

Epistemic Games Group Working Paper 2010-01

March 2010

The Bicycle Helmets of “Amsterdam”: Computer games and the problem of transfer

David Williamson Shaffer

Copyright © 2010 by David Williamson Shaffer

All rights reserved.

Readers may make verbatim copies of this document for noncommercial purposes by any means, provided that the above copyright notice appears on all copies.

Epistemic Games working papers are available on the Internet at www.epistemicgames.org.

Recommended citation:

Shaffer, DW. *The Bicycle Helmets of "Amsterdam": Computer games and the problem of transfer* (Epistemic Games Group Working Paper No. 2010-01). Madison: University of Wisconsin-Madison,

The Bicycle Helmets of “Amsterdam”: Computer games and the problem of transfer

David Williamson Shaffer

Abstract

This chapter looks at a central question in educational games: how do the things players do in a game help them do other things in the world outside of the game? The chapter frames the problem in both personal and literary terms, and then proposes a particular theoretical perspective on the issue. Empirical data and studies of computer games that use, explore, and expand this theory are presented; but the larger goal of the chapter is not to argue for one particular theory. Rather, the chapter provides an example of a broader class of theories that examine learning as simultaneously symbolic and situated, and makes the case that this theories of this kind are needed to develop to study learning that happens in games.

Introduction

Any discussion of games and learning has to address a fundamental question: *How do we know that players aren't just learning how to play the game?* That is, how do we know that what they do in the game will help them do other things in the world around them, the world outside of the game?

But answering such questions forces one to come dangerously close to—indeed, to step right on—the third rail in the study of learning today: the problem of *transfer*. No term, no word, no concept is as problematic, as debated, as contentious. Schema theorists say it is essential, sociocultural theorists say it doesn't exist, and never the twain shall meet, it seems. Debates are held in scholarly journals (see, eg, Anderson, Reder, & Simon, 1996; Cobb, 1999; Greeno, 1997; Edwin Hutchins, 2008; Packer, 2001; Seel, 2001). Graduate students and junior faculty avoid the term, knowing that where you stand on transfer determines where you sit in the field, and whatever you say will cast you in a dubious light somewhere.

But the problem of transfer is unavoidable in a field that examines the conditions under which, and the processes by which, people learn. As Dewey (1938) points out: “Every experience influences in some degree the objective conditions under which further experiences are had” (p. 37). A key question for the design of games for learning has to be: “What are the mechanisms by which one experience

influences another?” The premise of education writ large is that it is possible for experience to transfer—in this general sense—from one context to another. Otherwise there is no education, in games, in school or anywhere. There is no learning. There is no continuity, no culture, nothing beyond the immediate here and now.

To make matters worse, no currently-dominant theory accounts very well for this central problem of how experiences in one context (such as a computer game or a classroom) change the way learners think in other contexts (Sfard, 1998). At the microscopic level of analogical problem schemata (Anderson, 1993), theories of *learning as transformation of an individual* suggest that transfer occurs when a solution developed to address one problem is used to solve an analogically similar problem. The schema view is pervasive (formally or informally) in the current education system, in which isolated facts and problem-solving strategies are taught in classroom settings with the expectation that students will use those pieces of information and tools in other contexts. But nearly a half-century of research suggests that transfer in this sense is rare, difficult to achieve, and limited to problems that are very similar to the original context in which a solution is developed (the phenomenon of *near-transfer*). At the macroscopic level of membership in communities of practice (Lave & Wenger, 1991), theories of *learning as participation by individuals in larger practices* suggest that individuals learn by becoming members of a community, being mentored to do more tasks in different contexts within which the community operates.¹ In this view, however, skills do not transfer between settings; individuals can only learn to accomplish particular tasks within particular settings. Thus, neither view accounts very well for how students might be able to take an experience in a game, classroom or after school program and use it to accomplish meaningful ends in their lives outside of school; so neither view, by itself, is particularly useful for the development and analysis of new technologies and new environments for learning.

In presenting these “two views” I do not mean to suggest that there are no other extant theories of learning. I only argue that the two dominant paradigms for thinking about learning—as exemplified by

¹ And, of course, reflexively communities change through the participation of individuals.

particular theories—are mutually incompatible and individually incomplete as guides to the analysis of and development of learning environments.

This division is particularly problematic in the study of games for learning, a domain in which designers have little—too little—in the way of research and theory to guide them. They can follow one paradigm, and emphasize repetition of general principles. Or the other, and build richly immersive environments that encourage situated thinking, with little theoretical framework as to how learning in a virtual situation will lead to action in the real world. This makes it difficult to advance our understanding of how learning takes place in games—and thus what we can do as designers to make games for learning more effective, more powerful, more meaningful, and more useful in the broad curriculum of students' lives.

It does not have to be this way. A new generation of scholars has begun to work in the last decade, after the schism that divided sociocultural and symbolic approaches to the study of learning, trained to think about issues from both points of view. I deliberately include in this generation both those whose training incorporated symbolic and sociocultural perspectives and those who began working from one perspective and have taken up others later in their careers. A great deal of work went into distinguishing a coherent sociocultural science of learning from the orthodoxy of classical cognitive science, and the study of learning is deeply indebted to that work. But that hard-fought distinction having been made, some scholars are increasingly interested less in asking which view is right than in understanding how these views can inform one another—and how together they can advance our understanding of how to design environments for learning. In what follows, I argue that we can understand learning—the transfer of experience from one context to another—as something that can be discussed, analyzed, conceptualized and supported as both *transformation of an individual* and *participation in a practice*.

One way to start in this direction would be to provide an overview of extant theories, mapping the current landscape, highlighting the strengths and weaknesses of each approach. However, such overviews have already been written (see the particularly lucid and comprehensive Barab & Plucker, 2002), and in

any event tend to emphasize the existing dialectic rather than the possible synthesis. Instead, I follow Chi's (1997) example from some years ago, in which she mapped a complex theoretical landscape by providing one particular pathway through it.

In this chapter, I begin by framing the problem in both personal and literary terms. I then propose a particular theoretical perspective on the issue. In doing so, I will present some empirical data and refer to studies of computer games that use, explore, and expand this theory; but my purpose is primarily descriptive rather than demonstrative. My purpose is descriptive because my larger goal is less to argue for one particular theory than to provide an example of a broader class of theories that examine learning—and the mechanisms of learning—as simultaneously individual and cultural. I argue that these are the kind of theories that we need to develop to study learning that happens in games: learning that is simultaneously and authentically both symbolic and situated.

The Helmets of “Amsterdam”

I recently had the great pleasure and privilege of spending a year on sabbatical, living with my family in a city in Europe. This city—let's call it Amsterdam for convenience sake, although any similarity to the real city of Amsterdam is purely coincidental—is a city of some three quarters of a million people and, remarkably, over a million bicycles. Everyone in "Amsterdam", young and old alike, rides a bike.² Our 75 year old neighbor, for example, whose groceries we helped carry upstairs, rode her bicycle every day to the store.

² This is, of course, quite transparently a fiction, as anyone who knows me knows that in fact I was on sabbatical in the real European city called Amsterdam. I make the distinction between the Amsterdam in the Netherlands and the “Amsterdam” I discuss here because I am not presenting a rigorous empirical study of a particular city. Rather, I am making a point about the nature of culture and learning, which only requires that the details are illustrative. In this I am following Barthes, who in *Empire of Signs* (1982) writes in detail about the semiotics of Japanese culture but claims explicitly that he “can also—though in no way claiming to represent or analyze reality itself (these being the major gestures of Western discourse)—isolate somewhere in the world (*faraway*) a certain number of features (a term employed in linguistics), and out of these features deliberately form a system. It is this system which I shall call: Japan” (p. 3).

