

Live Performance and Post-Cinematic Filmmaking

Jeff Burke and Jared J. Stein

Introduction

The materiality of early cinema—for example, a certain number of rectangular images per second—has had a lasting impact on filmmaking across established and new media, and in turn, on live performances incorporating it. Even in the early twenty-first century, contemporary digital video production practices focused on content that was nonlinearly edited but still consisting of fixed, temporal progressions—whether representational or abstract. Like the media hardware platforms available to theatre-makers—e.g., monitors, projectors, and cameras—the content in live performances was often more or less congruous with audiences’ expectations and assumptions of content in cinematic and broadcast experiences: moving backdrops for location and mood, text for information and context, or sheer visual juxtaposition of abstract and recognizable elements. However, the steadily increasing use of computer graphics, even in nonanimated cinema, and the emergence of virtual reality (VR), with its nonrectangular presentation and audience point-of-view (POV) selection, are breaking down the vestiges of early cinema’s material affordances. Real-time editing and rendering and contemporary artificial intelligence (AI) techniques are changing how even the most conventional of cinematic images are being assembled. And notably, the increasing flexibility of *delivery* has been facilitating the divergence of media in live performance and cinema, as the former has been able to take advantage of, for example, cost-effective nonlinear playback opportunities—while unnecessary (for now) to the delivery of cinema, valuable to its production processes and thus heavily invested in. All of these evolutions parallel and impact new uses of digital cinematic media¹ in live performance and present new opportunities for live performance structures, including shifting roles for audiences, creators, and performers—during creation processes and performances, and in between each show.

UCLA’s Center for Research in Engineering, Media and Performance (UCLA REMAP) has explored the use of dynamically controlled digital video in live performance since 2004.² By “dynamic control,” we mean the use of software for triggering and/or manipulating media based on conditions within a given performance. Here, we focus on a set of the Center’s recent works that reflect parallel progress in three areas related to this control and the content it is controlling: *algorithms*, as embodied in software, used to select and render fragments of media; *cinematic media*, in

Jeff Burke is Professor-in-Residence of Performance and Technology at the UCLA Department of Theater and Associate Dean for Technology and Innovation at the UCLA School of Theater, Film and Television (UCLA TFT), where he has been a faculty member since 2001. Burke explores how technology can enable new relationships among media, performance, audience, and the built environment. He is co-founder of UCLA REMAP, a collaboration with the Samueli School of Engineering, which combines research, artistic production, and community engagement.

Jared J. Stein has directed and written plays, and co-created performance pieces incorporating drama, object, movement, and media-based approaches, for venues around the world. With UCLA REMAP and its predecessor, the UCLA HyperMedia Studio, since 2001, he has experimented with emerging media platforms within theatre-making processes as a writer, director, and dramaturge. He is currently a resident artist at Centro Hipermediatico Experimental Latinoamericano in Buenos Aires. He co-directed the Rhodope International Theatre Laboratory from 2005 to 2012.

terms of how it is conceived and shot relative to the performances to which the algorithms are ultimately applied;³ and the *performances* themselves, considering their evolution in narrative and dramaturgical structure and performative and scenographic form to become more tightly coupled with the other two areas.

In this paper, we do not trace many antecedents by name. The works discussed are influenced by and relying significantly on prevalent narrative, dramatic, and theatrical conventions—patterns of plot, some quite ancient—seen as clearly within work wherein those patterns are the focus of dramaturgical revolt. To some degree, the composition of dramatic action and choreography must always have resulted in algorithms for at least humans to follow. Endings, if nothing else, have to be prescribed by something. We have thousands of years of structures from which, dramaturgically, we can find parallels to power structures of new media technologies, e.g., *deus ex machina* of farce wrapping itself up farcically, the concert of form and improvisation within known Commedia dell'arte techniques, rituals of scheduled randomness such as happenings, and redefinitions of zeitgeist from the dawn of the Industrial Revolution through Modern Drama and the Symbolists to the patterns of Beckett.

We can also find parallels to countless explorations of relationships between manipulators and “puppeted” objects, and to different types of meaning that can be found in single, specific gestures within movement theatre traditions outside our own. For us, the structural points of reference intersect with investigations in database aesthetics—by media artists and theorists including Fabian Wagmister (at UCLA REMAP), Lev Manovich, Lynn Hershman Leeson, and George Legrady—while the corporeal role of the body and the object onstage intersect with many telematic performance and cybernetic explorations, such as works by Adriene Jenik and Stelarc. Approaches to nonlinear, respatialized film and video by Josef Svoboda and many contemporary projection designers also set up a post-cinematic approach. At the convergence of these two broad historical roads (though one is quite a bit shorter than the other!), we aim to explore how live performance, narrative drama in particular, with its evolving patterns and built-in expectations, can not only *react* to unique capabilities of digital media, computation, and networks but *drive* technology development specific to the theatrical stage.

By concurrently developing software systems, media, and dramaturgy (i.e., the algorithmic, cinematic and performative), we have been able to experiment with their intersections—iteratively moving forward on the pieces’ overall constructions as we attempt to make aspects of the three work in concert. With the projects described below, we explore how authoring on these three layers simultaneously is generating or uncovering for us an implicit post-cinematic concept of media. By articulating this concept here, and relating it to contemporary trends, we hope to identify questions that can be pursued in future work and generalize about opportunities for the evolution of the relationship between media and live performance.

In all of these processes, algorithms executed by software code have been employed for the real-time selection and display of media content based on aspects of a story-in-progress. Through this, the pieces address issues at the junction of storytellers’ need for authorial control and desire to provide agency to audiences and performers in response to fluctuating circumstances, in contrast to what normally coincides with more conventional media designs, such as uses of media per static sets of circumstances (e.g., scenes in a play) and predetermined events. We imagine in the future, machine learning (ML) and AI technologies will bridge the gap between composition of live experiences and agency within them—with new types of media within new types of human-authored

narratives, allowing for worlds that synthesize, in ways we have yet to imagine, new inputs (and/or contexts) and decisions to be made by artists. We will touch on this in the conclusion; first, it is useful to lay some groundwork, starting with common abstractions, and then with specific examples. Project by project, we have noticed some common threads that, to us, suggest some early fabrics of this synthesis.

When using “author” in its various forms, we mean to emphasize the notion of fixing elements within performance ahead of their use—within text, code, or other elements—accompanied by the general expectation those elements are to remain fixed during the life of the work even as they may be interpreted by others and contextualized differently in different productions. For our purposes here, even procedural or algorithmic compositions, or open-ended performance structures, are “authored” in this way, because the algorithms or procedures are fixed even if they describe an input-output relationship that is open and provides, for example, agency to the audience. We use “storytelling” to denote a more general creative act that may or may not involve authorship of fixed material. A general direction of this paper is to explore how new technologies are enabling authorship of more sophisticated relationships between preexisting, generative, and audience-supplied elements within storytelling, along with authorship of how those relationships themselves might evolve through artificial intelligence and machine learning.

