

A quantitative philology of introspection

C.G. Diuk^{*,1}, D. Fernandez Slezak^{*,2}, I. Raskovsky², M. Sigman³, and G. A. Cecchi⁴

¹Princeton Neuroscience Institute, Princeton University, USA
²Department of Computer Science, University of Buenos Aires, Argentina
³Department of Physics, University of Buenos Aires, Argentina
⁴Computational Biology Center, T.J. Watson IBM Research Center, USA

June 20, 2012

Abstract

The cultural evolution of introspective thought has been recognized to undergo a drastic change during the middle of the first millennium BC. This period, known as the “Axial Age”, saw the birth of religions and philosophies still alive in modern culture, as well as the transition from orality to literacy - which led to the hypothesis of a link between introspection and literacy. Here we set out to examine the evolution of introspection in the Axial Age, studying the cultural record of the Greco-Roman and Judeo-Christian literary traditions. Using a statistical measure of semantic similarity, we identify a single “arrow of time” in the Old and New Testaments of the Bible, and a more complex non-monotonic dynamics in the Greco-Roman tradition reflecting the rise and fall of the respective societies. A comparable analysis of the 20th century cultural record shows a steady increase in the incidence of introspective topics, punctuated by abrupt declines during and

^{*}Both authors (CGD and DFS) contributed equally to this work

preceding the First and Second World Wars. Our results show that (a) it is possible to devise a consistent metric to quantify the history of a high-level concept such as introspection, cementing the path for a new quantitative philology and (b) to the extent that it is captured in the cultural record, the increased ability of human thought for self-reflection that the Axial Age brought about is still heavily determined by societal contingencies beyond the orality-literacy nexus.

1 Introduction

The period in human history ranging from 800 BCE to 200 BCE marked a radical transformation in world civilizations, in particular those that constitute the Western tradition. Famously termed the Axial Age by K. Jaspers [1], this period produced world religions and philosophies that are still pillars of modern culture. Many scholars have argued that one of the most prominent consequences of this transition was a change in consciousness, from a “homeostatic”, living-in-the-present form, to a self-reflective and time-expanding inner life [2–4]. Studying two foundational texts of Western civilization, the Bible and the Homeric saga, Julian Jaynes [5] argued that the change in conscious life that took place during the Axial Age is part of a larger introspection-increasing arc explicitly expressed in the textual narrative. The most cursory analysis of the Bible shows how the wrathful, interventionist god of the Old Testament, who expels Adam and Eve from the Garden of Eden (Genesis 6:22), is followed by the rich inner life of the New Testament, where Jesus asks the accusers of an adulterous woman to examine their own guilt (John 8:1). The striking differences between the Iliad and the Odyssey in the way the characters’ behaviors are attributed, respectively, to divine intervention or to the individual’s volition, have been pointed out in numerous studies [6–9]. The impulsive, irreflexive heroes of the Iliad, driven by passions insufflated into them by the gods, give way to the wily, cunning Odysseus who outsmarts Polifemo and leads his men to Scylla with a bad conscience.

It has been argued that these changes may reflect not just artistic or even cultural tendencies, but profound alterations in the mental structure of those who wrote, collected and assimilated the stories. Marshall McLuhan, in his seminal work on the relationship between semiotics, media and thought structures, argued for a materialistic effect of the type of medium (the linearity of written language, the holistic nature of the moving image) on the

organization of thoughts (linear or integrative, respectively) [10]. Given that one of the defining features of the Axial Age is the widespread adoption of the written word, it is reasonable to consider that the transitions from orality to literacy and from lower to higher introspection are fundamentally linked [34]. Jaynes, moreover, proposed that the transition was also accompanied by a change from a mentality dominated by inner voices, issuing god-like commands necessary to maintain social cohesion, to one where the voices were replaced by a self-aware inner dialogue. Equally interesting and controversial from a neuroscientific perspective [11, 12], these ideas cannot be validated or refuted solely on the basis of philological and cultural studies. However, the ‘soft’ hypothesis, as Daniel Dennett termed it, that within the Judeo-Greco-Christian tradition at least, there exists an “arrow of time” in the cultural narrative of the Axial Age representing increased concern with *introspection* can be put to the test in a quantifiable framework [5, 13].

The increasing availability of large text databases has facilitated the study of the statistical regularities embedded in language, including literary subtleties such as the analysis of a novel’s plot, traditionally restricted to qualitative assessments [14]. Here we set out to develop a consistent analytic framework to measure the degree of introspection in extant texts spanning the Axial Age. With this tool, we tested the hypothesis of a “quantum leap” in the narrative during this transitional period, more generally, the possibility of measuring a *cultural history of introspection*.

