

Clusters and Graphs

Revealing and Modeling Documentary Indefiniteness¹

Michael Fox

ABSTRACT

The documentary record of interrelated, multimedia literary works is riddled with indefiniteness. The order of their materials is often uncertain, as is the medium that should serve as their point of reference. This article argues that such common and related indefiniteness, like ambiguity and genetic lineage, whether revealed through traditional scholarly methods or computational ones, can be modeled using graph technology. Graph technology, moreover, can be used to model much more about these literary works, from their atomic documentary features to their higher-order features. The article uses as its example Jaime de Angulo's Old Time Stories, a Modernist American masterpiece consisting of voice, text, and image. Graph modeling, or graph editing, the work results in a fine-grained, computationally accessible representation of it as it really is in all its indefinite and networked nature. Such a representation lends itself to typical hermeneutic investigations enhanced by the power of inferential queries, and it can even serve as an actually authentic source for more quantitative investigations. Most important, like all non-digital natural-language based artifacts, it can also be endlessly modified by future editors without ever giving up from its structure the history of its own making.

SCHOLARLY EDITORS OFTEN FACE THESE COMMON CHALLENGES, OF MANY, when we edit interrelated, multimedia literary works: how to order the materials and how to decide what medium should serve as a point of reference. When trying to address these challenges using either paper-based technology or technology that replicates the surface structures of paper-based technology like TEI, an editor might run into limitations.² What if the order of the materials can't be resolved decisively, given the available evidence? The editor would be stuck choosing one order and pointing out

1. My research for this article has been supported by a Frederick and Eleanor Kilgour Research Grant from the School of Information and Library Science at UNC-Chapel Hill.
2. See SCHLOEN AND SCHLOEN 2014, paras. 25–26 for a succinct explanation of how TEI fails to efficiently model the multidimensionality of texts.

possible others, making them known but perforce leaving them largely inaccessible. And what if each medium were equally important? Neither a book nor a hierarchical markup language can present multiple media at the same time without making one ancillary to another.

Jaime de Angulo's *Old Time Stories (OTS)*, a Modernist American masterpiece of ethnopoetics, is a case in point. It is an auidial work based on prompt texts, themselves produced from illustrated typescripts — typescripts that were further developed after the auidial work was published. And no one manifestation is subordinate to another. All, moreover, have been edited and re-edited over time. As a result, the original order of the auidial version of the work, in particular, has been obscured. That order may be recovered only partially from external evidence, and here I show that even evidence internal to the audio itself, when ultimately subjected to a sophisticated cluster analysis, is insufficient. All this and related documentary indefiniteness revealed about the work, however, may be modeled and made accessible by graph technology.³

Clusters

“i repeated them ovr the radio 3 times and each time i made some variation, transpositions, etc.”, de Angulo wrote in a letter to the Pound family (DE ANGULO 1950). By ‘them’ he was referring to live broadcasts that added up to about 22 hours’ worth of audio material, comprising the oral performance version of his *OTS*. The recordings of these broadcasts thus present a study in revision.

Born in Paris to Spanish parents, de Angulo moved to California at the start of the twentieth century, lived, in part, with Native American tribes during his career, and became a scholar of Native American cultures and languages. His *OTS*, based on this scholarship, prompted Pound to dub him the “American Ovid”. The composer Peter Garland called the work a “story-epic . . . unique in American literature” (GARLAND 1988, 10). The one-time performance version of it aired in 1949–1950 on the counter-cultural radio station KPFA FM, the first public broadcasting station in the United States. In mostly spoken English, it tells the stories of such characters as Coyote Old Man, Loon Woman, Fox Boy, Diniki, and Baby

3. The general documentary features of the *OTS* delineated in this paragraph are summarized from MCGANN 2022a and some of McGann’s unpublished research in MCGANN 2023.

Quail, and it includes songs in native California languages. Thoroughly Western and Modern in style but drawing its inspiration from Native American mythopoeic tales, it exhibits de Angulo's broad argument across his non-fiction writings that any philosophy of language must arise out of phonology.

