Mean-reverting self-excitation drives evolution: phylogenetic analysis of a literary genre, *waka*, with a neural language model

Takuma Tanaka
tanaka.takuma@gmail.com
Graduate School of Data Science, Shiga University

December 29, 2024

Abstract

To elucidate the evolutionary dynamics of culture, we must address fundamental questions such as whether we can interpolate and extrapolate cultural evolution, whether the time series of cultural evolution is distinguishable from its reverse, what factors determine the direction of change, and how the cultural influence of a creative work from the viewpoint of an instant is correlated with that from the viewpoint of a later instant. To answer these questions, the evolution of classical Japanese poetry, *waka*, specifically *tanka*, was investigated. Phylogenetic networks were constructed on the basis of the vector representation obtained using a neural language model. The parent–child relationship in the phylogenetic networks exhibited significant agreement with a previously established *honkadori* (allusive variation) phrase-borrowing relationship. The real phylogenetic networks were distinguishable from the time-reversed and shuffled ones. Two anthologies could be interpolated but not extrapolated. The number of children of a poem in the phylogenetic networks, the proxy variable of its cultural influence, evaluated at an instant, was positively correlated with that evaluated later. A poem selected for an authoritative anthology tended to have 1.1–1.5 times more children than a similar but nonselected poem, implying the existence of the Matthew effect. A model with mean-reverting self-excitation replicated these results.

Introduction

A quantitative analysis of cultural evolution is essential in understanding human history because the history of culture encompasses human history and beyond. Animals, including birds and fish, are known to transmit culture (Slater, 1986; Boesch and Boesch, 1990; Dugatkin and Godin, 1992). The oldest stone tools date back 3.3 million years (Harmand et al., 2015), nearly 1 million years before the emergence of the genus *Homo*, but already show signs of sophistication. Sophisticated stone tools as well as other creative works are a manifestation of cultural tradition and inherently evolve over time.

Naturally, there arise several questions to be addressed to understand the evolutionary dynamics of culture. Is the culture of a period the intermediate form of those of the preceding and following periods? Or, are they completely different? Can we predict the future direction of evolution from the history of the past? In other words, can we interpolate and extrapolate cultural evolution? Can the time series of cultural evolution be distinguished from its reverse? What factors determine the direction of cultural evolution? How is the cultural influence of a creative work judged at a time point correlated with that at a later time point?

Some of these questions have been tackled by a considerable amount of studies that investigate information transmission and replication in the various forms of creative works. Since the seminal works by Cavalli-Sforza and Feldman (1973a) and Cavalli-Sforza and Feldman (1973b), the evolution of music, scientific literature, social network posts, malware, images, and movies has been studied. Nakamura and Kaneko (2019) analyzed Western classical music data and found that the frequencies of dissonant intervals have steadily increased. A study on an English text corpus spanning about 400 years found that the authors in a similar period share the same literary style, which gradually changes over time (Hughes et al., 2012). The science of science (Fortunato et al., 2018) has elucidated the dynamics underlying scientific discovery and the scientific community itself: a rare combination of ideas (Foster et al., 2015), a combination of old and new ideas (Uzzi et al., 2013; Kim et al., 2016; Wang et al., 2017), and interdisciplinary collaboration (Larivière et al., 2015) result in higher citation rates; but these features are not advantageous in grant applications (Boudreau et al., 2016; Leahey and Moody, 2014; Lee and Bozeman, 2005); the number of citations of a paper is boosted by the total number of previous citations of the authors in the first years after publication (Petersen et al., 2014); the citation network has a long-tailed distribution (Price, 1965); and the number of citations of a paper is determined by its fitness, an obsolescence factor, and the number of previous citations (Wang et al., 2013; Eom and Fortunato, 2011). Scientific papers share a common pattern of a decaying collective memory with patents, songs, movies, and biographies (Candia et al., 2019). The rapidness of memory decay changes over time. An analysis of a 500 billion-word corpus revealed that inventions and people have been forgotten rapidly in recent years (Michel et al., 2011). Similar to the relationship between scientists and scientific papers, the relationship between users and their posts in social network services has also been studied (Wei et al., 2013; Weng et al., 2014).

One of the most powerful tools to study cultural evolution is phylogenetics. Phylogenetics, originally a branch of biology, is based on the theory of evolution. The evolutionary, *i.e.*, ancestor–descendant relationship of organisms estimated in phylogenetics is represented by a phylogenetic tree. However, the idea of the phylogenetic tree is older than Charles Darwin's theory of evolution; Friedrich Schlegel's language tree and Carl Johan Schlyter's manuscript phylogeny, or stemma, date back to the early nineteenth century (Atkinson and Gray, 2005). The method of phylogenetics has successfully been applied to the analysis of cultural transmission and evolution, forming a field of research called cultural phylogenetics (Mesoudi, 2011; Straffon, 2016). Cultural phylogenetics has elucidated the origin of a language family (Gray and Atkinson, 2003), the coevolution of livestock and descent rules (Holden and Mace, 2003), and the dynamics of political complexity (Currie et al., 2010). These previous works attempted to distinguish between correlations due to shared ancestry and convergent cultural evolution (Holden and Mace, 2003). This has been enabled by recent advances in computational methods.

Phylogenetic methods have also been applied to the creative works of individuals, such as Paleoindian projectile points (O'Brien et al., 2001) and Turkmen textiles (Tehrani and Collard, 2002). The rise of specialist forensic need and interest in Internet memes has stimulated techniques to reconstruct the phylogenetic trees of malware (Goldberg et al., 1998), images (Dias et al., 2010), movies (Dias et al., 2011), and audio files (Nucci et al., 2013). Extending the notion of classical stemmatics (Marmerola et al., 2016), Barbrook et al. (1998) employed computerized techniques to reconstruct the stemma of *The Canterbury Tales*, and Kanojia et al. (2019) used word embeddings to reconstruct the phylogenetic tree of a historical Sanskrit text, *Kāśikāvṛtti*. Most of the creative works investigated in these previous studies have been made by replicating existing ones (as in Internet memes and manuscripts) or carried an explicit indication of the influence between them as citations in research papers and reposts on social media. In other words, their phylogenetic networks are easily reconstructed.

However, not all kinds of creative works are endowed with such conditions favorable for reconstructing phylogenetic networks. Some kinds of creative works are not necessarily a direct replication of existing ones and contain no explicit citations, but nevertheless, are made under the influence of existing ones. To answer questions regarding the evolutionary dynamics of culture, we must extend the study of cultural evolution to these kinds of creative works and develop a method to estimate implicit influence.

Given this background, the present study aimed to answer these questions by investigating the evolution of a literary genre. Literary genres are easier to analyze than other cultural phenomena such as stone toolmaking, fashion, and rituals because databases of literary genres are publicly available, and as such, we can take advantage of recent neural language models in analyzing them. The literary genre studied in this paper is *waka*, the most authoritative poetic form in classical Japanese literature.

Waka has several advantages in studying evolution dynamics. First, *waka* is a fixed verse with 31 syllables, the brevity of which facilitates quantitative analysis. Using a vector representation of poems by neural language models, we can estimate the influence among poems and, consequently, their phylogenetic networks. Second, a comprehensive database of *waka* that contains poems ranging in date from the eighth to the sixteenth century is available. Third, as is the case with every type of classical poetry in the world, poems in the form of classical *waka* were written by poets who were traditionalists and had a thorough knowledge of the great poems of the past. Respect for tradition was so deeply rooted in the poets' hearts that borrowing words from past great poems (*honkadori*) had been an established method of poem writing for more than a millennium. Conducting a comparison with the known *honkadori* relationship is an ideal means to check whether the estimated phylogenetic network is accurate. Fourth, *waka* is so central to classical Japanese culture that a vast amount of research is available to shed light on the subject from another angle.

Thus, the present paper attempts to estimate the phylogenetic network of *waka* and, thereby, address the above questions. This paper is organized as follows. The Materials and Methods section describes the dataset, data preprocessing, and methods of analysis, and also reviews the basic properties and history of *waka*. The Results section reports an analysis of the phylogenetic networks estimated by using vector representations generated by BERT (Devlin et al., 2019), one of the most successful neural language models. First, the estimated phylogenetic network is compared with *honkadori* pointed out by previous studies. Second, the congruence between language models with different initial parameter values is quantified. Third, to characterize the time evolution of the real data, the phylogenetic network is examined by comparing it with those estimated from time-reversed and shuffled data. An index is shown to be able to distinguish the real data from the time-reversed data. Fourth, to examine the constancy of the evaluation of the cultural influence of a poem, the numbers of the phylogenetic children of the poem with the language models trained by using the full dataset and the dataset up to an anthology are compared. Fifth, the poems in two anthologies are classified to determine whether interpolation and extrapolation are possible. Sixth, the effect of selecting a poem for an anthology is measured. More specifically, the increase in the number of phylogenetic children after being selected for an anthology is observed. Finally, a simple model with mean-reverting self-excitation to reproduce these results is presented. The Discussion section summarizes and contextualizes the results. First, the relationship between *honkadori* of *waka* and other literary genres and, thereby, the applicability of the present results to other genres, are discussed. Second, the possibility of influence from other literary genres to *waka* is examined. Finally, the limitations of the present study are presented along with directions for future research.

Materials and Methods

Waka and its history

Waka had been the most authoritative form of Japanese poetry for more than a millennium. Although it is unclear when the form of *waka* was established, the earliest historically verifiable examples date back to the early half of the seventh century. The earliest *waka* anthologies were compiled in the eighth century. In this paper, *waka* refers to *tanka*, the most major poetic form, which consists of five lines with 5-7-5-7-7 syllables. This form has remained productive to date. The images, allegories, metaphors, and symbols of *waka* gave birth to *nō*, *haiku*, as well as novels such as *The Tale of Genji* (Brower and Miner, 1961; Kato et al., 1979; Konishi et al., 1984; Keene, 1999). The authority of *waka* comes from the fact that it was the most essential communication tool among the society of nobles in the *Heian* period and an indispensable part of education in the later periods. *Waka* was composed in the grammar and vocabulary of the early *Heian* period for a millennium (Keene, 1999).

