Mapping Narrative Space in Hollywood Film

James E. Cutting, Catalina Iricinschi, and Kaitlin L. Brunick

Abstract: This article presents a new method to create maps that chart changes across a cinematic narrative. These are unlike narrative spaces previously discussed in the literature—they are abstract, holistic, dynamic representations based on objective criteria. The analysis considers three films (*All About Eve, Inception,* and *MASH*) by counting the co-occurrences of main characters within scenes, and *12 Angry Men* by counting their co-occurrences within shots. The technique used combines the statistical methods of correlation, multidimensional scaling, and Procrustes analysis. It then plots the trajectories of characters across these spaces in *All About Eve* and *Inception,* regions for characters in *Inception* and *MASH,* and compares the physical arrangement of jurors with their dramatic roles in *12 Angry Men*. These maps depict the changing structures in the visual narrative. Finally, through consideration of statistical learning, the article explores the plausibility that these maps mimic relations in the minds of film viewers.

Keywords: characters, film, maps, multidimensional scaling, narrative space, scenes, shots, statistical learning

> What is narrative space? In both writing and in film, it refers to any space in which a narrative can take place. This space does not necessarily have to be always physical: it can be emotional, psychological. Many films take use of these intangible narrative spaces to create moods and tones.
>
> —Miss Yujin (9 October 2009)

This article offers a new way to map narrative structure and to show how it changes across the shots and scenes of a film. It does so by looking at four quite different Hollywood films—*All About Eve* (1950), a studio-era drama with a well-defined sequence of scenes that can be placed neatly within four roughly equal-length sections; *Inception* (2010), a much more complexly structured action film whose scenes we divide across six acts; *MASH* (1970), an episodic comedy without traditional film structure but whose scenes can be allocated to a setup and five major episodes; and *12 Angry Men* (1958), a drama with little traditional scene structure and with almost all of the story taking place over continuous time within one room with a constant set of characters.

Volume 7, Issue 2, Winter 2013: 64–91 © Berghahn Journals
ISSN 1934-9688 (Print), ISSN 1934-9696 (Online)

In the context of film studies our method of mapping narratives is unusual and only tangentially related to the broader concept of narrative spaces found in literature and film.[1] There are at least six reasons for this contrast. First, we really do mean maps —two-dimensional graphic representations of relations among main characters across the course of a film. Second, most of our maps have nothing to do with the physical spaces that are portrayed in films. Third, we ignore entirely the verbal content of the film—screenplay, dialog, plotline.[2] Fourth, our maps are objective; they are generated by computer algorithms from inputs based on the visual presentation of the characters as they co-occur within scenes and shots and accumulate over portions of film. Fifth, our maps are holistic. They consider relations among all major characters simultaneously, not just a select few. Finally, we create them multiple times across the length of the film. Comparisons across such maps reveal some of a film's dynamic narrative structure. We believe that these maps show the close relationship between how filmmakers compose shots and how the verbal content of the narrative unfolds. Narrative dynamics of almost any popular film can be portrayed in a sequence of two-dimensional maps.

Narrative dynamics of almost any popular film can be portrayed in a sequence of two-dimensional maps.

We create maps that are representations of the narrative structure of a film, not its structure in the mind of the viewer. Nonetheless, as the result of statistical learning, we claim that filmgoers encode the co-occurrences of characters and build dynamic representations of character relations that are consistent with these maps. We claim further that this knowledge, as it is acquired, helps them understand the story of the film.

Making Maps

The essence of any map is that it presents distances between things in a continuous space. Consider a concrete example. The top left panel of Figure 1 shows a schematic map of six major cities in Western Europe—Berlin, Geneva, London, Madrid, Paris, and Rome. The air distance in kilometers among these cities is noted. If these fifteen numbers are entered into a multidimensional scaling program, one can reconstruct the array of cities, shown in the top right panel. This is a two-dimensional display of the distances among the cities.

There are several peculiarities about the production of this scaling solution, or simply map. First, the "proper" orientation of the map is not known to the scaling program or represented in its output. In particular, the algorithm does not know about north, south, east, or west. Thus, with such a scaling solution, one is free to rotate the map however one wishes, although clearly the map shown at the upper right is oriented in the generally accepted way. The second peculiarity of this procedure is that it could be mirror reversed with Berlin to the left of London, and Rome to the left of Madrid. In other words, from the input distances the algorithm does not know that one usually looks

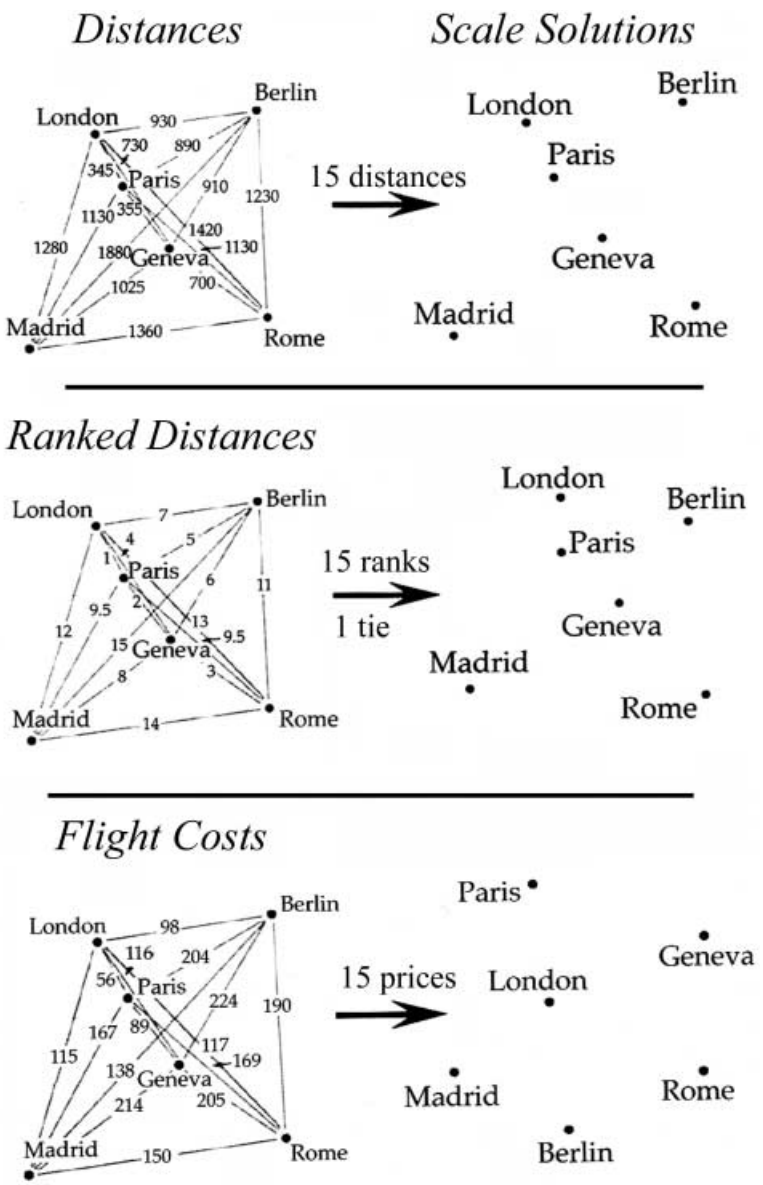

Figure 1. The creation of three maps through multidimensional scaling. The top panels use the distances in kilometers among Western European cities; the middle panels use the ranked order of distances among them; and the bottom panels use round trip flight costs from 2002. Adapted from Cutting (2006).

down at such a map, rather than up through it. Thus, one also has the freedom to flip the map over as if it were transparent. Despite these constraints, the program allows one to recover the map from the distances. In this case the map is essentially exact.

Surely, this does not seem very impressive: one inputs distances and gets back the map from which one obtained the distances, and one does not even

know its orientation or which side is face up. Distances in this first case were given in kilometers, which are measured on what is called a ratio, or metric, scale. The aggregation of their collective relations is called metric multidimensional scaling. But there are other ways to scale the data without this result, and these are much more interesting.

Consider: rather than mark intercity distances in kilometers, one might simply rank them from closest to farthest. This is done in the middle left panel of Figure 1. Here, London and Paris are ranked the closest, Paris and Geneva second, Geneva and Rome third, and so forth, with Madrid and Berlin the farthest apart. Notice that the distances between Paris and Madrid and Paris and Rome are very nearly the same, so these are given ranks of 9.5 (the average of rank 9 and rank 10). Using nonmetric multidimensional scaling, which uses only such ordinal ranks as input, one can also create a new map. This one is shown in the middle right panel of Figure 1.[3]

In this process of comparing all possible ranks, the algorithm tries to create a map with the distance between London and Paris as the shortest link, that between Berlin and Madrid the longest, and the rank order of all the distances between cities the same as given in the middle left panel. The scaling result is impressive. Notice that this map, despite its nonmetric origin, is almost the same as the metric one in top right of Figure 1. Oriented and flipped appropriately, it is very difficult to tell the difference. Scrutiny reveals that Berlin is perhaps a little bit too close to Rome, having slid somewhat south.