Common Threads

One of the authors’ key collaborative interests has been to develop new forms of dramatic literature and the software (or systems) they require in parallel, in order to explore new theatrical possibilities that emanate outward from the performance text via algorithmic control of media and other performance elements. Here, we omit some strands more related to nontextual aspects of the works, for example, the use of real-time sensing to create new engagements with the body on stage. As it may (accidentally) reinforce a dangerous indifference to bias and subjectivity in the real-world creation and applications of algorithms, we do not fully subscribe to the concept that algorithms are content-neutral or media-neutral, discussed recently in Pizzo, Lombardo, and Damiano (2019). However, with these threads, we emphasize those elements that can be considered within the creation of texts to be subject as much as possible to artists within worlds they create, with many different potential realizations, and with the real-world realities of the media or device being used as thematically related to their creation as they wish.

From Scripts and Trim Bins to Databases

Similar to others creating on digital platforms, we have long been organizing digital content, including text and cinematic media, in databases as opposed to according to a conventional linear list or timeline. This is an unsurprising approach for contemporary artists working on digital platforms, given the affordances of new digital media, including those identified by Manovich at the turn of the century, such as its numerical representation and modularity (Manovich, Malina, and Cubitt, 2001). Rather than triggering digital cinematic media segments in live performance according to a more linear and fixed schedule—e.g., stage managers’ cueing or scripted events—in much of our work, content objects need to be accessed nonlinearly and automatically based on algorithms and their applications to improvised events. While the trim bin of digital editing could be considered a database, we refer here to a collection of content that is ready and intended for real-time assembly and fluid delivery rather than raw materials used in preparing these final cuts.

Bridging Humans and Algorithms via Folksonomies

To organize these database media, we often apply semantic tags (keywords) following a taxonomy formulated for a piece and/or emergent within it. More complex metadata structures for organizing database media within our earlier works have given way to flat collections of such tags, which are easier for collaborators and audience members to understand and less arduous to apply to large corpora of content. All the tags used in a piece (the tagset) form a folksonomy, with a role similar to that of hashtags that create relationships on social media, yet more highly curated. Their use was influenced by other work at REMAP in community-generated digital artworks. When necessary, depending on their relationships to the content they are labelling, points in the story, possible audience inputs, any real-time digital manipulation required, and other elements of the experience, tags can be organized into hierarchical families for a more structured folksonomy. A group of characters attempting to create division among an audience, for example, could apply tags to individual audience members corresponding with their answers to certain questions and then create further division, or even reconciliation, depending on how those tags relate to tags applied to audiences of past performances. The tagset itself could be provided by a piece's authors, generated by the audience, or devised ahead of time and/or revised in real time by the performers, or any combination of these. We use "folksonomy" rather than "taxonomy" to convey this flexibility, as well as the potential for the public (audience-facing) availability of the tags for reuse, as with hashtags in social media. While extensive metadata systems, such as Dublin Core, have their place in organizing media, we have found that the use of folksonomies is an efficient way to bridge human understanding of media relevance with algorithmic decision-making.

Three Dimensions of Agency: Interpretation, Improvisation, and Participation

Implicit in the use of random-access databases to hold and organize media according to folksonomies is a view of media, including text to be delivered to performers or audience members, as a dynamic, evolving facet rather than a fixed overlay. This modular, random-access view of media, including cinematic elements, emerged from our interest in exploring new types of agency enabled by digital systems in live performance that parallel some of the societal promise of contemporary digital technology, whether it has been realized commercially or not.⁴ In an *a posteriori* look at the Center's work in live performance, we identify three dimensions of agency being explored:

Interpretation: the agency of ensembles, wherein the relationship between authorial and directorial control is delineated and the creators of a performance design a digital system—particularly the code embodying its algorithms—to work for their particular progressions (more traditional dramaturgical elements) and their interpretation of them, which can change as expressed in actions, dialogue, and media, as well as in their relationships to each other.

Improvisation: the agency of performers, choices made in response to varying circumstances within the performance including those executed by digital components, and in return potentially influencing changes in components or creation of new components that then influence other choices by performers or even audiences.

Participation: the agency of audiences, opportunities for public participation to impact circumstances and make choices, including those related to creation, changes or use of media content and taxonomy.

Cinematic media content has become a vehicle for these dimensions of agency in our work, as it has become easier to create, capture, select, remix and manipulate with technological advancements in digital media while retaining its visual power and cultural relevance. These dimensions of agency often relate closely to a thematic and technical focus on the *points of view* of performers and audiences, as fields of capture, as well as fields for providing customized content, as in augmented reality and its personalized presentations overlaid on individuals' points of view. This is described in more detail in our examples.

Metatheatrical Bridges

In this paper, we define a metatheatrical bridge as a technological system that is purpose-built for a given piece and is acknowledged in some form within the piece's own story or world. It is a technological construction with narrative functions that bridge fiction and reality, to varying degrees acknowledged directly or disguised by story elements—parallel to how, within more conventional dramaturgy, blackouts can be due to power outages, characters turning off switches, or simply the “hands of the storytellers.” With much of the work described below, the goal was for such systems' diegetic and nondiegetic presence to be interwoven and, in doing so, bridge the agency and identity of the audience and performers within both the fiction and the real world. One of our aims for using such bridges is to achieve intrinsic relationships between text and software (or system) that explore what is unique about digital computing among the many technologies incorporated into theatre and performance throughout history.

Live Performance Use Cases

Starting with two original works by students as part of a research and curricular program created and supervised by Burke and advised by Stein, and concluding with a series of three projects created by Burke and Stein together, we describe specific instances of these threads, and their evolving integration within what we have come to see as a single dramaturgical-technological project.

Grace Plains (2014)

Grace Plains, a student-authored, -directed and -designed interactive theatre piece at YouTube Space Los Angeles, integrated live-action role-playing, murder mystery, and morality-play tropes. Small groups of audience members wearing optical head-mounted displays (Google Glass), providing them with information about their identities as potential investors and other role-playing hints, were taken along a tour of a scientific facility by its high-ranking staff members, played by two actors. The audience was led along a continuously changing set of circumstances to decide what to do with an increasingly powerful (and human-seeming) AI. They were introduced to the AI, voiced by an actor, and got to know it. At the height of a conflict between the lab director and the AI inventor about the development and purpose of the AI, the potential investors came across a dead body. The AI was blamed, and the audience was forced to decide as a group if the AI should live or die.⁵

Real-time agency, of the ensemble, performers, and writers, was the most prominent exploration within this piece. A team of writers in a control room fed the audience participants and actors, also wearing the devices, suggestions for dialogue and motivations for action—giving them guided agency to improvise.



Control display from *Grace Plains*, showing each audience member's and actor's POV and current customized text, focused on their character within a fast-evolving experience.