Something like 83% of the population of Amsterdam is over the age of 14, which makes Amsterdam a city of more than 620,000 adult bike riders. But in a year of living among 620,000 bicycle-riding adults, I literally *never* saw a Dutch adult riding with a helmet. The 620,000 bicycle-riding adults of Amsterdam certainly know about helmets. They make their children wear them. They use helmets when riding racing bicycles through the countryside. But no adult could be seen—or would be seen—wearing a helmet while riding his or her commuter bike in the city.

This may be purely a practical matter, in the sense that Amsterdam is well-planned for bicycle traffic. There are separate lanes for bicycles on most streets, and many are actually physically separated from lanes for other vehicular traffic. Cars and busses—and even pedestrians—reliably give bicycles the right-of-way. So Amsterdam is arguably a safer city for bicyclists than many, and that certainly could be part of the reason adults do not wear helmets.

However, I argue that this phenomenon is better understood as an aspect of *culture*. In making this claim, I am using the term “culture” following Geertz (who is himself following Mead and others) as “a traffic in... significant symbols—words for the most part, but also gestures, drawings, musical sounds, mechanical devices like clocks [or, for that matter, bicycles and helmets] or natural objects like jewels—anything, in fact, that is disengaged from its mere actuality and used to impose meaning upon experience” (p. 45). But although Geertz was writing half a century ago (and Mead yet again a half a century before him) current scholarship continues to look at “shared cultural systems of meaning” (Lave & Wenger, 1991, p. 54, p. 54) and “cultural models” (Gee, 1999): at culture as a system of public symbols through which experience is understood (Sahlins, 2000).

It perhaps requires little convincing that a bicycle helmet—or rather, in this case, the choice someone makes about whether to wear a helmet or not—is a cultural phenomenon, but still I offer as evidence the fact that during my stay in Amsterdam, Dutch friends and colleagues almost without fail commented that wearing a bicycle helmet made me look like a foreigner. Their occasional laughter (good-natured, of course) marked the helmet as the token of an outsider. Carrying a helmet after riding to a meeting led to discussions of the relative skill of Dutch and American bicyclists. And, of course, the

pragmatics of cycling in Amsterdam—from the organization of the roads to the habits of drivers—are also part of the cultural *habitus* of Amsterdam’s social organization.³

I make this point about the cultural nature of bicycle helmets because I want to draw attention to a particular fact about life as an American living in Amsterdam: There were other Americans living there too, and *most of them did not wear bike helmets either*. Or, to be more precise, as in Dutch families, most American children wore bicycle helmets in Amsterdam and most American adults did not.

Being a curious ethnographer, I naturally conducted some informal interviews with a convenience sample of Americans who were living in Amsterdam at the time.⁴ In this small survey of American acquaintances, I discovered four interesting results:

1. All of the adults I spoke with said they wore helmets when riding a bicycle in the United States.
2. Among those who did not wear a helmet in Amsterdam, the reasons they gave were various: “I feel silly doing it here,” “No one else does,” “I feel safer here,” “I have to take it on and off so many times,” “It messes up my hair,” and so forth.
3. Those who continued to wear helmets while in Amsterdam gave the same reason every time: “It’s safer.”
4. The length of time someone had been living in Amsterdam did not seem to make much difference as to whether someone wore a helmet or not. After a few months (or even weeks) recent arrivals were just as likely to be wearing a helmet as those who had been in Amsterdam for longer periods of time.

Now I tell this little story—this small fly of personal experience now preserved in the amber of print—because although the issue is perhaps a small one (except, of course, for those Americans

³ I refer to these physical features of the city and social norms as *habitus* because they are part of the taken-for-granted infrastructure of Dutch society that shapes the meaning of experience (Bourdieu, 1977).

⁴ The study is, as yet, unpublished, but the convenience sample consisted primarily of other parents at my children’s school.

unfortunate enough to be in a bicycle accident while in Amsterdam), it illustrates a central dilemma in the study of learning:

If all the Americans living in Amsterdam wore helmets in the United States, why did some of them stop wearing helmets in Amsterdam while others continued to wear them?

Golding v Defoe

The wearing of bicycle helmets is a cultural practice: one that is a part of the culture in much of the United States and part of the culture in very little of Amsterdam. This means that the question of who does and does not wear a helmet is a matter of *enculturation*. By this I mean, following Brown, Collins, and Duguid (1989), that wearing (or not wearing) bicycle helmets is one of the ways in which “people, consciously or unconsciously, adopt the behavior and belief systems of new social groups” (p. 34). In other words, we have to ask why some Americans adopted the Dutch behavior and belief system regarding bicycle riding, while others did not.

Wertsch (1998) discusses this question of the adoption (or lack of adoption) of cultural practices in terms (following Bakhtin, 2002) of *resistance* and *appropriation*, where, in this case, Americans who continue to wear helmets are resisting the dominant local culture and those who follow local practice are appropriating the local culture. But this theme of selective enculturation is quite universal. Consider, for example, two well-known novels: William Golding’s (1954) *Lord of the Flies*, and Daniel Defoe’s (1996) *The Life and Strange Surprising Adventures of Robinson Crusoe of York, Mariner: Who lived Eight and Twenty Years, all alone in an un-inhabited Island on the coast of America, near the Mouth of the Great River of Oroonoke; Having been cast on Shore by Shipwreck, where-in all the Men perished but himself. With An Account how he was at last as strangely deliver’d by Pyrates. Written by Himself.*

In *Lord of the Flies*, a group of British schoolboys are marooned on an island. They quickly abandon the social norms and conventions of British society, and, led by the former choirboy Jack Merridew, engage in ritual torture and human sacrifice. Robinson Crusoe, in contrast, is a British sailor who is similarly shipwrecked. But unlike Merridew, Crusoe recreates in the wilderness a model of his

homeland (complete with his famous manservant, Friday) and defeats a tribe of native cannibals in battle before being rescued.

Leaving aside the obvious political overtones of these works (both imperialist and anti-imperialist), these novels both explore the stability of enculturation in the face of new social and material conditions. That is, they ask if, once we have become enculturated, practices can be reenacted in another setting, or whether they can only be expressed within the original cultural surround. Once we have learned to function in a particular context, can we recreate the conditions we need to continue using those practices elsewhere? Or when we leave the supporting structures of the original cultural context, do we lose the ability to enact the practices that came with it? Are we, in the end, more like Robinson Crusoe or the Lord of the Flies?

The answer is, of course, both... and neither. Or rather, that we are sometimes one and sometimes the other. Developing a theory that explains under what conditions we are like the British sailor Crusoe or the British schoolboy Merridew is central to understanding the nature of learning, and thus of learning in games and other technological environments.

Communities of Dysentery

To illustrate why Robinson Crusoe is so important to the science of learning, permit me to use one more piece of personal experience, this time from being a Peace Corps Volunteer in a country which I shall call “Nepal”.⁵

In the United States, all of the water that comes from a tap (in a sink, drinking fountain, or bathtub, for example) has been treated to remove parasites such as amoebae, bacteria, and other threats to human health and digestion. In Nepal, quite the reverse is true: there is, for all intents and purposes, no treated drinking water anywhere. As a result, part of the intensive training in Nepali culture given to Peace Corps Volunteers includes water purification practices that are the norm neither in the United States nor in Nepal.

⁵ See note 1 above for the relationship between the country I describe and the country of Nepal in Asia.