It is fun to make such maps, and one quickly realizes that aerial distances are not the only thing that might be considered. The lower panels of Figure 1 show the creation of a different kind of map. Shown on the left are costs in US dollars to fly roundtrip between the same fifteen pairs of cities according to Travelocity.com™. These data are old but instructive,[4] and all seasoned travelers will be familiar with these types of results. Notice that it was cheaper to fly from Paris to London to Berlin and back than to fly roundtrip between Paris and Berlin directly. The same is true from Geneva to London to Madrid, compared to roundtrip flights between Geneva and Madrid. Moreover, it is almost the same price to fly through London to most other cities from most other cities. This is a peculiarity of air hubs, routing, and heavily trafficked links between certain cities. And notice also that usual hubs of Frankfurt and Amsterdam were left out of this modeling.

Contemporary travelers understand these situations well, but rarely convert their intuitions into a map. Nonetheless, one can easily do so through nonmetric multidimensional scaling. If we start with the costs of travel as the "distances" between cities rather than the kilometers that separate them, we can create a new map of Western Europe. The scaling results are shown in the lower right panel of Figure 1.[5] This map shows why, in terms of monetary cost—although not total trip time or connection aggravations—London is on

the way from Paris to Berlin, and why London is between Geneva and Madrid. In other words, if airfare prices were the only consideration for the traveler, he or she could use such a map of Europe to plot business or vacation trips.

Deriving Character Distances and Making Maps of Film Narratives

As suggested in the cases above one can use many different sources of information as "distances"—kilometers, ordinally ranked distances, airfares, and so forth. One can also use correlation values as distance measures. That is, because correlations vary from 1.0 through 0.0 to –1.0, one can assert that items with correlations of 1.0 are very close to each other (indeed in the same location), those with correlations around 0.0 are at middling distance; and those near –1.0 very far apart. In this manner, one can use correlations as numbers on an ordinal scale, relating objects (in this case characters in a film) to one another, and then scale them nonmetrically to make a map.

Correlations come in several varieties and with different tweaks. The one we use is called Cohen's kappa, or κ. The distance measure we use is $1 - \kappa$, yielding a set of possible distances between 0.0 and 2.0.[6] Cohen's kappa is usually used to measure the agreement between two people making judgments in psychological experiments. Here we generalize it to the appearance of any two characters as they appear across scenes of a film.

Consider four cases, the first three from *Santa Fe Trail* (1940). First, Windy (Guinn Williams) and Tex (Alan Hale) are characters of comic relief. Like Tweedledum and Tweedledee they always appear together in a scene. More concretely, they appear together in twenty-six scenes, Windy never appears without Tex, or Tex without Windy, and neither appears in 109 other scenes.[7] The "distance" between Tex and Windy, then, is zero ($\kappa = 1.0$; distance = $1.0 - \kappa = 0.0$). Second, Jeb Stuart (Errol Flynn) and George Custer (Ronald Reagan) are main characters in the film and present a different relation. The film begins as a "bromance" between the two at West Point and then in Kansas in the 1850s, but shifts as Stuart begins to win over Kit Carson Holliday (Olivia de Havilland). Stuart and Custer appear together in sixty-nine scenes, Stuart appears without Custer in twenty-five, Custer without Stuart in only one, and neither appears in forty scenes. Thus, $\kappa = .64$ and thus the distance measure is .36. Third, Kit's father Bob Holliday (William Lundigan) and John Brown (Raymond Massey) appear together in only one scene, the last of the film where Brown is to be hanged. Holliday appears in twelve scenes without Brown, Brown in thirty-six without Holliday, and neither appear in eighty-seven scenes, yielding a distance of 1.15.

Finally, for a distance of 2.0 one must imagine a film in which two characters appear across all scenes, but they never appear together. There are likely a number of experimental films that have this composition, but the closest popular movie may be *Sliding Doors* (1998). It depicts Helen (Gwyneth Paltrow)

living in two separate realities. In one she catches a train and discovers her lover in bed with an ex-girlfriend. In the other she misses that train, is mugged, and comes home to her lover alone in the shower. The film goes back and forth between the two realities as they progress, but clearly the two Helens (one now with short hair having jettisoned her partner, and the other retaining long hair and having kept him) cannot appear in the same scene.[8] If one or the other appears in all scenes their "distance" would be 2.0 [$\kappa = -1.0$; distance = 1.0 – (–1.0) = 2.0].

Given *n* characters in a given film, there are $n^*(n-1)/2$ distances to consider. Thus, if there are four characters, there are six distances among them to scale; 8 characters, 28 distances; and 15 characters, 105 distances. Critically for this effort the more characters the more difficult it will be to fit them easily into a two-dimensional space.

Film Units, Small and Large

We had a number of viewers parse twenty-four films, three per film. Their instructions were to track "events" in the film (without having been given an overt definition of an event) and to mark frame numbers that began new events. Handily, these events generally obey continuities of space and time, particularly space. That is, viewers generally coded a new event as occurring when the film shifted to a new locale. Thus, these events map reasonably well onto scenes and subscenes in film structure (Cutting et al. 2012).[9] For three of the films we investigate here—*All About Eve*, *Inception*, and *MASH*—these parsed events are our "small" units. We started with the original scene-parsing data, going through each film again and generally tracking where at least two viewers agreed on an event boundary. We then recorded each character that was shown within each scene, entering a value of one for each into a large matrix. All characters not shown in that scene were given a value of zero. Because our fourth film, *12 Angry Men,* does not have a traditional scene structure, we shifted our "small" unit to the shot, analyzing the co-occurrence of characters in all shots of the film, with 1s and 0s in a large matrix for all shots. The constructed matrix was then the input for a Matlab computer program that computed the kappa-based distance measures among all possible characters across all scenes or shots in the film. It then generated a nonmetric scaling solution in two dimensions. We call this the *base map* and used it as a backdrop against which to compare other maps from that film.

Next new maps were computed independently for each of the "large" units of each film. We use here the term *act* as a technical term for *All About Eve* and *Inception,* which in our analyses contained four and six of these large-scale units, respectively. For *All About Eve* our acts seem to correspond to the structure promoted by Thompson (1999), who divided most films into roughly four equal-length acts, although she allowed longer films to have more. We

make no claim that she or others would parse these films in the same way as we have. What is important here is the serial nature of these units across the film, each one following directly on the previous, and that they are of roughly equal length. For a third film, *MASH,* we abandoned the notion of an act and substituted the term *episode*. Act structure implies coherence across the arc of a film, and this simply does not apply to a film like *MASH*. And finally, in *12 Angry Men* we divided the film according to juror votes. For lack of a better term, we call these larger units *sections*.

Our analytic method for the larger units is this: we first scaled the characters in each film portion, creating a new map, and then used a statistical technique called Procrustes analysis. This procedure rotates, shifts, expands or contracts, and possibly flips the newly scaled solution to best fit the base map. We followed this procedure for each larger portion of a given film—act, episode, or voting segment—matching its solution back to the base map so we could cross-register the separate maps and gain insight into character relations as they might change across the film.

All About Eve (1950): Tracing Paths in a Narrative Space

We chose *All About Eve* as the first film for our analysis. This drama presents the story of a small cross-section of mid-twentieth-century New York City theater society, taking place mostly within a flashback. The core of the story pits an established actress, Margo Channing (Bette Davis), against newcomer Eve Harrington (Anne Baxter). It also involves Margo's partner, Bill Simpson (Gary Merrill), a playwright and screen writer; another couple, Margo's director Lloyd Richards (Hugh Marlowe) and his wife Karen (Celeste Holm); a theater critic, Addison DeWitt (George Sanders); Margo's housekeeper, Birdie (Thelma Ritter); and the producer, Max Fabian (Gregory Ratoff), of Margo's play. We omit discussion of Birdie and Max, although their appearances were used in the construction of the base map and subsequent maps.

An Analysis of Acts

Following Thompson, the first act is a setup (here 31 minutes long). It introduces the characters at a banquet in honor of Eve, and then flashes back to a beginning with Eve lurking in an alley and having done so after every performance of a play in which Margo stars. The plot then progresses through to the point where Karen introduces Eve to the theater crowd, Eve charms them with her (fictitious) personal story, moves in with Margo, becomes her personal assistant, and arranges—against Margo's knowledge and remembrance—for her to have a midnight phone call with Bill for his birthday.

The second 31-minute act in our segmentation, the complication in Thompson's terms, has Margo growing increasingly fed up with Eve, while Eve insinuates herself among the others and into all matters of the theater pro-

duction in which Bill, Lloyd, Max, and Margo are involved. In the third act, the development (36 minutes), Margo learns that Eve is now her understudy, Karen arranges that Margo miss a performance, Eve performs well, and Addison writes a glowing review. In the climax act (33 minutes, including a 5-minute epilogue), the two couples meet at a restaurant, with Eve and Addison at a nearby table and, after an opening in New Haven and a behind-the-scenes tussle between Addison and Eve, the narration eventually flashes forward to the banquet that opened the film, celebrating the success of Eve. Then, in an epilogue an aspiring young actress, Phoebe (Barbara Bates), has talked her way into Eve's empty hotel room and, after Eve arrives, sets herself up to become a new Eve.

Maps

The four panels of Figure 2 lay out the co-occurrence relations of the players across the four acts of the film, with the different colors distinguishing them.[10] Again, each map is oriented and fitted to the base map (not shown) and the directions in these spaces are arbitrary: distances up the map are no different than distances to the right. These are abstract spaces that encode only the scaled distances among characters. Critical, however, are the distances among points, which represent eight of the characters.