Technicians, in a separate control room, manipulated media within the piece's surroundings. Online audiences could watch a composite POV cut live from participants' head-mounted feeds by a television director. On-the-fly choices were made by the writers (watching each feed) to instigate improvised theatrics, using a database of text providing them with prewritten phrases, and also to further instigate drama headed toward an eventual (known) finale. The audience participants had significant agency to make choices within their role-playing context, as guided by the writers, who were, in effect, also live directors. Unlike the online audience's single feed, consisting of its director's called shots (a multiple-camera live television shoot), the writers could feed any text to any or all participants, and could do this in any combination of pairs or groupings within the group as a whole. (See <http://graceplains.com/> for the students' explanation of the experience.)

In this piece, unlike the others to follow, there were no direct metatheatrical bridges between the technology and the story. The head-mounted displays existed outside of the narrative, functioning as both cue cards and hidden cameras. One of the most important experiences drawn from this experiment was the use of machine-supported real-time writing, by human authors, to support audience agency. The technology allowed aspects of oral storytelling, role-playing, and directorial techniques to be applied individually with each audience member, without the others knowing, and suggested a direction for experiences to be both personalized and collective, despite an inability to mitigate the distraction of each person having a head-worn display. It was also an energizing experience for authors to engage with the piece as it continued to evolve in the hands of each audience, similar to live-action role playing in other contexts—with a mix of events devised in real time and ones that were prewritten, organized by a simple and emergent folksonomy of their

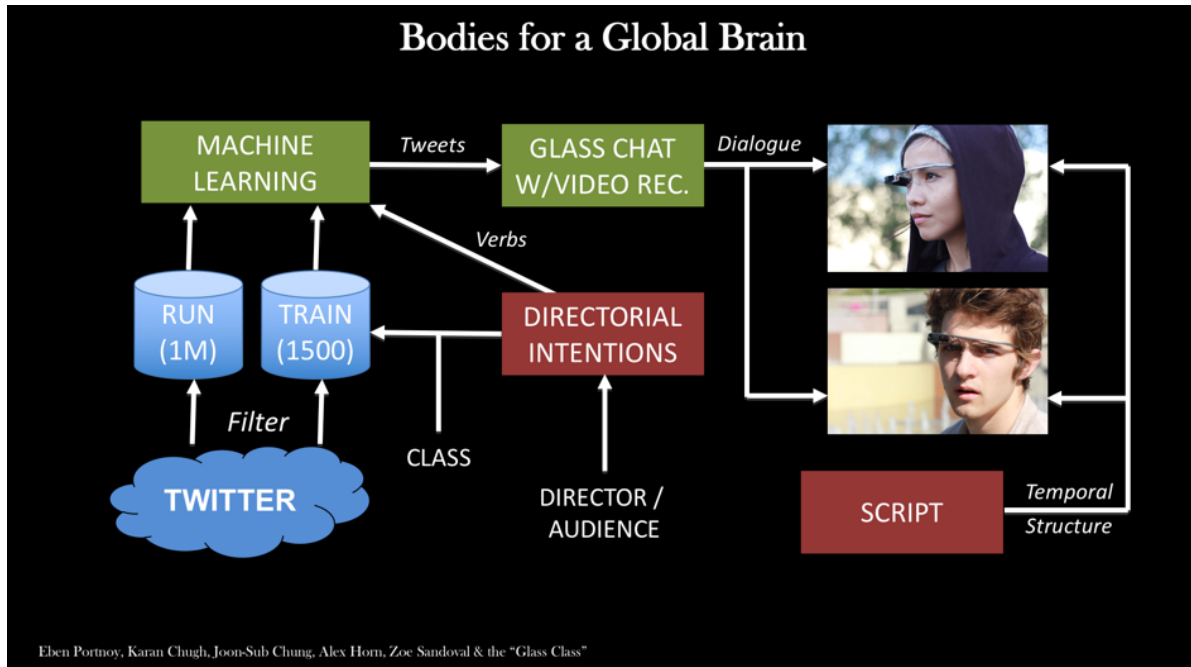
applicability to situation and character. In some ways, it was also similar to how stand-up comedians' acts are workshopped over time and bands revise their playlists for concerts, but far more intricate, having to incorporate actual actions of audience members rather than just changes in response to group laughter, applause, heckling, or calls for encores. The experience suggested considerations for creating dramaturgical structures that facilitate and incorporate evolution. We returned to these ideas in some of the later projects described below.

Bodies for a Global Brain (2014)

Bodies for a Global Brain, a short student web series created within the same research and curricular context as *Grace Plains*, used as a “script” text mined from social media that was selected dynamically for performance, using ML. It also explored real-time choice and reaction by the ensemble. A young couple, previously only having met online, decide to allow their thoughts to be controlled by the will of the Internet, via a new global consciousness system they designed. The pilot episode follows them as they meet and then begin to interact with the world and each other at the whim of the network. The result was a partially improvised performance based on a database text, selected via a writer/director-provided folksonomy that resulted in linear cinematic media.

While shooting, the actors, with a scripted set of circumstances and pre-rehearsed character arcs in mind, were fed the mined text via optical head-mounted displays (again, Google Glass). They were bound to this text to accomplish their characters' objectives—creating new obstacles, types of tactics, and plot twists; they had significant agency. Though the words forming the dialogue appeared as if completely random, the writer/director had several ways of sculpting the overall experience. A backstory was created to give the actors insight into their characters. The script itself was first constructed as a series of scenes with emotional arcs of the characters' relationship that could be rehearsed in a traditional way. Then, a large amount of Twitter data—a database of text—was captured, and a small set of that data was manually tagged with “intent” tags by a large group of students based on a folksonomy created by the director to match the “intents” within this arc. The tagged data were used to train the ML system, essentially generating an algorithm that provided a mapping between directorial intent keywords (tags) and matching text from a larger Twitter corpus. (Note here that the tags are *not* Twitter hashtags; they are keywords applied to existing tweets by the machine learning system.)

To deliver dialogue in real time to each performer during the performance, a software system was built enabling the director to select a set of intent tags and use them to recall a matching Tweet from the full database and deliver it to the appropriate performer via the head-mounted display.



Block diagram explaining the *Bodies for a Global Brain* system, challenging the two lead actors to react “in the loop” with a machine learning algorithm.

Unlike in *Grace Plains*, this entire *Bodies for a Global Brain* system was part of a metatheatrical bridge. In the story, the characters elect to receive instruction from the “global brain” of the Internet on what to do, thus embodying the technology’s real-world function more simply. The actual system does for the actors something very similar to what the system within the story does for the characters. (See <https://vimeo.com/119866019>, which includes an explanation of the circumstances.) It was the director—rather than a computer vision system—who observed the actors and their points of view and made decisions about what “type” of content to deliver. But the loop of control had a strong analogy connecting story and system, as did the presence of POV cameras. The characters choose to live by the word of a global brain not unlike the AI proposed by Google, Amazon, IBM, and others. This motivated the use of POVs and an envisioned but not implemented user interface for audiences to contribute to the control of the performers’ process. Even though the cameras were not directly in the loop of the system as realized—i.e., they were not influencing the AI and not diegetically described—and only existed in the story as unacknowledged parts of the head-worn displays, this work influenced bridges with even more automation and similarities between story and system created for the later pieces described below.

