (hence 'None').

Deep convolutional neural network for CBLV

The CNN consists of one input layer (of 400 and 1002 input nodes for trees with 50-199 and 200-500 tips, respectively). This input is then reshaped into a matrix of size of $201*2$ and $501*2$, for small and large trees, respectively, with entries corresponding to tips and internal nodes separated into separate rows (and one extra column with one entry in each row corresponding to the sampling probability). Then, there are two 1D convolutional layers of 50 kernels each, of size 3 and 10, respectively, followed by max pooling of size 10 and another 1D convolutional layer of 80 kernels of size 10. After the last convolutional layer, there is a GlobalPoolingAverage1D layer and a FFNN of funnel shape (64-32-16-8 neurons) with the same architecture and setting as the NN used with SS.

Neural network setting and training

For both NNs, we use the Adam optimisation algorithm^[55] as optimizer and the Mean Absolute Percentage Error (MAPE) as loss function. The batch size is set to 8,000. To train the network, we split the simulated dataset into 2 groups: [1] proper training set (3,990,000 examples); [2] validation set (10,000).

Preventing overfitting: Early stopping and Dropout

To prevent overfitting during training, we use: [1] the early stopping algorithm evaluating MAPE on a validation set; and [2] dropout that we set to 0.5 in the feed-forward part of both NNs^[56] (0.4, 0.45, 0.55 and 0.6 values were tried for basic BD model without improving the accuracy).

Neural networks for model selection

For model selection, we use the same architecture for FFNN-SS and CNN-CBLV as those for parameter inference described above. The only differences are: [1] the cost function: categorical cross entropy and [2] the activation function used for the output layer, that is, softmax function (of size 2 for small trees, selecting between BD and BDEI model, and of size 3 for large trees, selecting between BD, BDEI and BDSS). As we use the softmax function, the outputs of prediction are the estimated probabilities of each model, summing to 1.

The FFNN-SS and CNN-CBLV are trained on 8×10^6 trees in the small tree setting (4×10^6 trees per model, BD and BDEI). In the large tree setting, the FFNN-SS is trained on 12×10^6 trees (4×10^6 trees per model, BD, BDEI and BDSS) and the CNN-CBLV is trained on 9×10^6 trees (3×10^6 trees per model, BD, BDEI and BDSS), instead of 12×10^6 for GPU limitation purposes.

CONFIDENCE INTERVALS (95% CI)

Computation of 95% CI

We compute 95% CI using parametric bootstrap. To facilitate the deployment and speed-up the computation, we perform an approximation using a separate set of 1,000,000 simulations for calculation of CI. For each simulation in the CI set, we store the true parameter values (*i.e.* values with which we simulated the tree) and the parameter values predicted with both of our methods. This large dataset of true/predicted values is used to avoid new simulations, as required with the standard parametric bootstrap.

For a given simulated or empirical tree T , we obtain a set of predicted parameter values, $\{p\}$. The CI computation procedure searches among stored data those that are closest to T in terms of tree size, sampling probability and predicted values. We first subset:

- 10% of simulations within the CI set, which are closest to T in terms of size (number of tips), thus obtaining 100,000 CI sets of true/predicted parameter values.
- Amongst these, 10% of simulations that are closest to T in terms of sampling probability.

We thus obtain 10,000 CI sets of real/predicted parameter values, similar in size and sampling probability to T . For each parameter value p predicted from T , we identify the 1,000 nearest neighboring values amongst the 10,000 true values of the same parameter available in the CI sets, $R_{CI} = \{r_{i=1,1000}\}$, and keep the corresponding predicted values, $P_{CI} = \{p_{i=1,1000}\}$. We then measure the errors for these nearest neighbors as $E_{CI} = \{e_i = p_i - r_i\}$. We center these errors around p , using the median of errors, $median(E_{CI})$, which yields the distribution of errors for given prediction p : $D = \{p + e_i - median(E_{CI})\}$, from which we extract the 95% CI around p . Individual points in the obtained distribution that are outside of the parameter ranges covered through simulations are set to the closest boundary value of the parameter range. For example, for f_{ss} , if for a point in the distribution we obtain a value lower than 0.05, we set the value of that point to 0.05; and if we obtain a value larger than 0.20, we set it to 0.20. We undertake this procedure

for all parameters except for the time related ones, that is, infectious and incubation period as these depend on the time scaling. The width of our 95% CIs is defined as the distance between the 2.5% and 97.5% percentile.