Which is precisely why it is essential to establish the documentary facts about this audial version of the *OTS* as completely as possible before taking up the problem of editing the entire multimedia work — even if those facts turn out to reveal indefiniteness. As de Angulo's words to the Pounds indicate, the oral performance of the *OTS* was given in three original installments. The first installment of broadcasts was a major achievement that we know extended over thirteen hours. The second and third installments of the broadcasts were almost as important as the first and were revised and added to it. Various interested parties, over the decades, edited and re-edited the recordings of these installments, unfortunately blurring their distinction. And that distinction is necessary to know in order to accurately mark all those "variations, transpositions, etc." that de Angulo speaks of. Variations, transpositions, etc., after all, imply orderliness.

But because only some of the recordings are dated, an editor must turn to internal evidence to order the others. Is there something within the recordings themselves that might help us determine which recordings were made in which installments? Something about the overall sound of the recordings? Something like the sound quality, perhaps, or the signal-to-noise ratio? Or maybe we would need to expand the horizon by moving into the broadcast studio itself — into the studio through the recording — to consider background noise or ambient noise, which, to be precise, are not the same as each other, nor do they have much to do with the 'noise' of the signal-to-noise ratio since that 'noise' refers to the noise level of the recording equipment itself. What if a combination of all or some of these qualities distinguishes recording sessions and therefore recording installments? Or what if some combination of esoteric acoustic properties not perceptible to the average listener would accomplish the goal — qualities like the mean zero crossing rate over each recording, the mean spectral centroid, or the mean spectral rolloff?

Such questions concerning possible internal evidence of distinguishing qualities of the recording sessions are primed to be answered through the indefinite process of exploratory data analysis (EDA), the definition of which by John Tukey is worth repeating here: "Procedures for analyzing data, techniques for interpreting the results of such procedures, ways of

planning the gathering of data to make its analysis easier, more precise or more accurate, and all the machinery and results of (mathematical) statistics which apply to analyzing data” (TUKEY 1962, 2). In what follows, I describe the EDA undertaken to distinguish the three installments of the *OTS*, explain why it likely failed, and argue that the failure still bears fruit and, in the end, benefits an innovative graph model of the data.

The data, in this case, consist of tapes made in 1991 under the editorial supervision of de Angulo’s daughter. Though they have been taken to constitute the standard performance version of the *OTS*, they likely comprise a fractured, incomplete copy of the original one-hundred reel-to-reel tapes produced from the first broadcasts done in 1949–1950 — ‘likely’ because we can’t be certain that some of the 1991 tapes did not stem from any of the multiple re-broadcasts done between 1950 and 1991.⁴ The tapes come in twenty-two approximately hour-long parts, each divided into six to eight tracks. De Angulo made the original recordings of each part in about four fifteen-minute segments, but Pacifica Radio Foundation has provided the parts in this particular format, which totals one-hundred and sixty-five tracks. The greater number of divisions are fine, though, even desirable, we will see, when trying to separate them into the groups in which they were supposedly produced.

For the first foray into the EDA, the data were left as they came. They would only have to be cleaned and volume-normalized for an edition, but in their raw form, they could be analyzed unbiasedly. Repeated episodes, or repeated performances, were then culled from the data and separated into two groups Known and Unknown, Known containing those repeated episodes whose orderings are known and Unknown containing those whose orderings are unknown. Group Known would serve as the baseline against which the results of some procedure p from the EDA would compare. If p resulted in identifying the correct orderings of the pairings from Known, then we could be relatively confident in p ’s orderings of the pairings from group Unknown.

First, single variables were tested as determining factors for orderings, the variables I mentioned previously: sound quality, signal-to-noise ratio, background noise, and ambient noise. Since sound quality is a highly abstract notion, measuring it is a highly subjective procedure. But independent of each other, McGann and I were able to consistently group the episodes into three categories of sound quality: Not-so-clear,

4. The original reel-to-reel tapes were rediscovered in 2021 but they have not yet been made accessible.

Clear, and Very Clear. However, the results indicated that these categories did not correspond in any consistent way to stages of the orderings in group Known, so we concluded that sound quality did not determine any cleanly bounded clusters of the recording installments. The signal-to-noise ratio of each episode was then measured and tested. This ratio is related to the sound quality, but it's a more objective measure, comparing the strength of the audio signal to the power of the noise level. But as with the sound quality, the signal-to-noise ratio proved not to be helpful in the goal of clustering the purported recording installments of the *OTS*. Background and ambient noise were then measured and tested. Background noise, in this case, consisted of any sound not de Angulo's voice. It included airborne sound, or sound from the environment, as well as structure-borne sound, like static or sound arising from the audio system, the mic, the recording equipment, or otherwise. Ambient noise here consisted of only airborne sound. Samples of these types of noise were taken from every track, and we tried to group the samples into clusters according to common traits that we could discern, however subjectively. We failed to do so in any meaningful way that would bring us closer to our aim of establishing the orderings of the repeated episodes in group Unknown.