The special position occupied by *waka* in classical Japanese high culture is illustrated by the existence of the Imperial Anthologies (*chokusenshū*). The Imperial Anthologies were the official anthologies

Table 1: List of Imperial Anthologies

ID	Title	Year of publication
1	*Kokinshū*	905
2	*Gosenshū*	955
3	*Shūishū*	1005
4	*Goshūishū*	1087
5	*Kin'yōshū* 1	1124
6	*Kin'yōshū* 2	1125
7	*Kin'yōshū* 3	1126
8	*Shikashū*	1151
9	*Senzaishū*	1187
10	*Shinkokinshū*	1205
11	*Shinchokusenshū*	1232
12	*Shokugosenshū*	1251
13	*Shokukokinshū*	1265
14	*Shokushūishū*	1279
15	*Shingosenshū*	1304
16	*Gyokuyōshū*	1312
17	*Shokusenzaishū*	1320
18	*Shokugoshūishū*	1326
19	*Fūgashū*	1346
20	*Shinsenzaishū*	1359
21	*Shinshūishū*	1364
22	*Shingoshūishū*	1385
23	*Shinshokukokinshū*	1439

of the imperial court compiled on the order of the emperor or ex-emperor. They were compiled by the most prominent poets, some of whom were also the most renowned scholars of *waka* at that time. The Imperial Anthologies were compiled from the tenth to the fifteenth century (Table 1). Poets deeply revered and intensively studied past Imperial Anthologies and wished their poems to be selected for future ones. Consequently, this paper focuses on the Imperial Anthologies.

As is the case with classical poetry in other regions and periods, *waka* poets were encouraged to study and imitate great poems of the past. The rhetorical technique of *honkadori* (allusive variation) borrows material and phrasing from an older poem or poems (Brower and Miner, 1961; Bialock, 1994). An example of *honkadori* is found in a poem by Kiyohara no Fukayabu:

Mukashi mishi
haru wa mukashi no
haru nagara
wa ga mi hitotsu no
arazu mo arukana.

The original is one of the most famous poems by Ariwara no Narihira:

Tsuki ya aranu
haru ya mukashi no
haru naranu
wa ga mi hitotsu wa
moto no mi ni shite.

Sharing 14 syllables, these poems are so strikingly similar that Kiyohara no Fukayabu might have been accused of plagiarism according to present standards. However, *honkadori* was accepted rhetoric theorized in *Kindai Shūka* (1209) by Fujiwara no Teika, an anthologist of two Imperial Anthologies. *Honkadori* was regarded as a means to enrich and deepen the world behind the poem and represented a way to show respect to the great poems of the past. This is one of the reasons why the intensive study of old poems was encouraged in the periods during which the Imperial Anthologies were compiled. To understand the beauty of *waka*, the audience needs to have deep knowledge of the precedents on which the poems are based (Konishi et al., 1984). If a poem was a *honkadori* of an older poem, they could have easily pointed it out. The widespread practice of *honkadori* justifies assuming a phylogenetic network structure among poems. It also allows us to measure the cultural influence of a poem based on the number of children in the phylogenetic network.

Dataset

The present study used the *waka* database (https://lapis.nichibun.ac.jp/waka/index_era.html) created by Katsuhiro Seta and maintained by the International Research Center for Japanese Studies. Each poem in the database is included in an anthology, the date of publication of which ranges from ca. 700 to 1527. This database covers the periods when Japanese classical poetry was prolific and creative. Some anthologies lack a date of publication and were used as the validation set. Because the database contains other forms of poetry, poems with more or fewer than five lines were excluded. In addition, poems with lacunae were excluded. Regarding the data cleansing, line separators were deleted and all characters were replaced with *hiragana*, the most widely used phonetic lettering system in Japanese.

Several variants of poems may be found in an anthology, and some famous poems are included in multiple anthologies, resulting in multiple entries of a single poem in the database. To determine the best method to eliminate multiple occurrences of a single poem, the Levenshtein ratio l_{ij}, which measures the closeness of two sequences, of poems i and j was calculated. Figure 1 shows a histogram of the maximal Levenshtein ratio of each poem and all other poems, *i.e.*, $\max_j l_{ij}$. This histogram is bimodal, suggesting that a poem should be identified as an existing one if their Levenshtein ratio exceeds 0.8. Thus, newer poems satisfying this criterion were excluded from the dataset, except in the analysis shown in Fig. 7. The database contains a total of 206 965 poems in 496 anthologies. The data cleansing and elimination of the poems identified as existing ones resulted in 146 738 distinct poems, 6343 of which belonged to the validation set.

To examine the historical development of *waka*, a total of 24 training sets were used. Training set 0 contains all poems. Because *Kin'yōshū* has three versions, there are 23 versions of the Imperial Anthologies (Table 1). Training set i ($i = 1, \ldots, 23$) contains all the anthologies no later than Imperial Anthology i. The single validation set was used for all training sets.

The database of *honkadori* was constructed from a modern critical edition of *Shinkokinshū* (Tanaka and Akase, 1992), in which *honkadori* reached its highest sophistication. There are 418 poems with *honkadori* in *Shinkokinshū*. There are 450 *honkadori*-original pairs because some poems borrow phrases from more than one older poem.

Neural language model and distance metric

The vector representation of poems was obtained using BERT (Devlin et al., 2019), a Transformer-based language model (Vaswani et al., 2017). The language model and data loader were implemented with HuggingFace Transformers (Wolf et al., 2020) using default parameters unless otherwise stated. The dimensionality of the feed-forward layers was set to 768. The training sets and validation set were tokenized by SentencePiece with a token size of 5000 (Kudo and Richardson, 2018). Early stopping was used to prevent overfitting of the neural language model. If the validation loss had not improved for the last 10 epochs, the training was stopped, and the parameter values with the smallest loss were saved. The model was trained four times with different initial parameter values.

To measure the similarity between poems, the average vector of the intermediate representation of all tokens of the last hidden layer was calculated for each poem. The Euclid distance between the vectors was used as the similarity measure. Because BERT is not a metric learning model, for the average vectors $\boldsymbol{a}$, $\boldsymbol{b}$, $\boldsymbol{c}$, and $\boldsymbol{d}$, the comparability of the distances $|\boldsymbol{a}-\boldsymbol{b}|$ and $|\boldsymbol{c}-\boldsymbol{d}|$ is not necessarily guaranteed. This means that methods that depend on the axioms of metric space, such as linear regression and logistic regression, might not be trustworthy. Rather, methods that depend only on a comparison of the distances from a vector, such as $|\boldsymbol{a}-\boldsymbol{b}|$ and $|\boldsymbol{a}-\boldsymbol{c}|$, should be used, as these can be more reliable. Thus, in classification tasks, the k-nearest neighbor algorithm is used throughout this paper instead of logistic regression. k with the greatest validation accuracy was selected from $k = 1, 3, 5, 7, 9$ using leave-one-out cross-validation.

The phylogenetic network was estimated as follows. Each poem is a node in the directed network. The parent poem of a given poem is the poem closest to it among the poems older than it. Note that identifying the parent poem of the poem whose vector representation is $\boldsymbol{a}$ is done solely by a comparison of $|\boldsymbol{a}-\boldsymbol{b}|$ and $|\boldsymbol{a}-\boldsymbol{c}|$. Through this construction, the network comprises a set of directed trees. As the poems in the oldest anthology have no parent, they become the root nodes. Throughout this paper, the phylogenetic network estimated by the neural language model trained using training set i is referred to as phylogenetic network i. Phylogenetic network 0 is sometimes simply referred to as the phylogenetic network. An example of phylogenetic network 0 is visualized in Supplementary Figure 1.

The reasonability of the estimated phylogenetic network was evaluated as follows. In molecular phylogenetics, species with similar genotypes are placed close to each other. Hence, if a phylogenetic network is reasonable, each poem and its parent are sufficiently close. This suggests that the distance between a poem and its parent can be a measure of reasonability. However, as we have seen, the distance itself is unreliable. In other words, the meaning of the summation of distances such as $|\boldsymbol{a}-\boldsymbol{b}|+|\boldsymbol{a}-\boldsymbol{c}|$ remains unclear. Thus, the rank order of the distance is more reliable than the raw distance because it is calculated based on a comparison of the distances from only $\boldsymbol{a}$. Let us assume that poem i belongs to anthology a and that there are n_a poems before and m_a poems after anthology a. The parent of poem i is closest to poem i among the preceding n_a poems. Let us define $r_i = (k_i - 1)/(n_a + m_a - 1)$ if the parent is k_i-th closest to poem i among $n_a + m_a$ poems. $\bar{r}$ is defined as the average of r_i over all non-root poems. The smaller the value of $\bar{r}$, the more reasonable the phylogenetic network.

Let us note that $\bar{r}$ can distinguish divergence from and convergence to a poem. Supplementary Figure 2 shows older and newer poems in lighter and darker colors, respectively. In Supplementary Figure 2a, poems are diverging from the oldest one. The phylogenetic network shown by the arrows is constructed by connecting each poem (child) with the closest one among the older poems (parent). Here, all anthologies are assumed to have only one poem each. The parent of the poem marked by the asterisk (poem $*$) is the closest to poem $*$ among all other poems, *i.e.*, $k_* = 1$ and $r_* = 0$. If a phylogenetic network is constructed from the time-reversed data (Supplementary Figure 2b), the parent of poem $*$ is the fifth closest to poem $*$ among all other poems, *i.e.*, $k_* = 5$ and $r_* = 1$. For most of the other poems, k_i and r_i are greater in the time-reversed phylogenetic network. Thus, poems diverging from a poem (Supplementary Figure 2a) and poems converging to a poem (Supplementary Figure 2b) result in a low and a high $\bar{r}$, respectively. Note that time reversal does not affect $\bar{r}$ for continuously transitioning poems (Supplementary Figure 2c, d). Hence, $\bar{r}$ can be used as an index of divergence and convergence.