Acquiring those practices means acquiring a new Discourse, in the sense that Gee (2001) describes: a new way of “talking, listening, writing, reading, acting, interacting, believing, valuing, and feeling (and using various objects, symbols, images, tools, and technologies)” (p. 719). So, for example, it requires developing new skills, such as learning how to ask that food be served on a plate that has been dried with a towel, or learning the proper way to use iodine to purify a bottle of water.⁶ It means developing a fairly elaborate repertoire of cultural knowledge, including when fruit or eggs might be safe to eat, or under what conditions the rag being used to wipe the plate might be more contaminated than the water itself. It means acquiring the values that let you protect your personal hygiene in a way that is culturally sensitive. It means developing a new decision-making process regarding food and drink, a new set of concerns and criteria: a more complex calculus of hunger and thirst than either Americans or Nepalese practice. And the cumulative acquisition of this Discourse requires developing a new identity that is neither quite American nor quite Nepali. Anyone who doubts that Peace Corps Volunteers have a shared culture in this regard need only listen to the endless and intimate discussions volunteers have of the state of their intestines whenever they gather.

In this sense, acquiring the Peace Corp Volunteers’ Discourse of Water Purification in Nepal means becoming part of a *community of practice*. Lave and Wenger (1991) describe a community of practice as a group of individuals who share a repertoire of knowledge about and ways of addressing similar (often shared) problems and purposes. A community of practice shares a common body of knowledge and set of skills, but also a system of values that determines when and how those skills and that knowledge should be employed, and a set of processes through which such decisions are made. And, of course, such a community also has a shared identity that is manifest in both overt markers (such as the carrying of a water bottle or wearing of a bicycle helmet) and the shared enactment of the skills, values,

⁶ For those planning travel to Nepal, one way to ask for a dry plate is: “Plet rumalle putchepaachi khaana raknus.” But although that phrase may seem simple, actually getting a dry plate requires a far more complex demonstration of skills, usually including various forms of gesture and mime. The concern about the dryness of the plate is that if the plate is only washed clean, drops of contaminated water may remain on it. By contrast, using iodine to purify water requires only waiting 20 minutes, but you also have to remember to turn the bottle over and let some of the iodized water run over the screwtop, or the water trapped there can still contain parasites.

and decision-making processes. In fact, acquiring any Discourse involves becoming part of a community of practice, in the sense that a Discourse is the manifestation of a culture—a way of acting and being in the world—and a community is a group of people with a shared culture (Rohde & Shaffer, 2004).

I raise Peace Corps' attempt to bring volunteers into a community of practice regarding water purification because it illustrates the dilemma of learning environments more generally. Volunteers go through a scaffolded process of learning to purify water. They are, quite literally, newcomers, who move from easier and more peripheral participation in the practice, such as treating water with iodine, to more central and challenging ones: determining what food is safe to eat under which conditions, for example, or securing a clean plate from which to eat. By the end of training, all the volunteers are more or less a part of this community of shared practices. Then, volunteers head to their posts, where for the most part they are the only member of that shared community for miles around. They are surrounded by a different culture, with very different practices regarding the cleanliness of water.⁷

While I was in Nepal, an informal sampling of my fellow volunteers showed that some continued to treat their water at post, and some did not. And some of those who did not treat their water when they were at post continued to treat their water when they were around other volunteers in Kathmandu.⁸ Some volunteers, in other words, enacted the practices of the community of volunteers without the original cultural surround while others adopted the practices of their new surrounding context. Some were Crusoe and some were Merridew—and in saying this I want to be clear that I mean in no way to cast aspersions on Nepali culture, since Americans do not ordinarily purify their own water individually either. Rather the question is about the persistence of the use of a particular configuration of skills, knowledge, values, and decision-making processes—and thus also the display of identity—from one community when the original setting of practice is altered.

⁷ In many places in Nepal, the main criterion for the purity of water is the speed with which it moves. That is, flowing water is cleaner than standing water, no matter what its original source. So water flowing in a strong stream or tap is regarded as clean.

⁸ And not surprisingly, while everyone got parasites sometimes, volunteers who did treat their water got fewer serious parasites and got parasites less often than those who did not.

In what follows, I will argue that one productive way to think about conditions under which enculturation persists across context is by looking at the extent to which a person has become proficient in a Discourse within its original context. That is, we can understand the extent to which learning transfers by analyzing participation in (and thus enculturation into) a community of practice.

One Productive Way, Part I: Framing a Discourse

I want to draw particular attention to one key phrase in the preceding paragraph: “one productive way.” Quite specifically, I want to point out that I am offering this idea—that we can examine the extent to which a person has become enculturated into a community of practice and use it to understand whether they can mobilize the practices from that community in new situations—not as “the” solution to the problem of transfer, but as an example of an approach to the study of learning that draws on both individual and sociocultural perspectives. The approach I outline here, in effect, describes two theoretical constructs, each of which I have written about in more detail elsewhere, and brings them together to suggest a way of conceptualizing (and describing) the broader class of theories that the field of learning needs to develop.

The first, and more fundamental construct is an *epistemic frame*, which I have written about elsewhere as “the combination—linked and interrelated—of values, knowledge, skills, epistemology, and identity” that people have as part of a particular community of practice (Shaffer, 2007, p. 160). The term draws on Papert’s (1980) use of the concept of *epistemology* as a *particular way of thinking*, and on Dewey’s (1933) idea that *knowing is doing*, to define epistemology as a community’s norms for making decisions and justifying actions. The concept of a *frame* comes from Goffman (1974; see also Tannen, 1993), who argued that any activity is interpreted in terms of a frame: the organizational rules and premises, partly existing in the minds of participants and partly in the structure of the activity itself, that shape the perceptions of those involved. That is, a frame is something that is simultaneously individual and social: a set of norms and practices through which experiences are interpreted.

The terminology of epistemic frame is useful because one can think of such frames like a pair of glasses that colors the world in particular ways. They highlight some things as important and relevant and hide or elide other aspects of experience; in this sense, they are a description of Goodwin's (1994) concept of *professional vision*. And, of course, putting on such a frame marks one as the kind of person who sees the world in a particular way. At the same time, it is useful to describe such frames as "epistemic" because they constitute a way of knowing, and suggest that the practices of one community are no more or less fundamental than those of another. Thus, there is no reason, *a priori*, to privilege the educational value of Discourses of school subjects over extant practices in the world—or over the practices that are employed in a game. All communities have ways of knowing, and the relevant question is: How useful are those practices to solve problems and attain goals that matter to someone or to some group of people?

The concepts of epistemic frame and Discourse obviously share a lot in common, in the sense that "the combination—linked and interrelated—of values, knowledge, skills, epistemology, and identity" that people have as part of a particular community of practice can be seen as a more specific description of a way of "talking, listening, writing, reading, acting, interacting, believing, valuing, and feeling (and using various objects, symbols, images, tools, and technologies)." In this sense it can be quite useful to think of an epistemic frame as the grammar of a Discourse: a more formal description of the configuration of elements of a Discourse exhibited by members of a particular community. However, there are at least four properties of an epistemic frame that are distinct from the more general notion of a Discourse—although in making this comparison, I want to emphasize that one is not necessarily better than the other; rather, their different features are useful for different kinds of analyses.

Three of these distinctions are important to note, but not necessarily central to the utility of epistemic frames as a theoretical construct. First, describing the epistemic frame of a community requires specifying some set of elements of the Discourse of the community as being critical to the community's way of seeing and acting in the world. These are usually described as skills, knowledge, identity, values, and/or epistemology, but those categories are clearly guidelines rather than a rigid formula. There may be

settings in which only some of these elements are significant, or analytical circumstances where only some can be observed. The choice and definition of frame elements can come *a priori* from some theoretical or empirical analysis—that is, it can be an etic structure imposed on the data by the researcher—but more often (and more effectively) they arise from some ethnographic or other grounded examination of the community in action. In the same vein, the claim could never be that this listing of discourse elements is exhaustive, only that it is sufficient for the analysis at hand. Further, these discourse elements are often speech acts, but need not be limited to verbal utterances. As in a Discourse, doing is a form of speaking—that is, trafficking in meaningful symbols—just as speaking is a form of doing.