We understand meaning in language as arising from the mutual dependencies of concepts, in a holistic fashion [15–17]. Hence, our challenge is not merely to count the occurrence of a given word (for example, introspection) in a historical corpus but, instead, to measure the extent to which a concept, in its distributed semantic sense (as captured, for instance, in dictionaries and thesauri), is represented in the text. Several methods have been introduced to identify regularities and obtain a notion of semantic proximity [18–21]. One of the more widely used resources is Latent Semantic Analysis (LSA) [22]. LSA assumes that semantically related words will necessarily co-occur in texts with coherent topics. For a concept such as introspection to convey meaning, it must evoke a large array of related concepts or “images” in the mind. Conversely, semantically related words can convey the meaning of introspection without the word itself being present in a text. This is particularly relevant for our study of classical literature, as we do not expect words representing higher-order concepts such as introspection to be present in the Bible or Homer at all. That is, the use of semantic similarity analysis allows

us to measure the relatedness of a classical text to a concept with no explicit word at the time.

Our goal of the present study is to quantify the incidence of introspective topics in Axial Age literature, and uncover its cultural history. For this purpose, we decided to analyze two relatively self-consistent cultural traditions: the Judeo-Christian, anchored by the Biblical texts, and the Greco-Roman, whose foundational text is the Homeric saga, but include also the classical Greek and Latin writings up to the second century CE, encompassing among others Plato and Aristotle’s extant oeuvre. Further, we extended this analysis to modern times, by measuring the evolution of introspection in the Google n-grams database [23] during the 20th Century. While this database consists of an aggregate of word frequencies in the entire literary production per year, and thus qualitatively different from the book-based classical corpus, it allows us to explore to what extent the dynamics of introspection can be coupled to factors other than literacy.

2 Materials and methods

2.1 Texts

For the analysis of the Judeo-Christian tradition, we downloaded the Old and New Testaments, as well as St. Augustine’s *Confessions* from the Project Gutenberg database [24]. The Greco-Roman texts were downloaded from the MIT Classic archive [25], all in plain-text format (see Suppl. Mat. for details). Most documents had headers and footers with commentaries and information about the source database which were cropped by hand; only the title and full text of each book was considered for this study.

We preprocessed each text in this corpora using the NLTK (Natural Language Toolkit) library [26]. As a first step, we word-tokenized the text, discarding punctuation marks and symbols. The result is a list of words for each text, with repetitions. We then lemmatized the text. First, we tokenized each text into sentences, and POS-tagged (part of speech) them using the Treebank tagger supplied by NLTK. Once each word was tagged, we lemmatized each sentence using NLTK’s WordNet lemmatizer. The result of preprocessing is a list of lemmatized words, each one in a new line maintaining original order, in lowercase and without any punctuation marks or symbols.

For the analysis of the modern era, we used the Google n-grams database

[23]. This database consists of all the words found each year on the books digitized by Google, along with their frequencies for that year. In this case no lemmatization was done, as words are isolated from the sentence where they were written. Words without the sentence are impossible to tag into their Part of Speech, impeding correct lemmatization.

2.2 Semantic analysis

To establish a measure of semantic similarity between each pair of words, we used Latent Semantic Analysis (LSA) [22], a high-dimensional associative model that captures similarity between words. LSA generates a linear representation of the semantic content of words based on their co-occurrence with other words in a text corpus. If the corpus is sufficiently large and diverse, the frequency of co-occurrence of words across different documents represents the extent to which the words are semantically related [27].

The input to LSA is a word-by-document occurrence matrix X , with each row corresponding to a unique word in the corpus (N total words) and each column corresponding to a document (M total documents). An entry x_{ij} in this matrix indicates the number of occurrences of word i in document j . Using Singular Value Decomposition (SVD), the dimensionality of this matrix is reduced to a smaller number of columns, preserving as much as possible the similarity structure between rows. Formally, using SVD, we obtain a decomposition $X_{N \times M} = U_{N \times R} S_{R \times R} V_{R \times M}$, where R is the rank of matrix X . If we reduce dimensionality to k dimensions, $X_{N \times M} \approx \tilde{X}_{N \times M} = \tilde{U}_{N \times k} \tilde{S}_{k \times k} \tilde{V}_{k \times M}$, where $k < R$ and \tilde{U} , \tilde{S} , \tilde{V} are the same as the original decomposition (U, S, V) cropped to k dimensions. By reducing dimensionality, each word is projected into a space where semantic “meaning” is just its corresponding vector. The similarity in meaning between two words can be measured by calculating the cosine between the corresponding vectors. That is, similarity between two words is computed as $W(a, b) = \vec{a} \cdot \vec{b}$, where the vectors are the SVD representation of the words a and b . Since the vectors are normalized, the range of possible values for the similarity measure is $[-1, 1]$.

