Think Like a Machine (or not)

Katharine Coles
University of Utah

ABSTRACT

In the context of a larger discussion about how to incorporate Humanities values into evaluations of visualization projects in Digital Humanities, this position paper uses the Poem Viewer and POEMAGE collaborations between computer scientists and poets to consider a specific case in which working with computer scientists and computers on poetry visualization tools helped one poet/poetry scholar refine and deepen her thinking about poetry in both its sonic and figurative dimensions as she continued to engage in her primary work of close reading and textual analysis. Based on this case, it then argues that working with the computer may benefit close readers of poetry not only by showing them relationships in poems they might not otherwise have noticed, but also by teaching them new, more precise ways of approaching and thinking about the operations of poems, including metaphor. It concludes by suggesting a few concrete ways of incorporating these kinds of scholarly activities not subject to the usual kinds of measurement into evaluations.

Keywords: Digital Humanities, Poetry, Visualization, aesthetics, rhyme, sound, collaboration, methodology, assessment.

Index Terms: Digital Humanities, Poetry, Visualization, aesthetics, rhyme, sound, collaboration, methodology, assessment.

1 INTRODUCTION

In 2012, poet-scholar Julie Gonnering Lein and I entered into a collaborative project with Min Chen and Alfie Abdul-Ramen at the e-Research Centre at Oxford University in developing Poem Viewer [1], which sought to visualize a variety of relationships in poems. Later that year, we decided to collaborate with two visualization scientists at the University of Utah, Miriah Meyer and Nina McCurdy, on the development of a second tool, POEMAGE [2], specifically to identify and visualize complex sonic relationships in poetry. Gonnering Lein and I have written elsewhere about our initial concern in regards to both projects that the technology might distract us from our core work, close reading and intense engagement with poems, rather than enhance it [3], [4]. Because much Digital Humanities research at the time was driven by technology and the desire to extract data from or quantify texts or their features rather than to engage the kind of qualitative, aesthetic experience we considered central to our pursuits, we worried that the project would require us to subordinate our deepest values to accommodate ourselves to what the machine could already do. As I continue to consider the ongoing value of the tool to my own reading – in other words, to consider how I as a poet would "evaluate" the tool and its usefulness to my primary work [5] – I am increasingly led to think about how the project as a whole encouraged me, and continues to encourage me, to extend and refine my thinking (even my *imagining*) about how poetry itself actually operates. This process began with the thinking I had to do to help the computer scientists "teach" the machine what to look for in poems in general, and has continued as my colleagues and the machine have eventually "taught" me to think about poems in both a more precise and a more nuanced way. Given that I see this thinking as perhaps the most important result of our work, at least to me, I consider the question of how to evaluate it as we consider the overall impact of these projects pressing. Others on the technical side of both projects have offered both project evaluations and new perspectives on evaluation [6], [7]; the view I present here remains personal and rooted firmly in my own experience.

2 EXPOSITION

2.1 Project Overview

My colleagues and I have written and presented extensively, in group papers and individually, about both the Poem Viewer and POEMAGE projects and how the resulting tools differ from each other and from other tools available at the time. In both groups, we sought to create tools that open the possibility for the kinds of exploratory investigations experienced scholars of poetry most value by engaging the richness and ambiguity of a given poem and the reader's interaction with it. To this end, both tools permit users to load poems of their own choosing and visualize poetic elements and relationships within these poems in real time. Poem Viewer's strengths lie in its ability to identify a variety of poetic elements and in how it gestures in its glyphs toward the embodied nature of sound as it works in the poem, but this tool does not retain the poem's spatial and temporal integrity and, for this and other reasons, cannot show how these elements interact and play out across the space of the poem. The more recently-created tool, POEMAGE, visualizes only one element, sound. Rather than replacing or reconfiguring the poem, the visualization retains what we call the "poem space," honoring complex and evolving sonic relationships while also allowing the poem itself to remain visible and legible.