***Los Atlantis* (2015)**

A faculty-directed participatory stage piece created in collaboration with students, *Los Atlantis* followed a small band of travellers as they explored the futuristic archive of a historical city. The audience, invited to accompany the travellers, was offered agency to interact with various interfaces to explore and contribute to the archive’s contents, online at first and then in person. During each performance of the run, audience members, both in person and online, were guided through a series of vignettes integrating media from the archive, with their participation further expanding its holdings.



A scene from within one of the *Los Atlantis* vignettes, where the performers exist within media pulled from a cloud database.

We discuss this project in more detail in Burke et al. (2017). Here, we focus on specific facets that continued the exploration of how the cinematic, algorithmic, and performative intersect in the threads described above. In the previous two projects described, media content was one of the *outputs* of the performances—the live feed of *Grace Plains* and the edited episodes of *Bodies for a Global Brain*—but the story-related algorithms embodied in software code were concerned entirely with the text, and from that code emerged the mechanisms for audience and performer agency. In *Los Atlantis*, we pursued this fluidity and importance for media itself. Most of the software was concerned with the collection, categorization, and recall of cinematic media using similar techniques. “The Archive” was the story’s repository of the history of the city discovered by the travellers; in the real world, it was a database of media, hosted on YouTube, tagged with a folksonomy defined by the authors. It was a metatheatrical bridge and the focal point for contribution and agency: media was supplied to the repository and tagged by the ensemble, who gathered it on their cell phones throughout Los Angeles, and optionally by audience members who logged on to a website in advance of a performance. Media matching certain criteria within the folksonomy, and certain formal characteristics (length, etc.), was used in the performance.

The original intent was to have a large number of dynamically selected media cues, in which the media came from database queries with results changing depending on what was uploaded before, or potentially during, the performance. In practice, the desire for visual and temporal control drove only a small subset of cues to be truly dynamic, and with unpredictable, often shifting results, though all media for the show was retrieved in real time from YouTube. The limited dynamism may have arisen from the use of a fixed text that did not provide sufficient flexibility or dynamic prompting, as in the previously mentioned projects. Despite this, the Archive was a successful bridge between the

technological world of the production and the story, because of its library-like nature, which was a parallel intent. The holdings were intended to be as fluid and accessible to the characters within the story as they were to the storytellers' execution of it. Additionally, the piece expanded the earlier two projects' use of live POV streams as outputs, by allowing online audiences to pick between one of four characters to follow into different, sometimes parallel vignettes (see <https://vimeo.com/141879159>).

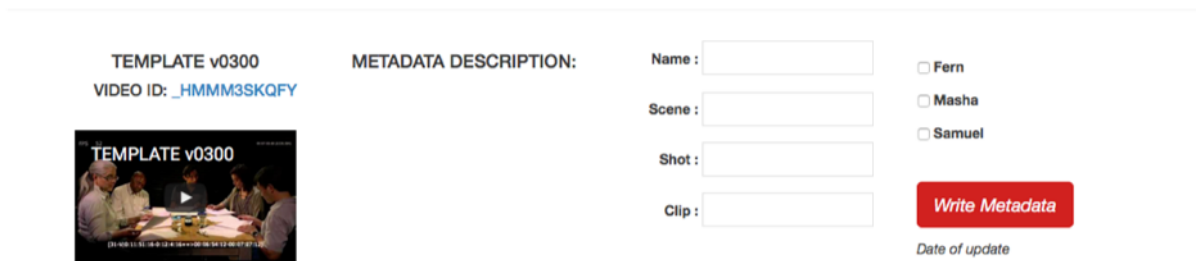
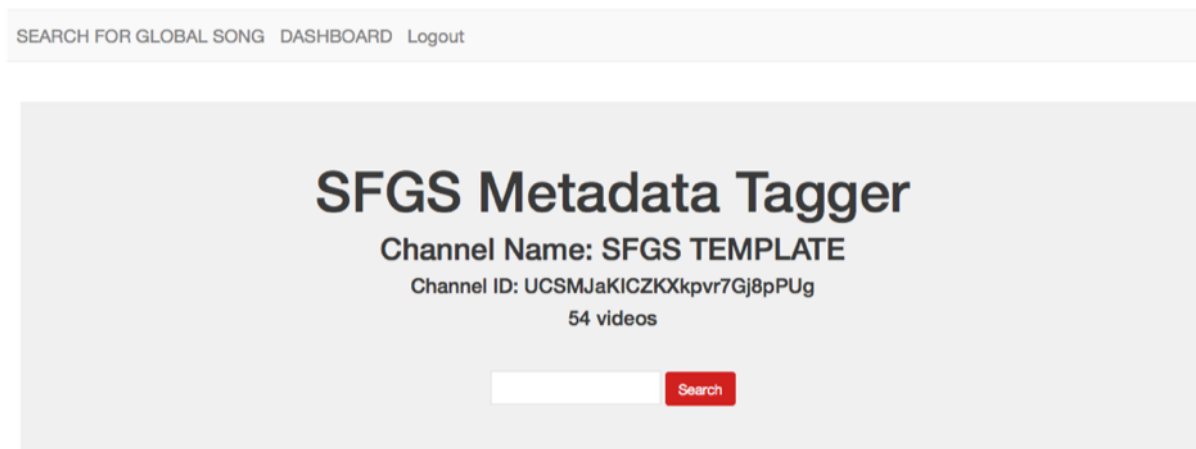
Our reactions to the realization of *Los Atlantis* led directly to *Search for Global Song*, described below. First, we wanted to focus on not just the selection but also the assembly of a sequence of media from a database—expressing control of that assembly within algorithms rather than manual editing, and creating the foundational footage with unknown, shifting assemblies in mind. Therefore, second, we wanted to make *all* media algorithmically selected while still being assembled into a cohesive narrative sequence, to enable a mix of control by text, code, and interpretation. Finally, we wanted to explore the relationship between text and interpretation more carefully because of the strong interrelationship of these elements and media-processing code. (In a piece that starts with a text, software that processes media typically affects the realization of a specific instance of that piece, not the text itself—though that is the subject of the final example.) The motivation for these choices was twofold: *Search for Global Song* was part of a general research effort to explore the joint authoring of code and text, and there was a specific objective to tell a familiar type of story—intersecting slice-of-life character studies—using a key facilitator of intersection today, albeit algorithmic media selection in Internet advertising. We aimed to present familiar input and output forms to the interpreters and the audience, respectively, by focusing on cinematic short-form drama, while creating an algorithmically driven system to select and remix interpretations within the limitations of those inputs and outputs.