Assessment of 95% CI coverage and width

To assess this fast implementation of the parametric bootstrap, we used the test set of 10,000 simulations (and 100 simulations for comparison with BEAST2 95% CI). We measured the coverage being defined as the fraction of simulations where the true/target parameter values are inside the obtained 95% CI:

$$95\% \text{ CI accuracy} = \frac{\# \text{ true values inside 95\% CI}}{\# \text{ simulations}}$$

We applied the same criteria for BEAST2. For comparison of all methods, we excluded BDEI and BDSS simulations for which BEAST2 did not converge after 10 million steps. To draw BEAST2 CIs, we discarded the burn-in, *i.e.* the first 10% of the MCMC, and calculated the CI on the remaining part of the chain. The CI width and coverage within the CIs obtained either by NNs or BEAST2 are reported in **Supplementary Table 4**.

There exists a plethora of approaches for assessment of uncertainty and CI estimation. For example, [1] in a similar ABC context, the use of neighboring trees (based on the Euclidean distance, not applicable to CBLV and questionable with SS) combined with a regression-based correction similar to that explained above^[19,20]; [2] a (non-approximated) parametric bootstrap^[56]; [3] predicting values from a distribution of trees reconstructed with Bayesian methods^[10]; *etc.* We chose an approximation of parametric bootstrap for its easy deployability, speed, coverage and width of produced CIs. The easy deployability comes from the fact that CIs are based on pre-calculated data stored in our CI set. The speed of the method comes from it not requiring simulations of new trees, and thus producing CIs within 2-4 seconds. The coverage and width are comparable to those of BEAST2 (**Sup. Table 4**), a Bayesian method, intended to estimate the distribution of parameters and the uncertainty of inferences, with high computational cost.

MODEL ADEQUACY

A priori check

We performed a sanity check using the SS of the test set simulations and the SS measured on the empirical HIV phylogeny. We reduced and centered the SS and performed a Principal Component Analysis (PCA) using the PCA function from the scikit-learn^[51] package.

We highlighted the datapoint corresponding to the Zürich HIV MSM phylogeny in **Supplementary Fig. 8**, for each model (BD, BDEI and BDSS). Dissemblance between the simulations and the HIV phylogeny would manifest by this datapoint being outside the distribution corresponding to the simulations.

A posteriori check

We performed a test analogous to the *a priori* model adequacy check. This time, instead of using the test set as representative of simulations, we simulated 10,000 additional simulations under the selected BDSS model. Parameter values were resampled from uniform distribution with boundaries given by the 95% CIs, and sampling probability fixed to presumed value of 0.25 (**Fig. 4**).

MODELS

The models we used for tree simulations are represented in the form of flow diagrams in **Fig. 1**. We simulated dated binary trees for [1] the training of NNs and [2] accuracy assessment of parameter estimation and model selection. We used the following three individual-based phylodynamic models:

Constant rate birth-death model with incomplete sampling

This model (BD^[8,9], **Fig. 1 a**) contains three parameters and three compartments: infectious (I), removed with sampling (R) and removed unsampled (U) individuals. Infection takes place at rate β . Infectious individuals are removed with rate γ . Upon removal, an individual is sampled with probability s .

For simulations, we re-parameterized the model in terms of: basic reproduction number, R_0 ; infectious period, $1/\gamma$; sampling probability, s ; and tree size, t . We then sampled the values for each simulation uniformly at random in the ranges given in **Supplementary Table 1**.

Birth-death model with exposed-infectious classes

This model (BDEI^[10-12], **Fig. 1 b**) is a BD model extended through the presence of an exposed class. More specifically, this means that each infected individual starts as non-infectious (E) and becomes infectious (I) at incubation rate ϵ . BDEI model thus has four parameters (β , γ , ϵ and s) and four compartments (E, I, R and U).

For simulations, we re-parameterized the model similarly as described for BD, set the ϵ value via $1/\gamma$ and incubation ratio ($=\epsilon/\gamma$). We sampled all parameters, including ϵ/γ , from a uniform distribution, just as with BD.