A second distinction between an epistemic frame and a Discourse is that an epistemic frame is always and explicitly a construct that is both internal and external: that is, a property of the community and of an individual, to the extent that an individual has appropriated or internalized the Discourse of the community. And finally, epistemic frames take epistemology as a central analytical focus, rather than identity, which is the organizing structure of a Discourse.

I suggest that these first three distinctions are worth noting but are not central because, in fact, these are shared properties of both a Discourse and an epistemic frame, and the differences here are matters of emphasis rather than inclusion or exclusion. Like an epistemic frame, a Discourse similarly is both a property of a community and something that can be appropriated by an individual. A Discourse naturally includes both forms of epistemology and of identity. And any Discourse analysis requires focusing on some key discourse practices and not others. In practice, these elements of the two constructs are highlighted differently, but this is a pragmatic rather than an ontological distinction. Pragmatic differences in emphasis can be important, and often sufficient to merit the introduction of a new theoretical construct. In this case, however, there is a more important distinction that marks epistemic frames as a distinct hypothesis about the nature of learning and enculturation.

The critical distinction of epistemic frame theory is its explicit—indeed, central—focus on the *linkages between frame elements*.

One Productive Way, Part II: Forming a Frame

What it means for two frame elements—that is, two discourse markers within a Discourse, or two sign tokens within a symbolic system—to be “linked” depends, of course, upon what the frame elements are. To take a simple example, the “skill of asking for a dry plate in Nepali” is linked to “knowledge of when the towel that will be used to wipe the plate is likely to be dirtier than the water it is wiping off” in the sense that experts only use the skill after taking this particular piece of knowledge into account. Thus, we might describe the nature of the linkage as “the knowledge mediating the application of the skill.” On the other hand, the decision to politely refuse a cup of tea served by a government official because the glass was not dried first could easily require caring about the cultural sensitivities of the moment, the ability to explain the decision clearly but respectfully, and the willingness to assert one’s identity as a quasi-outsider; that is, the decision-making process would depend on holding a particular value, having the skill to act in accord with it, and holding a particular view of oneself.

I will come back to the nature of the links in a moment. But first I would like to focus on the question of linkage itself. Strictly speaking, because discourse is a cultural artifact, elements of discourse always invoke some combination of skills, knowledge, values, and decision-making processes—as well as the identity or identities associated with them. That is, anything we say or do marks us as the kind of person who makes some set of assumptions, interprets a situation in particular ways, and decides on certain kinds of actions. The process of discourse analysis, writ large, is an attempt to understand how some action (or sequence of actions) is interpreted by the participants in their original context (Wood & Kroger, 2000). This necessarily means identifying the ways of knowing, doing, thinking, and being that the action signifies. Put another way, in practice discourse elements are always interpreted in terms of a larger Discourse or Discourses. Any display of skills invokes a particular identity. Any interpretation of a situation through the lens of a particular value requires knowledge of the context and its relevant features. Any decision is conditioned on the skills, knowledge, values and identity of the community.

Thus one might argue that it is impossible, in principle, to distinguish frame elements in isolation. Indeed, what would it mean to use a particular skill “on its own,” independent of any relevant knowledge,

values or identity? But this is the wrong question, because although skills are always linked to some form of knowledge, values, identity, and epistemology (and each of the other elements are, in turn, associated with all the others), they are *not always linked to the same ones*, or in the same ways. The purpose in marking certain frame elements as significant to a community of practice is to understand *how* they become linked in the ways characteristic of expertise within the community. Put another way, if we assume *a priori* that using some skill from a community is an assertion of identity in that community (for example), we are assuming the very thing that from the perspective of learning we are trying to find out: that is, how someone comes to see that skill as part of a particular identity in the first place.

The careful reader—and particularly the careful reader familiar with Discourse analysis—will no doubt recognize at this point that the concept of links in an epistemic frame plays an analogous role to the idea of *cultural models* in a Discourse (D'Andrade & Strauss, 1992; Gee, 1999; Holland & Quinn, 1987). Cultural models are memories of prototypical events in a Discourse, which could be real events experienced by an individual or stories told by or about events in the experience of members of the community. These prototypical event memories create images of expected patterns of behavior within the community. That is, they do the work of linking elements of the epistemic frame of the Discourse to one another in particular configurations: they show someone what choices are appropriate within the context of the community. In so doing, they provide a mechanism by which the norms of the community are enacted by individuals—and also a hypothesis about how the process of enculturation unfolds. Indeed, much of the work of Discourse analysis is the examination of moment-by-moment interactions to understand the cultural models at work, and thus the structure of the Discourses being employed. Cultural models are a “tool of inquiry” that guide the asking (and answering) of important questions about the nature of action in a community.

In a similar way, we can conceive of the identification of frame elements and the structure of the links between them as a tool of inquiry. The concept of an epistemic frame is agnostic as to the mechanisms by which those links are created within individuals or within a community. Some are surely encoded in prototypical stories (Nelson, 1996). Others in direct event memories (Donald, 1991; Dreyfus

& Dreyfus, 1986). Some are reduced to heuristic “rules of thumb,” or sayings or maxims. Some are enshrined in written or unwritten rules (Pinker, 1997). Some are maintained in the structure of the physical environment within which a community operates (Pea, 1993). But whatever the mechanism by which links between frame elements are created or maintained, modeling the structure of those links is a useful tool of sociological and psychological inquiry—and thus an important tool in the kit of the science of learning.

In saying that the structure of links in an epistemic frame models the structure of a Discourse, I want to emphasize the word *model*. I am not suggesting that a set of actual links exists, either within the mind of an individual or within the Discourse community, physically or even psychologically.⁹ The organization of links in an epistemic frame is a tool by which the cumulative structure of the norms of a Discourse can be described. That is, it is a technique for providing a *mechanic grip* on what Pickering (1995) referred to as “the mangle of practice.” Epistemic frames are a framework for describing the complex unfolding of events in the real world.

We can create a model of the epistemic frame of the Discourse (or more properly, of some community of practice that operates with a Discourse); but we can also create a model of the epistemic frame of any individual within that community.

And the similarity between the two provides a tool with which learning-as-enculturation can be analyzed.

One Productive Way, Part III: Fitting a Frame

The second already-existing theoretical construct that I want to introduce here builds on the first. *Epistemic Network Analysis* (Shaffer, et al., 2009), or ENA, is a technique for quantifying—and thus analyzing—an epistemic frame. The central insight of ENA borrows, broadly speaking, from Minsky’s (1988) notion of a *society of mind*: that is, of the mind as a set of interacting agents. Metaphorically, we

⁹ And even if they did, we are not currently able to observe them directly. Any representation of a Discourse is always a model, although recent advances in neuroscience suggest that perhaps a model of neural functioning could be similarly conceptualized as a set of links.

can think of the different elements of an epistemic frame (the skills, knowledge, identities, values, and epistemology of some community of practice) as agents in their own right, who are connected to one another through different patterns of association. The pattern of their interaction can thus be mapped using existing tools from Social Network Analysis. As one account of ENA (Shaffer, et al., 2009, pp.40-1) explains:

[U]sing social network analysis we might examine the relationships among a group of people meeting for the first time at a cocktail party. To do so, we might take a photograph of the party at appropriate intervals—perhaps every time the music changes, every time someone orders a drink, or at a fixed time interval, depending on the nature of the party and our hypotheses about the people and relationships involved and the social forces at work. If we make the assumption that people who spend more time in the same conversational group develop a closer relationship over time, we can quantify the social network being developed at the party by summing, for each pair of party-goers, the number of times they are recorded in the same conversational group during the party. Once quantified in this way, social network analysis provides a wide range of analytical tools for investigating the properties and processes at work in the social relationships of the party.