An important step in LSA is the generation of the word-by-document matrix from which semantic directions are extracted. We decided to use the TASA corpus, a collection of educational materials, compiled by Touchstone Applied Science Associates to develop The Educator’s Word Frequency Guide. The TASA corpus consists of 37,651 documents and 12,190,931 words, from a vocabulary of 94,796 distinct words. The materials in TASA consist of

general reading texts assumed to be common in the US educational system up to 1st year of college, including a wide variety of short documents taken from novels, newspaper articles, and other sources. This corpus is frequently used in LSA applications.

After generating the occurrence matrix for TASA, SVD was executed obtaining the decomposition. As the LSA method proposes, the SVD matrix may be cropped – reducing dimensionality – conserving the range of the original matrix. The choice of dimensionality is an important factor for success in measuring distance of concepts. Landauer and Dumais studied the effect of the number of dimensions retained in LSA and word-meaning similarities [27]. The maximum performance was obtained retaining around 300 dimensions, the number we used in this study.

For the analysis of the Judeo-Christian and Greco-Roman corpora, texts were lemmatized. Thus, the TASA corpus was also preprocessed in the exact same way, including lemmatization. This processing resulted in a vocabulary of 78,633 distinct words, 17% smaller than the un-lemmatized TASA corpus. For the analysis of the Google n-grams, the complete un-lemmatized TASA corpus was used.

3 Results

Judeo-Christian texts:

Due to the relative uncertainty of dating its different books, we collapsed the Old Testament into a single document, to be compared against the New Testament. Merely for the purpose of graphical depiction, we approximated the average dating of the Old Testament to the fourth century BCE given the wide spread of possible dating from the oldest books (second millennium to fifth century BCE, although in their written form they are more recent [28]) to the newest ones (Maccabees, second-first century BCE) [29]. The dating of the New Testament is comparably much more confined being certainly constrained to the second century CE.

In order to compare these results with a more modern text of the same tradition, we performed the same analysis on the *Confessions*, written by St. Augustine (or Augustine of Hippo in the Protestant nomenclature) around 398 CE. This book, an auto-biography, is considered a fulcrum between the classic and modern conceptions of the self, infused with Augustine’s idea of free will and self-determination [30,31].

For each book (Old Testament, New Testament and Confessions), we calculated LSA semantic similarity W_i between each single word and the word ‘introspection’, resulting in a series with values between -1 (semantically dissimilar to introspection) and 1 (semantically similar to introspection). We binarized these values to boost the weight of highly introspective words, using an arbitrary threshold of 0.15 , so that for each word i in the text we assign $S_i = 1$ if $W_i = \vec{w}_i \cdot \vec{I} > 0.15$, $S_i = 0$ otherwise, where \vec{w}_i, \vec{I} are, respectively, the LSA representation of word i and of *introspection*. We finally computed a moving average of this series with a window of $5,000$ words, representing an estimate of the local similarity to introspection of each fragment of the text. This series is shown explicitly for the Old and New Testaments in the inset of Fig.1, and represented by its mean and standard deviation for the three books in the main plot.

The Old and New Testament are significantly different in terms of their semantic similarity to introspection (t -test, $p < 10^{-12}$), in agreement with the transition hypothesis. The *Confessions* further shows a significant increase in introspection from the New Testament (t -test, $p < 10^{-12}$), consistent with its placement in the post-transition era, and with its qualitative scholarly assessment. As we mentioned above, we can only rely on broad, approximate dates for the compilation and writing of the books that conform the Old Testament, in particular for the oldest ones. The final *order* of the books in the canonical format we know today is as much a reflection of their temporal precedence as of multiple editorial interventions [32]. We therefore found highly suggestive that our semantic analysis term-by-term along the textual linearity of both Old and New Testaments also shows a statistically significant trend towards increasing introspection. This is represented in the inset of Fig. 1, where the moving average follows the textual precedence order of the Old (blue) and New (red) Testaments. Linear regression shows a significant positive slope (for Old Testament, $m = 4 \times 10^{-9}$, $p < 10^{-12}$; for New Testament, $m = 8 \times 10^{-8}$, $p < 10^{-12}$). While this trend is markedly positive on the latter, it is still significant in the former. A detailed analysis of the structure of the New Testament is beyond the scope of this paper; let us note, however, that the marked increase in the similarity to introspection corresponds approximately to the Pauline epistles.

Greco-Roman texts:

As it is the case with the oldest Judeo-Christian texts, it is difficult to precisely date the Homeric saga. Recent studies have dated with certain confidence the Trojan War and events related to the *Odyssey* around 1200

BCE [33], while the compilation into a coherent corpus probably dates between 900-800 BCE, and the final composition in the newly introduced Greek alphabet between 700-650 BCE [34]. While Jaynes dates the two epics roughly a century apart [5], we only assume some form of temporal precedence affected by the introspective transition hypothesis.