Birth-death model with superspreading

This model (BDSS^[5,10,11], **Fig. 1 c**) accounts for heterogeneous infectious classes. Infected individuals belong to one of two infectious classes (I_S for superspreading and I_N for normal spreading) and can transmit the disease by giving birth to individuals of either class, with rates $\beta_{S,S}$ and $\beta_{S,N}$ for I_S transmitting to I_S and to I_N , respectively, and $\beta_{N,S}$ and $\beta_{N,N}$ for I_N transmitting to I_S and I_N , respectively. However, there is a restriction on parameter values: $\beta_{S,S} * \beta_{N,N} = \beta_{S,N} * \beta_{N,S}$. There are thus superspreading transmission rates $\beta_{S,-}$ and normal transmission rates $\beta_{N,-}$ that are $X_{SS} (= \frac{\beta_{S,S}}{\beta_{N,S}} = \frac{\beta_{S,N}}{\beta_{N,N}})$ times higher for superspreading. At transmission, the probability of the recipient to be superspreading is $f_{ss} (= \frac{\beta_{S,S}}{\beta_{S,S} + \beta_{S,N}})$, the fraction of superspreading individuals at equilibrium. We consider that both I_S and I_N populations are otherwise indistinguishable, that is, both populations share the same infectious period $(1/\gamma)^{[5,10,11]}$. The model thus has six parameters, but only five need to be estimated to fully define the model^[5,10].

For simulations, we chose parameters of epidemiological interest for re-parameterization: basic reproduction number $R_0 (= \frac{\beta_{S,S} + \beta_{N,N}}{\gamma})$, infectious period $1/\gamma$, f_{ss} , X_{ss} and sampling probability s . In our simulations, we used uniform distributions for these 5 parameters, just as with BD and BDEI (**Supplementary Table 1**).

SIMULATIONS

For the parameters R_0 , $1/\gamma$, and s , that are common to all three birth-death models, the same value boundaries were used across all models (**Supplementary Table 1**). We considered two spans of tree size: ‘small trees’ with 50 to 199 tips, and ‘large trees’ with 200 to 500 tips. We then sampled parameter values uniformly at random within these parameter boundaries with standard Latin-hypercube sampling^[57] using PyDOE package. We created 3,990,000 parameter sets for training, 10,000 for validation and early stopping, another 10,000 for testing parameter inference and model selection (comparison with BEAST2 used a subset of 100, for computing time reasons), and 1,000,000 parameter sets for fast computation of CIs.

With these parameter sets, we simulated trees under each birth-death model using our implementation in Python of Gillespie algorithm^[58], based on a standard forward simulator. Comparable accuracies (as in **Fig. 3** and

Supplementary Fig. 2, both for BEAST2 and our methods) were reached on test simulations obtained with a well-established, but slower, simulator TreeSim^[4,5,7] (data not shown).

Each simulation started with one infectious individual (the class was chosen randomly under the BDSS model) and stopped when we obtained a tree with the given number of sampled individuals (tips). If the epidemic died away stochastically, that is, there was no more infectious tips left due to stochastic death before reaching the given tree size, we re-initialized the simulation up to 100 times. Only around 11% of simulations reached more than 2 iterations (20% for BDSS), and less than 0.5% reached more than 50 iterations for all models. If still no tree of given size was obtained after 100 iterations, we discarded the parameter set (less than 0.3% of all sets) and generated a new one to keep the desired number of simulations. This enabled us to maintain a nearly uniform coverage of parameter space, within selected parameter boundaries.

METHOD COMPARISON

Parameter inference with BEAST2

To assess the accuracy of our methods, we compared it with a well-established Bayesian method, as implemented in BEAST2 (version 2.6.2). We used the BDSKY package^[4] (version 1.4.5) to estimate the parameter values of BD simulations and the package *bdmm*^[12,13] (version 1.0) to infer the parameter values of BDEI and BDSS. Furthermore, for the inference on BDSS simulations, instead of BEAST 2.6.2 we used the BEAST2 code up to the commit nr2311ba7, which includes important fixes to operators critical for our analyses. We set the Markov Chain Monte Carlo (MCMC) length to 5 million steps for the BD model, and to 10 million steps for the BDEI and BDSS models. The xml files and command lines are available at:

https://github.com/evolbioinfo/phylodeep/tree/main/data_publication.