Finally, a *combination* of possible determining variables was tested. A set of eight features was extracted from these noise samples using a Python library called pyAudioAnalysis: zero crossing rate, energy, entropy of energy, spectral centroid, spectral spread, spectral entropy, spectral flux, and spectral rolloff.⁵ Their technical definitions are of no significance here. What is significant is that these features are more esoteric than sound quality and signal-to-noise ratio but more objective than the qualia of our experiences listening to the noise. And they constitute multiple, distinct, objective measures of different aspects of sound that we couldn't apprehend just by listening. We presumed that if measured in noise, together they might reflect imperceptible conditions within the studio and the recording equipment which might help distinguish one recording session from another and therefore one recording installment of the *OTS* from another.

But before these features were extracted, the decibel levels of the noise samples had to be normalized to control for volume. We thought some of these features were dependent on volume, and we were interested only in their relative differences across the entire set of *OTS*

5. See GIANNAKOPOULOS 2020 for precise definitions of these features.

tracks. Furthermore, pyAudioAnalysis first splits the input signal into frames, then computes the values of each of these features for each frame, so after analyzing the samples with it, averages of the values across the frames had to be calculated. This procedure resulted in a new set of data, a matrix consisting of a row for each track and a column for each audio feature, or variable. But such a high-variable, or high-dimensionality, data set is difficult to interpret. Its dimensionality had to be reduced using Principal Component Analysis (PCA), and the resulting PCA plot was then clustered using the *k*-means clustering algorithm.⁶ Below is the 3-means clustered PCA plot of the 165 tracks comprising the *OTS* — each track is labeled X*Y, where X is a part number and Y is a track number for that part (see Fig. 1). Note the overlapping, which immediately suggests this clustering won't do. Indeed, despite all this effort, we found that when compared against our baseline group *Known*, these clusters did not determine separate and distinct recording installments.

Overall, the EDA failed to help distinguish the three installments de Angulo mentions to the Pounds because, I surmise, the conditions of the recording sessions for each installment were basically uniform. The recording equipment doubtfully changed or degraded, so sound quality and signal-to-noise ratios across sessions would have stayed relatively consistent. And unless the recordings took place in different locations — which they didn't — background or ambient noise

6. PCA is an unsupervised machine-learning procedure that reduces a large set of variables, in this case eight, which explain a particular phenomenon, to so-called principal components. The first two of these principal components, PCA 1 and PCA 2, serve as variables orthogonal to each other, like *x* and *y* in a two-dimensional plot, and explain most of the variance in the data. But unlike *x* and *y* in a two-dimensional plot, PCA 1 and PCA 2 are unitless and in need of being interpreted, or defined. In this case, however, we were not concerned with what they meant. We only cared whether or not plotting them for each track would result in clusters of tracks corresponding to the stages of orderings in our baseline group *Known*. But we couldn't just eyeball such clusters on the PCA plot, so we applied *k*-means clustering to its underlying data to assist us. Like PCA, *k*-means clustering is an unsupervised machine-learning algorithm. It takes as input *k*, or number of clusters, and a set of plot points. Then, through an iterative approach, it statistically optimizes positions of centroids, or centers of the *k* clusters it is asked to find, yoking each plot point under one of the *k* clusters.