The central problem we faced in each group will be familiar to many: how to create a collaboration that would be productive for both the scientists and the poets on the team. Poem Viewer scientists primarily sought to create new tools that would visually present information in ways that were legible and informative to

- *Alfie Abdul-Rahman, a Research Associate at the Oxford e-Research Centre, can be reached at alfie.abdulrahman@oerc.ox.ac.uk.*
- *Min Chen, a Professor at the University of Oxford, can be reached at min.chen@oerc.ox.ac.uk*
- *Katharine Coles, a Distinguished Professor in the University of Utah's Department of English, can be reached at katharine.coles@utah.edu.*
- *Julie Gonnering Lein, an independent poet and scholar, can be reached at jkglein@gmail.com.*
- *Nina McCurdy, a PhD candidate in the University of Utah's School of Computing, can be reached at nina@cs.utah.edu.*
- *Miriah Meyer, an Associate Professor in the University of Utah's School of Computing and a researcher in the Scientific Computing and Imaging Institute, can be reached at miriah@sci.utah.edu.*

the users. This was also a goal for the POEMAGE scientists, though their interests evolved as the poets expressed their desire to use tools in their reading (and, it turns out, their theorizations) if and only if these tools did not become the point; did not, that is, come between them and the poems at hand and so replace the pleasurable work of reading. The poets were less interested in information and clarity than in developing a tool that would show the movements of poetic elements in their full complexity. The POEMAGE scientists saw the poets' priorities as providing potential opportunities to develop novel analytical tools. For this reason, the second team decided to focus only on a single element, sound, that, though complex, seemed potentially subject to computer analysis. To capture the progression of sonic clusters as they repeat in different and evolving combinations across syllables presented a computational problem that required the POEMAGE technical team to develop a new system, RhymeDesign [8], which allows users to query a broad range of sonic patterns within a given poem. RhymeDesign enables POEMAGE, which is built on top of it, to allow users to explore the interactions of sonic patterns within the poem space. In other words, the insistence of our poets that we visualize something complicated enough to be of interest to us in our work as we actually practiced it didn't impede but enabled the computer scientists in their quest to move their field forward into previously uncharted territory.

2.2 Teaching People Teaching Machines

Obviously, to quantify the operation of even a single poetic device (rhyme, meter, etc.) across a poem's space while fully honoring its complexity is not a simple task. In fact, the first challenge we as poets faced was simply describing any poetic feature at all in terms that a machine *could* quantify – a task that turned out to be interesting, productive, and highly theoretical in itself [9]. In truth, poets often give themselves over to ridiculously vague and inflated definitions of poetry, some of which (say, that poetry tells us "in so many words exactly where we are" [10]) are nonetheless true and useful at least to poets, others (that poetry is "emotion recollected in tranquillity" [11]) perhaps lovely but even less useful, not to mention not strictly accurate. The problem became clear in the first meeting of the Oxford/Utah team. Min, Alfie, and company had been listening patiently for some time as Julie and I waxed on about poetry. Finally, Min asked me if I could tell him "exactly" what I meant when I talked about "poetic time."

Well, no. I still can't, though these days I can come a lot closer. I realized in that moment that when he said "exactly," he meant in terms that could be *quantified*, an idea that was frankly overwhelming. For much of both Poem Viewer's and POEMAGE's conceptualization, design, and prototype periods, then, both Dr. Gonnering Lein and I wrote and presented extensively not about our work with the tools but about our work on the tools [3], [4], [9], [12-16], by which I mean that our

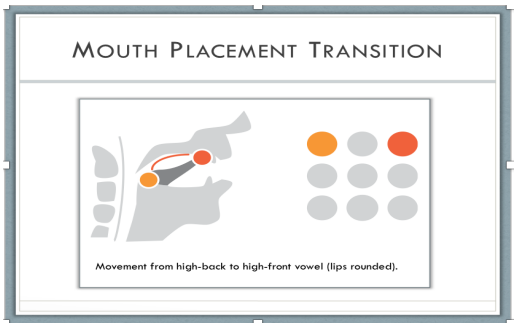

Figure 1: Mouth placement glyph from Poem Viewer.

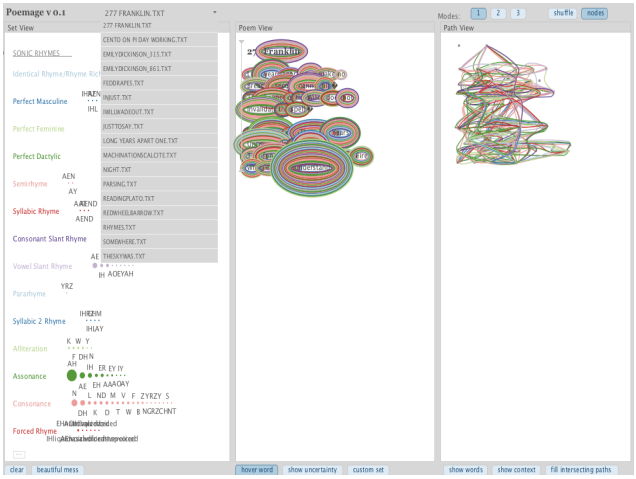

Figure 2: Screen shot of menus and Dickinson's #277.