***Search for Global Song* (post-production)**

Search for Global Song is an experimental short film (about twenty minutes long) that tells the story of three strangers who frequent the same café—a retiree, an unfulfilled middle-aged classical musician, and an exchange student electronic musician who has overstayed their visa—living in the same neighbourhood, and coincidentally connected by people in a single but unspecified, unseen place across the world. Their lives are sampled and interwoven by the electronic musician, resulting in a final musique concrète composition encountered in the café, for which the audience just happens at that moment to include the retiree and the classical musician. The script was written for interpretation by different directors to yield different sets of foundational footage to be algorithmically edited together in real time to produce different variations of the same film (see <https://vimeo.com/223828818>). The editing code also listens to the audio evolving from a given cut and inserts it back into the musical composition of the final performance within the film, providing a metatheatrical reflection on the film's structure via “found sound” from its corpus. Like the film's characters, interpretations collaborate with coordination to create a new piece of music each time it is cut. While not strictly a live performance by humans, the edit assembly is done live by software, resulting in a live music composition. As such, it has, for us, been an informative intersection of cinematic media creation that applies methodologies from and our theoretical interests in live performance.

We are finishing an initial implementation of the code that edits the film. It will operate on source media from four completed shoots—two productions of the entire script and two productions of short selected pieces (see <https://vimeo.com/223829046>). Shot with four different casts and crews

in different locations, *Search for Global Song* incorporates the most extensive notion of a database of media of any of the examples given here. All selected clips from all shots of all scenes are individually uploaded to YouTube, from each interpretation. They are associated with a given place and/or specific story requirement by a folksonomy of properties (tags with values, as described below). Expanding on *Los Atlantis*, the organization of the media takes into account formal requirements, such as shot length, by organizing this information via an edit decision list (EDL) that we created based on a human edit, as well as semantic tags that represent story milestones. These are embedded in the media through the use of timecode markers in the editing software. An assembly program uses these EDL “targets,” which express the writers’ intent, the directors’ interpretations, the story marker information and other metadata associated with each uploaded clip from each interpretation and attempts to assemble a cut following algorithms created by the authors and our collaborators. Algorithms progressively select and assemble shots in order to meet the story objectives, while also balancing other external parameters such as the amount of each interpretation to use on average across a given edit of the film. The algorithms’ parameters are varied by the film’s creators and production team to generate different edits and can change depending on events of the day, attributes assigned to the location of a viewer, viewer preferences, etc. (that is, the parameters of the algorithms can be varied manually or by other algorithms).




One metadata entry screen from the *Search for Global Song* system, used by contributing authors to assign folksonomy tags and other metadata linking their version’s media elements to the broader structure.

As the characters go about their daily lives, they also communicate with people in the same unspecified, unseen location across the world—via email, webchat, SMS, mobile phones, etc.—which motivates their actions and coincidentally leads to the events being recorded by the electronic musician, and why the electronic musician is recording them. Similarly, the audio system listens to the resulting edit and uses the composition of the audio mix to generate an algorithmic composition

of the final song in the piece, which is overlaid onto a character's performance. This composition bridges the assembly algorithms and the character's use of found sound throughout the story, in the final performance. The script and code together make the database a continuous (if hidden) metatheatrical authority within the experience and vice versa. Independently from how any final cut may purposefully or inadvertently—depending on the assembly—shed light on common threads between the interpretations, the characters within the story communicate across cultures, and this motivates sound collection and the electronic musician's composition.

The traditional, modern notion of montage, i.e., stemming from early Soviet directors such as Eisenstein, includes action being displayed while condensing space and time. It does, though, assume conventional Western notions of character, as if portrayed by the same actor within one story world, and the story world itself is absolute in terms of its locations. *Search for Global Song*, too, is defined by its script's absolutes—a progression of occupations and actions, types of locations and off-screen locations, and certain character descriptions, for example, age, and age in contrast to the other characters' ages. The script itself does not suggest or presume montage, nor even montage sequences in the contemporary, non-Soviet sense (for passages of time). Its plot is purposely arranged within real time, with jumps in time arranged in scenes over subsequent days. It does, however, explicitly allow for a montage of the directorial interpretations, especially as expressions of the locality of the places and the ensemble. An action portrayed by one actor playing a character could be concluded by another actor (in another interpretation) playing the same character; a character portrayed by one actor could be recording street noises in one location at the beginning and end of a scene while different actors portraying the same character could be recording noises from different streets in different locations in between. The piece condenses space and time between interpretations (the individual shoots) while following via the “unity of character” (the progressions of their actions), regardless of differing interpretations (how the same characters may appear differently, in terms of gender, ethnicity, fashion, etc.)—within one narrative frame that guides intercutting of media to allow the various interpretations to emerge for the audience.



Three directors' interpretations of one lead character in *Search for Global Song*.

Entropy Bound (current)

Digital technology is central to the audience *experience* in each piece above, yet it has a limited impact on the emotional trajectory of the characters, despite the presence of metatheatrical bridges, so was often perceived as not impacting the *story* as it unfolded. In *Los Atlantis*, the limit was likely the fixed text, preventing shifts in media from shifting the characters. In *Grace Plains* and *Bodies for a Global Brain*, very limited text structure hampered audiences' abilities to follow a story. This was not necessarily a problem for those works, but our desire in the works that have followed has been to balance agency and authored structure while making the technology-supplied fluidity of database media central to the characters' drama. This balance became a key challenge of interest.

For this reason, with the new work *Entropy Bound*, we are exploring how to integrate characters' objectives and reality with the technology more closely, so as to make that integration itself a metatheatrical bridge—that is, their actions are not only enabled by the system's placement within the story but also *required* for their very existence. There could be no progression of actions without it. The pursuit followed directly from our experience with *Search for Global Song*. In our opinion, as producers and early test audiences for the editing algorithms for that short film, we found that perhaps unsurprisingly, it was the characters that provided unity across very different interpretations. Unlike the plot, which of course, one doesn't know until one has seen the whole piece, parallel characters from parallel interpretations became immediately relevant and integrated quickly our ability to find coherence amidst the cross-cutting. The value of the unity of character, particularly across different and changing media, was what we chose to explore specifically within our most recent project.



Lead actor Rey Jarrell wearing a point-of-view camera in a workshop rehearsal of *Entropy Bound* that she felt explored “how much data you need to be sure who you are” (December 2018, UCLA).

Entropy Bound is an experimental comedy for the stage based on a metatheatrical bridge embedded within a character. It follows a highly determined, at times obsessive-compulsive urban professional who has suffered a traumatic brain injury and is now living life with a first-generation digital brain implant. Living with her parent and her best friend, who have moved in with her to help her recover, and participating in the best friend's web series about their cohabitation, she relies on the implant and its evolving base of media for all memory since her accident. Her last real memories end at the moment of impact, and, as enforced by the growing corpus of media from her POV of continuously trying to perfect that last moment, she assumes the evolution of the perfection is significant because it's the only new memory available to her, and she has, since her last real memory, been sharing it. Her best friend, since she came to, has been "broadcasting" her journey to the web series audience, that is, the audiences attending the play.