The sampling probability was fixed during the estimation. Since the BD, BDEI and BDSS models implemented in BEAST2 do not use the same parametrizations as our methods, we needed to apply parameter conversions for setting the priors for BEAST2 inference, and for translating the BEAST2 results back to parameterizations used in our methods, in order to enable proper comparison of the results (**Table 1**). More specifically, the BEAST2 parameters can be converted to those used in our methods, that is, instead of infectious period and incubation period, BEAST2 uses the inverse of these, namely the infectious rate and incubation rate, respectively; instead of superspreading

transmission ratio and superspreading fraction at equilibrium, it uses individual sub-component parameters $R_{0,SS}$, $R_{0,SN}$, $R_{0,NS}$ and $R_{0,NN}$, which we will collectively refer to as “partial R_0 ”. For BDSS, the BEAST2 prior was thus not the same as that of our simulations for BDSS (**Table 1** and **Supplementary Table 1**), since BEAST2 does not infer the same parameters. We used the range of all parameter values used in our simulations to set the boundaries of uniform prior distributions of parameters inferred by BEAST2. The initial values in the MCMC were set to the medians observed in the training set. During the inference, the parameter values were constrained in the same way as in the simulations, namely, we used the following constraint $R_{0,NN} * R_{0,SS} = R_{0,SN} * R_{0,NS}$ (equivalent to $\beta_{N,N} * \beta_{S,S} = \beta_{S,N} * \beta_{N,S}$) in the BDSS model inference. Furthermore, the effective frequency of superspreading individuals (parameter called “geo-frequencies” in *bdtmm*) was constrained to be between 5% and 20%. Due to the parameter conversions, and despite these constraints the inferred f_{ss} and X_{ss} can reach values outside the boundaries used for simulations, in which case we set them to the closest boundary for fair comparison with deep learning methods in **Fig. 3** (e.g. if the median *a posteriori* f_{ss} was estimated to be larger than 0.20, it was set to 0.20 and if inferred f_{ss} was less than 0.05, it was set to 0.05). The goal of this correction was to avoid penalizing BEAST2 when it converged to local minima outside of the parameter boundaries used for simulations, which are implicitly known to NNs since they were trained on simulations with parameters within these boundaries.

After we obtained the parameters of interest from the original parameters estimated by BEAST2, we evaluated the Effective Sample Size (ESS) on all parameters. We reported the absolute percentage error of the median of *a posteriori* values, corresponding to all reported steps (reported steps being spaced by 1,000 actual MCMC steps) past the 10% burn-in. For simulations for which BEAST2 did not converge, we considered the median of the parameter distribution used for simulations (**Fig. 3**, **Sup. Fig. 2**, **Sup. Tab. 2-3**) or excluded them from the comparison (**Sup. Tab. 2-3**, values reported in brackets, **Sup. Tab. 4**).

For the HIV application, the prior of infectious period was set to [0.1, 30] years (uniform). All the other parameters had the same prior distributions as used in simulations and shown in **Table 1**.

Table 1: BEAST2 priors and their relation to the parameters of interest.

Parameters	Name	BEAST2 parameter	Range	Initial value	Formula
γ	become uninfected rate	Yes	U(0.1,1.0)	0.55	
$1/\gamma$	infectious period	No	[1,10]	1.8	
s	sampling probability	Yes	fixed	true value	
R_0	basic reproduction number	Yes	U(1.0,5.0)	3.0	$= \beta/\gamma$ (BD & BDEI)
ϵ	incubation rate	Yes	U(0.02,5.0)	2.51	
$1/\epsilon$	incubation period	No	[0.2,50]	0.40	
$R_{0,SS}$	partial, within deme superspreading R_0	Yes	U(0.14,4.31)	1.25	
$R_{0,NN}$	partial, within deme normal spreading R_0	Yes	U(0.14,4.31)	1.44	
$R_{0,SN}$	partial, outside deme superspreading R_0	Yes	U(0.034,32.30)	9.0	
$R_{0,NS}$	partial, outside deme normal spreading R_0	Yes	U(0.034,32.30)	0.20	
R_0	basic reproduction number	No	[0.28,8.62]	2.69	$= R_{0,SS} + R_{0,NN}$ (BDSS)
X_{SS}	superspreading infectious ratio at equilibrium	No	$[4 \cdot 10^{-4}, 127]$	6.25	$= R_{0,SS}/R_{0,NS} = R_{0,SN}/R_{0,NN}$
f_{SS}	fraction of superspreading individuals at equilibrium	No	$[4 \cdot 10^{-3}, 0.99]$	0.12	$= R_{0,SS}/(R_{0,SS} + R_{0,SN})$ $= 1 - R_{0,NN}/(R_{0,NN} + R_{0,NS})$

This table shows parameters and their prior distributions used during inference with BEAST2. We display the parameters, their definitions and priors in BEAST2, that are common to all models (in yellow), common to BD and BDEI (in red), BDEI-specific (in purple) and BDSS-specific (in green). From these parameters, we deduce the values and distributions of parameters of interest as shown in the table. Note that the parameters of epidemiological interest are basic reproduction number and infectious period for BD, BDEI and BDSS, incubation period for BDEI, and superspreading infectious ratio and fraction of superspreading individuals for BDSS. We check convergence (ESS) on all parameters and extract median a posteriori and CI value exclusively for the parameters of epidemiological interest.