thinking at the time was rooted less in reading with the tools than in our consideration of what such tools needed to be and do in order to tempt us to read with them. Simultaneously, we needed to reconsider the operations of poetry in a way that might allow us to teach any tool to be an effective partner in reading poems. This re-theorization of close reading itself began as a necessary step in the design process, since figuring out both what elements and relationships to visualize and how to teach the computer to identify them involved a productive shift in the way we engaged poems. In developing POEMAGE, for example, this process led not to the narrowing that might have been expected and that we feared, but rather to a radical expansion of our working definition of "rhyme." Such constant and expansive reconsideration has continued apace in my work, as I have responded to visualizations not only with close readings but by extending and generalizing those readings into the theoretical sphere, allowing the tool to teach me how to think in new ways about poems rather than only, in my return to poems, following the same kinds of paths down the page I always had [14], [15]. My experience so far suggests that if we continue to insist on the relationship between machine reader and human reader as reciprocal and alive, the machine might help us to imagine new ways of thinking about metaphor that will let us, in the meantime, not "solve" metaphor through the machine but at least newly consider it in attentive collaboration with the machine. In other words, the kinds of theorizing about poetry and how it works that we must do to "teach" the machine how to show us something we might not otherwise see may become one of the machine's mechanisms for teaching us things we didn't previously know about how poetry works.

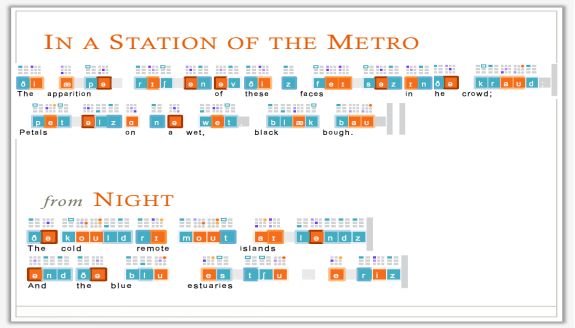

Figure 3: Poem Viewer visualizations of "In a Station of the Metro" and the first two lines of "Night."

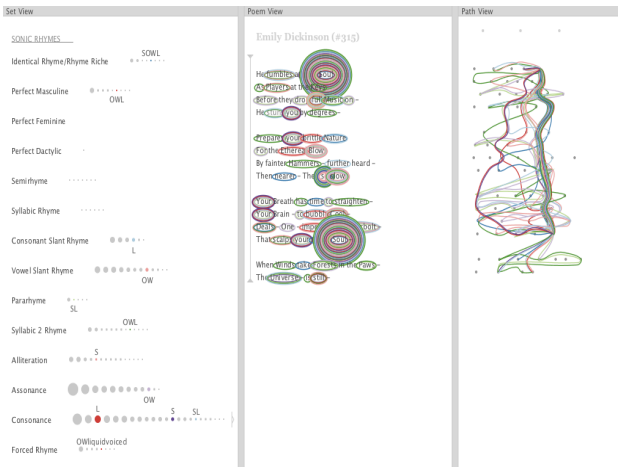

Figure 4: POEMAGE visualization of Dickinson's #315.

3 Discussion

3.1 Crossing Time

In response to Min's first question about poetic time, which engaged a mode of thinking about poetry to which I was entirely unaccustomed, I switched gears and began to talk about sound as one way in which a poem marks its time, by moving (like most poetic figures) simultaneously forward and backward, constantly retarding the poem's progress by recuperating and repeating sounds that have already occurred even as it drives the poem down the page through sonic resistance and change.

That early question also led me at the outset of this work to a serious consideration of how a poetic image emerges across the regulated space-time of the poem, during which I closely read the operation of the image in Ezra Pound's "In a Station of the Metro" (fig. 3) [4] among other poems. My argument questioned the way we think about the construction and operation of images in poems, especially in poems like Pound's that work to create a sense of immediacy in relation to the image. In fact, I argue, this sense is – must be – an illusion, given the way poems themselves work their temporal spaces through syntax across lines and down the page. We are not yet ready to consider ways of visualizing the development of images across syntax much less across poems, but I can hope that this theoretical work may eventually contribute to such an innovation.