Similarity to introspection for all Greco-Roman books was performed with the same approach used for the Judeo-Christian texts. The full figure shows the same analysis extended to the database of extant Greco-Roman texts (figure 2), for which dating is significantly firmer than for the Homeric saga. As suggested by Jaynes, and in line with the broader Axial Age idea, we observe that in the period up to 400-300 BCE, the literary record displays a dramatic divergence towards higher introspective value (fit to an exponential function: $R^2 = 0.92$, $p < 10^{-12}$, see Suppl. Mat. for details).

Following approximately the end of Greece’s Golden Era, however, introspective content in the literary record displays an abrupt fall, followed in its turn by a second rise, mimicking the previous transition, essentially with Roman writings. The Roman rise is less abrupt than the Greek one, but nevertheless clearly significant ($R^2 = 0.85$, $p < 10^{-12}$).

Modern Cultural Record: The finding of two “transitions” in the Greco-Roman tradition implies a possible link between introspection and societal upheaval. In order to test this hypothesis, we decided to investigate the semantic relatedness of introspection in the modern cultural record compiled in the *Google ngrams* database [23]. From all words included in this database, we filtered those included in the LSA dictionary – that is, all different words included in TASA. As mentioned in the Methods section, we used the unlemmatized TASA-trained LSA, as words in the ngrams database lack their sentence context.

For each year, we calculated the semantic relatedness of introspection as $f(y)$,

$$f(y) = \sum_{i=1}^{N(y)} \frac{W_{w_i(y)} \times A_{w_i(y)}}{T(y)} \quad (1)$$

where $N(y)$ is the number of different words during year y (intersected with those available in TASA), $W_{w_i(y)}$ is the introspective weight of word $w_i(y)$ as measured by its LSA distance to ‘introspection’, $A_{w_i(y)}$ is the number of occurrences of word $w_i(y)$ during year y and $T(y) = \sum_{j=1}^{N(y)} A_{w_j(y)}$ is the total words during year y . In order to assess the statistical significance of yearly changes, we computed a local measure of noise as the standard deviation of

the mean of $f(y)$ in a 6-year window centered in y . The results are shown in Fig. 3-top, where the 20th century CE was analyzed.

To verify the significance of the measurement for each year, we computed the expected similarity value for the null hypothesis that the results obtained can be simply explained by the distribution of similarity to introspection of the word ensemble for the year, regardless of the frequency of occurrence. For that, we generated 100 surrogate trials with randomized (uniform) frequencies for each word every year y , and calculated $f_{null}(y)$. In Fig3-bottom, we show the average and standard deviation of these surrogates. We confirmed the significance of the introspective value as each year surrogates are very distant to introspection compared to the n-gram values (all t-test's with $p < 10^{-8}$).

We observe a marked effect of WWI and WWII, which coincide with significant troughs, or at least punctuate the inter-war peak and a sustained increasing arc following WWII. While we can only speculate about what the significance of this arc means, the correlation between war and drops in introspection has obvious implications and is also consistent with our observation of the Greco-Roman cultural record.

It is possible that the growth of introspection observed in different cultural records merely reflects a cumulative effect of culture and knowledge. Similarly, it can be argued that the between-wars drop of introspective content in the cultural record may simply reflect a global trend that can be manifested by a broad number of concepts, not only introspection.

To investigate the specificity of this effect to introspective content, we repeated the same analysis on a list of 200 concepts deemed “fundamental” across 87 Indo-European languages [35]. A visual inspection clearly reveals that all the main effects reported here are specific to introspection, and do not replicate across these 200 other concepts (see Supplemental Materials for details and figures). By specificity we do not mean that this is the sole concept to show this dynamics, we simply argue that it is not a general trend observed in concepts which show less cultural dependence. First, these words did not show a systematic growth as compared, for example, to the growth of introspection in the New Testament. In the modern cultural records, control words did not show any systematic trend revealing, as we observe for introspection, a rise and fall between wars and a consistent progression after the second world war. To quantify this observation we performed two independent statistical tests. We measured the quadratic dependence of each word during the period 1914-1939 (from the beginning of the first to the beginning of the second world war) and the linear dependence in the

1945-2000 period (from the end of the second war to the end of the century). The quadratic term for introspection (1914-1939) was negative, reflecting a convex dependence (rise and then fall) with a significance of $p < 0.01$, considering as the null-hypothesis the distribution of the quadratic terms obtained from all control words. Similarly, the post-war growth was testified by a positive linear correlation which was also within the one percentile of all the control words. The growth of introspection in the New Testament was also significant within the distribution of control words ($p < 0.01$).