If we think of the “party” not as a collection of individuals in a room, but rather as a collection of elements from the epistemic frame... then we can use the same analytical tools to conduct epistemic network analysis.

I want to draw out two important points from this brief summary of the relationship between Social Network Analysis and ENA.¹⁰

The first is rather evident, but nonetheless significant: Social Network Analysis is a simplification of the complex dynamics of a social situation. As in any analysis of data, it is based on an assertion about the relationship between the observed data and some phenomenon of interest. In the hypothetical case above, there is an inference made that “being in the same conversational group” tells us something about a developing relationship. It is possible to posit more complex inferences (perhaps based on more

¹⁰ A third point worth making, although not central to the argument here, is that early empirical work with ENA suggests that epistemic networks, in fact, have mathematical properties that differ somewhat from the networks analyzed by traditional social network analysis. (Social networks tend to contain large numbers of nodes with sparse connections. Epistemic networks tend to have fewer nodes and are more densely connected.) Thus, although the same principles apply, in practice the use of ENA requires slightly different algorithms.

complex observations), including all manner of possible relationships between people using Social Network Analysis. The subsequent mathematics may become more complex (depending on the relationships hypothesized), but the principle remains the same: as in any model, Social Network Analysis simplifies the observed situation, and the test of whether the simplification is appropriate lies in the nature of the question being asked and the subsequent results obtained.

In a similar way, then, ENA is a simplified representation of the structure of the practices of a community—although one has to quickly admit that a cultural model (and any other model, qualitative or quantitative) is as well. I will provide an example of ENA being used on real data in a moment, but most ENA analyses use a relatively small number of frame elements (6-20), and take “co-occurrence in discourse” as a proxy for “linkage.” This is a simplification, but one of practical convenience, not theoretical necessity. The same technique could account for more elements, different kinds of linkages (including directional links), and more complex decision-rules for determining linkage.

The second issue to draw out is actually a point of omission: namely, that the account above makes no more mention of an epistemic frame being the property of a single individual than a social network is—although, of course, an individual has collections of both epistemic frames and social networks.¹¹ In fact, the description above does not even suggest that an epistemic frame is uniquely the property of people, in the sense that the artifacts with which people interact could similarly provide links between frame elements. In this sense, a theory of epistemic frames mimics Pask’s (1975) conversation theory or a theory of distributed mind (Shaffer & Clinton, 2006), where persons and machines contribute to the co-creation of systems of meaning. This is a central point, in the sense that one can use ENA to describe, quantify, and compare the epistemic frame exhibited by various combinations of actors in the system (including individuals and groups of persons and technologies).

What’s in the game?

¹¹ To be completely fair, the original text at this point actually *does* refer to the frame elements as belonging to an individual. In the original presentation this was for clarity of the description. Later in the same treatment of ENA, there is a discussion of quantifying the frames of collections of individuals.

By way of example, I'd like to present empirical work that, again, has been reported in more detail elsewhere (Hatfield & Shaffer, 2006, 2008; Shaffer, 2007): this time, a study by my colleague David Hatfield on an educational game he developed with Alecia Magnifico. In the game, called *science.net*, players become science journalists and produce an online science newsmagazine. *Science.net* is an example of an *epistemic game*, in the sense that it is a game designed to develop the epistemic frame of a particular socially-valued community of practice (Shaffer, 2007). It recreates in game form what cub reporters experience in their professional training.

This particular claim—that the game recreates what cub reporters experience in their professional training—provides an excellent example of the kind of analysis I have been describing. It asserts that the practices in one community (the game) model the practices in another (the professional practicum through which journalists are trained). If ENA can be used to empirically test Hatfield's claim about the structure of the game, then it provides a theoretically-grounded metric for comparing Discourses—and thus, a means to assess enculturation.

The computer game *science.net* was designed based on two ethnographic studies of journalism capstone courses, in which undergraduate journalism students were apprenticed into the community of practice in a class designed to mimic key practices of a real newsroom. One of those key practices was *copyediting*: line-by-line corrections to a news story that are copious, detailed, and blunt. The copyedits in *science.net*—a game designed for middle school students—were somewhat less blunt, but they were supposed to retain the other salient features of the original. Graduate students working on the project (none of whom were themselves journalists) were trained to give feedback in the game, and Hatfield's question was whether these mentors in the game setting were modeling the same practices as the journalists who were mentors in the professional training.

To assess whether the mentors in the game and in the journalism practicum on which it was modeled were engaging in the same practices, Hatfield created a coding scheme from the ethnographic data, cataloging some 18 skills, forms of knowledge, values, identity markers, and epistemological stances: everything from “knowledge of the reader” to the “skill of narrative storytelling,” the “value of

making a difference” and the “epistemology of accuracy.”¹² Each of the copyedits from each of the stories filed was coded as to whether the mentor referenced each of the 18 frame elements, as in the example below:

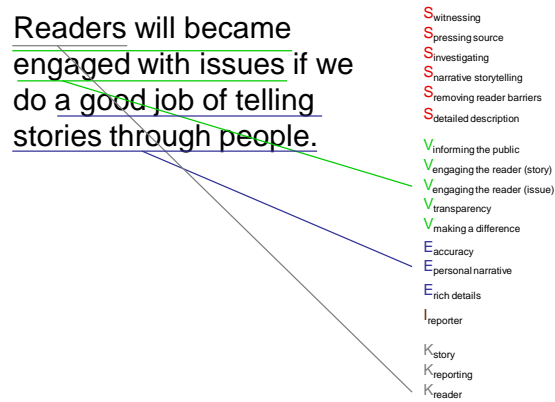


Figure 1: A copy edit coded with elements of the epistemic frame of a journalism practicum

Hatfield then used ENA to construct a model of the epistemic frame the journalist was enacting. He treated the individual copy edits across all stories as if they were snapshots of people (that is, frame elements) interacting at a party (that is, being linked in discourse), to model the *enacted epistemic frame* of the journalist: that is, to create a model of the journalist’s epistemic frame as reflected in the copy edits he made. This model of the Discourse of journalism copy editing served as a *reference frame*, since it was a model of the practices (or part of the practices) that the game *science.net* was trying to replicate.¹³

Hatfield then used the same coding scheme on all of the copy edits from the game *science.net*, and as a point of comparison, also on copy edits made on papers in a graduate psychology course, which were of similar style and form to the journalism edits. He then computed the epistemic frame for each of

¹² There are, perhaps, those who might wonder why accuracy is coded as a form of epistemology and not a kind of value. In fact, it could be coded as either, since, as above, the categories of skill, knowledge, and so forth are guidelines only. In the case of journalism, however, epistemology may be the better choice for accuracy, in the sense that the practice of journalism is inextricably bound up with “scientific” concepts such as accuracy and verification (Kovach and Rosenstiel 2001), or as some suggest, the fundamental commandment of journalism is: “Thou shall not fabricate” (Gardner, Csikszentmihalyi et al. 2001).

¹³ More details on the mathematics of ENA can be found in (Shaffer, et al., 2009). More details on Hatfield’s specific methods can be found in (Hatfield & Shaffer, 2010).

the stories in each of the venues: Journalism Practicum (JP), *science.net* game (G), and Psychology Course (PC). Using a technique called *relative centrality distance* (which is described in more detail in his work) he computed the “distance” between the reference frame and each story in the data set.