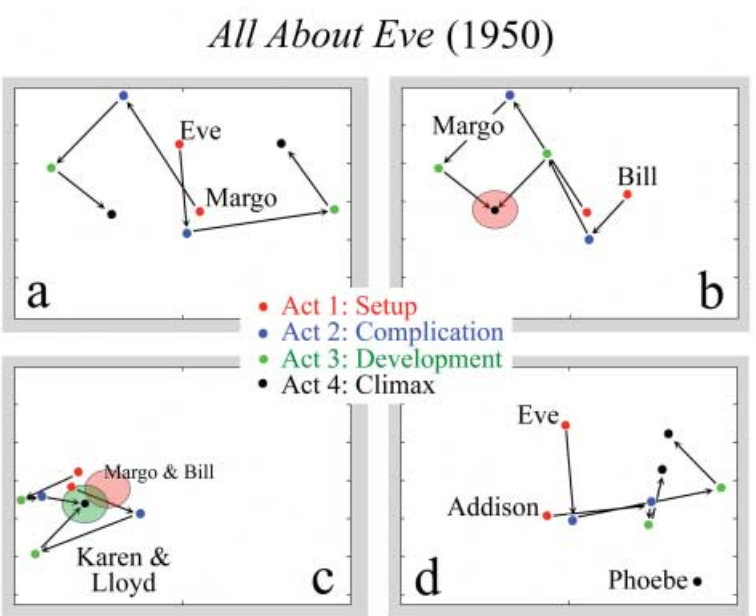

Figure 2. Four maps showing the trajectories of characters in a narrative space across the four acts of All About Eve. *Different colored dots represent the characters in the different acts. The layouts are based on the co-occurrences of characters in scenes (measured by Cohen's* κ*), with those values multidimensionally scaled and brought into register by Procrustes analyses.*

The relations of Eve and Margo, to each other and to all of the rest of the company, are the focus of* All About Eve, *and indeed they move most through the computed narrative space.

Figure 2a represents two paths across the four acts of the film. The relations of Eve and Margo, to each other and to the rest of the company, are the focus of *All About Eve,* and indeed they move most through the computed narrative space. Normalized to the distances moved across acts by the other characters (that is, giving them an average value of 1.0, the movements of Eve (1.72) and of Margo (1.87) are nearly twice as large. This suggests that their characters have changed most across the course of the film. Notice that initially Margo is in the center of the space, and indeed the flashback in the film starts centered on Margo. In terms of the data generating this position, Margo appears about equally often in scenes with each of the other characters. Alternatively, Eve generally co-occurs with only a subset of the characters, is computationally an outsider, and indeed a social outsider as well.

During the second act, however, Eve insinuates herself into the center of the space and Margo is ejected. Margo is in crisis as she begins to realize, across this act and the development to follow, that she is no longer suited to play the roles of younger women. Most important, through the filmmakers' choice of depicting the various characters in the scenes, great distance is placed between Eve and Margo. This is the case both in the computation of co-occurrences and in the narrative itself. Finally, some of the tension—at least for Margo—is relaxed in the last act. This is, in large part, due to her improved relations with Bill.

The Margo-Bill spatial relations are shown in Figure 2b. Initially, Margo and Bill are quite close in space, and indeed quite close in their relationship as depicted in the film. As Eve drives Margo into crisis, the distance between Bill and Margo dramatically increases during the complication and development. Indeed, Bill goes to Hollywood for a while and returns to a party in which sparks fly in all directions around Margo. In a memorable quote Margo warns: "Fasten your seatbelts, it's going to be a bumpy night." In the last portion they come close together. They occupy the same spot in the map. Bill is going to marry Margo, and Margo has accepted her inappropriateness for ingénue parts.

Figure 2c shows the relative constancy of two more minor characters, Lloyd and Karen. They too end in the same spot, and in close proximity to Bill and Margo, but they have never moved very far. This reflects not only their co-occurrences in the scenes but also in their relationship within the narrative. Figure 2d shows the co-occurrence mappings of Addison and Eve. The narrative mirrors these relations: Addison is initially at some distance from Eve, tracks and helps her progress in the theater, both increasingly separated from Margo and the other characters.[11] Finally, and in the same panel, Phoebe arrives in the

epilogue, a young outsider hard-bitten on the ladder of success, much like Eve at the beginning of the story.

We claim that the spatial layouts of each of the four panels make sense in terms of the narrative in *All About Eve*. We hasten to add that these panels were constructed without consideration of any verbal content within the film. Instead, they were compiled based only on visual appearance in the film's scenes. Also, to be clear, we make no claim that this is the only way to create a narrative space, but we do think that many similar schemes will converge on these same results. For example, if one weights the character co-occurrences by the duration of each scene, the results are almost identical to those in Figure 2.[12] Most important, we claim that this set of results makes narrative sense.

Inception (2010): Character Regions and Paths across a Narrative Space

Tracking the co-occurrences of characters in *All About Eve* yields an interpretable set of maps across the four portions of the film. One could argue, however, that dramatically this is a fairly straightforward older film and that perhaps those attributes contribute to the interpretability of the maps we have generated. Thus, we next look at *Inception,* a much more complex film with a much more complex structure. We divided *Inception* into six larger units, which we also call acts.

An Analysis of Acts

The first act is a prelude (21 minutes). The action that takes place mostly in two dream environments—a very large, Asian wooden building at one dream level, and a hotel room with a revolution going on outside at another. It then erupts into a nondream environment with the main characters on a fast train. It continues until Saito (Ken Watanabe), an industrial magnate in his private helicopter, drops off Cobb (Leonardo DiCaprio), the main dream controller, and Arthur (Joseph Gordon-Levitt), his major assistant. It is then that Saito offers Cobb a job in exchange for his heretofore-impossible repatriation to the United States.

The second act is a first setup (23 minutes) in which Cobb assembles his team. The team includes Arthur, Ariadne (Ellen Page) as the dream architect in Paris, Eames (Tom Hardy) as the impersonator, and Yusuf (Dileep Rao) as the chemist (or druggist) both in Mombasa, and with Saito tagging along, "protecting his investment." The third act is a second setup (20 minutes) during which the team, back in Paris, plans the inception —the insertion of an idea into the controlled dream of Robert Fischer (Cillian Murphy), the son of Saito's chief global competitor. This portion ends as the inception team and Fischer drop off in a drugged sleep in first-class compartment of a transcontinental flight from Asia to Los Angeles.

The fourth act is the complication (20 minutes), where the team and Fischer arrive in the first dream level in an urban setting to rain, a train careening through cars down an avenue, a small enemy army consisting of Fischer's mental "projections" firing at the inception team, and Saito being shot. It continues in a warehouse, through a continued battle with projections, until all members drive off in a van and Cobb hatches the revised plan called "Mr Charles," a gambit in which Fischer is informed he is dreaming within a dream.

The fifth act is the development (21 minutes). It begins at the second dream level in a hotel bar and continues upstairs while at the first dream level Yusuf dodges the small army and continues to drive the van with its sleepers. This act ends when the van drives through the bridge railing, triggering the beginning of the sixth and final act, the climax (35 minutes).

At the first dream level the climax shows the slow motion of the van dropping to the water. Because time is increasingly dilated at successively deeper dream levels in the film, this provides different amounts of time to the same deadline at which events must come to fruition at dream levels two (the hotel), three (the winter fortress and hospital), and four (limbo, and the city designed by Cobb and his dead wife Mal, Marion Cotillard). At the end of the climax, the epilogue (three minutes) has the group returning to cinematic reality in the airplane, Saito assuring Cobb's repatriation with a phone call, Cobb passing through immigration, and Miles (Michael Caine) taking him to see his kids.

Maps

By our counts *Inception* has 2,755 shots and 355 scenes and subscenes, most of the latter are quickly cut back and forth in parallel action, particularly in the climax. These numbers vastly outstrip *All About Eve,* which we count as having 775 shots and 90 scenes and subscenes. Thus, rather than simply tracking the progress of characters across acts as we did in Figure 2, we will set up local narrative regions in which most of the characters exist, and then plot the trajectories of three other characters through these narrative spaces.[13] Also, unlike the arrays in Figure 2, we have left in place the large array of across-act, colored points, circles, and crosses for each of ten characters. This presents what might appear to be chaos, but to create order we have drawn in colored ellipses around all the locations of a given character across all the acts in which he or she appears. These are shown in Figure 3a for Mal, Cobb's children, Miles, Arthur, Eames, Yusuf, and Saito.[14]

Notice that Mal, the children, and Miles all appear on the one side of the map in quite distinguishable regions. Contrasting with them and on the opposite side are Eames, Yusuf, Saito, and Arthur. They form highly overlapping set of regions with Arthur, the only team member to appear in the prelude, extending considerably downward. The high degree of overlap of these four sug-