4 Discussion

In light of the results obtained analyzing the historical literary record of the two root traditions of the Western civilization, a number of questions are left open. From a neuroscientific perspective the most prominent is, whence comes the transition? Numerous scholars have pointed to the consolidation of alphabetic writing during the Axial Age as one of the main causes of a change in the cultural dynamics of introspection [5, 34]. The Greek alphabet, in particular, developed around 700 BCE, had finally been internalized in the culture by Plato's time. Notably, however, Socrates, in Plato's *Phaedrus*, excoriates it on the basis of its corruption of the faculty of memory [25]. It is precisely this effect on memorization that has been regarded as introducing a dramatic change in mental life: by easing the onerous task of memorizing the narrative that defines a culture, it allows the mind to wander with new freedom, and releases it from the constant anchor of 'internalized external voices' [34]. As mentioned in the Introduction, McLuhan similarly relates the linearity and constrained structure of the written text to the development of rational thinking and the abandonment of "magical" or unconstrained form of thought, typical of orality [10].

There is evidence that literacy is associated with functional as well as anatomical changes in the brain of literate as opposed to illiterate adults, even when literacy is acquired later in life [36–38]. While these observations do not have a direct bearing on the possible link between literacy and introspection, recent findings indicate an association between brain symmetry and commanding hallucinations and psychosis [39–41]. It is therefore reasonable to hypothesize that the strong lateralization that accompanies literacy can influence the mental processes required for increased introspection. Moreover, introspection, at least in the restricted sense of the meta-cognitive ability

to judge your own performance in a task, also has measurable physiological correlates [42,43], and therefore cannot be dismissed as an ephemeral cultural by-product.

A typical argument leveled against the hypothesis of a historical evolution of introspection is that even in the relatively backward civilization of the Homeric epics and Ancient Israel, people were organized in complex social structures that required equally complex behaviors including planning, assessment and a form of “theory of the mind”, among other cognitive features inconsistent with a non-introspective mind. While it would indeed be nonsensical to assume that pre-Axial Age humans were mere automata, incapable of any form of reflexion, a growing body of research suggests that goal-oriented behavior in individuals is highly determined by social imposition, and can occur largely outside awareness [44].

Our analysis of the Greco-Roman and modern cultural records suggests that the dynamics of introspection are heavily determined by, or at least correlated with, the overall state of the society and liable to dramatic changes over short periods of time. We cannot consider this finding as a refutation of the hypothesis of a link between the orality-to-literacy transition and increased introspection, as a relatively thriving societal context may be a necessary pre-requisite but not a sufficient cause. Moreover, we need to consider that the relationship between outward cultural manifestations and the mental makeup of individuals in a culture is highly complex and devoid of a clear causal directionality. In fact, the rapid changes we observe around the World Wars may be interpreted precisely as indicating that cultural productions have a dynamics of their own, not reducible to a simple reaction to dramatic “external” factors.

To the extent that the texts we studied represent Axial Age civilizations, our results provide a framework on which hypothesized cultural changes affecting cognitive structures can be quantifiably measured. By the same token, our analytic approach can be extended to characterize other similarly high-level cognitive features in the cultural record as a whole, literary styles [45] and psychiatric dysfunction [46, 47].

Acknowledgements

The TASA corpus appears courtesy of Tom Landauer and Touchstone Applied Science Associates. The authors would like to thank critical readings by Pablo

Meyer, Stanislas Dehaene, Jean Pierre Changeux and Jean Remi King.

References

- [1] K. Jaspers. *The origin and goal of history*. Yale University Press, 1953.
- [2] S. N. Eisenstadt, editor. *The origins and diversity of axial age civilizations*. SUNY Press, 1986.
- [3] B.I. Schwartz. The age of transcendence. *Daedalus*, 104:1–7, 1975.
- [4] K. Armstrong. *The Great Transformation: The Beginning of Our Religious Traditions*. Alfred A. Knopf / Random House, 2006.
- [5] J. Jaynes. *The origin of consciousness in the breakdown of the bicameral mind*. Mariner Books, 2000.
- [6] R.B. Onians. *The origins of European thought about the body, the mind, the soul, the world, time, and fate: new interpretations of Greek, Roman and kindred evidence also of some basic Jewish and Christian beliefs*. Cambridge University Press, 1988.
- [7] E.R. Dodds. *The Greeks and the Irrational*. Berkeley: Univ. of California Press, 1951.
- [8] A.W.H. Adkins. *From the many to the one: a study of personality and views of human nature in the context of ancient Greek society, values and beliefs*. Constable & Company Limited, 1970.
- [9] I.J.F. De Jong and J.P. Sullivan. *Modern critical theory and classical literature*. Brill Academic Pub, 1994.
- [10] M. McLuhan. *The Gutenberg galaxy: The making of typographic man*. Routledge & Kegan Paul, 1962.
- [11] M. Kuijsten *Reflections on the Dawn of Consciousness*. Julian Jaynes Society, 2006.
- [12] A.E. Cavanna, M.R. Trimble, F. Cinti, and F. Monaco The “bicameral mind” 30 years on: a critical reappraisal of Julian Jaynes hypothesis. *Functional Neurology*, 22(1):11–15, 2007.