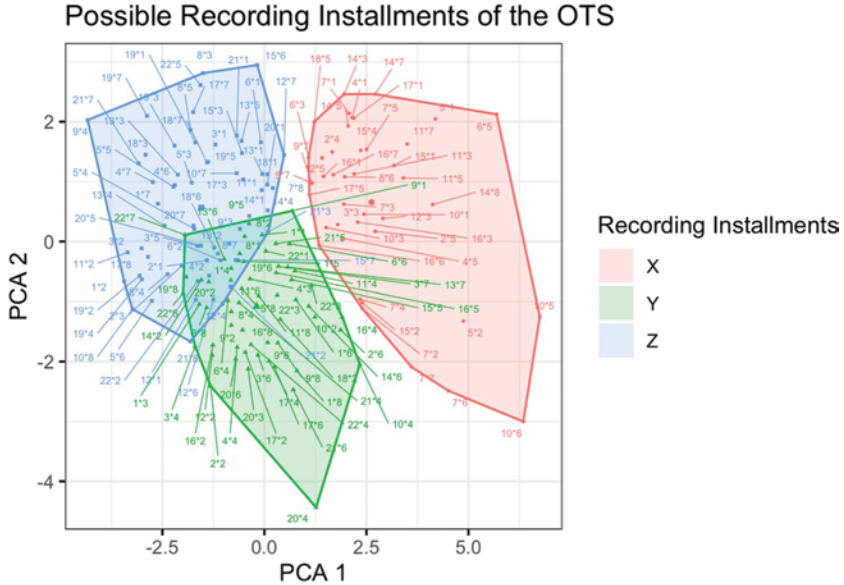


Figure 1. 3-means clustered PCA plot produced with R from measures of eight acoustic properties of noise samples taken from the 165 tracks of the *OTS*. The eight properties are: zero crossing rate, energy, entropy of energy, spectral centroid, spectral spread, spectral entropy, spectral flux, and spectral rolloff. Each track is labeled $X*Y$, where X is a part number and Y is a track number for that part.

would've been similar from session to session. Why the more esoteric properties analyzed in the noise samples using machine learning didn't indicate distinct installments is difficult to say. Being sensitive, objective measures of sound quality, they seemed promising. Perhaps by nature they vary too much within their relatively limited ranges to be determinants in separating sessions.

Regardless, the failure was still productive for two reasons: first, it ruled out all kinds of properties and combinations of properties intrinsic to the audio as determining factors; and second, it generated a great deal of data that may be rolled into a full-fledged graph representation of the entire multimedia *OTS* and its related materials. Those data include not only measures of acoustic properties but also possible orderings of episodes that haven't been dismissed. In a graph representation, they may be interpreted, despite not being definitive, in relation to each other and to the rest of the graph-modeled work, or they may await further analyses in light of newly discovered evidence.

Graphs

Indefiniteness is the very attribute that lends the graph structure to literary works because they themselves are indefinite, especially those, like the *OTS*, that come in multiple media, each as important as the other.

A graph of a literary work consists of nodes and edges, the nodes, at their basest level, referring to atomic, documentary phenomena and the edges representing relationships between those phenomena. These nodes and edges, as phenomena and relationships, can be combined into more nodes and edges referring to and representing more abstract phenomena and relationships, respectively. Theoretically, they can be combined into any documentary abstraction and its relationship to others. Of course, these very abstractions would be interpretive, but so is any edition, no matter how diplomatic it is, and an editor must proceed. To be sure, though, an editor must do so as responsibly as possible, trying to be as objective as possible. Responsibly graphing a literary work recreates, in all its dimensions, the work itself in a computationally accessible universe like our own. That is the goal: a computationally accessible universe like our own. ‘Computationally accessible’ implies a model. (Think of the difference between a video of a 3-D object and a topological model of it — the former does not mathematically represent the structure of the 3-D object, so cannot be computationally analyzed, while the latter does, so can.) And a ‘universe like our own’ means it exhibits ambiguity, the overriding distinctive feature of literary material, and any information, for that matter. Such a universe also represents multiple competing ideas at once. And no less important, it is flexible. It can be emended or added to without it ever giving up from its structure the history of its own making. And it will always be emended and added to in a never-ending process because a literary work is autopoietic.

A graph of a literary work also essentially comes in two layers. The first layer models the documentary foundation of the work. The second models all the editions of the work: the editor’s edition, the previous editions, and users’ editions.