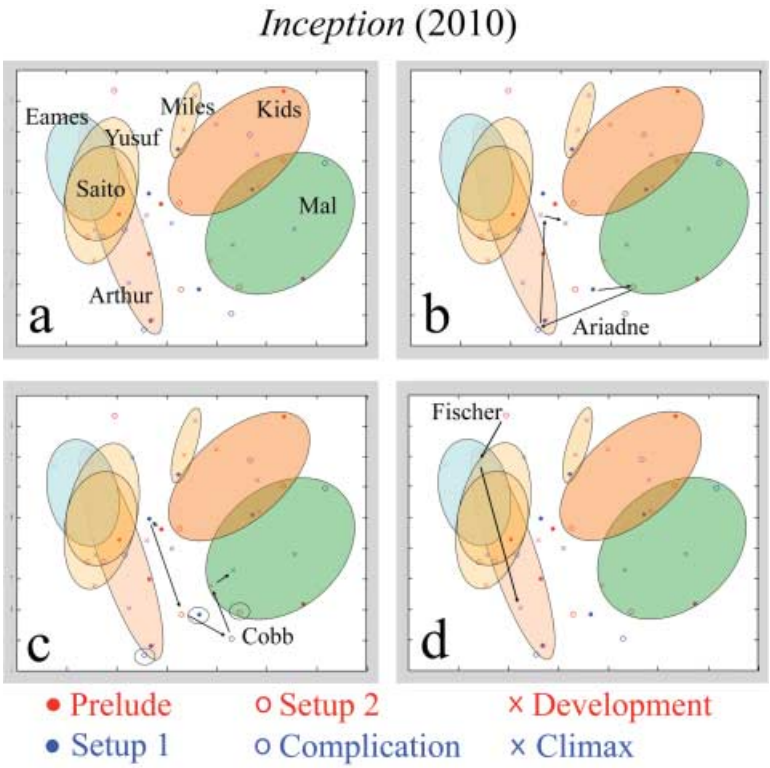

Figure 3. The upper left panel shows the scaled narrative space for seven characters as they co-occur in the scenes of Inception. Each ellipse encloses the scaled locations for each character across the acts of the film in which he or she appears. The locations of these ellipses show the nonoverlap between Cobb's family and his co-workers. The other panels show the trajectories of Ariadne, Cobb (who is near Ariadne, shown as small unfilled ellipses in three acts of the film), and Fischer.

gests that, as is true in the narrative, these characters are a team working together and are largely intersubstitutable in the larger story structure. The multidimensional scaling solutions of characters within and across scenes clearly separates the overlapping work of the *Inception* team from the separable narrative niches of Cobb's family.

Figure 3b shows the clockwise trajectory of Ariadne across the last five acts of the film, at and beyond her recruitment at the lower center of the map in the first setup, and then entering the sphere of Mal in the second setup, when she accompanies Cobb on his dreams/memories of her. Ariadne then becomes more integrated into the team and in the last two acts winds up at the center of the computed configuration. Indeed, she *is* the center of the film when dream levels collapse and the van crashes into the water in the climax. We follow Ariadne, but not the other characters, up through the dream levels. Figure 3c shows the counterclockwise

The multidimensional scaling solutions of characters within and across scenes clearly separates the overlapping work of the* Inception *team from the separable narrative niches of Cobb's family.

trajectory of Cobb, who starts out at the center of the space as he organizes everyone, then moves downward towards Ariadne (the points surrounded by open ellipses) and Mal in the last four acts of the film. His movement across the mapped space is 1.32 times greater than the normalized average (1.0) of the other characters. Finally, Figure 3d shows the trajectory of Fischer: he starts as the target dreamer outside the team in the second setup; is held captive in the complication, essentially joins the team in the development at dream levels two, three, and four; and then drifts to the outside during the climax. His across-act average movement is the greatest of all the characters (1.5). These normalized movement values suggest that Cobb and Robert are the characters that have changed the most across the film.

Again, these maps were constructed based only on visual character co-occurrence within scenes. Nonetheless, we claim they make narrative sense and generally depict more than a few aspects of plot relations among characters. But again, we make no claim that these are the only possible narrative maps.[15]

MASH (1970): Regions of Narrative Space across Film Episodes

Our analyses of *All About Eve* and *Inception* demonstrate that objective narrative maps can be constructed for quite different films from quite different eras in the face of both content and structural dissimilarities. Although our analysis gives more acts to *Inception* than most accounts might for Hollywood films, we claim that both films have a tightly coherent narrative structure, with cause and effect relations pervading them. This is not true for *MASH,* an episodic film that takes place in and around a remote *M*obile *A*rmy *S*urgical *H*ospital in a war zone in Korea, standing in for Vietnam. In such a film, there is less temporal coherence and only a smattering of cause and effect relations across episodes.

An Analysis of Episodes

We divided the film into a setup and five episodes. In the setup (26 minutes), we are introduced to the main characters as they arrive in camp—the surgeons, Hawkeye Pierce (Donald Sutherland), Duke Forrest (Tom Skerritt), and finally Trapper John McIntyre (Elliott Gould). Those already in camp include surgeon Maj Frank Burns (Robert Duvall), nurse Lt Dish (Jo Ann Pflug), Father John Mulcahy (René Auberjonois), Col Henry Blake (Roger Bowen), and Cpl Radar O'Reilly (Gary Burghoff). Episode 1 (23 minutes) begins with the arrival of head nurse Maj "Hot Lips" Houlihan (Sally Kellerman) with ideas of enforcing a strict regimen in camp and in the operating room. She begins a short passionate affair with Frank Burns, the sounds of which Radar and Trapper John broadcast over the camp's loudspeaker system. The next day Hawkeye

taunts Frank, who then attacks him, and the episode ends with Frank taken away in a straightjacket.

We consider episodes 2 and 3 together (27 minutes), with the former centered on the Last Supper event and the latter on the Hot Lips shower event. Troubled with manliness issues after failing to consummate an ad hoc fling with a nurse, the camp oral surgeon Capt Painless Waldowski (John Schuck) decides to commit suicide. He is thus the protagonist of a staged false suicide scene that humorously replicates Leonardo's Last Supper ensemble. Accommodating a request from Hawkeye, Lt Dish helps Painless regain his manly confidence on her last night in the camp. Her departure the following morning closes the episode. The shower sequence begins with camp helper Ho-Jon's (Kim Atwood) and Radar's endeavors over a device that successfully lifted the women's shower tent and revealed naked Hot Lips.[16]

In episode 4 (12 minutes) Hawkeye and Trapper John are summoned to Japan for surgery. They befriend another doctor, Me Lai Marston (Michael Murphy), and outrage then dupe Col Merrill (James Douglas). Episode 5 (22 minutes) sets up and enacts a football game between the MASH team and that of Gen Hammond (George Wood). Hawkeye and Trapper John arrange for the transfer of a heart surgeon and former professional football player, Spearchucker Jones (Fred Williamson) to the MASH unit. Spearchucker then leads the MASH football team to victory. A short epilogue (4 minutes) after the game declares the war suddenly over for the main characters.

Unlike most feature films, the episodic structure of a film like *MASH* can afford to relax the cause-effect plot connections. Such weak interepisodic causal connections allow (and account) for the continuous shift in the cast of characters. Frank Burns disappears after episode 1, Lt Dish after episode 2, and Painless is not seen after episode 3. With Hawkeye and Trapper John in Japan, the rest of the MASH crew does not appear in episode 4 and two new characters appear. Episode 5 introduces Spearchucker Jones in the context of the football game. The only characters who traverse all episodes are Hawkeye and Trapper John, the latter being introduced only three minutes from the end of setup.

Unlike most feature films, the episodic structure of a film like* MASH *can afford to relax the cause-effect plot connections.

Maps

By our analysis *MASH* has 962 shots and 153 scenes and subscenes. As for other movies in our study, we created a base map, this time involving 15 characters. We then fit to the base map the scale solutions of the setup and the episodes, interpret the resulting structure.[17] Again, we left the markers for each of the characters on the map displays. Figure 4a shows the layout of the characters in the setup. The three main characters form a small coherent

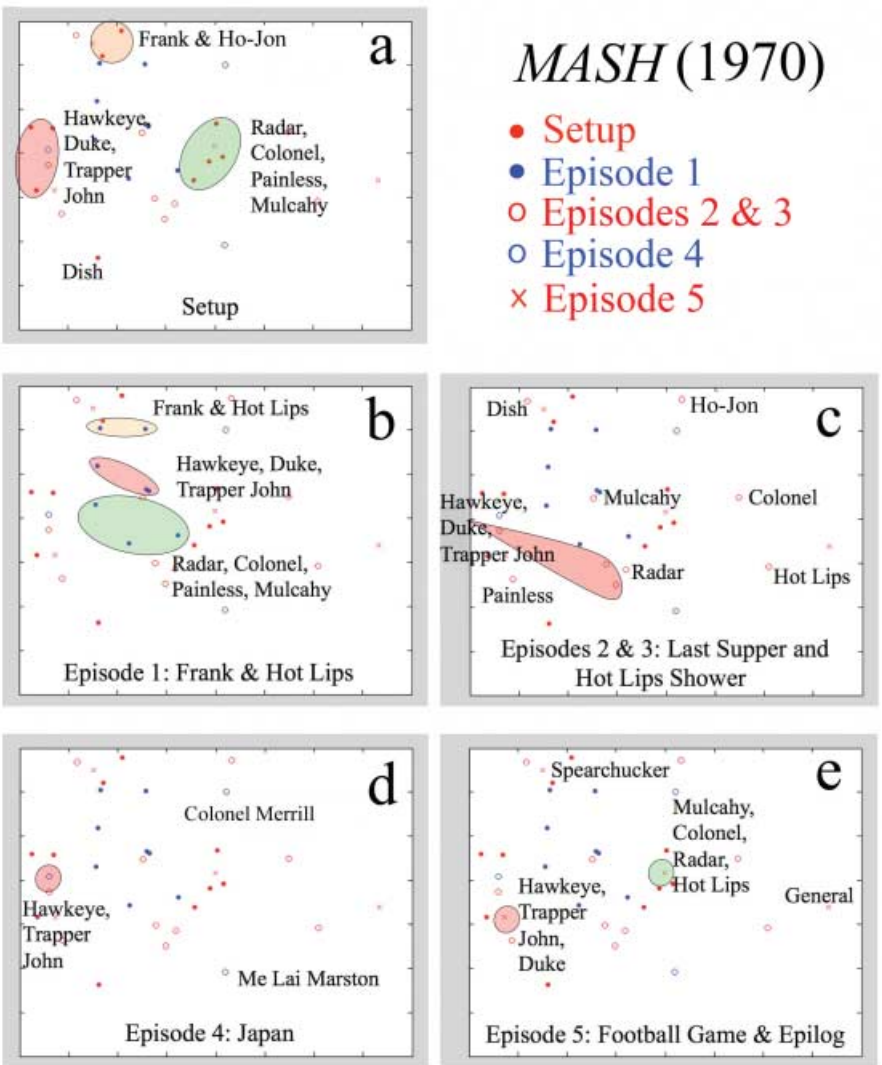

Figure 4. Narrative regions for characters in scenes of MASH *across the setup and five episodes as determined by multidimensional scaling.*