Likewise, Poem Viewer's use of glyphs to show where in the mouth a given sound is made (fig. 1) caused Gonnering Lein and me to think (and write) about the body as site and instrument for sonic play in the poem and also about poetic "turbulence" (as marked by the shifting location of sound in the mouth) as one marker of a poem's temporality. The idea of turbulence in the poem led us to think of the poem and its temporality in terms of "flow," which we defined as *a fluid (or fluids) moving via its linguistic devices and figures through a defined space* [9]. Of course, this definition is metaphorical. A poem is not a fluid, but may behave through its syntax and devices as a fluid does—as might also traffic, or neurons, or patterns of migration, all of which can also be described in terms of "flow."

This summer, I have presented two papers at poetry symposia abroad [17], [18] about how temporality behaves differently in lyric than narrative. This thinking is firmly rooted in the theorization I did for and from these digital projects, though for the time being, as with the image, questions of temporality in lyric, like those about metaphor, are questions I expect the human, rather than the machine, to address.

3.2 Sounding Metaphor

Still, there has been a lot of talk recently, some of it from me, about how we might get a machine to "solve" metaphor, by which I mean simply to identify it reliably. As difficult (perhaps impossible) a problem as this is to tackle algorithmically, the machine will have to address metaphor before it can address poetic time, if only because the temporality that emerges in any poem does so as a function of every one of its poetic elements, perhaps especially including metaphor, operating with and inextricably enmeshed with every other.

Metaphor has generally been considered computationally intractable because of its semantic and syntactic complexities. As I noted above, a poem is not a fluid, and to say it is reveals both the similarities between the two and the differences. In fact, metaphor's revelations inhere not in the similarities, where the metaphor succeeds, but in the differences, where it fails. An additional difficulty with metaphor lies in its grammar, which can be slippery even in simple instances. Getting the machine to understand why "Hope is a bird" (or, far more problematically, "'Hope' is the thing with feathers") is a metaphor but "Juliet is a Shepherd" and "Karen is a Carpenter" may be either metaphors or similar or different statements of fact is not straightforward. The difficulties compound when poets as different as Dickinson and Shakespeare (who tells us "Juliet is the sun") play metaphors out across poems or passages in elaborate and shifting figural structures. To develop a tool that can reliably identify metaphoric relationships as POEMAGE identifies sonic relationships—in real time across the entire poetic field—would require the solution of multiple, and hard, open problems in computer science.

My recent theoretical work, however, suggests that our understanding that every poetic element interacts with and informs every other as the various elements play out temporally through a given poem's space may provide one way of getting at, if not solving, metaphor through the element we already have "solved," sound. Combining sonic analysis with the kinds of syntactic analyses that are inadequate on their own may provide us with new insights.

Part of my extension of sonic patterning into metaphoric relationship occurred because I allowed the machine to "teach" me to think of rhyme even more broadly than, as a 21st-century poet, I was already accustomed to doing. POEMAGE can identify over a dozen different kinds of "rhymes," including not only those an experienced reader can quickly identify herself like perfect masculine and perfect feminine rhymes, but also more subtle relationships like assonance and consonance, and even less frequently noted rhyme types like syllabic rhyme, eye rhyme, and consonant or vowel slant rhyme. The user can identify as few or as many rhyming relationships to be visualized as she wishes, or even create what the POEMAGE group calls a "beautiful mess," in which all the relationships are visualized at once (figs. 2 and 4). Though the scientists expected this screen to be visually confusing and so useless, it shows precisely where the most intense rhyming activity occurs in the poem (the sites that are most sonically "interesting"), and which words are most heavily implicated in the poem's overall sonic patterning. Because the tool allows us always to see the poem alongside the visualization, the fact that the "beautiful mess" visualization obscures the poem doesn't matter – and that it is indeed "beautiful" matters to poets.

In close reading, we have long noted in passing that places the machine marks as being sonically "interesting" are also sites of metaphorical action. More recently, I have noticed that this action often inheres in, rather than simply existing alongside, the sonic relationships being indicated by the tool. This inherence can emerge through various kinds of sonic relationships, including but not limited to homonyms like "knot" and "naught," which

POEMAGE shows in Bradstreet's "Prologue," and eye-rhymes like "blood" and "mood," which it picks up across Pelizzon's "Blood Memory," about menstruation. In presenting these words as related, even conjoined, the machine opens a space for us to tease out figural connections between a loop in a rope and nothingness, or menstrual blood and emotional pain.