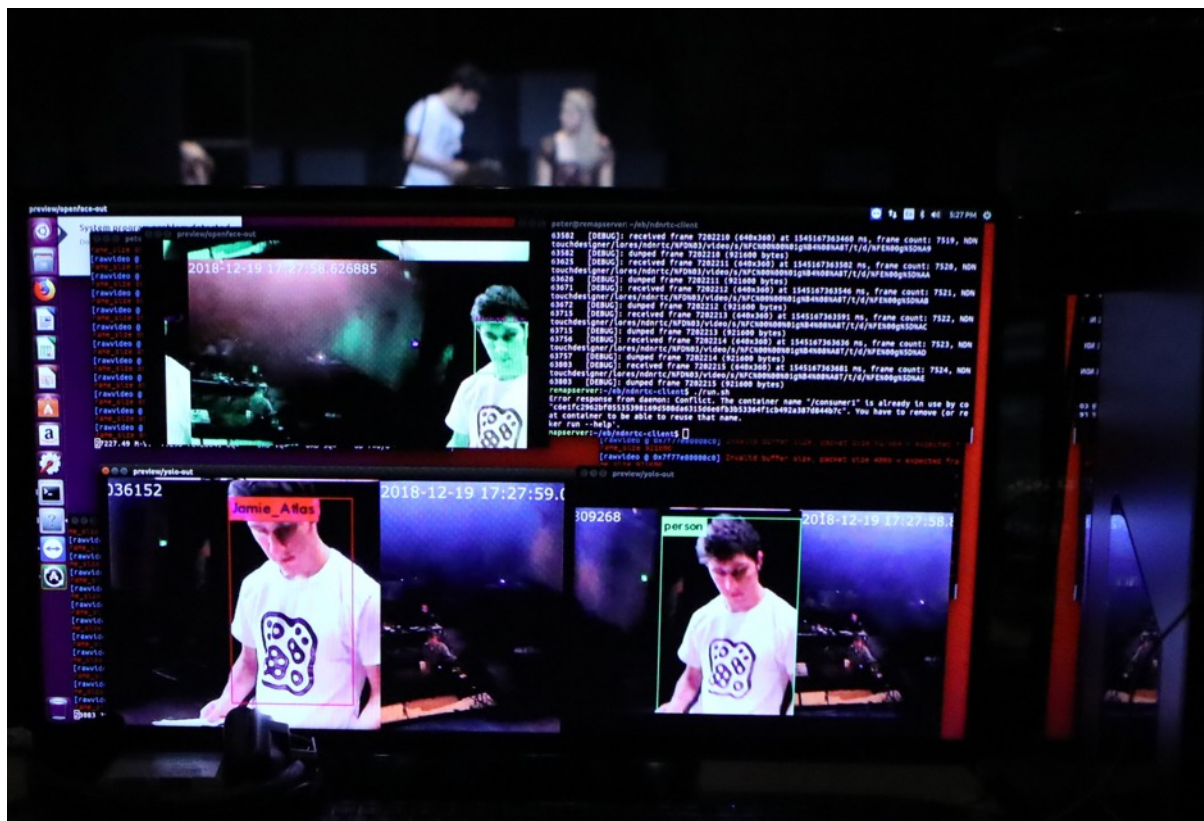
The composition of *Search for Global Song*'s presentation—a montage of interpretations—changes according to the choices of the algorithms applied; the narrative core will not. But algorithmic editing, and manipulation more broadly, are applied live in *Entropy Bound* to media content to enable real-time human reactions to shift the plot itself, further contributing to the corpus of media, and then again, the plot, and so on. For the authors, this circular yet progressive pattern mirrors increasingly prevalent cycles of media use in everyday life and numerous dramaturgical possibilities. To reflect this within a stage story and a theatre-making process, we anchored the central character's entire memory in a database of expanding media content, with the individual memories being recalled according to the whim of an algorithm. And similar to algorithms in social media systems, the algorithm in *Entropy Bound*'s system enables its user, the lead character, to react in real time to events within her POV, using tags—thus training the system to further enable her and placing the actor in the position of training the system on which her character must depend.

In an initial workshop of *Entropy Bound*, we learned from the actor playing the lead character that, for her, the distinction between her instrument as an actor and her character became blurred through the implant's role as a metatheatrical bridge and the use of rehearsal footage mixing the real world and the fiction. This was further reinforced by a need in the workshop for her movements to be tuned to frame images properly within the POV camera worn above her ear. She envisioned the same would happen to anything the implant could record and redisplay. Going forward, we use the term "actor/character" to convey this blurring. We also suggest a similar shift within the audience; they are already an audience both inside (the best friend's webcast) and outside the fiction, and are recorded by the implant.

Just as in *Los Atlantis*, the database in *Entropy Bound* has a defined role within the drama, as does the algorithm, and the audience and each performance impact the evolution of that role. It contains the actor/character's POV—an archive in her head, each frame and video segment annotated with its semantics and importance. This semantic folksonomy is applied in real time by machine learning⁶ analysis, as is the segmentation of the running POV video into discrete clips. Within the fiction, in addition to an implied analysis by the implant of her circumstances, the main character has been able to add tags within her field of view since coming to with the implant having been installed. This is actually what happens within the main actor's process through rehearsals and the run, as the first rehearsal is when she came to, when the implant starts recording while she is in character, and the actor/character can begin adding tags.

Objects—for example, different hams she prepares for the web series (and actual) audiences—are the primary focus of the character's POV, and these key objects are understood by the system. Each

group creating a new interpretation (production) must train the ML system to map objects from the real world to those required by the script, codifying their interpretation of key items and embedding it into the collection of algorithms that drive the piece—including the main actor/character’s reactions to (tagging of), for example, the hams and the other characters. Here, the tags of the folksonomy are envisioned as being defined by the authors of text and code, but with significant flexibility for each production to choose to what physical objects they are applied. These choices filter through the overall aesthetic approach to the production and yield media to which the character responds via improvisation. The actor portraying her not only is charged with being a dramatist in the moment, but also in that production’s dramaturgical evolution by continuous in-the-moment training, adding reactions/tags to what she sees, so the system can understand how she feels about what is around her, and what she needs to know in lieu of memory. In the script, [brackets] denote objects, for example, [ham], that the system must be trained to recognize, and *italics* represent opportunities for improvisation by the actor in response to the media associations retrieved by the system based on what is happening on stage. Further, the character directly interacts with the audiences, getting collective and individual reactions to her object, and thus within her POV, captures these for the database.⁷ This permits each audience’s interpretation to be incorporated by the character into the show. Their reactions to her [ham] are interpreted and tagged by her, further training the system.