group in the middle left of the space, the regulars of the MASH unit (Col Blake, Radar, Painless, and Mulcahy) are to the right, and Frank Burns and Ho-Jon, whom Burns teaches to read English by reading the Bible, are near the top of the space. Ho-Jon quickly disappears with a soft-porn magazine given to him by Hawkeye, and he does not reappear until the end of episode 3. The separation of these three groups of characters makes narrative sense; the main surgeons are distanced from Burns, and the camp regulars are off to the side and generally undistinguishable from each other as they each provide idiosyncratic color.

Figure 4b shows the layout for episode 1, where Frank and Hot Lips get together. Notice that they now form a tight spatial group, that the camp regulars (Col Blake, Radar, Painless, and Mulcahy) form another tight group, and that the three surgeons (Duke, Trapper John, and Hawkeye) who provoke the episode intercede between these two. This arrangement is consistent with the content of the narrative.

The separation of episodes 2 and 3, both shown in Figure 4c, is revealing. It allows a contrast between the Last Supper and Hot Lips shower episodes. In the former, Painless is now separated from the remainder of the camp regulars and tucked behind the space of the major surgeons. Spread around the surgeons are Dish, who ends the episode; Father Mulcahy, who has doubts about issuing last rights for a suicide and Radar, who helps the surgeons arrange the Leonardo tableau. More striking, this Last Supper space contrasts sharply with that of the shower episode. The major surgeons generally remain relatively stable in their positions. However, Ho-Jon and Radar, who make arrangements for the lifting of the tent sides around the shower, drift away. Hot Lips, outraged, rushes to Col Blake, who is in bed with one of the nurses. Most important is that Hot Lips has moved from her remote spatial position with Frank Burns in episode 1, to an even more remote position from the major surgeons and the camp regulars in episode 3.

Episode 4 takes place in Japan with only two of the surgeons, Hawkeye and Trapper John, participating. They befriend Me Lai Marston and run afoul of Col Merrill, and the triangular distances among them are consistent with the plotline. Notice also that Hawkeye and Trapper John have become a single point in Figure 4d. The reason for this, as with Windy and Tex in our discussion of *Santa Fe Trail,* is that they appear in all scenes in this episode, and all scenes together, never separately.

Episode 5 returns to the MASH camp and preparations for a football game against the team of Gen Hammond. The three surgeons appear together in nearly all their scenes and thus again become a single point in Figure 4e, representing the players in the game; Col Blake, Radar, Mulcahy, and Hot Lips appear in nearly all of their scenes together on the sidelines and also become a point. Spearchucker appears at equal distance from the surgeons and the camp regulars since he appears about equally often with both, and the General appears at some distance from all of this since he is on the opposite sidelines. Thus, Figure 4e is a schematic representation that accurately captures the four entities involved in the football game episode: Spearchucker, the MASH players, and the fans on the two sidelines. The General's team consists of individuals who are unknown and therefore not really characters.

The narrative movement of Eve and Margo in *All About Eve* and of Cobb in *Inception* is quite different than that for Hawkeye and Trapper John. The pro-

Our analysis shows that Hawkeye and Trapper John haven't changed much across the film; instead they have orchestrated the people in the episodes as they take place around them.

tagonists of *MASH* move less across episodes, 0.67 and 0.69, respectively, compared to the normalized average of 1.0 of the other characters. Our analysis shows that Hawkeye and Trapper John have not changed much across the film; instead they have orchestrated the people in the episodes as they take place around them. This idea fits well with the narrative of the film. Finally, even with 15 characters across disjointed episodes, the plot of the narrative space of those episodes in *MASH* seems reasonable based on the co-occurrence of characters that come in and out of those episodes. In this manner, we claim increasing generalization of this technique across different kinds of films—classic standards, complex contemporaries, and episodic gambits.[18]

12 Angry Men (1957): From Diegetic to Narrative Space

Our co-occurrence analyses of the characters in *All About Eve, Inception,* and *MASH* were carried out within scenes and then compared across all acts or episodes of the film. Most films have more or less readily defined scenes and subscenes, and characters change as the narrative shifts across time and space. Some films, however, take place in continuous time—for example, *My Dinner with Andre* (1981), *12 Angry Men*, *High Noon* (1952), and *Rope* (1948). This constraint allows few chances for standard scene change. Moreover, the central core of first two films takes place in a single location with a constant set of characters. Thus, with no time jumps, no location shifts, and no character changes one could argue that there are no real changes of scene. The latter fact might suggest that our type of analysis could not be used for such films. Thus as a strong test of our method, we attempted to construct a dynamic narrative space for *12 Angry Men,* shifting methodological gears and looking only at the characters seen in consecutive shots. In our analysis of the film we found 362 shots and all but 10 of those take place in the jury room and an adjoining bathroom.

An Analysis by Votes

Ignoring the prologue in the courtroom followed by jurors milling around the jury room (eight shots, nine minutes) and epilogue outside the courthouse (two shots, one minute), we divided the whole of the jury deliberation sequence by the votes that were taken. Section 1 (22 minutes) begins with a preliminary vote. This reveals that Juror 8 (Henry Fonda) is on the short end of an 11-1 guilty verdict. Unanimity would entail a death sentence for a young Puerto Rican accused of killing his father. Discussion follows with almost everyone sitting around the table focusing on the reticence of Juror 8. After a long while and in desperation he calls for another vote.

Section 2 (13 minutes) begins with a surprise: the vote is 10–2. The oldest member, Juror 9 (Joseph Sweeney), has joined Juror 8. That outcome is unpredicted by the previous discussion or by the camerawork. More discussion and displeasure accrue, and the jurors start to walk around, and in some cases simulating events of the murder.

Section 3 (15 minutes) begins with a new vote. Juror 5 (Jack Klugman), who like the defendant grew up in a slum, changes his mind. Juror 11 (George Voskovec), an immigrant who knows the difficulties of being out of the mainstream, soon follows. The vote is then 8–4 and more discussion, anger, and perambulations ensue.

Section 4 (11 minutes) begins with Jurors 2 (John Fiedler) and 6 (Edward Binns) changing their votes, deadlocking the jury at 6–6. This is followed by a pause in the proceedings, the start of rain and cooling temperatures outside, then further discussion, and a demonstration of a switchblade fight.

In Section 5 (17 minutes) the deliberation begins with Juror 7 (Jack Warden), and then Jurors 1 (Martin Balsam) and 12 (Robert Webber) changing their votes without the apparent conviction of the others. This tips the vote to 3–9, and further discussion ensues.[19] Juror 10 (Ed Begley) disqualifies himself for bias, Juror 4 (E. G. Marshall) changes his vote on a reconsideration of evidence, and Juror 3 (Lee J. Cobb) finally relents after a dramatic protestation.

Maps

The layout for the whole film and the maps of our five sections are shown in Figure 5.[20] The base map is shown in Figure 5a. In the jury room the 12 jurors sit at a rectangular table and they are numbered clockwise around it; the foreman (Juror 1) at the head and Juror 7 at the other end. If they all remained at the table and shots simply showed a discussant and his near neighbors the scaled arrangement would likely be an ellipse. However, during the course of deliberation all jurors get up and walk around. This dissolves their diegetic arrangement placing jurors nonadjacent at the table within the same shot while sometimes omitting the jurors in between. Nonetheless, some aspects of the table arrangement remain. We have placed a dot at the origin of the plot (the center of gravity around which all jurors are placed) and drawn lines from the origin to each of the jurors.

Notice that the clockwise order of the jurors perfectly mimics the order of the jurors as they sit around the table. The diegetic space in *12 Angry Men* is apparent in our base map, merged with its computed narrative space. As to the former, despite juror perambulations over the course of the film there are a sufficient number of co-occurrences of adjacent jurors in the shots to preserve their ordinal relations at the table. Yet overlaid on this order there are deviations that are consistent with the narrative. Obvious

The diegetic space in* 12 Angry Men *is apparent in our base map, merged with its computed narrative space.