Model selection with BEAST2

We performed model selection under BEAST2 using Akaike's information criterion through MCMC (AICM)^[32,33].

The AICM is based on the following formula:

$$AICM = 2s_l^2 - 2l$$

where l and s_l^2 are the sample mean and variance of the posterior log-likelihoods. The AICM is an equivalent of AIC

and the model with lowest AICM value is selected.

For 100 simulations obtained with each model (BD, BDEI and BDSS for large trees, BD and BDEI for small trees), we performed parameter estimation with BEAST2 under each model, computed AICM considering the whole MCMC, but excluding 10% burn-in (*i.e.* 9,000 log-likelihood values for BDEI and BDSS considered in total, 4,500 for BD). The results of model selection are shown in **Supplementary Fig. 3**. The BDEI and BDSS simulations for which BEAST2 did not reach an ESS of 200 for all parameters were excluded from the computation of model selection accuracy for all methods.

Linear regression

For each model, linear regression was trained using reduced and centered summary statistics (using scikit-learn package, as with FFNN). Its bias and accuracy were assessed using the same criteria as for the NN approaches (**Supplementary Tables 2-3, Supplementary Fig. 4**).

FFNN-CBLV

We trained an FFNN on CBLV representation. The FFNN architecture was close to the one described in *Architecture* with one extra hidden layer, so 5 layers in total, organized in a funnel shape with 128-64-32-16-8 neurons and 1 output layer of size 2-4 depending on the number of parameters to be estimated. The setting during the training and the sizes of training, validation and testing sets were the same as for the CNN-CBLV. Its bias and accuracy were assessed using the same criteria as for other NN approaches (**Supplementary Tables 2-3, Supplementary Fig. 5**).

TreePar

We used TreePar^[5] for MLE. With BD, we obtained results close to estimates under BEAST2, which is consistent with former studies^[58]. TreePar^[5] uses an exact analytical formula of likelihood for BD and thus these (and BEAST2) results are theoretically optimal.

We also performed several trials to do parameter inference for the more complex models, *i.e.* BDEI and BDSS, but in a large number of cases we encountered numerical problems, *e.g.* underflow or overflow issues, which resulted in infinite negative log-likelihood values, and eventually failed runs. When the calculations did not fail, we found that many estimations under BDSS and BDEI had lower likelihood than estimations performed with (nested) BD on the

same input data. These numerical issues, without available solutions at the moment, were confirmed by the authors of the TreePar package.

Null models

To assess how much information was learned on given problem, we compared FFNN-SS and CNN-CBLV to two null models.

The first null model was the FFNN trained for each model on 4,000,000 simulations using SS, but with randomly permuted target values (*i.e.* the initial correspondence between the SS and underlying parameter values was lost, while the range of values was conserved). We then predicted parameters for 10,000 test simulations (100 for comparison with BEAST2) and measured the mean absolute relative error (MRE, equivalent to MAPE; **Supplementary Table 2**). In such a case, the FFNN always predicted values close to the value with the lowest value of the cost function, *e.g.* 2.2 for parameter values uniformly sampled between 1 and 5 (and not random values inside the range). The MRE of this approach represents then the lowest MRE that machine learning approaches can have in the absence of information, but the knowledge of the parameter distribution. This can be used to get an idea of how well the trained approaches perform and how much information regarding each parameter they can extract from the data.

The second null model was a set of random values sampled from the parameter ranges that were used for simulations (**Supplementary Table 2**). In this model, as opposed to the previous null model, there is no training phase and we do not learn the best compromise in the absence of information.

PERFORMANCE ASSESSMENT

Mean relative error MRE

To compare the accuracy of parameter estimation, we used 100 simulated trees per model. We computed the mean absolute relative error (MRE, **Fig. 3, Supplementary Table 2, Supplementary Fig. 2**) between [1] the true (or target) parameter values and the predicted values for machine learning approaches; and [2] the true (or target) parameter values and the median *a posteriori* values obtained with BEAST2, which are more stable and accurate than maximum *a posteriori* values:

$$MRE = \frac{1}{100} \sum_{i=1}^{100} \frac{abs(predicted_i - target_i)}{target_i}.$$

We plotted individual absolute relative errors (RE) of predictions (**Fig. 3, Supplementary Fig. 2**) for each simulation i , calculated as:

$$RE_i = \frac{abs(predicted_i - target_i)}{target_i}$$

Not being limited by the computational cost for machine learning approaches, we computed the same metric but on 10,000 simulations (**Supplementary Figs. 4-6**; results from 1,000 simulations plotted in **Supplementary Fig. 7**).