I do something of an injustice to Hatfield’s study in describing it so briefly here, but suffice it to say that he triangulated his findings using both further qualitative investigation of the different stories, as well as further quantitative analysis of the frames he constructed, which both confirmed his initial findings, and led to some refinements in the methods for computing the distances between frames. But the overall result was clear and straightforward, and can be seen in the graph below, which shows that the distances from the reference frame to the frames from the game stories (G) were not significantly different than the distances from the reference frame to the stories in the journalism practicum (JP). The papers in the psychology course (PC), on the other hand, were significantly further away from the reference frame:

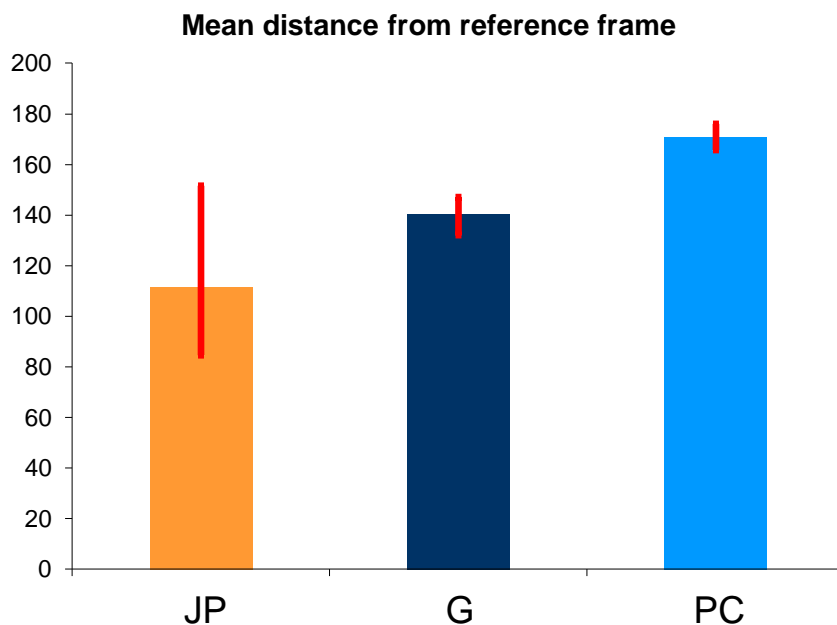


Figure 2: Graph showing mean distance from reference frame for stories in three conditions: journalism practicum (JP), game (G), and psychology course (PC)

The differences are even more dramatic when the frames are actually visualized, as in this projection from the high-dimensional space in which frames are analyzed into a 2-dimensional plane, where the blue network represents the journalism practicum, the green the game, and the red the psychology course:




Figure 3: A projection of the links between frame elements in a journalism practicum (blue), the game science.net (green), and a psychology course (red)

In both cases, the results are the same. The model suggests that the copy editing practices used by mentors in the game *science.net* working with middle school-aged players reflected the same practices used in the training of journalists during one capstone course. In other words, Hatfield used ENA to show that the epistemic game *science.net* was... well, epistemic.

Another colleague, Elizabeth Bagley, has taken a similar approach to looking at the epistemic game *Urban Science*, which recreates the training practices of urban planners (Bagley & Shaffer, 2009, 2010). Like Hatfield, she studied the training of professionals to model a reference frame for the game, this time from the practice of planning rather than journalism. In her study (again, reported in more detail elsewhere), she used ENA to show how, from one week to the next during the practicum, the enacted frame of novice planners converged with the frame being modeled by their mentor. That is, the Discourse

of the novices became more like the Discourse of the profession. Interestingly, she was able to show that this happened despite the fact that novices had every appearance of trying to resist the perspective the mentor was modeling.

ENA, then, and with it the concept of epistemic frames, appears to be a useful tool in assessing learning as the reflection of an individual's enculturation in a community of practice.

ENA and Crusoe

To explain how ENA might provide such a model, let us imagine, for a moment, superimposing two epistemic networks, one upon another. Like social networks, epistemic networks can be visualized in a number of ways, including as a *network graph*, similar to those in Figure 2, where frame elements are represented by nodes and links are represented by lines, with strength of association between two elements represented by the length and/or thickness of the line that connects their associated nodes in the graph. Networks can also be represented by adjacency matrices, where each cell in the matrix quantifies the association between two elements, or even as points in a very high dimensional space, where points in the same neighborhood represent networks of similar configuration.

There are thus several ways of imagining the superposition of networks, and perhaps Figure 3 below is as good as any:




Figure 3: Superposition of two epistemic networks. The set of all nodes is a reference frame from some community of practice. The orange (solid) nodes and orange (solid) lines represent the epistemic frame of a novice at some point in time

Let us assume, again for the sake of simplicity, that the full set of nodes and links in Figure 3 is the reference frame for some community of practice: that is, it is the set of frame elements and links between them that experts in the community exhibit when taking action. A novice's epistemic frame might be represented by the set of orange (solid) nodes and links: it is a subset of the reference frame, including only some of the nodes and links from the original. There could also be extraneous or over-weighted links, although this diagram does not show that case.

The key point is that the orange frame alone does not appear to be “stable.” Key connecting nodes and links that are part of the reference frame are missing. Some of the frame elements that are present (the node in the upper right, for example) are not connected to other elements. Thus, it is easy to believe that our Orange novice would not be able to enact practices of the community in an unfamiliar setting with much success. Too many of the conditions for successful practice are missing in his or her epistemic network. He or she is not sufficiently enculturated because he or she has not (yet) internalized enough of the frame of the community. On the other hand, some links are in place, so with appropriate support from the setting—whether from peers, mentors, or from the material conditions of practice—it appears that the Orange novice could participate in the practices of the community, even if perhaps at this stage still only peripherally.

As the Orange novice progressed through his or her enculturation, more of the skills, knowledge, values, epistemology, and identity of the community would become part of his or her epistemic frame. If the reproductive practices of the community were effective, the appropriate kinds of links would be formed between them. At some point, the configuration would become stable, in the sense that it was able to operate with minimal support from the setting in which our Orange (now) expert finds him- or herself. If we take the liberty of describing the reference frame at work here as the Peace Corp Volunteers' Discourse of Water Purification in Nepal, we might expect that the Orange links and nodes represent the epistemic frame of a volunteer who treated his or her water when around other volunteers in Kathmandu, but not when he or she was alone at post. If the reference frame is, in some broad sense, “British Culture,” our Orange novice might be Jack Merridew.

MSA and games

There is, of course, much to be said (some of which, as above, has already been said elsewhere) about the details of CodingConstructingQuantifyingAnalyzingVisualizing epistemic frames using ENA. But I began this meditation on the nature of learning in games with a claim that my intent was to describe a theory of learning that is simultaneously individual and situated and to suggest that it provides an example of the kind of theory that we need to develop to study learning in games. My argument is thus that the specific details of epistemic frames and ENA are less important (in this context) for their own sake, than in the way that they exemplify a class of theory that I propose to call *Models of Situated Action* (MSAs). MSAs are representations of participation in a practice that, like epistemic frames, are amenable to comparison between and across individuals, groups, and situations. They are post-schism theories that are both symbolic and sociocultural.

The term MSA is meant to recognize that any such theory is necessarily a *model* of action, in the sense that it does not necessarily propose that a specific physical or psychological structure exists in the minds of individuals or in the world around them. Rather, it proposes that some construct is useful for analyzing situated action. Further, any such theory must make a link between individuals, groups of individuals, and the context around them: that is, it must be *situated* in the sense that it provides a formal language that accounts for the way in which individuals take certain actions in certain situations. And finally, any such theory must focus on *action* in the sense that it can provide not only descriptive but predictive accounts of the discourse individuals and groups use (including, of course, but not limited to, particular speech acts). Put in simpler terms, any MSA has to be the kind of theory that could account for which Americans will wear bicycle helmets in Amsterdam and which Peace Corps Volunteers will treat their water at post in Nepal.