Let’s take the *OTS*, which consists of voice, text, and image. The first layer of a graph for the work would refer to its basic, atomic, documentary phenomena: audio words of the audio and the indivisible visual events that still carry meaning on the prompt texts and typescripts.⁷ The

7. One could theoretically break down the audio words into phonemes, but audio words are easier to delineate, and if phonemes are necessary to represent — say,

graph would refer to them using start and stop times and pixel regions, respectively. Segmenting the recorded broadcasts into audio words, however, is difficult, especially in an oral performance done in multiple languages and incorporating both singsong and song. Not to mention the fact that the notion of an oral word in any language as a unit is fraught with complication.⁸ But an editor can execute the task pragmatically by taking the *commonsense* view of a spoken word, which, of course, is itself an ambiguous notion, one arising from the idea that one simply knows a spoken or sung word when one hears it. Extracting atomic visual events from the typescripts is similarly difficult given problems of visual interpretation, but an editor may again proceed pragmatically by taking a *commonsense* view of various units of writing and drawing, such as letters and distinct parts of illustrations. Thankfully, machine-learning tools can help accomplish both these tasks. Google Cloud's Speech-to-Text service, for example, can provide an initial set of audio words, along with their start and stop times, that can then be manually corrected against one's own transcript. OpenCV can determine the pixel regions of every letter and every distinct part of every illustration on the typescripts.

After establishing the graph's documentary foundation, the editor can lay over it — or *into* it, since a graph can be both hierarchical and flat at once — a higher-order ontology, one that shapes the editor's own edition of the work. The pixel regions can be combined into more abstract entities relevant to the *OTS*. Written words, typographic designs, songs, illustrations, for example. The audio words, defined by start and stop times, can, too, be combined on the graph into more abstract entities pertaining to the work. Songs, for example. Below is a subgraph that illustrates such elemental and non-elemental entities across media along with their relationships to each other (see Fig. 2).

This subgraph implicitly shows letters combined into words — letters and words which would have their own nodes in the graph — and words combined into a Song node and a Typographic Design node along with distinct graphic units, which, too, would have their own nodes in the graph, combined into an Illustration node. And it explicitly shows audio words combined into another Song node. Note that each node comes with a set of properties. In the Property Graph Model of modern graph database implementations, such as Neo4j, nodes can have properties, as can edges,

if their expressed sound and therefore meaning is ambiguous — then they could be represented on an individual basis.

8. See HOCKEY 2000, 20 and MCGANN 2002, 11–12.

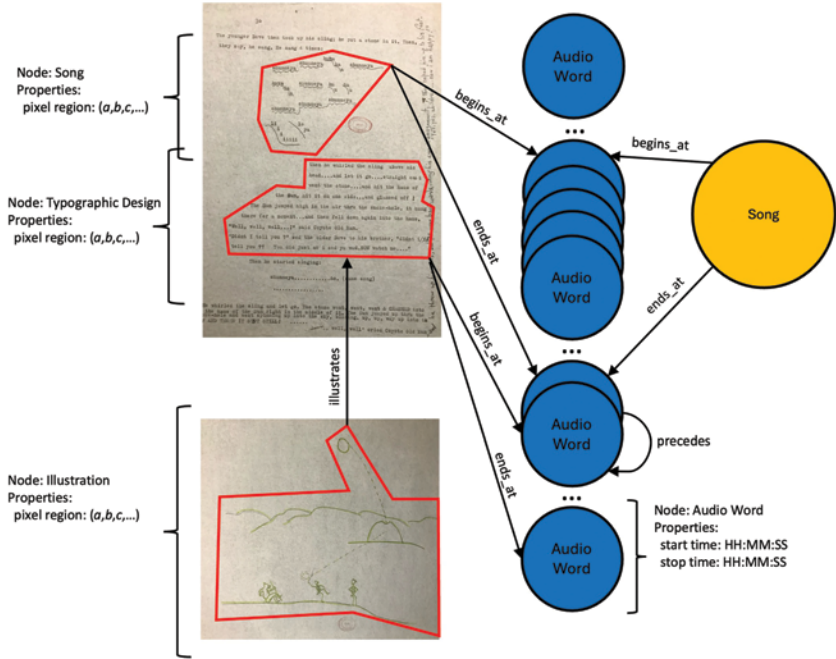


Figure 2. A subgraph that illustrates some elemental and non-elemental entities across the media of the *OTS* along with their relationships to each other. The visual material (visual text and illustration on the typescript pages) is shown explicitly for effect. In the graph, of course, the relevant parts from them would be represented with nodes that have properties like the ones shown.

or relationships. The operative properties here, again, are pixel regions for the visual-based nodes and start and stop times for the audio words. These tie the graph to the physicality of the materials, the digital versions of which wouldn't actually be stored in the graph implementation or database but on the file system.⁹ Other properties not shown here would refer the nodes to the filenames of those materials.