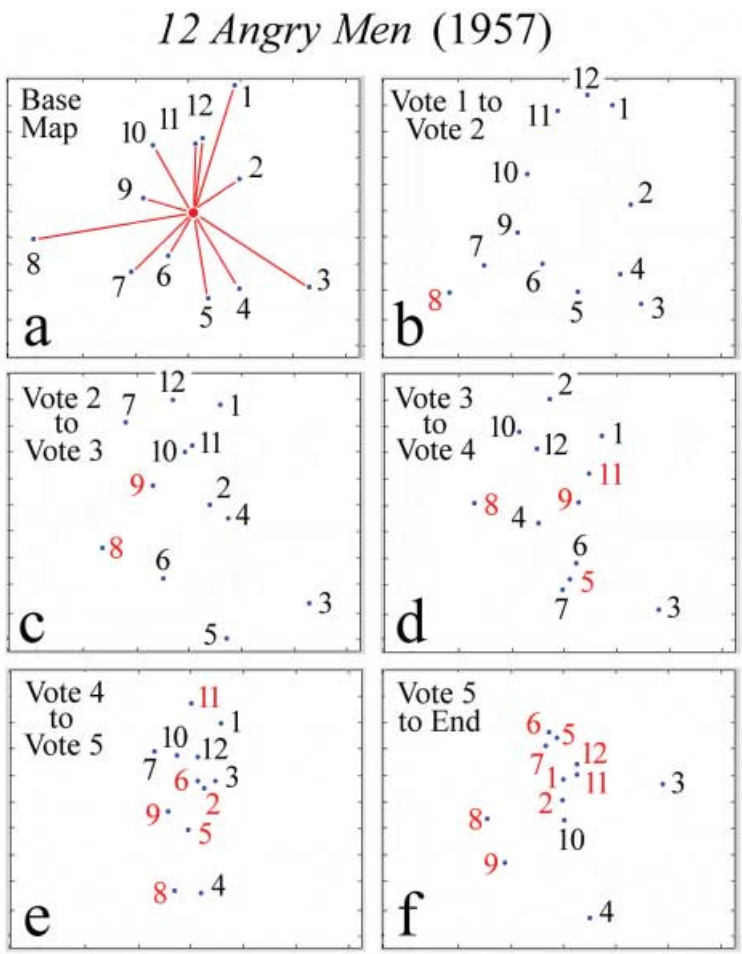

Figure 5. Narrative spaces for 12 Angry Men *as determined by multidimensional scaling co-occurrences of characters in shots of the film.*

outliers are Jurors 8 and 3. Juror 8 (Fonda) is the first to vote for acquittal and Juror 3 (Cobb) is the last. There is an important sense in which these two jurors act as polar opposites in the film. The reason for their separation from the others and from each other is the large number of shots in which they are isolated, appearing without others in the shot. We call these *one-shots*—shots depicting only one character. These reduce co-occurrence measures with all other jurors, increasing their distance from them. There are 54 one-shots of Juror 8, mostly near the beginning of the film; and 58 one-shots of Juror 3, mostly near the end.

Figure 5b depicts the scaled solution from the first vote to just before the second vote. The number of Juror 8 is colored in red because he has voted for acquittal. Notice that he is an outlier (23 one-shots in this segment), but that the general, ordinal arrangement around the table has been preserved. Across

this section of the film there are only a few shots where individual jurors have walked around. Notice too that at this point Juror 3 is not so extreme (only four one-shots throughout the segment), and in the narrative neither he nor the others is convinced of the innocence of the defendant.

Figure 5c depicts the film portion between the second and third votes. At the beginning Juror 9 (three one-shots) has joined Juror 8 and their relative positions could be predicted by proximity alone. Juror 3 (five one-shots) and Juror 5 (two two-shots and appearing with Juror 3 in all others) become more extreme. Juror 5 will soon change his mind and vote for acquittal. His isolation in the scaling solution seems to anticipate his new view based on his knowledge of switchblades, the implement of the murder, and how they are used.

Figure 5d depicts the portion between the third and fourth votes, showing that Juror 5 has changed his mind and now migrating more towards the center of the solution, no longer an outlier (and with only two one-shots in the rest of the film). In addition Juror 11, the immigrant (six one-shots in the remainder of the film), joins the still-minority view. During this sequence the jurors have walked around quite a bit and the ordinal relations of the jurors at the table is now no longer apparent. Notice that the jurors voting for acquittal generally occupy more central positions in the solution.

The period between the fourth and fifth votes (see Figure 5e) is spatially chaotic as jurors have continued their perambulations. There is no residual coherence of order around the table, no apparent visual structure in the co-occurrences of characters in shot, and no coherence of opinion as Juror 2 (convincingly) and Juror 6 (unconvincingly) have joined the emerging trend. Overall opinion is unclear in the narrative, and co-occurrences in shots reflect that lack of clarity.

Figure 5f shows the shot co-occurrences between the fifth vote and the end of the film. Shown in black are the last three jurors to change their minds. Juror 10 is the first to recant his position, essentially through embarrassment at his own racist monolog. It was a little surprising to us that his position is so central, but the depiction in this scene is of other jurors removing themselves from the table as he speaks, facing away from him. The cinematography retains all of jurors within the same camera purview in a long shot. Juror 4 is the most reason-driven holdout and is shown in five one-shots as others counter his claims. Eventually he is convinced that one of the key eyewitnesses could not possibly have seen the murder. Juror 3 soon self-destructs in a tirade against his son, himself, and the world. As he does so he is shown in 24 one-shots, many of them tight close-ups. Notice the cluster of votes for acquittal is bound together, but has completely lost its spatial order around the table.

Taken as a group, we think these scaling solutions built on the character co-occurrences in shots show the dynamic pattern of the narrative as rela-

tions develop and change over the course of the film. Moreover, their diegetic arrangement in the jury room—as might be predicted by the base map—systematically changes and is incrementally replaced by a more abstract narrative space based on votes.

A powerful numerical trend concerns the distances of the characters from the center of the space as each vote is taken. The center of the space can be taken as neutral, particularly as the spatial layout around the table disintegrates. When Juror 8 casts the only vote for acquittal (Figure 5b), his distance from the center is 1.79 times the average distance of those of the guilty votes. He is a true outlier and this distance is due to the large number of one-shots. But after being joined by Juror 9 (Figure 5c) the mean distance of the two dissenters is only 1.25 times the average of the majority. And after Jurors 5 and 11 join making the vote 8–4 (Figure 5d), the average distance from the center of the characters with votes for acquittal is *less* than for those with guilty votes, with a proportion of 0.83. When vote is deadlocked (Figure 5e), the distance of those voting not guilty continues to decrease and is only two-thirds (0.68) that of those jurors voting guilty. And with the vote 3–9 (Figure 5f) the mean distance of those voting not guilty is about half (0.57) that of those voting guilty. Thus, the sequence of mean spatial distances of characters voting for and against acquittal anticipates the outcome of a unanimous verdict of not guilty. And again, all of this stems from what is shown in the camera; no dialog is considered. What the visual narrative shows in its shots very much captures the dynamics of what the film is about.

Discussion: The Visual Narrative and the Viewer

> We are getting closer and closer to the techniques of *narrative* . . . and to the description of association networks ... In this fusion of qualitative literary qualities and the power of quantitative treatment, we expect a renewal of methods and explanations in the humanities. There is no more powerful explanation than the analysis of the contingent circumstances of association networks.
>
> —Geneviève Teil and Bruno Latour (1995)

What is the import of our results? Consider four intellectual gains. The first three are concrete but modest, and the fourth more tenuous but we believe is of considerable importance. First, in reading the wider literature, we have found previous accounts of narrative space in films and novels to be somewhat vague and unhelpful. Some are tied to an analogy between linear perspective in Renaissance art and narrative space in film, with separate shots contributing to a mental construction of a physical space. Others abandoned aspects of this idea but resuscitated it through an appeal to cognitive maps.[21] Such accounts seem incomplete, overly tied to physical space, and static even

when discussing a moving camera and shot/reverse-shot depictions. Our first contribution allows such spaces to be objective, explicitly mapped out, and largely freed from physical space. Moreover, our analysis over separate portions of films allows for the interpretation of dynamic changes in those spaces. Nonetheless, we admit that clearing up some of the problems or vagaries in ideas found elsewhere is not earthshaking science. What we have found here does not really add to our understanding of the narrative in each film. Instead, it adds objective depictions that are coherent with it.

Second, by bringing statistical rigor to the analysis of shots and scenes across larger units of film, we show a coupling of layout and changes in the visual content of the film to the relationships and changes in its story line. That is, our methods demonstrate that filmmakers sculpt the presentation of characters in a way that comports amicably with the narrative itself. Again, it is good to know that this can be done from objective data, but it is hard to imagine that the result could be any other way. Surely the visual narrative of a film must be tightly connected to the narrative itself.

Third, a major purpose of this article is to introduce a new method for the visual analysis of films. The mappings we have offered are concrete and largely deterministic.[22] One possible use might be in making narrative comparisons across a film. If it were important to an argument to assert that the relative amount of narrative change in the relationship of one particular pair of characters was greater than that of another pair, ours would be a method that could affirm or disconfirm of such an assertion. Similarly, if it were important to suggest that one character is more distant from the general flow of the narrative than another, this too could be tested with our method. And finally, as we have done here, if it is important to measure the amount of spatial change in a character compared to other characters as an index of narrative development, our method allows this too.

