## Visualizing Online Social Interaction

A group, according to the American Heritage Dictionary, is an assemblage of persons or objects gathered or located together. In a social context, this usually implies a shared psychological goal or physical state, and in the real world, often both. The set of students in the Sociable Media Group at the MIT Media Laboratory is such a group, sharing the academic goal of studying social interaction in a computer mediated world, as well as inhabiting the same physical location in the SMG common area.

In a society, the group is the atomic element. Indeed, the American Heritage Dictionary defines society only within the context of a group of people. Fundamentally, with only one person, there would be no one else to interact with, socially or otherwise. When we discuss any aspect of a society, then, we must necessarily include at least an implicit reference to the group context within which that aspect has meaning. Certainly there are interesting characteristics of the individual within the society, however all of those characteristics derive their significance from the others within or outside of the group to which that individual belongs.

As social beings in a social world, we learn from a very early age how to interact with groups of people. Especially when dealing with larger groups of people (more than five or ten), much of our understanding of how such groups work and how we can interact with them derives from our viewing of them. Imagine you are walking through a city square on a nice summer day. As you walk, you encounter people strolling and windowshopping. Ahead, you see a large gathering of people, arranged in a circle, who seem to be paying attention to something within the group. You walk up to the edge of the crowd, realizing that there must be something of interest, perhaps a street performer. How do you know that there is someone or someone beyond your view? You have watched the people in the group interact: there are people who are happily watching that thing in the center, people who are chatting to those around them, and people who are drifting into and out of
the crowd at a leisurely pace. All of these visual signals tell you that there is something not only of interest, but also entertaining at the center of the group. If the people ahead of you were not bemused, or were running into and out of the crowd quickly, then your impression of what was going on in that space might be completely different.

Similarly, imagine you were walking through a crowded subway station, with most of those around you walking hurriedly in the opposite direction. Your goal is the set of stairs beyond the crowd. You weave your way in between people ahead of you, taking the circuitous path of least resistance. You reach the staircase fairly quickly, not much slower than you would in an empty station. How were you able to navigate the sea of people so efficiently? You were able to see ahead of you the flow of traffic, and the patterns of conglomeration that you need to avoid, and the paths that seem to offer a lighter density of people.

These situations are ones where the viewing of people is of utmost importance for accomplishing the task at hand. When you are at a party, and people form cliques, you can see at a glance those groups of people that are having the animated, dynamic, and perhaps therefore interesting conversations that you would like to join. You would also easily see dead conversations: those in which participants are looking about in search of another conversation, or are sipping their drinks quietly, barely saying anything to those around them. There are the people watchers, people like yourself who watch the unfolding of the social situations around them, reading the groups and the interaction within them with interest, and perhaps learning about the participants in those groups.

Imagine trying to accomplish any of these tasks without sight. You may still be able, after some time, to hear the groups around you, and sense which are the interesting dynamic ones. But without sound, you would be at a loss to even do that. The visual portrayal of groups in the real world is the most concise and rich encoding of social interaction information that is available to us. We read groups of people like we read a book, assembling collections of objects into a greater whole, gaining new and different information as we our perception moves from small groups to larger ones. Without the
visual dimension, groups become mostly incomprehensible: we may be able to sense some disjoint information about a group, perhaps that person A is talking to person B , but the overall picture, and how each person fits into that picture, is gone.

The Internet, through various communications channels such as email, MUDs and MOOs, Usenet, and IRC, has become a social medium. We as a society have begun to interact with each other in a way similar to the real world. We form groups online, have discussions, and gather socially. Yet though the ideas behind our social interaction online mimic those of real world social interaction, their expression is a bit different. When we interact in real time through online chat, we play out the same social roles as we might at a cocktail party, but without the visual feedback that would allow us to navigate social space with ease.

As previous interfaces have shown, it is not good enough to merely put static pictures of people on a screen and have them talk to each other [refs]. We cannot understand the larger scale behavior of the group through such interfaces. We must create interfaces that are dynamic, and that provide convincing representations of the underlying human interactions taking place.

Creating an interface that provides some of the social cues given by the dynamic image of a crowd in real life should therefore be of utmost importance. Current chat interfaces do not prevent us from interacting socially online, but they do block many of those cues that allow us to understand the structures and dynamics of groups. This paper, therefore attempts to explore how we might begin to create such interfaces. Understanding the social psychology of groups in the real world is of utmost importance in order to understand how such phenomena can work online. Also, a thorough grounding in both the cognitive process of visual perception and the visual display of human forms is necessary if we are to build visualizations of online group interaction that allow us to read the patterns of such behavior from that visualization, and make clear the underlying social interaction.

Only through careful application of the principles learned from real world study, and through understanding if and how social interaction online is both superficially and fundamentally different from online social interaction can we build successful interfaces. This paper doesn't attempt to address all of the problems of social interaction online, but does provide an initial foray into understanding how we can portray online group interaction in a meaningful way.