Results

Validity of the estimated phylogenetic networks

First, the validity of the estimated phylogenetic networks was examined by comparing it with the ground truth. Specifically, the estimated child–parent relationship was compared with the original–*honkadori*-poem relationship described in a critical edition of *Shinkokinshū* (Tanaka and Akase, 1992).

Figure 2 shows the relationship between 418 *honkadori* poems in *Shinkokinshū* (child) and their original poems (parent). The parent–child relationship in the estimated phylogenetic network is classified into seven categories: the original poem of *honkadori* is the parent (parent), an ancestor but not the parent (ancestor), a node within 4 hops (close relative), a node more distant than 4 hops but in the same connected component (remote relative), in a different connected component (unconnected), found in an anthology later than *Shinkokinshū* (anachronism), and not found in the database (not found). If a poem is a *honkadori* poem of more than one old poem, the best one was used. This figure shows that 6.9% of *honkadori* pairs are captured as the parent and child in phylogenetic network 0 and that 72.9% of pairs belong to different connected components. This might not seem to be a very accurate estimate.

However, this is a statistically significant result if it is compared with a hypothetical random phylogenetic network (null hypothesis), in which the parents of poems in *Shinkokinshū* are randomly drawn from 48 057 poems older than *Shinkokinshū*. Three original poems at most have been identified for a *honkadori* poem in *Shinkokinshū*. If one of these poems is the parent in the estimated phylogenetic network, the *honkadori* relationship is regarded to be guessed right. Thus, the parent of a poem in the random phylogenetic network is identical to one of the original poems of *honkadori* with a probability of $p_h = 3/48\,057$ at most. The number of parents identical to the original, n_h, obeys the binomial distribution $P(n_h) = p_h^{n_h}(1-p_h)^{n-n_h} n!/\{n_h!(n-n_h)!\}$, where $n = 4\times 418$ because Fig. 2 is the average of the four phylogenetic networks. The probability that $n_k \geq 17$, *i.e.*, more than 1%, which is much lower than 6.9%, of the parents are identical to the original is less than $p = 5 \times 10^{-32}$. Hence, the phylogenetic network estimated by using the vector representation of a neural language model succeeds in identifying at least some of the *honkadori* relationships. This conclusion is robust for change in the model parameter values of BERT (Supplementary Figure 3).

Consistency and reasonability of phylogenetic networks and the arrow of time

Second, the cultural influence of poems judged by the models with different initial parameter values was quantified to examine the consistency of the estimation. If a poem has a large number of children, it is potentially the original poem of a large number of *honkadori* poems, implying that it should be judged to be a poem with a great influence on the following poems. Hence, this paper uses the number of children of a poem in the estimated phylogenetic network as a proxy variable of its cultural influence. However, because, by construction, an earlier poem tends to have a larger number of children than a later one, it is not fair to compare the raw number of children. Thus, by standardizing the number of children for each anthology, we can quantify the cultural influence of each poem relative to other poems in the same anthology. The Spearman correlation coefficients of the standardized number of children averaged for all six pairs of the four neural language models are shown in Fig. 3. Most exhibit a positive correlation coefficient. Thus, the estimated cultural influence of a poem is consistent among the neural language models with different initial parameter values.

Third, the reasonability of the structure of the estimated phylogenetic network and the existence of the arrow of time were examined. In molecular phylogenetics, closely related species are placed close to each other in the phylogenetic tree. Similarly, if a phylogenetic network of *waka* is reasonable, each poem and its parent are expected to be sufficiently close. To test this expectation, the phylogenetic networks estimated from the dataset in which the order of anthologies is reversed (reversed) and that in which the order of anthologies is randomly shuffled (shuffled) were made along with phylogenetic network 0 (real). Twenty shuffled networks were made for each neural language model. $\bar{r}$ measures the average distance of the parent–child relationship. Figure 4 shows that the real data exhibit the lowest value of $\bar{r}$. This means that the phylogenetic network estimated from the real data is more reasonable than those estimated from the reversed or shuffled data. The present analysis is consistent with and extends the results of Hughes et al. (2012), which showed a gradual change in literary style. Particularly, by using $\bar{r}$, we can distinguish the real data from the time-reversed data. A low $\bar{r}$ of the real data shows diversification, rather than continuous transition, in *waka*. This result indicates that the arrow of time is present and observable in the evolution of *waka*.

Cultural influence of poems

Fourth, the cultural influence of poems judged from the dataset up to a certain time point was compared with that judged from the whole dataset. Specifically, the number of children in phylogenetic network i was compared with that in phylogenetic network 0. In other words, the congruence of the cultural influence of a poem measured by a language model trained with a limited corpus of poems and that with the whole corpus in the database is examined. It is likened to asking poets in the *Shinkokinshū* (1205) era "Which do you think are most influential among *Kokinshū* (905) poems?" and comparing their answers with ours. The standardized number of children of poems up to Imperial Anthology i in phylogenetic networks i and 0 was positively correlated in most cases (Fig. 5a). Moreover, the standardized number of children until Imperial Anthology i in phylogenetic network i was positively correlated with the number of children after Imperial Anthology i in phylogenetic network 0 (Fig. 5b). Therefore, the cultural influence of a poem from the viewpoint of a certain time point is correlated with the cultural influence of the poem from the viewpoint of a later time point. Analogically, the influence of *Kokinshū* poems evaluated by poets in the *Shinkokinshū* era is positively correlated with our estimation of the influence of *Kokinshū* poems on poems later than *Shinkokinshū*. However, the correlation was not so strong presumably because of the collective memory decay of a poem (Candia et al., 2019).

Interpolation and extrapolation

Fifth, to investigate whether a culture in a period is an intermediate form of those in the preceding and following periods and whether the past history allows us to predict future culture, the interpolation and extrapolation of the Imperial Anthologies were examined using classification. The k-nearest neighbor algorithm discriminating the first Imperial Anthology *Kokinshū* (class label 0) and the last Imperial Anthology *Shinshokukokinshū* (class label 1) was applied to all Imperial Anthologies. Figure 6a shows the average of the class label predicted by the model for all Imperial Anthologies. The validation accuracy of leave-one-out cross-validation is used as the average class label of the labeled anthologies. This figure shows that the average class label of the anthologies in between exhibit intermediate values, that is, this k-nearest neighbor classifier can interpolate the anthologies between the two anthologies. However, the k-nearest neighbor model discriminating *Kokinshū* and *Shinkokinshū* exhibits no signs of extrapolation (Fig. 6b). The anthologies following *Shinkokinshū* are no more *Shinkokinshū*-like than *Shinkokinshū*. Taken together, these results indicate the presence of detectable and gradual, albeit unpredictable, change over time.

Effect of being selected for an Imperial Anthology

Sixth, the effect of being selected for an Imperial Anthology was examined. The Imperial Anthologies were so authoritative that being selected for them is expected to increase the number of children. This expectation was tested by comparing a pair of poems, the first of which is contained in Imperial Anthology i, referred to as x, and the second of which is not contained in Imperial Anthology i but regarded to be similar to the first, referred to as y. All poems that are included in Imperial Anthology i but first appear in an earlier anthology were classified as x. For each x, y was sampled under the condition that it appears first in the same anthology as x and that it gives birth to the same number of children as x in phylogenetic network i in the period preceding Imperial Anthology i. The numbers of children after Imperial Anthology i in phylogenetic network 0 were compared for the pairs of x and y (Fig. 7).

For most Imperial Anthologies, the average number of children of x after the Imperial Anthology tended to be greater than that of y, meaning that the effect of being selected is positive. In most cases, a poem selected for an Imperial Anthology gains 1.1–1.5 times more children than a poem that is not. However, this effect is inconclusive for some anthologies, particularly for the later ones. This may be the result of a decline in quality of, or loss of interest in, these Imperial Anthologies (Keene, 1999). The positive effect of most Imperial Anthologies can be interpreted in two ways. The first is that entry into the Imperial Anthologies boosted its fame and increased its number of children. This is a form of the Matthew effect (Merton, 1968). The second is that x is closer to the taste of a later period than y, and thereby had more children in later periods. At any rate, these results indicate that the number of children is largely affected by chance.

The difference among the three versions of *Kin'yōshu* (ID 5–7) is worth noticing (Supplementary Figure 4). Of these, being selected for the first and third versions (ID 5 and 7) does not seem to affect the number of children. In other words, it exhibits a weaker effect than the second version (ID 6). This is consistent with the fact that the second version was the most circulated. Particularly, the third version had been forgotten until the nineteenth century (Keene, 1999). The difference among them supports the existence of the Matthew effect.

Model

Both the chance factor's role in the number of children and the impossibility of future prediction suggest that randomness is a key feature of the evolution of *waka*. This leads us to a model that qualitatively replicates these results. This model assumes that a poem is generated in the vicinity of existing poems. Specifically, poem 0 is generated on $\boldsymbol{x}_0 = \boldsymbol{0}$, where $\boldsymbol{x}_i$ is a d-dimensional vector. Poem $t \geq 1$ is randomly drawn from the Gaussian mixture

$$p(\boldsymbol{x}_t) = \sum_{s=0}^{t-1} \frac{k^{t-1-s}(1-k)}{1-k^t} \frac{1}{(2\pi)^{d/2}} \exp\left(-\frac{|\boldsymbol{x}_t - \alpha \boldsymbol{x}_s|^2}{2}\right),$$

where k is the decay constant and α is a positive constant less than one. α makes this process a mean-reverting self-excitatory process. A poem facilitates the generation of another poem in the vicinity of itself and the origin. This is a self-exciting stochastic dynamical model (Golosovsky and Solomon, 2012). Because the influence of a poem decays at rate k, the poems as a whole can exhibit a long-term drift. If a poem is generated in the close vicinity of another poem by chance, the number of children of the latter can be boosted.