The machine also led me to connect the word "soul" in Dickinson's #313 with the intensifying "so," with "slow," and with "you," sonic connections that all point to the poem's evolving metaphorical complex (fig. 4). Even better, the machine can show me, in a way that my ear cannot tell me, how, through its dense sonic web, the poem connects "soul" through "slow" with the poem's last word, "paws"—a big enough stretch that I am not sure I would make it without the tool's suggestion. Through this link, and given the (violated) expectation of a rhymed couplet at the end of the sonnet, the poem arguably creates a situation in which "paws" metaphorically "rhymes" with "still" by suggesting to the attentive reader an absent but implied homonym, "pause," creating a semantic rhyme through an odd but dazzling indirection of sound.

Also instructive are the machine's productive misreadings, based on statistics, of words like "wind," which in Bogan's "Blue Estuaries" the machine reads as "wined." In insisting on "the restless [long-i] wind of the inlets," the machine marks as early as line four the beginning of an extended metaphor connecting the landscape with the winding human circulatory system, a metaphor that doesn't finally emerge until the last line of the poem, fifteen lines later, which invokes the movement of "blood in the heart."

The consideration of metaphor through sound also sent me back to Pound's "In a Station of the Metro." POEMAGE reminded me that "Metro" is a rhymed couplet (loosely pentameter, no less): though this is not what we generally attend to in the poem, "crowd" and "bough" make a perfectly respectable assonant rhyme. Though of course I had noticed the assonance before, I had done so only in passing. It was only as the machine taught me to consider how sonic resemblance might gesture toward or even contain other kinds of figuration that I came to think of it as important in the figural sense. In this case, focusing my attention on the near rhyme led me to reconsider how the image and its residual metaphor emerge: we see not only pale faces ("petals") against a black metro tunnel ("bough"), but also, as the poem substitutes "bough" for "crowd," those individual, luminous faces emerging from the collective darkness of the crowd itself.

4 Conclusion

Individual readings of specific poems mostly interest only poets and poetry scholars who are invested in such readings and the particular poems under examination. By using such specific examples here, I mean not to extend this branch of scholarship (I have published and presented far more detailed and extended readings of most of these poems elsewhere) but rather to provide examples of how not only the actual use of these tools but also the very framework and demands of the work itself, the constant needs of the machine and the people teaching the machine to understand the operations of poetry as precisely as possible, continually cause me to rethink how poetry actually works and so to extend my own theoretical practice in unexpected directions. As we consider the value of this kind of work, and perhaps especially as we consider how to evaluate its usefulness to humanities fields, we may find that such a product, which may seem to a computer scientist to be peripheral to the central work of developing a tool, is of at least equal value to the project overall and surely of greater value to the humanities scholars involved. Likewise, this kind of thinking may suggest new angles for attacking previously unsolved problems in the work of visualization science. In addition, for example, to trying to develop algorithms that might directly address metaphor's semantic complexities, visualization scientists working alongside poets might work to identify sonic and other movements and changes that are most likely to indicate the parallel development and extension of metaphor and other complex figures working across the poem's space. This could lead to ways of statistically evaluating the chances that a specific syntactic construction in a poem may indicate the presence of metaphor.

In my field, we would mark strong theoretical work, however developed, as a success in its own right – and we decide how powerful ideas are only after seeing the extent to which they contribute to and further the ongoing conversation within the field, a test not subject to immediate quantification. With this in mind, I would like to suggest that teams engaged in similar work fold some Humanities-based methods of evaluation into the process alongside those considered acceptable in the sciences.

First, evaluation should invite and include narrative and theoretical accounts by the humanities scholars of the impact of the digital project on the thinking that is at the heart of their humanities work. Groups should consider this separately from any interest or engagement a humanities scholar might have developed in the technology itself (shifting the attention of a humanities scholar away from core research and onto technology may or may not be counted as a success), and should include impacts that are both directly related to the digital work, and so expected, and indirectly related, even surprising, as in the "spun" poems developed by some of our younger collaborators [19]. Such evaluations may include surveys and questionnaires with answers that can be compared across groups, but for many scholars accustomed to thinking through writing, open-ended written accounts may be more revealing, or at least revealing in very different ways we should also count as important.

Second, while evaluation of research products should of course include conference presentations, which are highly valued in the sciences, evaluation of humanities work must emphasize essay and perhaps especially book publication, since these are the forms of scholarly dissemination most emphasized in humanities-related disciplines, including in DH. Since many of the services meant to provide metrics on scholarly publication and citation have arisen to serve scientific disciplines and so don't adequately represent humanities fields, groups should work together to document these products and their influence and to argue for their importance. Likewise, scholars (including computer scientists) engaged in DH all stand to benefit by working with their institutions to advocate for and encourage more effective inclusion and representation of humanities-based research products in such metrics.