## Social Interaction in the Real World

In the real world, people form groups and interact with each other effortlessly. We, as social beings, learn early on the proper social etiquette and behaviors that allow us to do so. Part of this process is to learn how groups of people interact so that we can "read" them and interact within them. By "read" I mean the process of watching a group of people and deducing psychological characteristics of the individuals and of the group as a whole, allowing inferences to be made, such as determining the leader of a conversation.

This knowledge, however, is anecdotal. To make this knowledge explicit, we must look to the fields of sociology, sociometry, and social psychology. Essentially, we are looking for a set of measurable qualities that are interesting and useful statistics for describing how a group of people behaves. Unfortunately, much of the work done in this are deals with goal directed groups. Such groups are collections of people that are studied as they undertake the solving of a problem or task. Though people interact in a social manner when performing in these studies, it is not altogether clear that the findings from goaldirected group studies are directly translatable to the seemingly goal-less social interaction that takes place in purely socially driven interaction in the real world and online. Still these studies may shed light on the types of measures of group behavior that might be interesting to look at in purely social situations.

Hemphill and Westie were the first researchers to study the dimensions of groups. In their 1950 study titled "The Measurements of Groups Dimensions", they describe fourteen measurable characteristics of group behavior [1]. Borgatta et al. explain these dimensions in their critique "On the Dimensions of Group Behavior." [2]

1. Autonomy is the degree to which a group functions independently of other groups.
2. Control is the degree to which a group regulates the behavior of group members.
3. Flexibility is the degree to which a group's activities are marked by informal procedures rather than by adherence to rigidly structured procedures.
4. Hedonic Tone is the degree to which group participation is accompanied by a general feeling of pleasantness of agreeableness.
5. Homogeneity is the degree to which members of a group possess similar characteristics.
6. Intimacy is the degree to which members of a group are familiar with the personal details of one another's lives
7. Participation is the degree to which a group permits ready access to membership.
8. Permeability is the degree to which a group permits ready access to membership.
9. Polarization is the degree to which a group is oriented and works towards a single goal which is clear and specific to all members.
10. Potency is the degree to which a group has significance for its members.
11. Size is the number of members of the group.
12. Stability is the degree to which a group persists over a period of time with essentially the same characteristics.
13. Stratification is the degree to which a group orders its members into status hierarchies.
14. Viscidity is the degree to which members of the group function as a unit.

These are surely only a small set of the measurable characteristics, though for Hemphill and Westie, these fourteen were the ones they deemed most important to their research. Borgatta et al. applaud Hemphill and Westie for carefully choosing pertinent measurements: "[The] study contributes an important corrective to the tendency occasionally found to 'through everything in the hopper' of factor analysis with the apparent expectation that meaningful structures will somehow result." Clearly there are going to be those measurements that are useful when studying group behavior, and those that are not even applicable, depending on the task at hand.

Perhaps surprisingly, almost all of the fourteen measures above seem to be applicable to purely social groups. The one characteristic that seems not to fit—polarization-can be rephrased to better fit our particular domain of study. In place of working towards a single clear and specific goal, we might describe polarization as the degree to which a group maintains a single clear and specific subject of conversation. It is important to distinguish polarization from viscidity. The former describes how the group interacts with
itself internally, while the latter measures how the group interacts with its surroundings, including other groups.

In a later work, Cattell defines three general classes of measures of group and individual behavior: [3]

Population variables or dimensions are merely means (or other statistical parameters) of the measured characteristics of the component individuals, such as the mean I.Q., mean structure, etc., of the component members. These are clearly distinct from the characteristics of the group as a group, which arise by interaction, for they can be measured in the individuals before they become a group. By structural variables or dimensions we mean the descriptions of the internal behavior of the group, such as the status gradients, the clique relations as revealed by sociometry, the reciprocal role relations, the form of leadership structure, and, in organized groups, what Stogdill has called "the sociometry of working relations in formal organizations." These are inferred from observations on the internal interactions, processes, and procedures of the group and they are often quite high level abstractions involving such complex concepts as status and leadership structure. The third category comprises the true syntality variables, which represent the performances of the group acting as a whole and commonly through its executive, e.g., its decision in a committee-like situation, its constructive performance on a building task or its acts of aggression or assistance towards other groups.

These three classes, population variables, structural variables, and syntality (meaning the personality of the group) variables would seem to cover all of the measures we can make about a group and the individuals that belong to it. Cattell makes no qualification as to how important each of these classes of variables are to the description of a group. Such a rating is dependent upon the types of questions or goals we have about a particular group.

In our case, the goal of meaningfully portraying groups online, the first two classes seem most important. The third class-syntality-seems to be an emergent property of the group's behavior as a whole. With any luck, the characteristics that fall under this last class will become apparent when we portray the first two classes. Indeed, since the syntality of a group is the most qualitative class of measure of the set, it is the most difficult to measure numerically. Through these classes of group and individual measurement, we can direct our thinking about interesting qualities of group social interaction.

Both Cattell and Borgatta et al. present frameworks for thinking about the measurable qualities of a group interaction. Theories which use such frameworks for studying social psychology are few, and those that do exist are conservative in their predictions of how group social interaction should work. This is understandable; the first research in this field took place less than forty years ago, and there has not been enough time to establish anything more than rudimentary work on this front.