Figure 8 shows the results of the model with 24 000 poems and the following parameter values: $d = 100$, $k = \exp(-1/5000)$, and $\alpha = 0.6$. To obtain poems in the steady state, a total of 74 000 poems were generated. The first 50 000 poems were discarded, and the last 24 000 poems were divided into 24 anthologies containing 1000 poems each. The k-nearest neighbor algorithm with the first anthology as class 0 and the 12th anthology as class 1 shows a steady increase up to the 11th anthology, but plateaus thereafter (Fig. 8a). This is consistent with the possibility of interpolation and the impossibility of extrapolation. The phylogenetic network made in the same way as *waka* has a lower $\bar{r}$ than those made from the time-reversed or shuffled anthologies (Fig. 8b). This is because poems diverging from and converging to a point exhibit a low and high $\bar{r}$, respectively. The standardized number of children before anthology i is positively correlated with that after anthology i (Fig. 8c). Mean reversion is essential because the model without mean reversion ($\alpha = 1$) exhibits a much weaker correlation (Supplementary Figure 5). Thus, the model succeeded in replicating the results qualitatively.

Discussion

The present study attempted to elucidate the evolutionary dynamics of culture in estimating the phylogenetic network of *waka*. The results are summarized as follows. First, the phylogenetic network reflects a significant part of the *honkadori* relationship. Second, the vector representation obtained using BERT gives reproducible results. Third, the estimated phylogenetic network is distinguishable from the phylogenetic network constructed from the time-reversed data. That is, the arrow of time is observable in a literary genre. Fourth, the cultural influence of a poem, which is measured by the number of children at a certain time point, is correlated with that at a later time point. Fifth, we can successfully perform the interpolation, but not the extrapolation, of the poetic style. In other words, we cannot predict a style in the future. Sixth, the number of phylogenetic children increases after being selected for an anthology. Last, a mean-reverting self-excitation model replicates these results. If not complete, these results are at least partial answers to the questions raised at the beginning of the paper.

It is quite natural to ask whether these results are universal to other literary genres and creative works. Although *honkadori* is a rhetorical technique characteristic of *waka*, borrowing and imitating phrases from old great literary works are prevalent in a diverse range of classical literary genres. Kato et al. (1979) pointed out that the *xīkūn* style in the early Sung dynasty and "Waste Land" by T. S. Eliot can be seen as a parallel to *honkadori*. Gahan (1987) presented a detailed analysis of a tragedy from the Silver Age of Latin literature and illustrated the abundance of *imitatio* and *aemulatio*, that is, the technique of borrowing phrases and ideas. Hence, the technique of borrowing is universal to classical literary genres. An analysis similar to the present study can also shed new light on these genres.

To analyze other creative works, the method of estimating the phylogenetic structure in this paper should be extended. The present study has assumed that a poem is a child of another poem and ignored exogenous factors. This assumption is justifiable because *waka* had been the most authoritative genre in Japanese literature and, consequently, influence from *waka* to other genres in Japanese literature exceeds *vice versa*. However, because the influence of Chinese literature is indisputable (Konishi et al., 1984), this should be taken into account in future analyses. In addition, although the present study has assumed that there is only one parent poem for a given poem, a poem can be a *honkadori* poem of multiple poems. Thus, phylogenetic networks allowing multiple parents should be examined in the future.

The effect of being selected for an Imperial Anthology is greater in earlier Imperial Anthologies but diminished in later ones. There are two possible explanations for this result. First, the Imperial Anthologies after *Shinkokinshū* were regarded to be of low quality and thus less intensively studied (Keene, 1999). Second, poems in older anthologies tend to have a larger number of children than those in newer ones. Thus, poems in the later Imperial Anthologies tend to have a smaller number of children, deteriorating the signal-to-noise ratio.

The present paper has proposed a model that replicates the results qualitatively. In this model, a poem is an event in a self-exciting point process. Although the timing of poem generation was not formulated in the present model, the spatiotemporal Hawkes process might be hopeful. Identifying a poem with an individual in population genetics, we can regard this model to be closely related to the neutral theories, which stress the importance of the interaction of selection and drift (Kimura et al., 1968; Ohta, 2002; Akashi et al., 2012). Testing whether the word frequency obeys the distribution predicted from the neutral model will be of interest (Bentley and Shennan, 2003; Bentley et al., 2004). Although there may sometimes be fixed directionality in evolution, the results of the present study indicate that the cultural evolution of *waka* is approximated by mean-reverting self-excitation.

However, this study has some limitations. First, the training set may not have been sufficiently large. Including other literary genres in the corpus might improve the performance. In particular, the poems in *The Tale of Genji* had a substantial influence on the history of classical Japanese literature. Taking the influence from and to *The Tale of Genji* into account could enable us to create a more holistic picture of the historical development of Japanese literature. However, other literary genres, such as novels, lack dating more often than *waka*. In fact, many of the poems that were excluded from the training sets and included in the validation set because of missing dating were from novels. Thus, this may be difficult to implement. Second, using the chronology of the *waka* anthologies as the chronology of the poems in them may not be appropriate in some cases. This is because an anthology can contain a poem by a poet from an older generation. Thus, the order of poem writing and publication may be reversed. Incorporating information from the poem description (*kotobagaki*), author, and volume name in the anthology might also improve the results. Third, as stated above, the influence of Chinese literature was not taken into consideration in the present analysis.

There are several possible future directions. First, the present paper has left a detailed analysis of the specificities of each anthology for future work. What is and is not influential, what determines the strength of influence, and what are the long-term trends in the strength of influence remain questions that need to be addressed. Second, examining whether the phylogenetic network reflects the schools of *waka*, such as *Nijō*, *Kyōgoku*, and *Reizei*, would also be of interest. Third, a method to measure the speed of evolution needs to be developed. This was difficult to measure in the present analysis, which utilized an intermediate vector representation of a neural language model. If a large number of nearly identical poems are in the training set, they might take a large volume in the vector space. Conversely, if we introduce Chinese poetry into the training set, the volume covered by *waka* would be compressed. This means that the distance between successive poems can be affected by both the speed of evolution and the relative abundance of similar poems. Thus, the development of a measure of evolution speed that is insensitive to the relative abundance is needed. Whether the dynamics follow biased cultural transmission (Henrich, 2001) would also be of interest. Because the styles of English authors are similar among contemporaries and differ from preceding generations (Hughes et al., 2012), conformity bias in a generation and anti-conformity bias between generations could also be observed in *waka*. In addition, we might be able to measure conformity in a school and anti-conformity between schools. Fourth, applying this method to other literary genres and creative works would be of interest. Specifically, whether similar results can be obtained for longer works such as novels should be examined. Other forms of creative works, such as music (Nakamura and Kaneko, 2019), fine arts (Cetinic et al., 2019; Sandoval et al., 2019), and Internet memes, should also be subjects of the present method. Self-supervised representation learning, such as SimCLR (Chen et al., 2020) for images, gives a vector representation of input without labeled data. This can be used in the same way as the vector representations produced by language models. Fifth, the present study may provide hints about how to train generative artificial intelligence (AI) with output of another generative AI. This procedure has been reported to deteriorate the quality of output (Alemohammad et al., 2023; Shumailov et al., 2024). If this is the result of an unbounded random walk of the generated data, it can be mitigated by mean reversion, which led to the persistent influence of poems in the present model. Similarly to the model's mean reversion, selecting the output that is closer to the mean of existing creative works could prevent deterioration. Furthermore, in multimodal generative AI, natural photos and sounds could not only serve as the "mean" but also prevent model collapse (Shumailov et al., 2024).

References

Akashi H, Osada N, Ohta T (2012). Weak Selection and Protein Evolution. *Genetics* 192(1):15–31

Alemohammad S, Casco-Rodriguez J, Luzi L, Humayun AI, Babaei H, LeJeune D, Siahkoohi A, Baraniuk RG (2023). Self-consuming generative models go MAD. *arXiv preprint arXiv:2307.01850*

Atkinson QD, Gray RD (2005). Curious parallels and curious connections–phylogenetic thinking in biology and historical linguistics. *Systematic Biology* 54(4):513–526

Barbrook AC, Howe CJ, Blake N, Robinson P (1998). The phylogeny of the canterbury tales. *Nature* 394(6696):839–839

Bentley RA, Hahn MW, Shennan SJ (2004). Random drift and culture change. *Proceedings of the Royal Society of London. Series B: Biological Sciences* 271(1547):1443–1450

Bentley RA, Shennan SJ (2003). Cultural transmission and stochastic network growth. *American Antiquity* 68(3):459–485

Bialock DT (1994). Voice, text, and the question of poetic borrowing in late classical japanese poetry. *Harvard Journal of Asiatic Studies* 54(1):181–231

Boesch C, Boesch H (1990). Tool use and tool making in wild chimpanzees. *Folia Primatologica* 54(1-2):86–99

Boudreau KJ, Guinan EC, Lakhani KR, Riedl C (2016). Looking across and looking beyond the knowledge frontier: Intellectual distance, novelty, and resource allocation in science. *Management Science* 62(10):2765–2783

Brower RH, Miner E (1961). *Japanese court poetry.* Stanford University Press

Candia C, Jara-Figueroa C, Rodriguez-Sickert C, Barabási AL, Hidalgo CA (2019). The universal decay of collective memory and attention. *Nature Human Behaviour* 3(1):82–91

Cavalli-Sforza L, Feldman MW (1973a). Models for cultural inheritance I. group mean and within group variation. *Theoretical Population Biology* 4(1):42–55

Cavalli-Sforza LL, Feldman MW (1973b). Cultural versus biological inheritance: phenotypic transmission from parents to children. (A theory of the effect of parental phenotypes on children's phenotypes). *American Journal of Human Genetics* 25(6):618