We assessed the statistical significance of MRE differences using paired z-test. The two NN approaches were also compared using the same test, but no significant differences were found.

Mean relative bias MRB

To compare the bias in parameter estimation, we used 100 simulated trees per model. We computed the mean relative bias (MRB) between [1] the true (or target) parameter values and the predicted values for machine learning approaches; and [2] the true (or target) parameter values and the median *a posteriori* values obtained with BEAST2 (**Supplementary Table 3**):

$$MRB = \frac{1}{100} \sum_{i=1}^{100} \frac{(predicted_i - target_i)}{target_i}$$

Model selection accuracy

We performed model selection with CNN-CBLV, FFNN-SS and BEAST2 on 100 simulations obtained with each model (10,000 for a sub-comparison of CNN-CBLV and FFNN-SS). Results are shown in **Supplementary Fig. 3** in the form of confusion matrices, where the columns represent the true/target classes, and the rows are the predicted classes.

We then computed the accuracy of each method:

$$accuracy = \frac{\# \text{ true predictions}}{\# \text{ total predictions}}$$

For BEAST2 model selection and large trees, the chain did not converge (displayed as “ESS<200” in **Supplementary Fig. 3**) for 24.3% simulations of large trees and 4.5% simulations of small trees. We did not consider these in accuracy measurements, for all the methods.

Comparison of time efficiency

For FFNN-SS and CNN-CBLV, we reported the average CPU time of encoding a tree (average over 10,000 trees), as reported by NextFlow workflow manager^[60], a pipeline software that we used. The inference time itself was negligible.

For BEAST2, we reported the CPU time averaged over 100 analyses with BEAST2 as reported by NextFlow. For the analyses with BDEI and BDSS models, we report the CPU time to process 10 million MCMC steps, and for the analyses with BD, we report the CPU time to process 5 million MCMC steps. To account for convergence, we re-calculated the average CPU time considering only those analyses, for which the chain converged and the ESS of 200 was reached across all inferred parameters.

HIV DATASET

We used the original phylogenetic tree reconstructed by *Rasmussen et al.*^[25] from 200 sequences corresponding to the largest cluster of HIV-infected men-having-sex-with-men (MSM) subpopulation in Zurich, collected as a part of the Swiss Cohort Study^[24]. For details on tree reconstruction, please refer to their article.

ADDITIONAL REFERENCES

50. Zarestkii, K. Reconstructing a tree from the distances between its leaves. (in Russian) *Uspehi Matematicheskikh Nauk* **20**, 90-92 (1965).

51. Pedregosa, F. et al. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* **12**, 2825–2830 (2011).

52. Abadi, M. et al. TensorFlow: Large-scale machine learning on heterogeneous systems. Preprint at <https://arxiv.org/abs/1603.04467> (2015).

53. Chollet, F. K. <https://keras.io>. (2015).

934 54. Clevert, D.A., Unterthiner, T., Hochreiter, S. Fast and Accurate Deep Network Learning by Exponential Linear
935 Units (ELUs). *ICLR* (2016).

936 55. Kingma, D.P. & Ba, J. Adam: A Method for Stochastic Optimization. *ICLR* (2015).

937 56. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R. Dropout: A simple way to prevent
938 neural networks from overfitting. *Journal of Machine Learning Research* **15**,1929–1958 (2014).

939 57. Efron, B. *Breakthroughs In Statistics, Ch. Bootstrap Methods: Another Look at the Jackknife*. (Springer, New
940 York, 1999).

941 58. McKay, M., Beckman, R., Conover, W. A Comparison of Three Methods for Selecting Values of Input
942 Variables in the Analysis of Output from a Computer Code. *Technometrics* **21**, 239-245 (1979).

943 59. Gillespie, D.T. Exact stochastic simulation of coupled chemical reactions. *The Journal of Physical Chemistry* **81**,
944 2340-2361 (1977).

945 60. Di Tommaso, P. et al. Nextflow enables reproducible computational workflows. *Nat. Biotechnol.* **35**, 316–319
946 (2017)