- [13] D. Dennett. Julian jaynes’s software archeology. *Canadian Psychology/Psychologie canadienne*, 27:149–154, 1986.
- [14] F. Moretti. *Graphs, Maps, Trees: Abstract models for a literary history*. Verso Books, 2005.
- [15] W.V. Quine. Two dogmas of empiricism. *The Philosophical Review*, 60:20–43, 1951.
- [16] R. Ferrer i Cancho and R.V. Sole. The Small World of Human Language. *Proc. Roy. Soc. Lond. B*, 268: 22612265, 2001.
- [17] M. Sigman and G.A. Cecchi. The global organization of the Wordnet lexicon. *Proc. Natl. Acad. Sci. USA*, 99(3): 1742–1747, 2002.
- [18] C. Fellbaum. *Theory and Applications of Ontology: Computer Applications*, chapter WordNet, pages 231–243. Springer, 2010.
- [19] T. Pedersen, S. Patwardhan, and J. Michelizzi. Wordnet: similarity-measuring the relatedness of concepts. In *Proceedings of the National Conference on Artificial Intelligence*, pages 1024–1025, 2004.
- [20] S. Patwardhan, S. Banerjee, and T. Pedersen. Using measures of semantic relatedness for word sense disambiguation. *Lecture notes in computer science*, pages 241–257, 2003.
- [21] K. Lund and C. Burgess. Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods*, 28(2):203–208, 1996.
- [22] S. Deerwester, S.T. Dumais, G.W. Furnas, T.K. Landauer, and R. Harshman. Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, 41:391407, 1990.
- [23] Jean-Baptiste B. Michel, Yuan Kui K. Shen, Aviva Presser P. Aiden, Adrian Veres, Matthew K. Gray, The Google Books Team, Joseph P. Pickett, Dale Hoiberg, Dan Clancy, Peter Norvig, Jon Orwant, Steven Pinker, Martin A. Nowak, and Erez Lieberman L. Aiden. Quantitative analysis of culture using millions of digitized books. *Science (New York, N.Y.)*, 331(6014):176–182, January 2011.

- [24] Michael S. Hart. Project gutenber. <http://www.gutenberg.org/>, last accessed: Oct. 4th, 2011.
- [25] Daniel C. Stevenson. Mit classics. <http://classics.mit.edu/>, last accessed: Feb. 27th, 2010.
- [26] S. Bird, E. Klein, and E. Loper. *Natural language processing with Python*. O'Reilly Media, 2009.
- [27] T.K. Landauer and S.T. Dumais. A solution to plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological review*, 104(2):211, 1997.
- [28] W. Schniedewind. *How the Bible Became a Book: the Textualization of Ancient Israel*. Cambridge University Press, 2004.
- [29] M.D. Coogan (Ed.). *The Oxford History of th Biblical World*. Oxford University Press, 1998.
- [30] P. Brown. *Augustine of Hippo*. University of California Press, 1967.
- [31] J.D. Caputo and M.J. Scanlon (Eds.). *Augustine and Postmodernism*. Indiana University Press, 2005.
- [32] R. Scholes and R. Kellogg. *The nature of narrative*. New York: Oxford UP, 1966.
- [33] C. Baikouzis and M.O. Magnasco. Is an eclipse described in the odyssey? *Proc. Nat. Acad. Sci. USA*, 105:8823–8828, 2008.
- [34] W.J. Ong. *Orality and Literacy: the Technologizing of the Word*. Methuen & Co., 1982.
- [35] Mark Pagel, Quentin D Atkinson, and Andrew Meade. Frequency of word-use predicts rates of lexical evolution throughout Indo-European history. *Nature*, 449(7163): 717–720, October 2007.
- [36] K.M. Petersson, C. Silva, A. Castro-Caldas, M. Ingvar, and A. Reis. Literacy: a cultural influence on functional leftright differences in the inferior parietal cortex. *European Journal of Neuroscience*, 26(3): 791799, 2007.