Cetinic E, Lipic T, Grgic S (2019). A deep learning perspective on beauty, sentiment, and remembrance of art. *IEEE Access* 7:73694–73710

Chen T, Kornblith S, Norouzi M, Hinton G (2020). A simple framework for contrastive learning of visual representations. In *International conference on machine learning* pages 1597–1607. PMLR

Currie TE, Greenhill SJ, Gray RD, Hasegawa T, Mace R (2010). Rise and fall of political complexity in island south-east asia and the pacific. *Nature* 467(7317):801–804

Devlin J, Chang M, Lee K, Toutanova K (2019). BERT: pre-training of deep bidirectional transformers for language understanding. In Burstein J, Doran C, Solorio T, editors, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)* pages 4171–4186. Association for Computational Linguistics

Dias Z, Rocha A, Goldenstein S (2010). First steps toward image phylogeny. In *2010 IEEE International Workshop on Information Forensics and Security* pages 1–6. IEEE

Dias Z, Rocha A, Goldenstein S (2011). Video phylogeny: Recovering near-duplicate video relationships. In *2011 IEEE International Workshop on Information Forensics and Security* pages 1–6. IEEE

Dugatkin LA, Godin JGJ (1992). Reversal of female mate choice by copying in the guppy (*Poecilia reticulata*). *Proceedings: Biological Sciences* 249(1325):179–184

Eom YH, Fortunato S (2011). Characterizing and modeling citation dynamics. *PLOS ONE* 6(9):e24926

Fortunato S, Bergstrom CT, Börner K, Evans JA, Helbing D, Milojević S, Petersen AM, Radicchi F, Sinatra R, Uzzi B, Vespignani A, Waltman L, Wang D, Barabási AL (2018). Science of science. *Science* 359(6379):eaao0185

Foster JG, Rzhetsky A, Evans JA (2015). Tradition and innovation in scientists' research strategies. *American Sociological Review* 80(5):875–908

Gahan J (1987). *Imitatio* and *aemulatio* in seneca's phaedra. *Latomus* 46(2):380–387

Goldberg LA, Goldberg PW, Phillips CA, Sorkin GB (1998). Constructing computer virus phylogenies. *Journal of Algorithms* 26(1):188–208

Golosovsky M, Solomon S (2012). Stochastic dynamical model of a growing citation network based on a self-exciting point process. *Physical Review Letters* 109(9):098701

Gray RD, Atkinson QD (2003). Language-tree divergence times support the anatolian theory of indo-european origin. *Nature* 426(6965):435–439

Harmand S, Lewis JE, Feibel CS, Lepre CJ, Prat S, Lenoble A, Boës X, Quinn RL, Brenet M, Arroyo A et al (2015). 3.3-million-year-old stone tools from Lomekwi 3, West Turkana, Kenya. *Nature* 521(7552):310–315

Henrich J (2001). Cultural transmission and the diffusion of innovations: Adoption dynamics indicate that biased cultural transmission is the predominate force in behavioral change. *American Anthropologist* 103(4):992–1013

Holden CJ, Mace R (2003). Spread of cattle led to the loss of matrilineal descent in africa: a coevolutionary analysis. *Proceedings of the Royal Society of London. Series B: Biological Sciences* 270(1532):2425–2433

Hughes JM, Foti NJ, Krakauer DC, Rockmore DN (2012). Quantitative patterns of stylistic influence in the evolution of literature. *Proceedings of the National Academy of Sciences* 109(20):7682–7686

Kanojia D, Dubey A, Kulkarni M, Bhattacharyya P, Haffari G (2019). Utilizing word embeddings based features for phylogenetic tree generation of sanskrit texts. In *Proceedings of the 6th International Sanskrit Computational Linguistics Symposium* pages 152–165

Kato S, Chibbett DG, Dore RP, Sanderson D, Etiemble, McClellan E (1979). *A History of Japanese Literature*. Kodansha International

Keene D (1999). *Seeds in the heart: Japanese literature from earliest times to the late sixteenth century.* Number 1 in A history of Japanese literature. Columbia University Press

Kim D, Cerigo DB, Jeong H, Youn H (2016). Technological novelty profile and invention's future impact. *EPJ Data Science* 5:1–15

Kimura M et al (1968). Evolutionary rate at the molecular level. *Nature* 217(5129):624–626

Konishi J, Miner ER, Gatten A, Teele NJ, Harbison M (1984). *A History of Japanese Literature*. Princeton University Press

Kudo T, Richardson J (2018). SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. In Blanco E, Lu W, editors, *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations* pages 66–71 Brussels, Belgium. Association for Computational Linguistics

Larivière V, Haustein S, Börner K (2015). Long-distance interdisciplinarity leads to higher scientific impact. *PLOS ONE* 10(3):e0122565

Leahey E, Moody J (2014). Sociological innovation through subfield integration. *Social Currents* 1(3):228–256

Lee S, Bozeman B (2005). The impact of research collaboration on scientific productivity. *Social Studies of Science* 35(5):673–702

Marmerola GD, Oikawa MA, Dias Z, Goldenstein S, Rocha A (2016). On the reconstruction of text phylogeny trees: Evaluation and analysis of textual relationships. *PLOS ONE* 11(12):1–35

Merton RK (1968). The matthew effect in science. *Science* 159(3810):56–63

Mesoudi A (2011). *Cultural evolution: How Darwinian theory can explain human culture and synthesize the social sciences*. University of Chicago Press

Michel JB, Shen YK, Aiden AP, Veres A, Gray MK, Team GB, Pickett JP, Hoiberg D, Clancy D, Norvig P et al (2011). Quantitative analysis of culture using millions of digitized books. *Science* 331(6014):176–182

Nakamura E, Kaneko K (2019). Statistical evolutionary laws in music styles. *Scientific Reports* 9(1):15993

Nucci M, Tagliasacchi M, Tubaro S (2013). A phylogenetic analysis of near-duplicate audio tracks. In *2013 IEEE 15th International Workshop on Multimedia Signal Processing (MMSP)* pages 099–104. IEEE

O'Brien MJ, Darwent J, Lyman RL (2001). Cladistics is useful for reconstructing archaeological phylogenies: Palaeoindian points from the southeastern united states. *Journal of Archaeological Science* 28(10):1115–1136

Ohta T (2002). Near-neutrality in evolution of genes and gene regulation. *Proceedings of the National Academy of Sciences* 99(25):16134–16137

Petersen AM, Fortunato S, Pan RK, Kaski K, Penner O, Rungi A, Riccaboni M, Stanley HE, Pammolli F (2014). Reputation and impact in academic careers. *Proceedings of the National Academy of Sciences* 111(43):15316–15321

Price DJDS (1965). Networks of scientific papers: The pattern of bibliographic references indicates the nature of the scientific research front. *Science* 149(3683):510–515

Sandoval C, Pirogova E, Lech M (2019). Two-stage deep learning approach to the classification of fine-art paintings. *IEEE Access* 7:41770–41781

Shumailov I, Shumaylov Z, Zhao Y, Papernot N, Anderson R, Gal Y (2024). Ai models collapse when trained on recursively generated data. *Nature* 631(8022):755–759

Slater P (1986). The cultural transmission of bird song. *Trends in Ecology & Evolution* 1(4):94–97

Straffon LM (2016). *Cultural phylogenetics: concepts and applications in archaeology* volume 4. Springer

Tanaka Y, Akase S (1992). *Shinkokinwakashū*. Number 11 in *Shinnihonkotenbungakutaikei*. Iwanami Shoten

Tehrani J, Collard M (2002). Investigating cultural evolution through biological phylogenetic analyses of turkmen textiles. *Journal of Anthropological Archaeology* 21(4):443–463

Uzzi B, Mukherjee S, Stringer M, Jones B (2013). Atypical combinations and scientific impact. *Science* 342(6157):468–472

Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Ł, Polosukhin I (2017). Attention is all you need. *Advances in Neural Information Processing Systems* 30

Wang D, Song C, Barabási AL (2013). Quantifying long-term scientific impact. *Science* 342(6154):127–132

Wang J, Veugelers R, Stephan P (2017). Bias against novelty in science: A cautionary tale for users of bibliometric indicators. *Research Policy* 46(8):1416–1436

Wei X, Valler NC, Prakash BA, Neamtiu I, Faloutsos M, Faloutsos C (2013). Competing memes propagation on networks: A network science perspective. *IEEE Journal on Selected Areas in Communications* 31(6):1049–1060

Weng L, Menczer F, Ahn YY (2014). Predicting successful memes using network and community structure. In *Proceedings of the international AAAI conference on web and social media* volume 8 pages 535–544

Wolf T, Debut L, Sanh V, Chaumond J, Delangue C, Moi A, Cistac P, Rault T, Louf R, Funtowicz M et al (2020). Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations* pages 38–45

Acknowledgments

This work was supported by JSPS KAKENHI Grant Number JP22K18526.

Competing interests

The author declares no competing interests.

Data availability

All data used in this paper can be found at https://lapis.nichibun.ac.jp/waka/index_era.html. The program files for data preprocessing, model training, phylogenetic network reconstruction, and figure generation can be found at https://github.com/tanaka-takuma-lab/.

Ethical approval

Not applicable.

Informed consent

Not applicable.

Author contributions

Not applicable.