Finally, since humanities scholarship is generally driven not by specific, provable hypotheses or the desire to obtain information but rather by the generation and development of ideas in conversation, often over periods spanning years, decades, and even centuries, groups should find ways of documenting how the ideas produced as a result of digital work express themselves not only in the work immediately connected to a project but also over time, at least on the scale of careers. As a part of this effort, humanities scholars might consider using citations and acknowledgements to explicitly link any DH project in which a line of thinking originated to later expressions and extensions of that thinking.

Acknowledgments

I wish to thank my collaborators as well as the organizations that supported this work: The Office of Digital Humanities and the Digging Into Data project at the National Endowment for the

Humanities, the University of Utah's Seed Grant Program, and the SCI Institute, especially Chris Johnson.

REFERENCES

[1] Abdul-Rahman, A., J. Lein, K. Coles, E. Maguire, M. Meyer, M. Wynne, A.E. Trefethen, C. Johnson, and M. Chen. (2013). Rule-Based Visual Mappings—With a Case Study on Poetry Visualization. *Computer Graphics* Forum 32, pp. 381-390.

[2] McCurdy, N, J. Lein, K. Coles, and M. Meyer. (2016). Poemage: Visualizing the Sonic Topology of a Poem. *IEEE Transactions on Visualization and Computer Graphics*. 22:1, pp. 439-448.

[3] Lein, Julie Gonnering. Sounding the surfaces: Computers, context, and poetic consequence. Western Humanities Review, pages 84–109, Fall 2014.

[4] Coles, Katharine. Slippage, spillage, pillage, bliss: Close reading, uncertainty, and machines. Western Humanities Review, pages 39–65, Fall 2014.

[5] Coles, Katharine. Show Ambiguity: Collaboration, Anxiety, and the Pleasures of Unknowing. #Vis4DH, InfoVis 2016.

[6] A. Abdul-Rahman, E. Maguire, and M. Chen. Comparing Three Designs of Macro-Glyphs for Poetry Visualization. In N. Elmqvist, M. Hlawitschka, and J. Kennedy, editors, EuroVis - Short Papers. The Eurographics Association, 2014.

[7] McCurdy, Nina, J. Dykes, and M. Meyer. Action Design Research and Vidualization Design. BELIV 2016.

[8] McCurdy, Nina, V. Srikumar, and M. Meyer. (2015). RhymeDesign: A Tool for Analyzing Sonic Devices in Poetry. In Proceedings of Computational Linguistics for Literature, pp. 12-22.

[9] Coles, Katharine and J. Lein. (2014). Turbulence and Temporality: (Re)visualizing Poetic Time. *Things My Computer Taught Me About Poems*. MLA2014. Chicago, IL.

[10] Strand, Mark. 'On becoming a poet' in *The weather of words: poetic invention*, New York: Alfred A Knopf, 33-44, 2000.

[11] Wordsworth, William. Observations Prefixed to the *Lyrical Ballads*. 1800. *The Poetry Foundation*. Web. 13 July, 2017.

[12] Lein, Julie Gonnering. Digital Humanities and Dickinson's 'Tell': Recounting Poetic Encounter. New Work on Dickinson: Flash Talks. Modern Language Association. Vancouver, BC Jan. 2015.

[13] Lein, Julie Gonnering. Computers in my Classes: A Pedagogy Round-Table on Workshopping (With) the Digital. Panel Discussion. AWP2015. Minneapolis, MN. April 2015.

[14] Coles, Katharine. Ghost (in the) Machine. Keynote lecture. Australasian Association of Writers and Writing Programs Annual Conference. Melbourne. Dec. 1, 2015.

[15] Coles, Katharine. In Motion in the Machine. Invited lecture. Poetry on the Move/International Poetry Studies Institute. Canberra. Sept. 2015.

[16] Lein, Julie. Seeing the Sonic: Aesthetics, Poetry, and Data Visualization. Aesthetics Reloaded. Aarhus, Denmark. Dec. 2012.

[17] Coles, Katharine. The Poem in Time. Inside-Outside Poetry Symposium. Winchester University, Winchester, UK. June 2017.

[18] Coles, Katharine. Lens. Great Writing Conference. Imperial College, London. July 2017.

[19] McCurdy, Nina, J. Gonnering Lein, A. Hurtado. Deep in Poetry: Improvisations in Tech, in Text, in Time. IEEE VISAP2-15. Chicago, IL. Oct. 2015.