Thus, epistemic frames are but one example of an MSA. Other researchers are similarly working on models—often, like ENA, quantitative and computational—of situated action. For example, Shute and her colleagues have been using Bayesian Networks to model player choices in a game (Shute, Ventura, Bauer, & Zapata-Rivera, 2009). Mislavy has proposed Evidence Centered Assessment Design as a

framework for constructing, describing, and validating assessments of complex practices (Mislevy & Riconscente, 2006). Hickey has looked at models that link “proximal” situated data with “distal” standardized assessments to understand the efficacy of learning activities, particularly in the setting of educational games (Hickey, A., & Jameson, 2009). Suthers and colleagues have looked at quantifying the distance between forms of discourse (Constantino-Gonzalez, Suthers, & Escamilla de los Santos, 2003). Each of these is a different quantitative and theoretical framework for looking at situated action. And each is thus, I claim, a kind of MSA.

Which brings us to a key question for the future of learning : Why are MSAs so important for understanding learning in games?

My answer starts with the recognition that a game is fundamentally a culture (Gee, 2003; Shaffer, 2007; Steinkuehler, 2006). Playing a game—and certainly playing a game well—requires a certain way of thinking about and being in the world: a set of skills, knowledge, values, ways of making decisions, and ways of seeing oneself and being seen by others that “work” in the game. The players of any game form a community of practice, complete with newcomers and oldtimers, with rites of passage and increasingly central participation. Thus, one powerful way to analyze the experience of players of a game is through some model of enculturation. That is, using an MSA.

A second aspect of computer games that makes MSAs such useful conceptual tools is that, compared to other interactive settings, games record a relatively large subset of activity. No game can record everything that happens with, by, and to players, since some actions that matter in the game always take place outside of the computer system, and even away from the game interface itself. For example, players talk with one another or with onlookers while playing, multiple players design or play the same character, and players talk about games at the schoolyard or water cooler. But these issues notwithstanding, games record an enormous amount of discourse in the form of players’ choices, actions, and even (in good games) dialog with other players and non-player characters. There is, in fact, an overwhelming stream of cultural information available about players in a game. Too much to be managed, in most cases, without some formal model of participation. That is, without an MSA.

Third, although some games are more “open-ended” than others—a multiplayer online role playing game has more possible interactive contexts than a game like Tetris, for example—the cultural context of a game (particularly a computer game) is bounded in predictable ways. That is, relative to social interactions writ large, the structure of the game situation constrains the set of discourse moves that players may undertake. The game provides a restricted semiotic system (Gee, 2003) which, while not finite, is more formally tractable than discourse in many interactive settings. The organization of the game into set levels, scenes, non-player characters, and other structural features—not to mention formal rules through which these game elements operate—provides a degree of standardization across contexts for different players. The game system consists of arenas in which particular kinds of discourse are likely to recur.

Games worlds are thus fundamentally cultural in nature. They provide a rich record of players’ discourse. And they have the potential to provide standardized elements of situations in which players interact. All of these features, in turn, suggest that MSAs are particularly powerful—and perhaps essential—tools for understanding learning in games. Contrariwise, games are also particularly powerful—and if not essential, certainly quite valuable—contexts in which to develop and understand MSAs and learning through enculturation in practices more generally.

Consider, for example, one of the players in Hatfield’s *science.net* game. Let’s call her Sarah. At any point in the game, we could, in theory, construct the epistemic frame of journalism that Sarah has enacted during the game: that is, we can quantify the structure of links between frame elements in Sarah’s recorded discourse. We could also construct the frame enacted by the group or groups with which Sarah has worked. In fact, we could construct the frame enacted by all of the elements of the game (including other players) that have been visible to Sarah during the game. This would constitute her *observed frame*, which would represent all of the discourse within her *horizon of observation* in the game (E Hutchins, 1995). As discussed above, we could, from studies of expert performance—either in real newsrooms, journalism practica, or by having journalists play the game—construct a *reference frame* (or reference frames) that represent an ideal set or sets of expert performance.

These frames would constitute an MSA, and using ENA these different frames could be compared quantitatively. Thus, we might ask, for example, whether during the game Sarah's frame converged with one of the reference frames of expert journalists. We could ask whether her frame was more likely to converge toward the reference frame if her observed frame also converged to the reference frame. That is, we could quantify one of the most elusive concepts in education, *opportunity to learn* (Darling-Hammond, 2006), by examining the progress of players relative to the qualities of the context in which they were embedded.

We could determine the trajectory over time of Sarah's frame development in the game, and compare it to successful trajectories of experts, or of other players. We could determine whether certain learning progressions are more or less likely to lead to significant enculturation. We could determine which kinds of responses from mentors are most and least effective in promoting enculturation at different points in the game. We could even, with sufficient data, potentially replace the cognitive model in an automated tutoring system with an MSA to produce realistic mentoring in an ill-formed domain like journalism (Shaffer, 2009).¹⁴

To the extent that any of these have been done, can be done, or will be done with ENA, they could also be done with any MSA that similarly provides a model of participation in a community of practice. In other words, MSAs provide a mechanism for assessing enculturation by determining the extent to which—and the particular ways in which—an individual has or has not adopted the discourse practices of a community: the community of players of a game, or a community of practice in the world outside the game. MSAs do so by explicitly connecting events at the sociocultural level (participation in a Discourse) with events at the individual level (actions of a player). They do so by explicitly modeling the relationship between discourse markers (symbols) and the structure of a Discourse (practices).

¹⁴ I use the term ill-formed domain here because problems in journalism—as in many domains of complex practice—cannot be solved using a set of procedures specified in advance. (This is in contrast, for example, to work in a factory, or a game like tic-tac-toe.) Well-formed domains can be represented computationally using cognitive models, and automated tutoring systems have had some success in supporting learning in well-formed domains. The same cannot be said (yet) for ill-formed or open-ended problem spaces (Dreyfus & Dreyfus, 1986; Graesser, et al., 2004).

In this sense, an MSA is a theoretical construct that looks at learning—at transfer in the general sense of how one experience can impact another—from both a sociocultural and symbolic perspective. In so doing, an MSA is a model of the authentic practices of a community in action, providing an assessment of learning over time and *in situ*. A model of learning that has the potential to distinguish Robinson Crusoe from the Lord of the Flies. A model for the kind of learning that happens in games, and that ideally happens in effective learning environments of any kind.

Acknowledgements

This work was funded in part by the Macarthur Foundation and the National Science Foundation through grants REC-0347000, DUE-091934, DRL-0918409, and DRL-0946372. The opinions, findings, and conclusions do not reflect the views of the funding agencies, cooperating institutions, or other individuals. Thanks to these funders, and to David Hatfield, Elizabeth Bagley and the Epistemic Games Research Group, and to Jim Gee and Katie Clinton for their suggestions on the manuscript.

References

- Anderson, J. R. (1993). *Rules of the mind*. Hillsdale, NJ: Erlbaum.
- Anderson, J. R., Reder, L. M., & Simon, H. A. (1996). Situated Learning and Education. *Educational Researcher*, 25(4), 5-11.
- Bagley, E., & Shaffer, D. W. (2009). When people get in the way: Promoting civic thinking through epistemic game play. *International Journal of Gaming and Computer-Mediated Simulations*, 1(1), 36-52.
- Bagley, E., & Shaffer, D. W. (2010). *The epistemography of urban and regional planning 912: Appropriation in the face of resistance*. Paper presented at the ICLS, Chicago, IL.
- Bakhtin, M. M. (2002). *The dialogic imagination : four essays*. Austin: University of Texas Press.
- Barab, S. A., & Plucker, J. A. (2002). Smart people or smart contexts? Cognition, ability, and talent development in an age of situated approaches to knowing and learning. *Educational Psychologist*, 37(3), 165-182.
- Barthes, R. (1982). *Empire of signs* (1st American ed.). New York: Hill and Wang.
- Bourdieu, P. (1977). *Outline of a theory of practice*. Cambridge ; New York: Cambridge University Press.
- Brown, J. S., Collins, A., & Duguid, P. (1989). Situated Cognition and the Culture of Learning *Educational Researcher* 18(1), 32-42.