- [37] M. Carreiras, Mo.L. Seghier, S. Baquero, A. Estevez, A. Lozano, J.T. Devlin, and C.J. Price. An anatomical signature for literacy. *Nature*, 461: 983–986, 2009.
- [38] S. Dehaene et al. How learning to read changes the cortical networks of vision and language. *Science*, 330: 1359–1364, 2010.
- [39] S.H. Laval et al. Evidence for linkage to psychosis and cerebral asymmetry (relative hand skill) on the X chromosome. *American Journal of Medical Genetics*, 81: 420–427, 1998.
- [40] R. Olin. Auditory hallucinations and the bicameral mind. *The Lancet*, 354: 166, 1999.
- [41] M.C. Corballis. The evolution and genetics of cerebral asymmetry. *Phil. Trans. Royal Society B - Biol. Sci.*, 364: 867–879, 2008.
- [42] R. Kiani and M.N. Shadlen. Representation of confidence associated with a decision by neurons in the parietal cortex. *Science*, 324: 759, 2009.
- [43] S.M. Fleming et al. Relating introspective accuracy to individual differences in brain structure. *Science*, 329: 1541–1543, 2010.
- [44] R. Custers and H. Aarts. Social effect on unconscious goal setting, also references therein. *Science*, 329:47–50, 2010.
- [45] F. Moretti. Network theory, plot analysis. *New Left Review*, 68:80–102, 2011.
- [46] P.J. McKenna and T.M. Oh. *Schizophrenic Speech: Making Sense of Bathrooms and Ponds that Fall in Doorways*. Cambridge University Press, 2005.
- [47] N.B. Mota et al. *Speech graphs provide a quantitative measure of thought disorder in psychosis*. *PLoS One*, 7(4): e34928, 2012.




Figure 1: Introspection in the cultural record of the Judeo-Christian tradition. The New Testament as a single document shows a significant increase over the Old Testament, while the writings of St. Augustine of Hippo are even more introspective. Inset: regardless of the actual dating, both the Old and New Testaments show a marked structure along the canonical organization of the books, and a significant positive increase in introspection.




Figure 2: Introspection in the cultural record of the Greco-Roman tradition. Dotted lines represent exponential fits for each of the traditions.

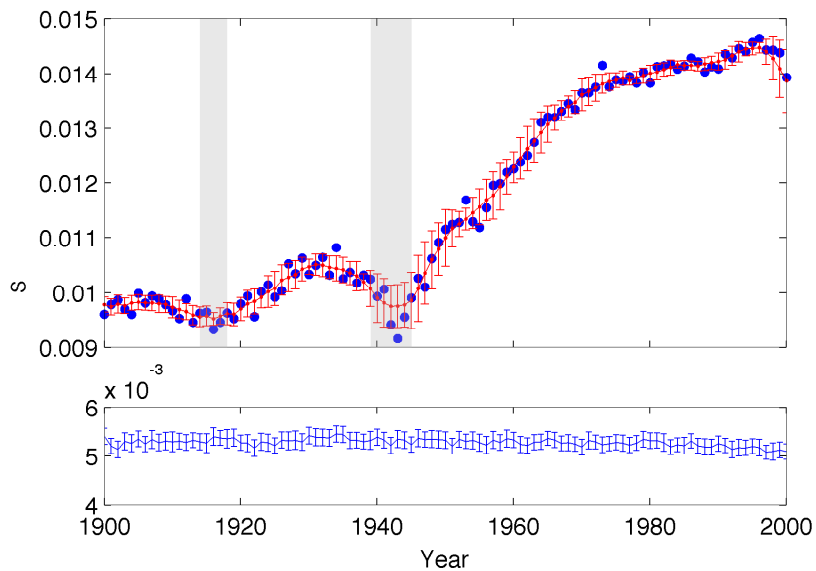


Figure 3: Evolution of introspection in the modern cultural record. Upper panel: the blue circles correspond to semantic relatedness to introspection, and the red trace to the mean and standard deviation of these values on a 6-yr. running window (i.e. running average and noise). The gray shades correspond to the two World Wars, and the red vertical dashes the standard deviation of the null hypothesis. Lower panel: mean and standard deviation of the null hypothesis corresponding to randomized word frequencies.