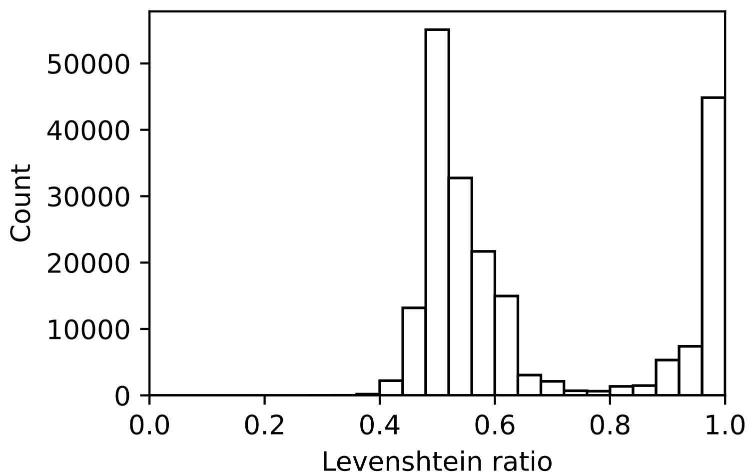

Figure 1: Histogram of the maximal Levenshtein ratio, $\max_j l_{ij}$, for all i.

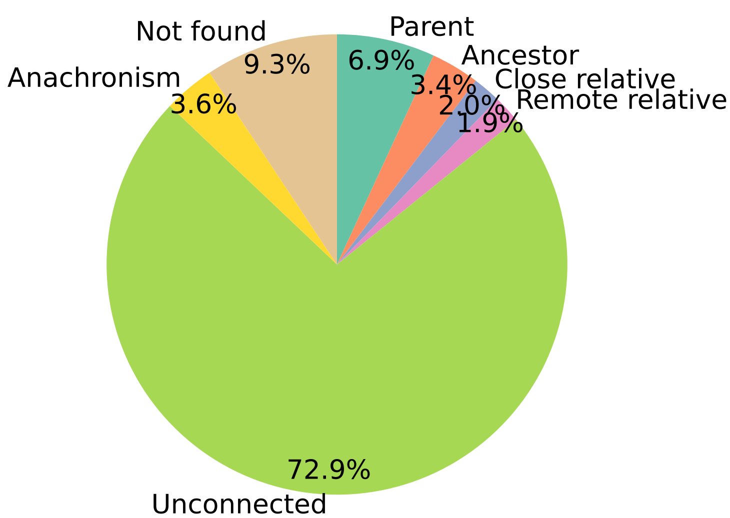

Figure 2: Classification of *honkadori* relationships in the phylogenetic network. The average for four neural language models is shown.

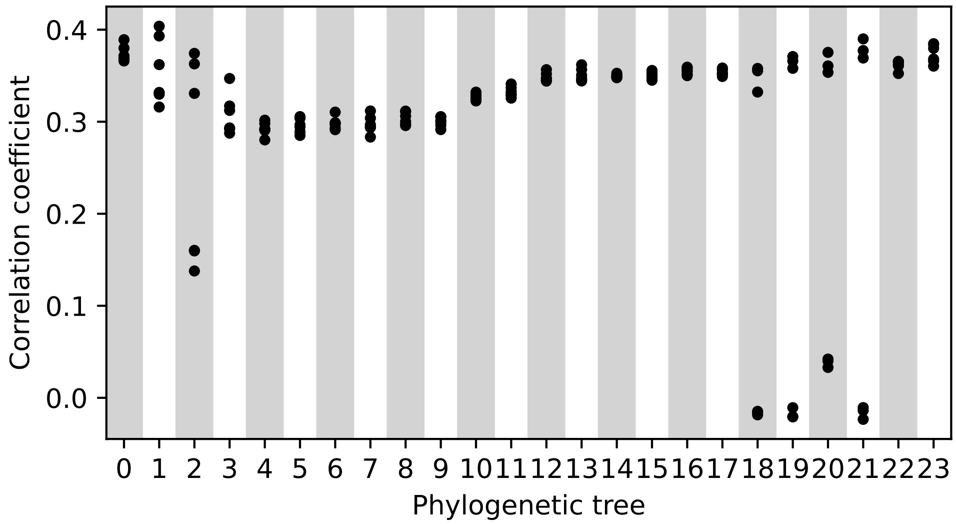

Figure 3: Correlation coefficients of the standardized numbers of children in phylogenetic network i between neural language models with different initial conditions.

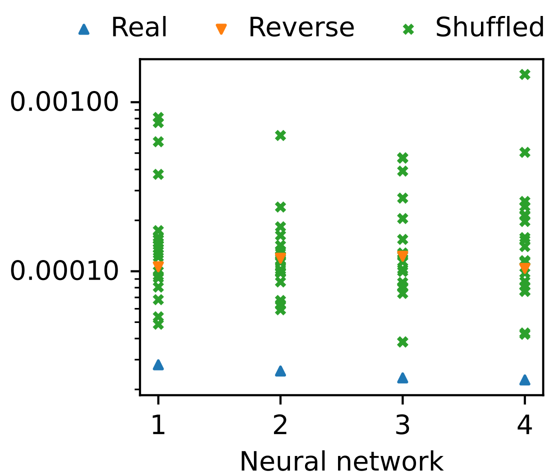

Figure 4: The $\bar{r}$ values for four neural language models with different initial conditions.

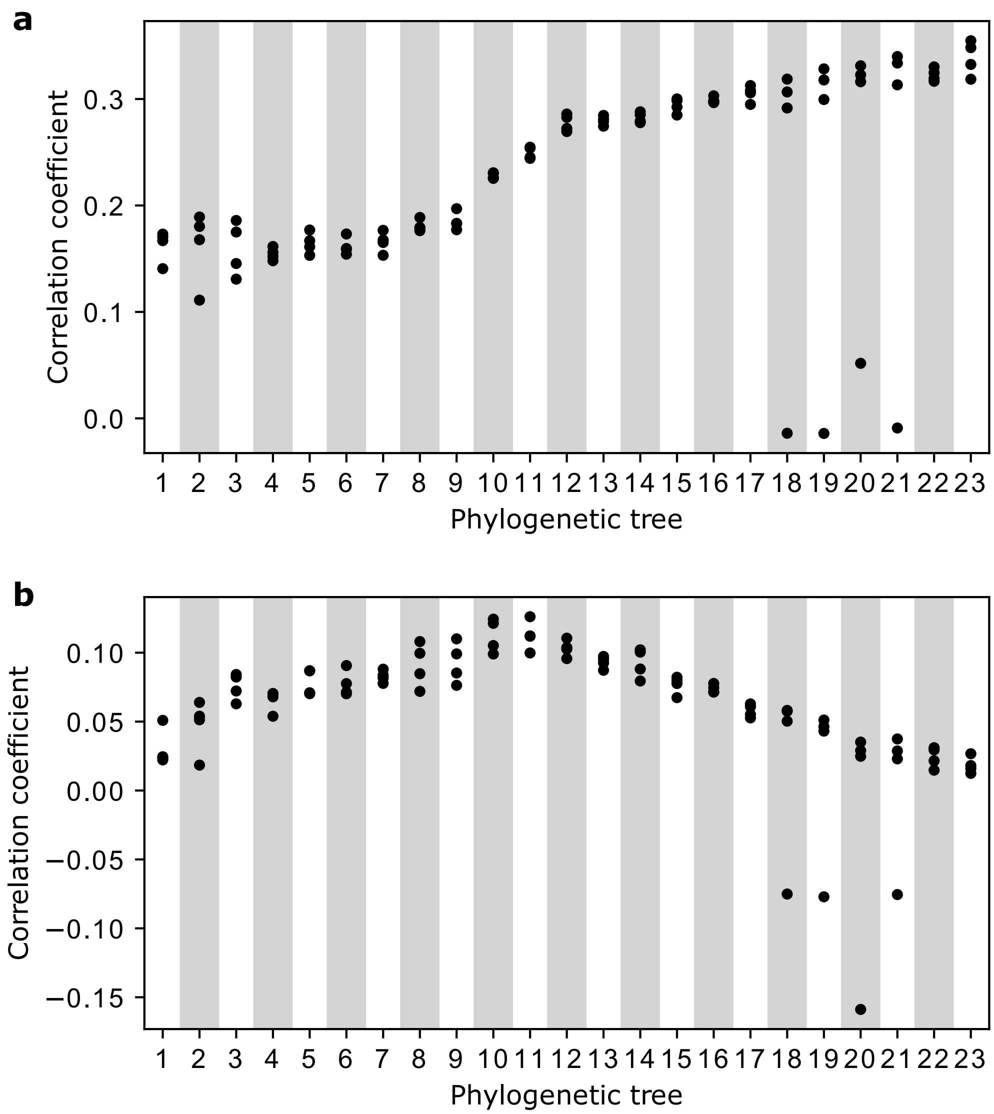

Figure 5: Correlation coefficients of the standardized numbers of (a) children before Imperial Anthology i in phylogenetic network i and phylogenetic network 0 and (b) children before Imperial Anthology i in phylogenetic network i and children after Imperial Anthology i in phylogenetic network 0. The panel labels indicate i.

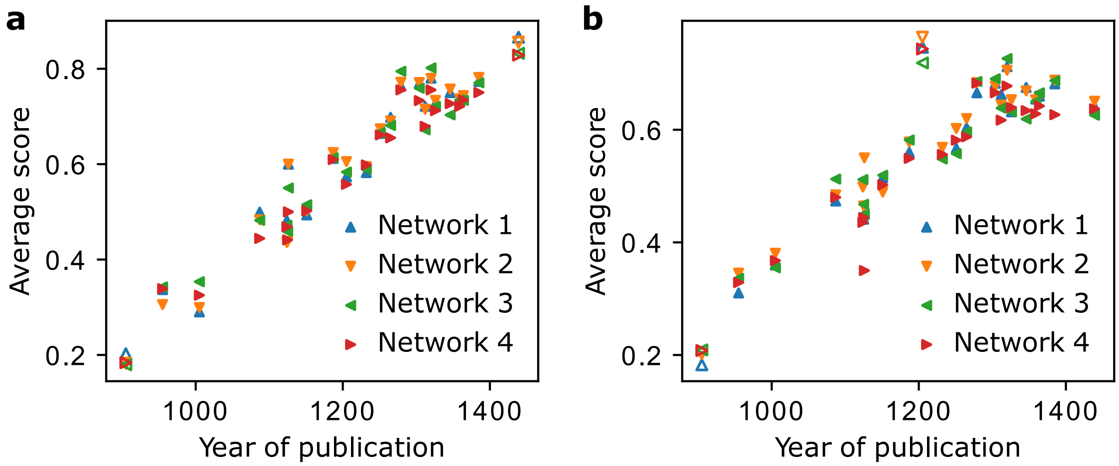

Figure 6: The k-nearest neighbor classification discriminating (a) the first (*Kokinshū*) and last (*Shinshokukokinshū*) Imperial Anthologies, and (b) the first and tenth (*Shinkokinshū*) Imperial Anthologies. The results for the anthologies used as the training set of the k-nearest neighbor classification are indicated by open triangles.