- Chi, M. T. H. (1997). Quantifying qualitative analyses of verbal data: A practical guide. *Journal of the Learning Sciences*, 6(3), 271-315.
- Cobb, P. (1999). Individual and collective mathematical learning: The case of statistical data analysis. *Mathematical Thinking and Learning*, 1, 5-44.
- Constantino-Gonzalez, M. d. I. A., Suthers, D. D., & Escamilla de los Santos, J. G. (2003). Coaching Web-based Collaborative Learning based on Problem Solution Differences and Participation *International Journal of Artificial Intelligence in Education*, 13(2-4), 263-299.
- D'Andrade, R., & Strauss, C. (Eds.). (1992). *Human Motives and Cultural Models* Cambridge: Cambridge University Press.
- Darling-Hammond, L. (2006). Securing the Right to Learn: Policy and Practice for Powerful Teaching and Learning *Educational Researcher*, Vol. 35, No. 7, 13-24 (2006), 35(7), 13-24.
- Defoe, D. (1996). The life and strange surprizing adventures of Robinson Crusoe, of York, mariner who lived eight and twenty years all alone in an un-inhabited island on the coast of America, near the mouth of the great river of Oroonoke; having been cast on shore by shipwreck, wherein all the men perished but himself. With an account how he was at last as strangely deliver'd by pyrates Available from http://gateway.proquest.com/openurl?ctx_ver=Z39.88-2003&xri:pqil:res_ver=0.2&res_id=xri:ilcs-us&rft_id=xri:ilcs:ecf:Z000001109
- Dewey, J. (1933). *How we think, a restatement of the relation of reflective thinking to the educative process*. Boston, New York [etc.]: D.C. Heath and company.
- Dewey, J. (1938). *Experience and education*. New York: Collier Books.
- Donald, M. (1991). *Origins of the modern mind: Three stages in the evolution of culture and cognition*. Cambridge, MA: Harvard University Press.
- Dreyfus, H. L., & Dreyfus, S. E. (1986). *Mind over machine : the power of human intuition and expertise in the era of the computer*. New York: Free Press.
- Gee, J. P. (1999). *An introduction to discourse analysis : theory and method*. London: Routledge.
- Gee, J. P. (2001). Reading as situated language: A sociocognitive perspective. *Journal of Adolescent & Adult Literacy*, 44(8), 714-725.
- Gee, J. P. (2003). *What video games have to teach us about learning and literacy*. New York: Palgrave Macmillan.
- Goffman, E. (1974). *Frame analysis : an essay on the organization of experience*. New York: Harper & Row.
- Golding, W. (1954). *Lord of the flies, a novel*. London,: Faber and Faber.
- Goodwin, C. (1994). Professional vision. *American Anthropologist*, 96(3), 606-633.
- Graesser, A. C., Lu, S., Jackson, G. T., Mitchell, H., Ventura, M., Olney, A., et al. (2004). AutoTutor: A tutor with dialogue in natural language. *Behavioral Research Methods, Instruments, and Computers*, 36, 180-193.
- Greeno, J. G. (1997). On claims that answer the wrong questions. *Educational Researcher*, 26(1).
- Hatfield, D., & Shaffer, D. W. (2006). *Press play: Designing an epistemic game engine for journalism*. Paper presented at the ICLS, Bloomington, IN.
- Hatfield, D., & Shaffer, D. W. (2008). *Reflection in professional play*. Paper presented at the ICLS, Urecht, Netherlands.
- Hatfield, D., & Shaffer, D. W. (2010). *The epistemography of journalism 335: Complexity in developing journalistic expertise*. Paper presented at the ICLS, Chicago, IL.
- Hickey, D. T., A., I.-G., & Jameson, E. (2009). Designing assessments and assessing designs in virtual educational environments. *Journal of Science Education Technology*, 30, 837-861.
- Holland, D., C., & Quinn, N. (Eds.). (1987). *Cultural Models in Language and Thoughts*. New York: Cambridge University Press.
- Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.
- Hutchins, E. (2008). The role of cultural practices in the emergence of modern human intelligence. *Philosophical Transactions of the Royal Society*, 363(1499), 2011-2019.

- Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge, UK: Cambridge University Press.
- Minsky, M. L. (1988). *The society of mind*. New York: Simon and Schuster.
- Mislevy, R. J., & Riconscente, M. M. (2006). Evidence-centered assessment design: Layers, concepts, and terminology. In S. Downing & T. Haladyna (Eds.), *Handbook of Test Development* (pp. 61-90). Mahwah, NJ: Erlbaum.
- Nelson, K. (1996). *Language in cognitive development: emergence of the mediated mind*. Cambridge, UK: Cambridge University Press.
- Packer, M. (2001). The Problem of Transfer, and the Sociocultural Critique of Schooling. *Journal of the Learning Sciences*, 10(4), 493-514.
- Papert, S. (1980). *Mindstorms: Children, computers, and powerful ideas*. New York: Basic Books.
- Pask, G. (1975). *Conversation, cognition and learning*. Amsterdam: Elsevier.
- Pea, R. (1993). Practices of distributed intelligence and designs for education. In G. Salomon (Ed.), *Distributed cognitions: Psychological and educational considerations*. Cambridge, UK: Cambridge University Press.
- Pickering, A. (1995). *The mangle of practice : time, agency, and science*. Chicago: University of Chicago Press.
- Pinker, S. (1997). *How the mind works*. New York: W.W. Norton & Company.
- Rohde, M., & Shaffer, D. W. (2004). Us, ourselves, and we: Thoughts about social (self-) categorization. *Association for Computing Machinery (ACM) SigGROUP Bulletin*, 24(3), 19-24.
- Sahlins, M. D. (2000). *Culture in practice : selected essays*. New York: Zone Books.
- Seel, N. M. (2001). Epistemology, situated cognition, and mental models: 'Like a bridge over troubled water'. *Instructional Science*, 29(4-5), 403-427.
- Sfard, A. (1998). On Two Metaphors for Learning and the Dangers of Choosing Just One. *Educational Researcher*. *Educational Researcher*, 27(2), 4-13.
- Shaffer, D. W. (2007). *How Computer Games Help Children Learn*. New York: Palgrave.
- Shaffer, D. W. (2009). AutoMentor: Virtual Mentoring and Assessment in Computer Games for STEM Learning. University of Wisconsin-Madison: National Science Foundation.
- Shaffer, D. W., & Clinton, K. A. (2006). Toolsforthought: Reexamining thinking in the digital age. *Mind, Culture, and Activity*, 13(4), 283-300.
- Shaffer, D. W., Hatfield, D., Svarovsky, G., Nash, P., Nulty, A., Bagley, E., et al. (2009). Epistemic Network Analysis: A prototype for 21st Century assessment of learning. *The International Journal of Learning and Media*, 1(2), 33-53.
- Shute, V. J., Ventura, M., Bauer, M. I., & Zapata-Rivera, D. (2009). Melding the power of serious games and embedded assessment to monitor and foster learning: Flow and grow. In U. Ritterfeld, M. Cody & P. Vorderer (Eds.), *Serious games: Mechanisms and effects* (pp. 295-321). Mahwah, NJ: Routledge, Taylor and Francis.
- Steinkuehler, C. A. (2006). Massively multiplayer online videogaming as participation in a Discourse. *Mind, Culture & Activity*, 13(1), 38-52.
- Tannen, D. (1993). *Framing in discourse*. New York: Oxford University Press.
- Wertsch, J. V. (1998). *Mind as action*. New York: Oxford University Press.
- Wood, L. A., & Kroger, R. O. (2000). *Doing discourse analysis : methods for studying action in talk and text*. Thousand Oaks, Calif.: Sage Publications.