Basic one-way relationships are depicted here as well, but in the actual graph for the comprehensive edition of the *OTS*, there could be any number of directed edges between any two nodes that would model more fine-grained relationships, even unreconcilable ones, should the editor choose to supply them. Take, for example, the relationship between the illustration and the

9. Storing the actual digitized materials in the database would amount to storing binary data in the database, which would be unnecessary and isn't normally done, no matter the type of database.

typographic design. The illustration indeed ‘illustrates’ the typographic design, but the same can be said the other way round if the typographic design is taken as the reference point. And because of this, each one’s connotations and denotation of its rhetoric act upon their counterparts of the other’s rhetoric, forcing substitutions back and forth, causing the whole visual system, the illustration and typographic design together, to be in superposition between the two. Notice the glaring enjambment in the typographic design: “The Sun jumped high in the air thru the smoke-hole, it *hung*/there for a moment . . . and then fell down again into the haus” [*my emphasis*]. This self-referential enjambment of ‘hung’ denotes the illustration and vice versa. The enjambment as an enjambment also denotes a pause, which in turn implies a future event that takes place on the next line — in this case, the sun falling back down. But the illustration does not denote the pause; there is no visual sign of the sun’s future motion back into the house. Nonetheless, because the illustration ‘illustrates’ the typographic design, it must relate in *some* way to that denoted pause. It must *connote* that pause. And thus the pause becomes both denoted and connoted by the superposed visual system, modeled by a multiplicity of relationships and sub-relationships. When Barthes says “the discontinuous connotators are connected, actualized, ‘spoken’, through the syntagm of the denotation, the discontinuous world of symbols plunges into the story of the denoted scene as though into a lustral bath of innocence,” he is thinking of only one image in itself (BARTHES 1977, 51). But here, with illustration and typographic design relating to one other, the denotators become actualized in a connotated scene as well. The result is like some M. C. Escher of lustral baths plunging into each other from different dimensions.

Also not shown here would be relationships linking the expressive nodes of letters and audio words — again, defined by their properties, pixel regions and start and stop times, respectively — to a plane of standardized content nodes in the full graph. After all, the same letter can be written in countless ways, just as the same word can be spoken or sung in countless ways. And the graph need not store this content redundantly, only the content once along with the countless unique expressions of that content, each expression meaning something different and exposing the ambiguous condition of all communicative exchange.¹⁰

Even higher-order data can be laid over or into such a graph, data like those generated by the previously described EDA: numerical measures of acoustic properties and ordinal information indicating possible sequences

10. In Fox 2024, I explain the content plane of a graph for a literary work more fully but as applied to a manuscript’s revision stages.

of episodes. The first type of data would simply be stored as values for properties of audio words or combinations of audio words. The second would be stored in the graph's structure itself. Below, for example, is a subgraph that illustrates two performances of the same episode, marked as such by the reciprocal edges, or relationships, lying between them and labeled 'is_a_repeated_performance_of'. In this case, because both relationships are marked as active via their 'active' properties, the two performances remain unordered (see Fig. 3).

To order the episodes, the editor would just need to mark one of the reciprocal relationships as inactive. The subgraph also shows corresponding dialogue between these episodes, and different annotations describing the tone of two equivalent segments of the dialogue. Note that in one instance, the tone is marked as 'assertive', and in the other 'submissive'.