We regard the fourth gain, although more tenuous than the first three, to be quite deep. It concerns our analysis of the visual film structure and the viewer. As suggested by Teil and Latour (1995), the human mind—among its many other attributes—is an engine driven by associations. Indeed, the mind can be described as a Humean machine (after David Hume); it builds up knowledge through associations among things that co-occur in the world across our individual lifetimes. Moreover, co-occurrence commonalities within a culture create the consistency of cultural knowledge across individuals as each of us absorbs them. This is not the only means to knowledge gain but it is a very powerful one. Moreover, units of the nervous system—often called Hebbian units (after Donald Hebb)—function this way as well. The oft-quoted phrase is that "neurons that fire together, wire together."[23] In this manner, from basic neurophysiological processes to the attainment of knowledge and culture, our brains and minds seem to follow associative mechanisms.

For filmgoers we claim that part of their comprehension process is the same: to understand the narrative over the course of a film a viewer must build up representations of the relations among the characters and their goals. Understanding how this is done and how films manipulate this build-up are important tasks within cognitive film theory. As a preliminary sketch of this process, Bordwell (1992: 184) suggested that:

> *the process of understanding many things in films is . . . likely to draw upon ordinary, informal reasoning procedures . . . Presented with a set of circumstances (flat tire, man opening trunk), you categorize it (Driver Changing Flat Tire) and draw an informal, probabilistic conclusion, based on a structured piece of knowledge about what is normally involved in changing a tire.*

For our purposes the key ideas in Bordwell's analysis concern probability and the making of meaning out of what one sees. Although surprise is always a possibility in life and in movies—think of the endings in *The Usual Suspects* (1995) and *Atonement* (2007)—the filmgoer's aggregation of probabilistic information from a lifetime of experience with what she has seen in a film up to the moment creates an evolving set of mental representations of the characters and story. These accruing structures allow for reasonably accurate predictions about what will happen next. Again, violations of our filmic expectations can be delightful, surprising, even horrific, and they can keep one riveted to the narrative, but the ability to predict a good bit of what will happen over the intermediate course of a film is central to its comprehension. Moreover, this ability to use the comparative probabilism from pieces of experience is the essence of what is called statistical learning.[24] It occurs over lifetimes and also over minutes.

More formally, statistical learning is a branch of machine learning and cognitive science. It explores how learning can occur in computers and in humans, and in the latter has been particularly used to explore language learning and visual learning by infants (Kirkham et al. 2002; Saffran et al. 1996). Particularly interesting and most relevant to our analysis is unsupervised statistical learning. In this domain, a computer or human is typically exposed to repeating sequences of nonsense items and passively learns about the structure of the co-occurrences that are presented. Thus, to use an overly simple example, if a red star is followed by a green triangle 60 percent of the time in a long sequence of varied items of various shapes and colors, adults and infants would likely come to expect that, given the presentation of a red star, the next item is likely to be a green triangle. Most interesting about this expectation is that it is typically learned without awareness. In more complex situations, viewers often have no insights concerning their predictions and regard them as just guesses. This learning is essentially nonconscious. More-

over, this learning occurs not only for sequential items, but also for features that co-occur in the same arrays and for complex aspects of static visual scenes (Brady and Oliva 2008; Fiser and Aslin 2001).

We are quite sure that no one has experimentally studied the statistical learning of co-occurrences of characters in a film. Nonetheless, given the wealth and breadth of demonstrations of statistical learning, we have no doubt that this occurs for filmgoers and that, for the most part, it also occurs nonconsciously. People fascinate us, we pay attention to them to a greater degree than anything else, and we do so from a very early age. Moreover, they are the prime objects that we look at in movies (Mital et al. 2011; Smith et al. 2012). We follow them, remember them, and that we likely encode their appearances with other characters seems uncontroversial. Otherwise we could not understand movies.

We make no direct claim that the diagrams shown in figures 2 through 5 are actually "in the head" of a viewer who watches these films. These diagrams are representations of the character relations in the films. Nonetheless, our data demonstrate that character co-occurrence relations by themselves can yield narrative insights, a fact heretofore not known. We claim that these maps should bear some recognizable similarity to the relations among characters in the minds of viewers, and by themselves could guide some coarse understanding of the film without a necessary comprehension of verbal events across the film.

James E. Cutting is Susan Linn Sage professor and chair of the Department of Psychology at Cornell University where he has taught courses on perception for over 30 years. He has written two books, *Perception with an Eye for Motion* (1986) and *Impressionism and it Canon* (2006) and published more than 100 scientific articles on form, motion, depth, and related areas of perception.

Catalina Iricinschi is a graduate student in the Department of Psychology at Cornell University. Her interests include psycholinguistics, narrative processing, and language development.

Kaitlin L. Brunick is a graduate student in the Department of Psychology at Cornell University. Her interested include film, TV, and children's cognitive development.

Notes

[1] The discussion of narrative space in film studies begins with Heath (1976), who described its creation through camera work. This space is the implied physical space within which action takes place in the film, specifically carved by the filmmakers to promote the story in its specific location. With modifications Cooper (2002) reviewed and continued this idea. Zoran (1984) discussed several types of narrative space, including a *chronotopic level* representing space and time within a film. Similarly in literature, Bjornson (1981: 59) suggested that narrative spaces are created by "piece-by-piece construction of an image which maps the imaginary space described in the text" and Ryan (2003) discussed cognitive maps that are built up over the course of a story. O'Toole (1980) came closest to our approach, even using two dimensional matrices to depict relationships among characters in episodes of stories but he never attempted further statistical analysis. Multidimensional scaling, which we employ here, has been used to generate cognitive maps (Kosslyn et al. 1974), but these map back to physical spaces. Throughout empirical psychology, however, multidimensional scaling has been used to create conceptual spaces, and following Miss Yujin (2009), our maps are abstract and do not typically reflect the depicted physical space in a film.

[2] Film viewers intuitively claim that the verbal behavior displayed onscreen carries more of the plot than visual information. In addition, research on content analysis in film regards the plotline as a combination of dialog and action, where action is a mixture of verbal and nonverbal visual input. See, for example, Rosenfeld et al. (2003).

[3] With multidimensional scaling solutions one typically computes the *stress* of the solution, sometimes called "badness-of-fit." The lower the stress value the better the solution captures the relational patterns in the data. It is not surprising that the stress for the solution of the data for six European cities is very low—0.004.

[4] These costs were gathered on 10 September 2002 for flights to be taken on 1 October 2002. These are the average of round trips each way: for example, Paris-London-Paris and London-Paris-London. A more expansive discussion of these data can be found in Cutting (2006). We also note that, as its name suggests, multidimensional scaling is not confined to two dimensions. If there are more objects to scale, the algorithm can easily handle the data in three or more dimensions. We have confined ourselves to two dimensions both because it yields more understandable maps and because the two-dimensional solutions generally provide reasonable and interpretable fits to these data. Notice that the more typical hubs of Amsterdam and Frankfort were left out of this analysis. Finally, taking flight costs for June 2013 into account, we found that London had lost its centrality in the mapping, but that strong deformations from veridical maps still exist.

[5] Nonmetric multidimensional scaling routines automatically convert metric information into nonmetric ranks. Again, the stress is reasonable for this solution—0.095.

[6] Cohen's κ in this context corrects for chance co-occurrences, which Pearson's r does not. Also, unlike other correlations, it is typically confined to a range between 0 and 1. Nonetheless, the calculation easily generalizes to a range of –1.0 to 1.0, which we use here.

[7] The number of scenes was determined by aggregating the results of three different viewers who parsed *Santa Fe Trail* (1950) for a different project (Cutting et al. 2012).

[8] One could also argue that such a film would need two narrative spaces that are quite separate from one another. We have no stake in either interpretation, and use this as an example only.

[9] Bellour (1976) called scenes and subscenes by the terms segments and subsegments.

[10] The stress of the base map for the nine characters in *All About Eve* was rather high—0.144—indicating a less than optimal fit. The stresses for maps of the act 1 and act 3 maps

were also high, 0.127 and 0.13, respectively. The stress for act 2 was more reasonable—0.08, and that for act 4 very good—0.0001. Despite the relatively high stress in these scaling solutions, we claim it is their utility in making sense of the narrative that matters in this context. Interpretability is an important criterion in multidimensional scaling (Kruskal 1964).

[11] The climax representations in Figure 2 of the two heterosexual couples—Margo and Bill, and Karen and Lloyd—contrast with Eve and Addison. It is worth reconsidering an argument concerning this film and classical Hollywood's treatment of homosexuality. White (1999) and Corber (2005), for example, proposed that Eve and Addison, and particularly Eve, are portrayed as characters of diminished fulfillment and happiness because of their homosexuality. Most provocatively in one short scene, Eve enlists a female companion to assist her with a phone call in a plot concerning Lloyd. After that call, she and Eve walk up a staircase in bathrobes, arms around each other's waists. Thus a lesbian reading seems not inappropriate. Nonetheless, it is worth noting that Eve also tries to seduce all three male leads in the film, and late in the film Addison confronts her with her past, which involved a relationship with married man and a quick, paid departure from her hometown. We think that Eve is essentially omnivorous, rather than simply lesbian. In addition, Addison seems asexual rather than homosexual. He asserts to Eve in a climactic scene: "After tonight you will belong to me. . . . [We both] have a contempt for humanity, an inability to love and be loved. . . . We deserve each other."

[12] After Procrustes analysis the correlation of coordinates is extremely high, $r = .993$, $t(14) = 31.5$, $p < .0001$.

[13] The stress for the base map for the 10 characters in *Inception* was adequate—0.099, and that for the six separate acts was 0.005, 0.048, 0.076, 0.032, 0.107, and 0.099, respectively.