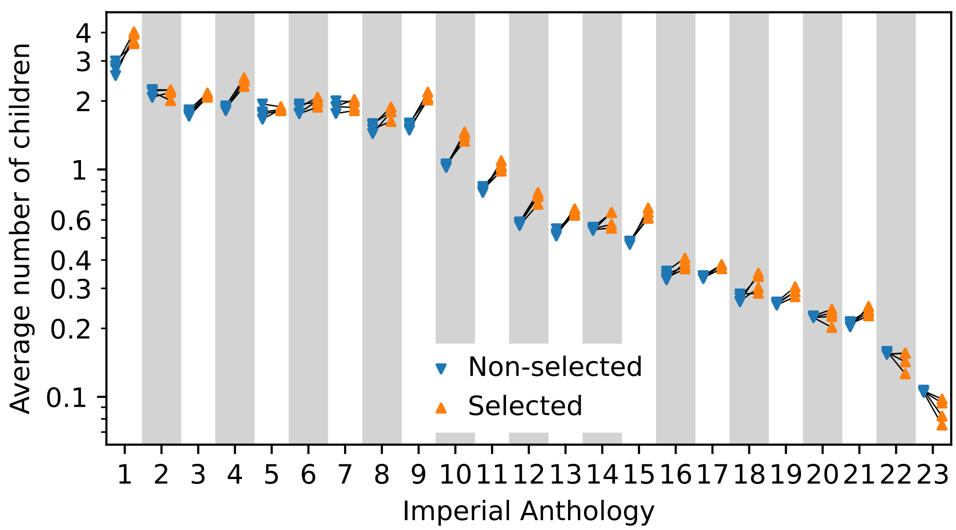

Figure 7: Effects of being selected for an Imperial Anthology. The average numbers of children of poems selected and not selected for Imperial Anthology i after this anthology are compared in panel i. The phylogenetic networks estimated from the neural language models with the same initial conditions are connected by lines.

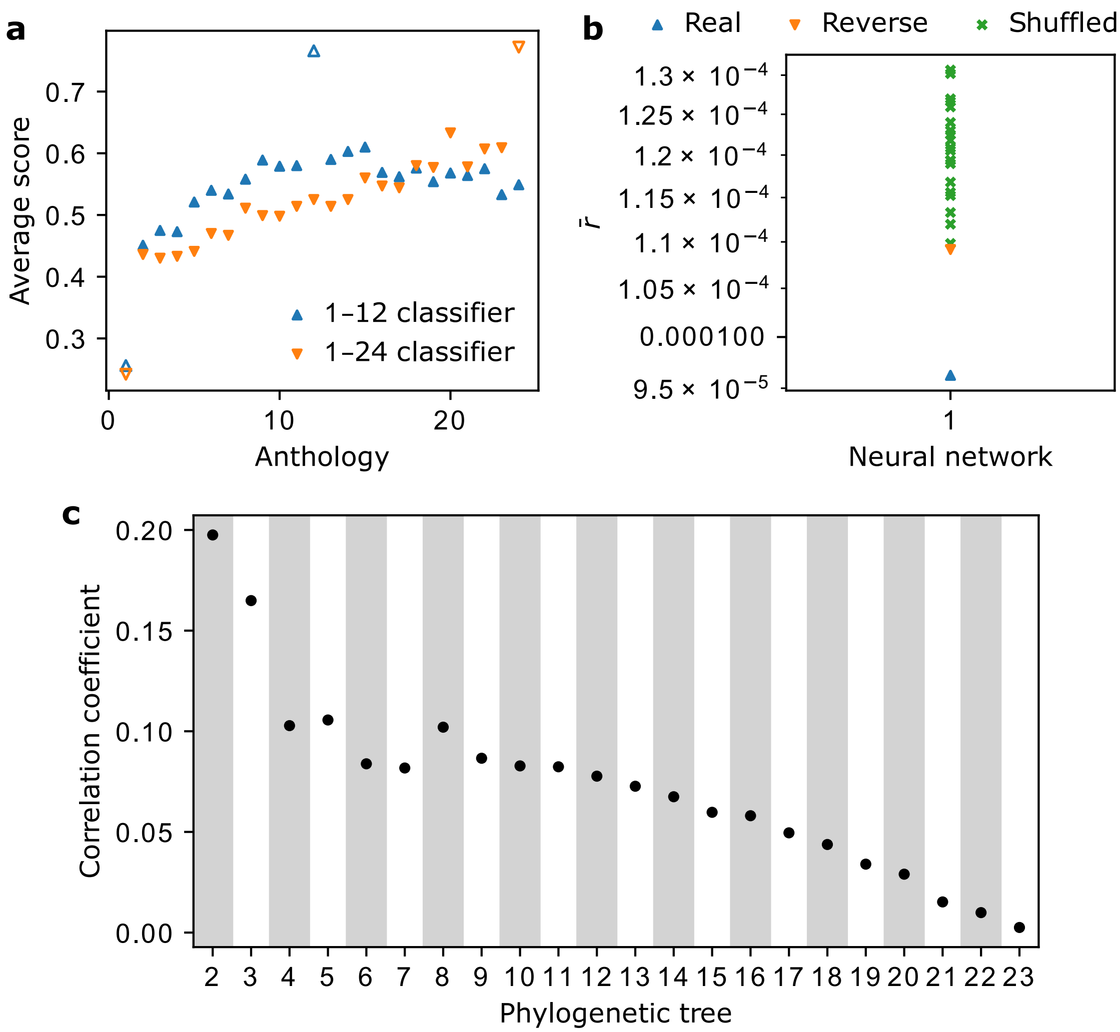

Figure 8: Model results. (a) Performance of k-nearest neighbor classifiers discriminating anthologies 1 and 12 and anthologies 1 and 24. (b) Values of $\bar{r}$ for the real, time-reversed, and shuffled data. (c) Correlation coefficient between the standardized numbers of children before and after anthology i.

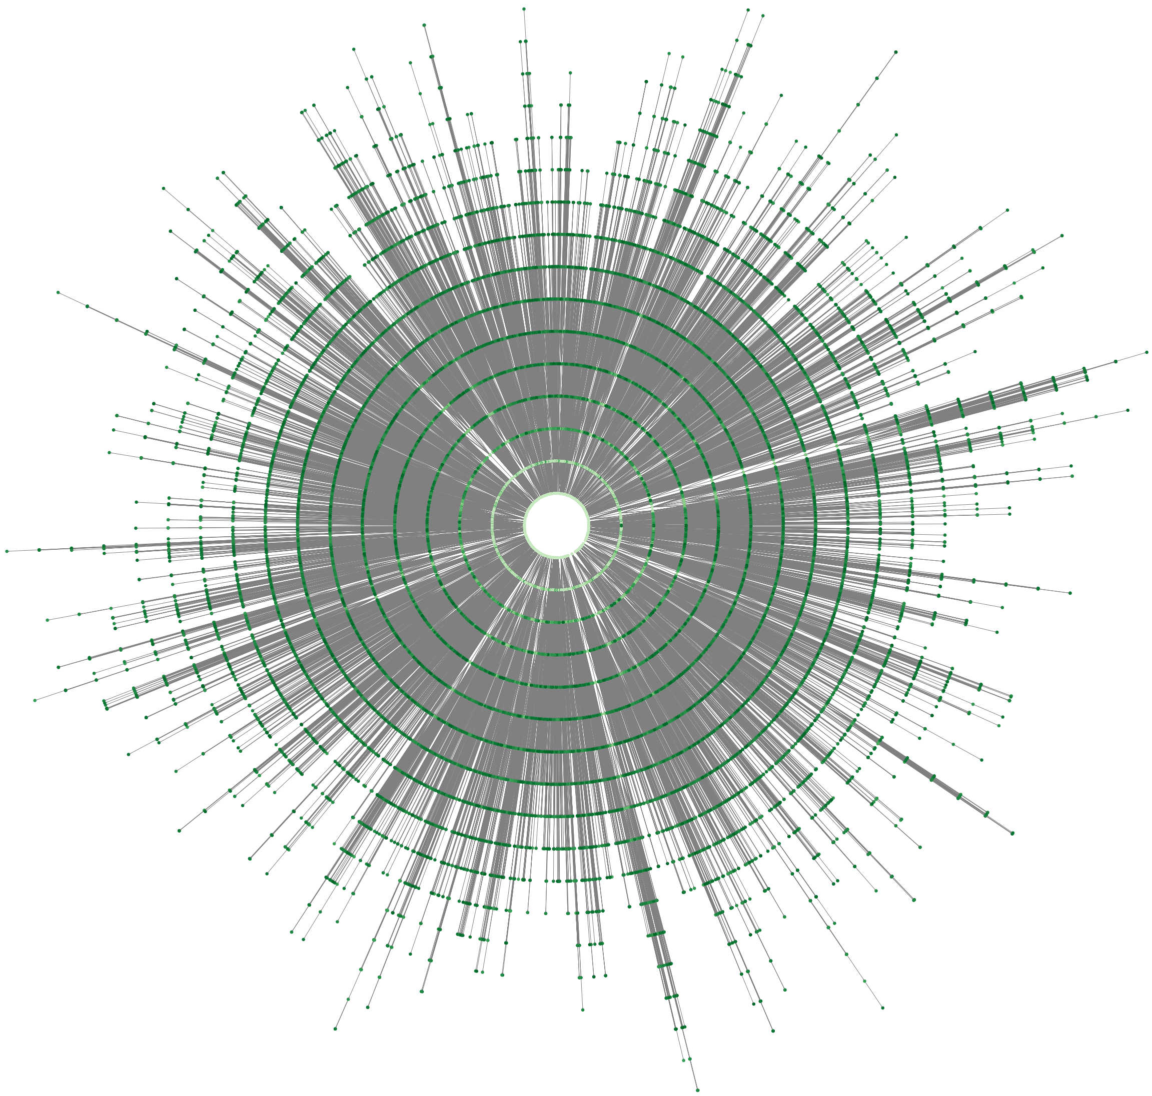

Supplementary Figure 1: A reconstructed phylogenetic network. Poems and parent–child relationships are indicated by colored disks and arrows, respectively. Light and dark green indicate earlier and later poems, respectively. Each of the concentric circles contains one generation of poems, the disks on the innermost circle being the poems in the oldest anthology, *i.e.*, the first generation poems.