All of this would constitute an editor's own edition of the work built upon the foundational layer of the graph. But the editor can incorporate previous editions as well, simply by adding new edges to the foundational layer's nodes and labeling the edges with the name of this or that edition. What has been taken to be the standard performance edition made in 1991 is only the last in a long line of audial editions. It includes, the one-hundred reel-to-reel tapes culled from the original 1949–1950 broadcasts, twenty-five cassettes made from a re-broadcasted version in 1976–1977 and given to Peter Garland; and a small selection of cassettes made by Andrew Schelling from a second re-broadcast in 1983. And there is a history of textual versions, too, stemming from typescripts and prompt texts. This list is even more extensive. It includes, among many others, the often-reprinted *Indian Tales*, first published in 1953 by A. A. Wyn, Inc., and the 1976 multi-volume *Old Time Stories* published by Turtle Island Foundation. The graph can include every extant edition, both performative editions and textual, standing in equal relation to each other at the most granular level, and the graph's content plane would serve as a nexus for all shared content between the editions, ensuring that no edition across time, thus no medium either, serves as the only point of reference.

And now imagine a scholar comes to this collection of editions of the *OTS* seeking to create yet another edition, ordering the work, defining its documentary structures and their relationships, and annotating it however he or she wants. The scholar can do so on top of the foundational layer, or even on top of another edition, adding the new edition in the same manner to all the previous ones, without erasing any part of them.

Such a comprehensive graph edition, one that contains many, would afford a scholar great analytical power in the form of inferential queries.

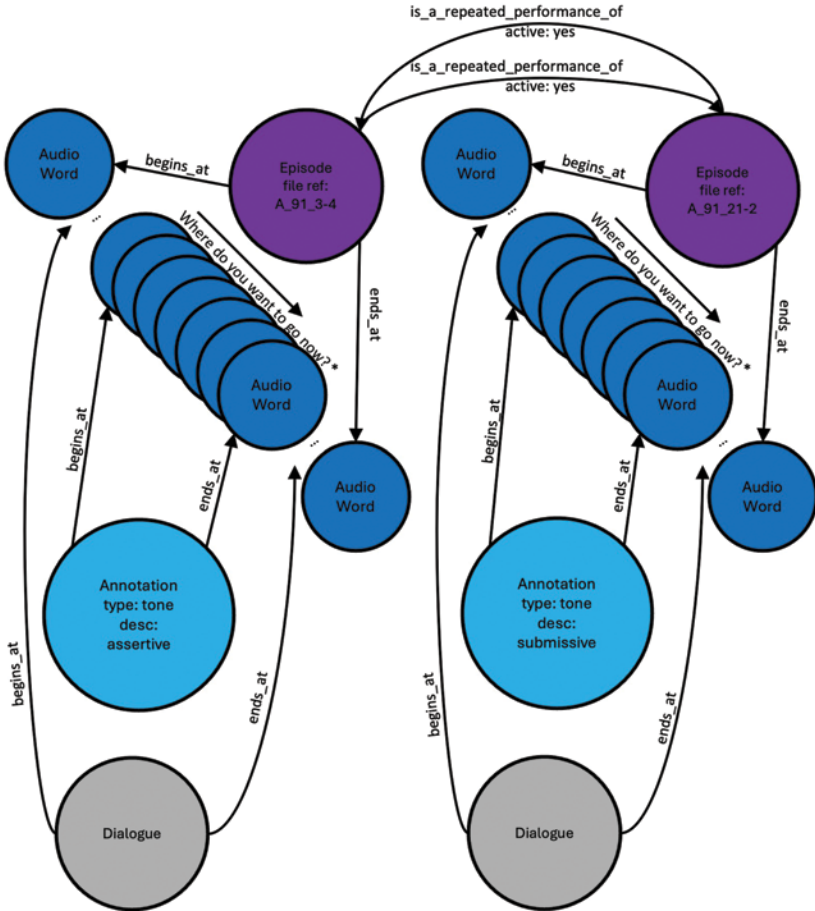


Figure 3. A subgraph that illustrates two performances of the same episode, marked us such by reciprocal edges, or relationships, between them labeled ‘is_a_repeated_performance_of.’