Live machine learning analysis of the lead actor’s POV camera in *Entropy Bound*. System by Peter Gusev.

The folksonomy has three types of tags: 1) those corresponding to bracketed items as recognized by the ML system; 2) additional tags from the ML system that provide general semantic analysis of its observation (for example, labelling person-like elements in the POV camera as “person” with a confidence value); and 3) emojis⁸ that are used by the lead character (as improvised by the actor) as

diegetic labelling of her environment (as seen by the POV camera). The actor/character selects and applies these latter tags within the action of the play, while the first two are selected automatically and invisibly, although a particular production may choose to reveal them to the audience. The tags applied to any given moment are unique within each run of the performance. The available tags in the folksonomy are defined by the code and used to organize the media within the database (and the character's memory). The metatheatrical bridge (the implant/system) uses this database of semantically tagged media segments as the artificial memory within the story. Implicit within the action are overlapping hierarchies of the media segments.

The hybrid folksonomy is eventually explicitly revealed to the audience in a way that it is not in the other projects. It is the audience's reactions to the events of the performance, including the current state of the algorithm and its results, that motivates the character's participation in the use of the folksonomy. The main character's memory of everything since the production began is limited by what the database contains, and the database contains only actions obtained since the process of the production began. The audience provides the mechanisms for expansion of the database through the opportunity to do another performance, their presence in the POV media, and—still under discussion—the expansion of the folksonomy itself by providing new tags. With each performance, the audience's existence and participation within the piece motivate the next one—with each audience gradually aware of past audiences' inclusion and impact, then eventually, theirs. (See <https://vimeo.com/324643293/02d628f91b> for details about the recent UCLA workshop of the piece-in-progress.)

Dramaturgy of Robots

Based on the work discussed above and that of many other contemporary artists, we anticipate that instructions for humans to create performances will soon be regularly accompanied by instructions for computers to participate in the realization of those performances. Just as scripts allow for interpretation, so too will these algorithms that are embodied in code. Unlike a traditional text, code used in performance will continue to enforce authorial decisions in real time on an ongoing basis, including the management and manipulation of cinematic media. They may even maintain connections to their authors—in real time or at any customized intervals.

After a talk on robotics (Bell 2018), anthropologist and Intel Senior Fellow Genevieve Bell cites a colleague's phrase, "any algorithm surrounded by an object is a robot." This can be extended to the mechanisms and rules of dramaturgical segments—actions—and how they are executed. Whether tactics of characters portrayed by actors, gestures within traditional dance drama forms, movements by puppets or objects, or progressions of frames, drama is generally crafted by maps of these actions. There are already some robotic roots advancing narratives—"puppeted" progressions—regardless of the psychological, emotional, or spiritual reality being created, or the cultural and human connections. The maps themselves are what motivates and elicits these intangibles. For thousands of years, there have been endless rearrangements of actions, within and across cultures, centred on events. The complexity and options of continued rearrangements, and their ability to incorporate cinematic media fluidly, now grow by the operational possibilities of code, which can not only be applied to the map for the performers, but also to the nature of the live events, the rituals justifying them—show by show for the performers, the system and the use of media—each with a new audience.

Digital media, including the cinematic image, are the inputs and outputs of this emerging robot dramaturgy, and its memory. Its expressive properties are malleable based on algorithms and how these function within the creative process. One could envision algorithms in counterpoint or opposition to expressive properties authored within text being layered in for a given interpretation, analogous to other types of interpretation. In a way, to allow for this more sophisticated, and perhaps interesting, direction, we are seeking to establish the role of algorithms embedded within the text, as opposed to merely accompanying it—and to find a way to invite and even celebrate interpretation by others of digitally driven, media-rich forms of performances that are often conceived/written and also realized by the same group of people.

Previously, the manipulation of media has often been left to directorial and design interpretation rather than being a core aspect of traditional playwriting. There is an arc in the works described here toward both input and output media processing being described by code during the writing process—code that will ultimately manipulate actual media, as in real-time filmmaking. A significant challenge—with some analogy to writing plays without code—is the practicality of creating code that expresses media relationships with sufficient generality for it to apply to any interpretation, yet with enough specificity for there to be “a piece.” Conversely, there must be code that implements those generalities for a given production’s interpretation and media specifics. (A simple example is how a character’s “social media channel” in a script and its corresponding code might be manifested in a projection within one production and through an application on the audiences’ cell phones in another. The same general manipulations of media apply to both, but what controls pixels on screens is quite likely different.) A full treatment of this challenge and how it might relate to typical techniques in large-scale software systems is outside the scope of this paper, but we address some new tools that could be applied to make it more feasible in the conclusion below.

In filmmaking, it seems that yesterday’s post-production process is today’s real-time, on-set capability. In developing use for algorithmically manipulated media, we have borrowed and been inspired by real-time techniques, taking us further and further away from historically cinematic processes and structures, though not necessarily from the visual formalisms that result from them. As conceivers of new works of text and code that allow for interpretation while we actively (via code) continue to participate in those interpretations, our challenges are structural: what to show when, what to capture, to whom to yield control or offer agency, and within what ruleset. In addition to Manovich, we might look at Eco’s *The Open Work* (1989) and its notion of an artistic work completed in collaboration with, for example, its audience, as it intersects with the procedural art of Lewitt, in which a detailed set of instructions can be used anywhere to create the same piece of art. In our case, instructions executable by a computer are provided to implement the authors’ objectives with respect to media. Lewitt said, “the idea becomes a machine that makes the art” (Cox, 2007). Such allusions sample just a few precedents; in fact, they omit decades of critical scholarship and foundational work from many digital artists, such as Lynn Hershman Leeson’s use of POV and robotics in many of her works. There is a rich history to draw from as we consider how to write code that accompanies text, embodying authorial constraints while trying to offer mechanisms for agency. For us, the post-cinematic is the algorithmic, which leads to the question: What kind of algorithms do we have today that progress the possibilities for both digital media and performance?

Conclusion—A New Thread: Narrative (and Media) Extension by Machine Intelligence

The quickening of the interconnectivity of the world and the pervasiveness of media and their content is culturally unavoidable. The significance and utility of customs built into dramaturgical structures are perpetually shifting as that connectivity and our human place in it evolves. Performance, from its earliest roots, started as an entirely localized ritual—in form, context, and content. Generally, dramaturgical structures contain justifications for that ritual of sharing space, as well as shared vocal, physical, and visual languages and taxonomies (from shared culture and as created specifically for particular performances). Generally, individual components of the structures—actions—progress toward further justification and further contextualization of their rituals and languages. One specific gesture—within Kathakali, for example—or grand realization—for example, Oedipus—can have concrete relationships with the justifications for the performances themselves.