Figure 2.TIF

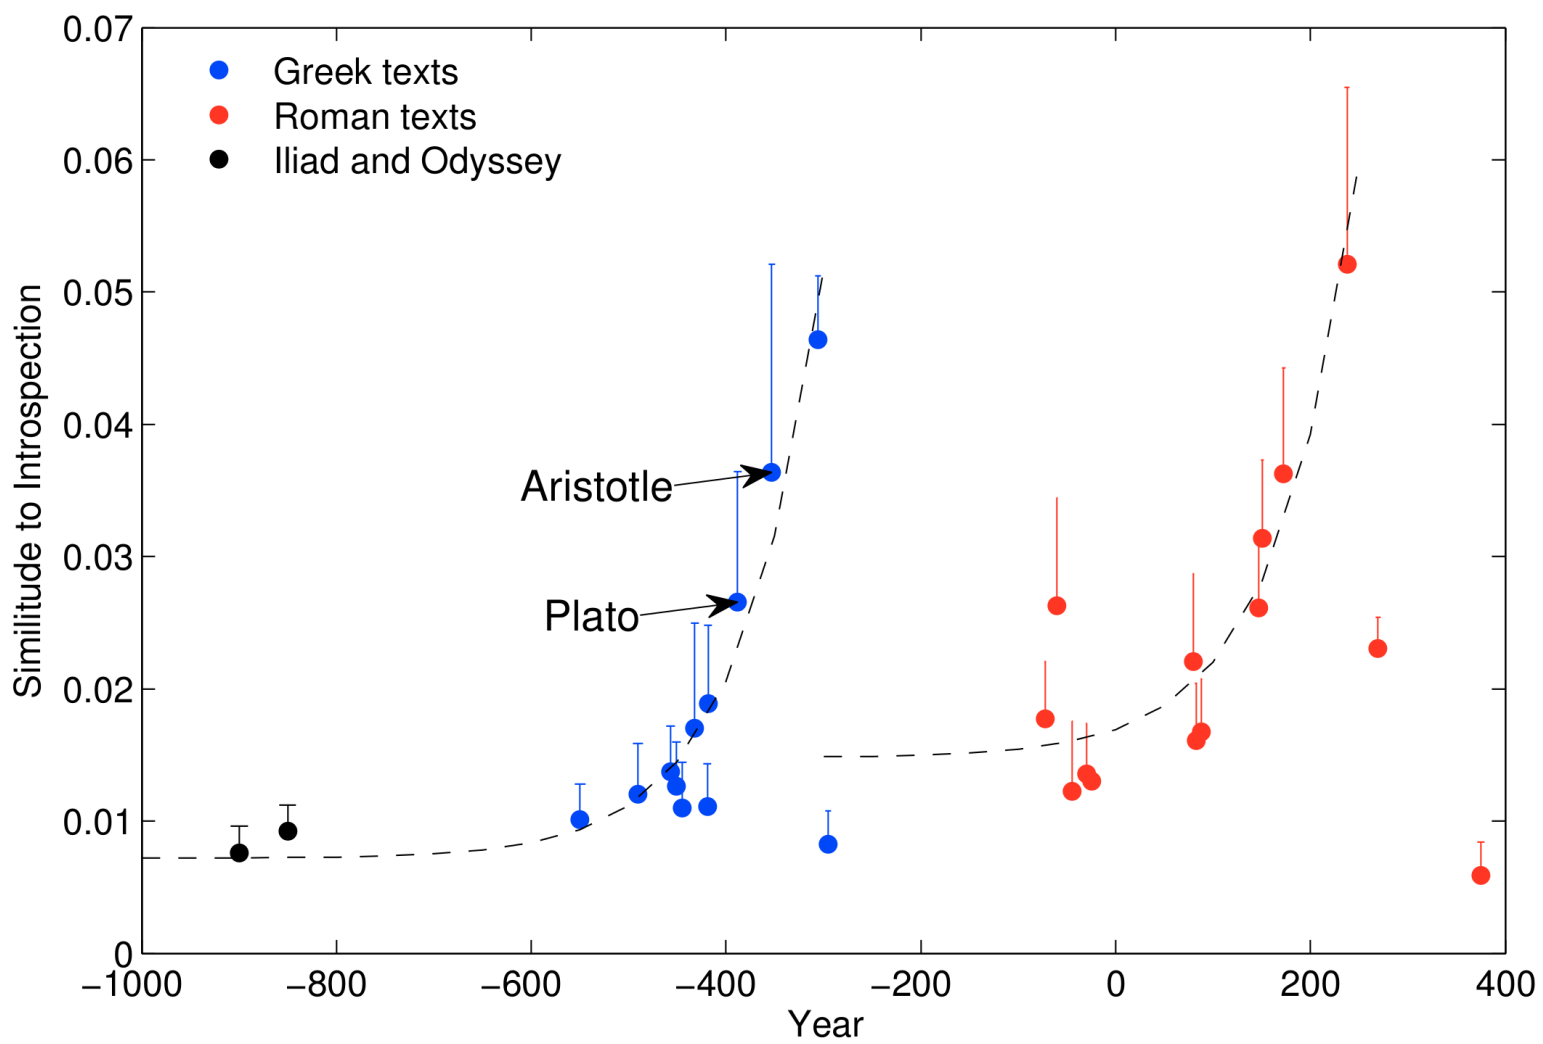


Figure 3.TIF