Suppose, for example, one were to ask the full-fledged subgraph of any particular edition to return the place and time of each tale being told, then list all the characters present during each tale-telling along with the name of the character who tells the tale. Given different native language sources for different tales, such a query would expose inferential knowledge about the varying nature of storytelling across native language sources. Or, consider further inferences that could be drawn if a scholar were to query the graph to return relationships between a character such as Baby Quail, who never speaks, only laughs, and other characters. What do these other

characters say in her presence? And with the graph's answer to that, one could answer this most basic question: What does what they say tell us about who they are? Often, indeed, as Horatio does for Hamlet, to recall a salient example, a character's relatively silent interlocutor serves to make the character most self-revealing.

These examples apply to only a couple of the work's specific informational horizons, its narratological and dialogic rhetorics. A scholar might focus elsewhere on the graph, say, on the multimedia documentary evidence for the three installments that comprise the audial version of the work, discussed at length in the previous section. Then, the scholar might ask the graph to return the subgraphs of the broadcasts and typescripts whose linguistic content is largely equivalent, and to list the acoustic and textual differences within each set from edition to edition. Even before the graph might serve as a richer source for quantitative or statistical approaches to literary studies — richer because it explicitly models the multidimensional expression of the work — the hermeneutic possibilities are endless, especially given that the graph allows for endless modification.

This is key. For editors who edit natural language-based historical artifacts, whether they're oral, textual, or visual, the fundamental problem of communication concerns neither linearity nor finality but what McGann calls the textual condition, which is dialogic.¹¹ Artifacts are born and exist in constant flux from a process of exchange or transaction across space and over time — transaction that is agent co-dependent and that works upon, or changes, through many mechanisms of interpretation, however minutely, without end. Such an indefinite domain is primed to be revealed and modeled by computational methods that are themselves characterized by indefiniteness.

University of North Carolina at Chapel Hill

Works Cited

- ANGULO, Jaime de. 1950. "To Dorothy and Ezra Pound". July.
 BARTHES, Roland. 1977. *Image-Music-Text*. New York: Macmillan.
 Fox, Michael. 2024. "Graphing VALA, or The Four Zoas: Toward a Dynamic Edition".
 In *William Blake's Manuscripts: Praxis, Puzzles, and Palimpsests*, edited by Mark CROSBY and Josephine A. MCQUAIL, 205–30. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-031-47436-1_10.

11. See MCGANN 2020, 3 and MCGANN 2022a, 54–55, where he in turn cites KIRSCHENBAUM 2014 and KIRSCHENBAUM 2021, 69–73.

- GARLAND, Peter, ed. 1988. *Jaime de Angulo: The Music of the Indians of Northern California*. Stockton: Soundings Press.
- GIANNAKOPOULOS, Theodoros. 2020. "3. Feature Extraction". GitHub. July 3, 2020. <https://github.com/tyiannak/pyAudioAnalysis/wiki/3.-Feature-Extraction>.
- HOCKEY, Susan M. 2000. *Electronic Texts in the Humanities: Principles and Practice*. Oxford: Oxford University Press.
- KIRSCHENBAUM, Matthew. 2014. "Operating Systems of the Mind: Bibliography After Word Processing (The Example of Updike)". *The Papers of the Bibliographical Society of America*, 108.4: 380–412. <https://doi.org/10.1086/681565>.
- . 2021. *Bitstreams: The Future of Digital Literary Heritage*. Philadelphia: University of Pennsylvania Press.
- MCGANN, Jerome. 2002. "The Gutenberg Variations". *Text*, 14: 1–13.
- . 2020. *The Textual Condition*. Princeton University Press.
- . 2022a. "Editing and Curating Online: Beginning Again". *Textual Cultures: Texts, Contexts, Interpretations*, 15.1. <https://doi.org/10.14434/tc.v15i1.34497>.
- . 2022b. "Breakthrough into Performance: A Touchstone Work of Late Modernist American Poetry". *PAJ: A Journal of Performance and Art*, 44.1 (130): 16–29. https://doi.org/10.1162/pajj_a_00594.
- . 2023. *Untitled*. Unpublished manuscript.
- SCHLOEN, David, and Sandra SCHLOEN. 2014. "Beyond Gutenberg: Transcending the Document Paradigm in Digital Humanities". *Digital Humanities Quarterly*, 8.4.
- TUKEY, John W. 1962. "The Future of Data Analysis". *The Annals of Mathematical Statistics*, 33.1: 1–67. <https://doi.org/10.1214/aoms/1177704711>.