The potential difference between the drama of the past and the drama of the present is the difference between literature and the sort of literary robot described above, which can be distributed to enforce its abstractions and rules (algorithms) at any time in performance, including real-time execution of a given production. This opportunity shifts the demands on imagery from the cinematic and televisual worlds that have been a part of theatre since at least the 1950s to the real-time filmmaking (e.g., Watercutter 2020) of today: media can be accessed on demand from databases; tags can provide easy descriptive power over that media; algorithms can influence, curate and embody various types of agency and generative behaviour (and, then, may guide their further embodiment by performers or audiences); performers can closely influence not only media but the semantic structures organizing that media, or be tightly coupled to both, embodying the outputs of the algorithms; authors can continue to influence interpretations via shared code that becomes a platform for future stagings of a work. The social media-powered evolution of remix culture (Lessig 2008) motivates choices within dramas that can ask audiences to contribute, by adding, tagging, commenting on, or changing what they are given—all on top of new types of materiality and very different physical affordances—but we are also interested in the machine within a given work that enables the remix itself.

To conclude with a direction in which to head, we suggest ML and AI more generally offer ways to make increasingly sophisticated algorithmic mappings between inputs and outputs that can embody the layered interpretation that exists in any live performance. Specifically, we are interested in mechanisms that more easily map between the computational abstractions in code accompanying a text versus that used to implement media manipulations in a specific interpretation, as discussed above. More fundamentally, new types of cause and effect relationships between actions and media, with the technologies enabling them as metatheatrical bridges, and new opportunities to define cause and effect through algorithms with broad-reaching inputs and outputs, will result in new levers of control, interpretation, and agency. In *Entropy Bound*, we provide a system that analyzes an actor's (and character's) POV based on semantics we as authors have decided on in advance and then use to model and execute how that media is to be incorporated in a given production's interpretation. Crafting this new type of system is exciting, yet comes with it for authors new unknowns—about how a given work will change, incrementally and overall. It yields aspects of authorship to the lead actor, within a context that requires they make decisions in the moment, similar to how one psychologically, emotionally and even spiritually interacts with machines on a daily basis (e.g., what

to react to on social media, what to ignore, what to share and with whom, how to categorize connections, who to mute or unfollow) and what is done by the machines behind the scenes that orchestrate these rituals.

Compellingly for authors, ML offers an opportunity to train systems by example—of how to map complex inputs to complex outputs, to drive media selection and assembly. As we hope to have shown, there are many mappings central to the common threads we described: for example, from the world of the authors to the world of the directors, and from the story to the meta-theatrically relevant real world. AI, more generally, is an approach that models the world (or *a* world), uses inputs to refine that model, and generates decisions and outputs based on the model. Authors are already model makers, and for example, folksonomies represent a distilled model of a story world; the algorithms that manipulate media based on them describe relationships between elements in that model. These ideas, familiar in gaming, hypertext, and interactive media art, become more relevant to cinematic media in live performance as we gain capabilities to relate media to story abstractions in real time via ML and AI.

The cinematic image is a certain kind of memory for this robot dramaturgy—algorithms realized onstage with a traditional text—that borrows from the audience’s visual familiarity with film. Each piece becomes a sort of purpose-built robot with digital and human parts, fluidly digesting and regurgitating the previously inviolable chain of twenty-something cinematic rectangles per second alongside each other element of the theatrical event.

Acknowledgements

All of the projects discussed were collaborative works with many participants; full credits can typically be found on the REMAP website, <http://remap.ucla.edu>. In particular, we thank: Eben Portnoy, creator of *Bodies for a Global Brain*; Cole Baker, Pierre Finn, and Phillip Powers, creators of *Grace Plains*; J. Ed Araiza, director of *Los Atlantis*; and Rey Jarrell, lead actor in the UCLA *Entropy Bound* workshop; as well as the entire cast and crew of each project. Each of these projects was supported in some way by Google, through the Glass Creative Collective, a Faculty Research Award, and Focused Research Award. *Entropy Bound* has also been supported by the 2019 Hearst Senate Faculty Award for Playwriting at the UCLA School of Theater, Film and Television.

Notes

1. By “cinematic media,” we mean visual media content that, in some substantive way, incorporates optical capture of the physical world to distinguish it (imperfectly) from purely computer-generated imagery. We sometimes use this two-word term interchangeably with “video” in this essay. Also, we mix plural and singular uses of “media” and “data” as needed to keep the language familiar within a given context.
2. Except where noted, in this paper, by “live performance,” we mean a live event seen by some number of physically present audience members and perhaps also through Internet streaming or other means.
3. Our work has considered other types of media as well, such as sound and 3D graphics; here, we focus on cinematic media.
4. It’s also relevant that media in theatrical performances is often already modular—broken into cues that are distinctly triggered—even if their effect on stage is that of a smooth transition. This already relates more clearly to the cinematic shot or edited sequence than to the unbroken experience of an entire film.

Algorithmic control of modular media is not that big a step; the bigger question is who controls the algorithm and how predictable it is.

5. For more on this piece and the one following, see Burke and Stein (2017).

6. A specific sub-area of AI in which algorithms learn to improve their performance over time.

7. Is the audience a metatheatrical bridge? In some sense, yes, as they are also the best friend's audience for his streaming show on the main character's life. That said, for this paper, we have focused only on technologies that have an author-controllable function as being such bridges.

8. Our goal with this is to use a language that is visually expressive and can be used in a variety of language contexts.

References

- Bell, Genevieve. 2018. "Making Life: A Brief History of Human-Robot Interaction." *Consumption Markets & Culture* 21 (1): 22–41.
- Burke, Jeff, Peter Gusev, Zoe Sandoval, Jared J. Stein, and Zhehao Wang. 2017. "The Storytelling Systems of Los Atlantis." In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (1046–54). New York: ACM.
- Burke, Jeff, and Jared J. Stein. 2017. "Location-Based and Audience-Aware Storytelling: Grace Plains and Bodies for a Global Brain." <https://howlround.com/location-based-and-audience-aware-storytelling>.
- Cox, Geoff. 2007. "Generator: The Value of Software Art." In *Issues in Curating Contemporary Art and Performance*, edited by Judith Rugg and Michèle Sedgwick, 147–62. Bristol: Intellect.
- Eco, Umberto, and Anna Cancogni. 1989. *The Open Work*. Cambridge, MA: Harvard University Press.
- Lessig, Larry. 2008. *Remix: Making Art and Commerce Thrive in the Hybrid Economy*. New York: Penguin.
- Manovich, Lev, Roger F. Malina, and Sean Cubitt. 2001. *The Language of New Media*. Cambridge, MA: MIT Press.
- Pizzo, Antonio, Vincenzo Lombardo, and Rossana Damiano. 2019. "Algorithms and Interoperability between Drama and Artificial Intelligence." *TDR/The Drama Review* 63 (4): 14–32.
- Watercutter, Angela. 2020. "ILM Used *Fortnite* Tech to Make *The Mandalorian*'s Virtual Sets." <https://www.wired.com/story/fortnite-mandalorian-filmmaking-tech/>.