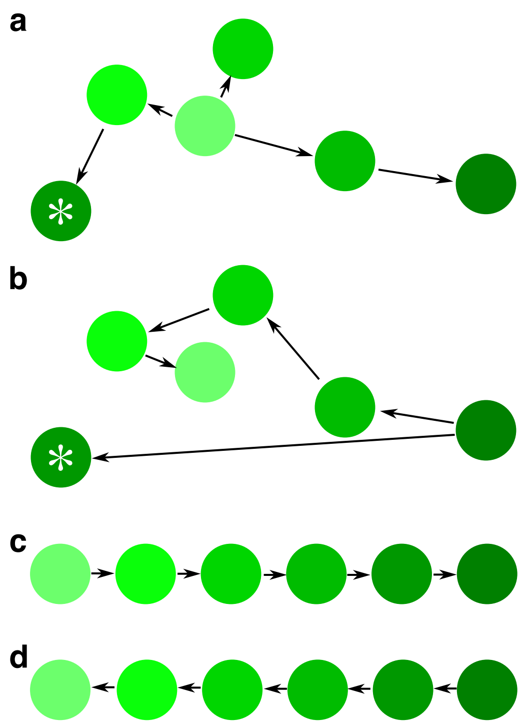

Supplementary Figure 2: Phylogenetic networks diverging from a poem (a), converging to a poem (b), and transitioning continuously (c, d). b and d are the time-reversed networks of a and c, respectively.

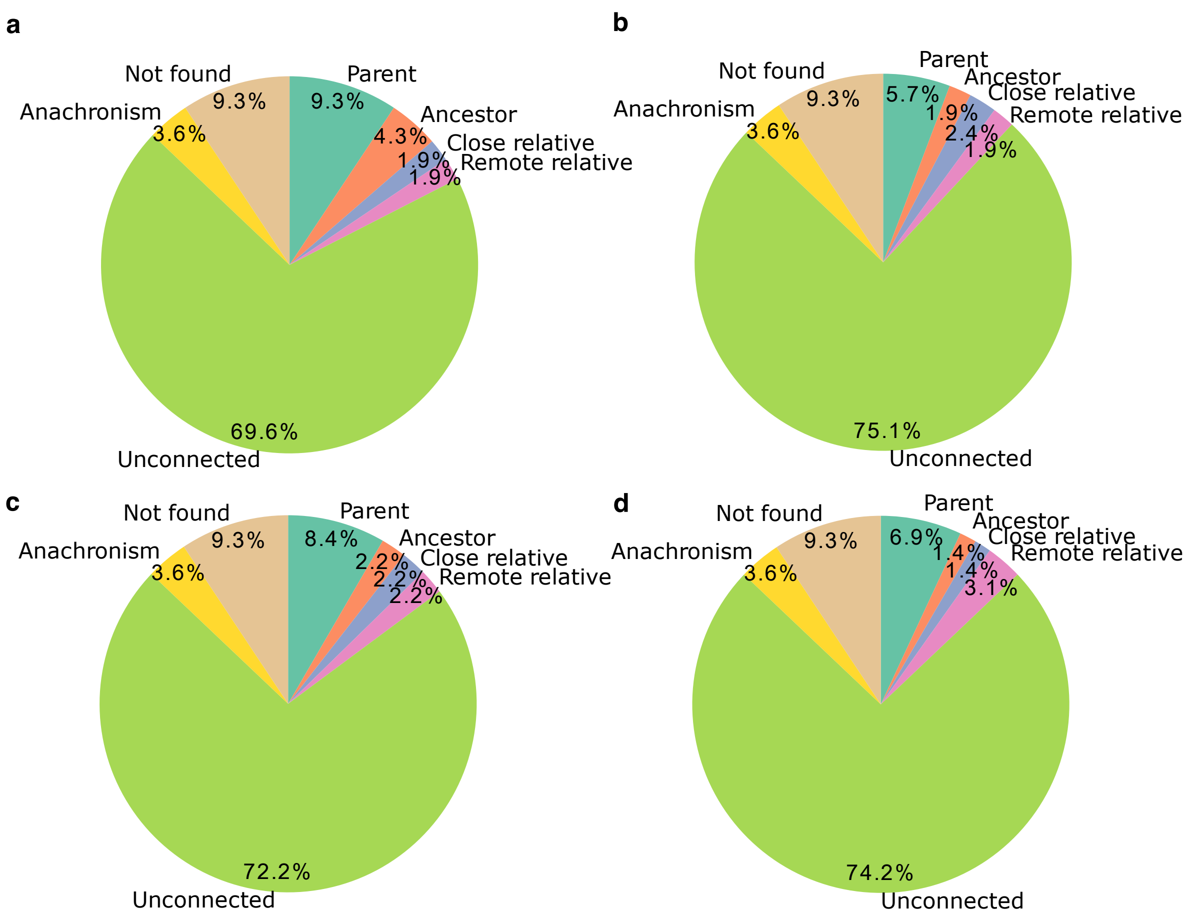

Supplementary Figure 3: *Honkadori* relationships in the phylogenetic network generated with (a) a token size of 3000, (b) a token size of 10000, (c) the dimensionality of the intermediate layers of 256, and (d) the dimensionality of the intermediate layers of 1024. The same model parameter values were used unless otherwise stated.

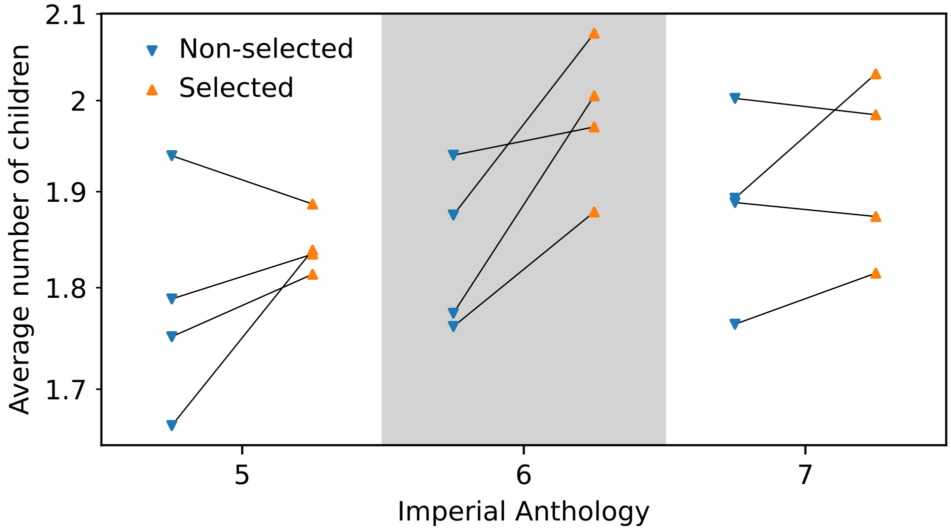

Supplementary Figure 4: Effects of being selected for *Kin'yōshū* 1, 2, and 3.

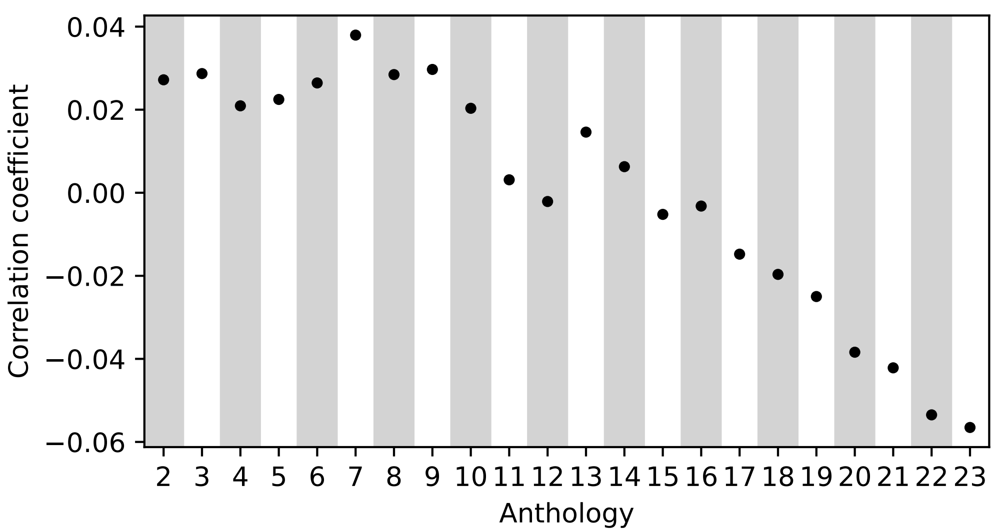

Supplementary Figure 5: Correlation coefficient between the standardized numbers of children before and after anthology i in the model with $\alpha = 1$.