[14] We did not consider Robert Fischer's dying billionaire father, Maurice Fisher, nor his executive assistant Peter Browning as sufficiently important characters to include in this analysis or in constructing the base map. The father appears in only 13 shots across the film and, although Browning appears in 38 shots, Eames is impersonating him in 24 of them. In our assessments Eames was always counted as present when impersonating any character. By contrast Miles, the character with the smallest role in our analysis, appears as himself in 31 shots.

[15] The correlation between the two base map solutions, one for simple scene coding versus another that weighted co-occurrences by the duration of each scene, was again extremely high, $r = .9993$, $t(18) = 113.3$, $p < .0001$. Thus, the two procedures yield the same results.

[16] Tacked on to the end of the shower sequence is a short episode of Ho-Jon and Hawkeye at a Korean clinic (two minutes), but we did not consider it separately.

[17] The stress of the base map for *MASH* was generally high—0.19, but this value is understandable given as many as 15 characters scaled in only two dimensions. The stress for the setup was 0.126, and in the subsequent plots of episodes 0.058, 0.175, 0.001, and 0.001, respectively.

[18] The correlation for *MASH* between the solutions of the unweighted co-occurrence codes versus those weighted by scene duration was not nearly as impressive as those for *All about Eve* and *Inception*: $r = .799$, $t(28) = 7.05$, $p < .0001$. This is undoubtedly due to the relatively high stress in both two-dimensional solutions.

[19] Juror 12 actually changes his vote back to guilty and then later back to not guilty.

[20] Stress values of the six panels of Figure 5 for *12 Angry Men* are 0.082, 0.051, 0.130, 0.065, 0.089, and 0.062, respectively. Moreover and again, weighting the shots by their duration does not change the overall layout of the base map. The correlation between the character positions in the map generated weighting the shots equally versus weighting them by their duration was very high, $r = .989$, $t(22) = 31.4$, $p < .001$.

[21] Heath (1976), in comparing photography and cinema, suggested that space was built up in a manner not unlike Renaissance perspective. Cooper (2002) suggest that this view was passé, but resuscitated aspects of the approach in an analysis of physical space in *Sleepless in Seattle* (1993), and how it might be lit. Ryan (2003) switched gears towards the observer and suggested that a mental space was built up in the form of a cognitive map.

[22] The running of algorithms of multidimensional scaling yields a certain amount of variability. As these programs iteratively run, cycling through possibilities, shifting points in space, and remeasuring the goodness of fit, they can often reach what is called a local minimum and stop. This process is somewhat analogous to water running downhill and getting caught in a puddle that will not allow it run down to a lower elevation. To keep this from happening one typically starts the algorithmic process with a number of different random positionings of the to-be-scaled points and selects the outcome "that runs downhill the farthest," yielding the lowest stress. When several different starting configurations reach the same stress level, one can be satisfied that what is called a global minimum has been reached.

[23] See Cutting (2003, 2006) for an analysis of how an artistic canon might be built up from shared experiences within a culture. And Doidge (2007: 427) attributes a first version of the "wiring together" phrase to Carla Shatz, a neurobiologist at Stanford.

[24] Prior to the evolution of statistical learning, script theory (Schank and Abelson 1997) relied on the accumulation of probabilities for the buildup of information as in Bordwell's changing-tire example. See also Gernsbacher (1990).

References

Bellour, Raymond. 1976. "To Analyze, to Segment." *Quarterly Review of Film Studies* 1(3): 331–353, doi: 10.1080/10509207609360959.

Bjornson, Robert. 1981. "Cognitive Mapping and the Understanding of Literature." *SubStance* 10(1): 51–62.

Bordwell, David. 1992. "Cognition and Comprehension: Viewing and Forgetting in *Mildred Pierce*." *Journal of Dramatic Theory and Criticism* 6(2): 183–198.

Brady, Timothy, and Aude Oliva. 2008. "Statistical Learning Using Real-World Scenes: Extracting Categorical Regularities Without Conscious Intent." *Psychological Science* 19(7): 678–685, doi: 10.1111/j.1467-9280.2008.02142.x.

Cooper, Mark Garrett. 2002. "Narrative Spaces." *Screen* 43(2): 139–157, doi: 10.1093/screen/43.2.139.

Corber, Robert J. 2005. "Cold War Femme: Lesbian Visibility in Joseph L. Mankiewicz's *All about Eve*." *GLQ: A Journal of Lesbian and Gay Studies* 11(1): 1–22.

Cutting, James E. 2003. "Gustave Caillebotte, French Impressionism, and Mere Exposure." *Psychonomic Bulletin & Review* 10(2): 319–343.

Cutting, James E. 2006. *Impressionism and Its Canon*. Lanham, MD: University Press of America.

Cutting, James E., Kaitlin L. Brunick, and Ayse Candan. 2012. "Perceiving Event Dynamics and Parsing Hollywood Films." *Journal of Experimental Psychology: Human Perception and Performance* 38(6), 1476–1490

Doidge, Norman. 2007. *The Brain That Changes Itself*. New York: Viking Press.

Fiser, József, and Richard N. Aslin. 2001. "Unsupervised Statistical Learning of Higher-Order Spatial Structure from Visual Scenes." *Psychological Science* 12(6): 499–504.

Heath, Stephen. 1976. "Narrative Space." *Screen* 17(3): 19–75.

Gernsbacher, Morton A. 1990. *Language Comprehension as Structure Building*. Hillsdale, NJ: Lawrence Erlbaum.

Kirkham, Natasha Z., Jonathan A. Slemmer, and Scott P. Johnson. 2002. "Visual Statistical Learning in Infancy: Evidence for a Domain General Learning Mechanism." *Cognition* 83(2): B35–B42, doi: 10.1016/S0010-0277(02)00004-5.

Kosslyn, Stephen M., Herbert L. Pick, and Griffin R. Fariello. 1974. "Cognitive Maps in Children and Men." *Child Development* 45(3): 707–716.

Kruskal, John B. 1964. "Multidimensional Scaling by Optimizing Goodness of Fit to a Nonmetric Hypothesis." *Psychometrika* 29(1): 1–27.

Mital, Parag K., Tim J. Smith, Robin L. Hill, and John M. Henderson. 2011. "Clustering of Gaze during Dynamic Scene Viewing is Predicted by Motion." *Cognitive Computation* 3: 5–24, doi: 10.1007/s12559-010-9074-z.

O'Toole, Lawrence. 1980. "Dimensions of Semiotic Space in the Narrative." *Poetics Today* 1(4): 135–149.

Rosenfeld, Azriel, David Doermann, and Daniel DeMenthon, eds. 2003. *Video Mining*. Boston: Kluwer Academic.

Ryan, Marie-Laure. 2003. "Cognitive Maps and the Construction of Narrative Space." Pp. 214–242 in *Narrative Theory and the Cognitive Sciences,* ed. David Herman. Stanford, CA: Publications of the Center for the Study of Language and Information.

Saffran, Jenny R., Richard N. Aslin, and Elissa L. Newport. 1996. "Statistical Learning by 8-Month-Old Infants." *Science* 274(5294): 1926–1928, doi: 10.1126/science.274.5294.1926.

Schank, Roger C., and Robert P. Abelson. 1977. *Scripts, Plans, Goals, and Understanding: An Inquiry into Human Knowledge Structures*. Hillsdale, NJ: Lawrence Erlbaum.

Smith, Tim J., Daniel T. Levin, and James E. Cutting. 2012. "A Window on Reality: Perceiving Edited Moving Images." *Current Directions in Psychological Science* 21(2): 142–145, doi: 10.1177/0963721412437497.

Teil, Geneviève, and Bruno Latour. 1995. "The Hume Machine: Can Association Networks Do More Than Formal Rules?" *Stanford Electronic Humanities Review* 4(2). http://www.stanford.edu/group/SHR/4-2/text/teil-latour.html (accessed 16 October 2012).

Thompson, Kristen. 1999. *The Way Hollywood Tells It*. Cambridge, MA: Harvard University Press.

White, Patricia. 1999. *The Uninvited: Classical Hollywood Cinema and Lesbian Representability*. Bloomington: Indiana University Press.

Yujin, Miss. 2009. "Narrative Space in *Sleepless in Seattle* and *Suspicion*." October 9. http://missyujinfilm.blogspot.com/2009/10/narrative-space-in-sleepless-in-seattle.html (accessed 16 October 2012).

Zoran, Gabriel. 1984. "Towards a Theory of Space in Narrative." *Poetics Today* 5(2): 309–335.

Filmography

Altman, Robert. 1970. *MASH*. USA

Curtiz, Michael. 1940. *Santa Fe Trail*. USA.

Ephron, Nora. 1993. *Sleepless in Seattle*. USA.

Hitchcock, Alfred. 1948. *Rope*. USA

Howitt, Peter. 1998. *Sliding Doors*. UK and USA.

Lumet, Sidney. 1957. *12 Angry Men*. USA.

Mankiewicz, Joseph L. 1950. *All About Eve*. USA.

Nolan, Christopher. 2010. *Inception*. USA and UK.

Malle, Louis. 1981. *My Dinner with Andre*. USA.

Singer, Bryan. 1995. *The Usual Suspects*. USA and Germany.

Wright, Joe. 2007. *Atonement*. UK, France, and USA.

Zinnemann, Fred. 1952. *High Noon*. USA