Newcomb et al., for example, suggest that there is a limit to the amount of interaction that can take place between people within a group [4]. With a large number of members (with no note as to how large is "large"), interaction within a group will take place between subgroups, and not between individual members. For very small groups, with perhaps five members, all interactions are person to person. Newcomb et al. goes on to say that within these large groups, "there tended to be an 'inner circle' and an 'outer fringe.'" This prediction, however, has more to do with the process of social interaction, defining people either as active interactors or as onlookers, than with the actual physical locations of the group members.

Milgram presents a similar finding his "The Individual in a Social World" [5]. Much of what makes Milgram's work so powerful and concrete are the images he includes as evidence for his findings. He defines a visual inner circle of activity-a center of a large group-within which most of the activity takes place. Around this center is the outer fringe - a ring of people who attend mostly to the goings on in the center. Milgram claims that for this type of group configuration, there are only a few defining qualities: the diameters of the center circle and of the ring, the permeability of the boundary of the ring, the sharpness of the edge of the ring, and the inter-crowd space or the distance between rings if more than one exists. For more general types of crowds, there is the polarization of the crowd, which is a measure of where the members of a crowd are looking, and the size of the crowd, which Milgram claims is a very important absolute measure that will determine if the crowd is massive enough to support certain phenomena such as looting.

After considering this research, we still do not have a concrete set of worthwhile measurements of this type of interaction. However, we do have a framework from which to begin to describe those measures, and a sample set of real world measures to start from. Perhaps the most important lesson to learn from Borgatta et al. and Newcomb et al. is that each type of group is going to have distinct characteristics and dimensions that will define it, and that the set of characteristics that are appropriate for one group situation may not be appropriate for another.


Milgram's work is perhaps the most useful to us in determining how a group of people will look, based on a small set of criteria. His research points to certain fundamental defining qualities of groups that are independent of the qualities of the individuals within them. With more research, more of these defining characteristics will be discovered, leading eventually to a usable method for simply and reliably describing crowd behavior.

## Social Interaction in the Online World

Now that we have seen how group social interaction can be described, measured, and to a very small part predicted in the real world, we must ask how similar types of interaction take place in the online world, and how these real world techniques and understandings translate. In the real world, group social interaction takes place in real time, with a set of people who share the same space and some of the same psychological beliefs or goals, and who interact with each other within that group. Online, this type of interaction takes place most obviously through chat.

There are many forms of chat in the online world. IRC is one of the most widely used chatting technologies encompassing a number of worldwide chat networks supporting tens of thousands of people simultaneously. Other chat systems do exist, such as chat in MUDs and MOOs, myriad web based chat systems, and instant messaging systems
(though there is usually not a perceived "group" when people interact over instant messaging systems), however many of these systems are functionally the same as IRC. Such systems support real time messaging among members, as well as the manual grouping of people who wish to talk with each other about particular topics.

It is important now to acknowledge the differences between social interaction in the real world and through online chat, as they are key to applying real world social psychology and sociometry techniques to an online environment. First and foremost, there is no physical space on the Internet. Though people can share a place in the virtual world, this place has traditionally been an imagined one with limited visual expression. IRC, for instance, presents a chronological, textual display of messages to a channel. People are represented only by their user names, and the place that people share is expressed solely that textual window. Since traditional research into group social interaction usually assumes a shared space for that interaction, the very limited existence of a common space online may limit applications of those frameworks and theories presented in the previous section.

Secondly, social groups in real life usually emerge naturally from the gathering of people who share both a space and psychological goals (namely to interact with others in a social way). Online, however, the groups are more artificial-we often create the group and set a subject for conversation, and people then join in the discussion. While new groups do emerge out of older, larger ones (such as the creation of a new IRC channel to provide a place for a subset of participants to talk to each other), many IRC channels are longstanding ones that now exist independently of their members. The technology forces us into this view: we can create a new channel within IRC effortlessly, and that channel exists as such even if there is only one person there. While the channel in IRC really represents a space in which people can interact, users often refer to the channel as the group (or set of groups) of people interacting, pointing to a blurring of the line between a place of interaction and the groups of people within that place.

Thirdly, in text chat systems the groups that form are difficult to see. Previous work has shown that conversational groups in social settings form within and around the "space" limits imposed through the system interfaces [6,7]. Since the only grouping of people that we see when using these text systems are the artificial ones created through IRC channels and MUD rooms, these conversational groupings are invisible. Research by both Reid and Cherny point out that the actual social interaction that takes place online is similar to such interaction in real life, so it is safe to assume that there would be some utility in actually seeing this interaction.

Graphical chat systems, such as Chat Circles and The Palace, are a step in the right direction [8]. These systems make explicit and visible the space in which the social interaction is situated, as well as provide users the ability to group themselves in accordance with the conversations they are having. However it is not yet proven that users utilize this feature of the system when close visual proximity is not a requirement for interacting with other users. Chat Circles is a counterexample since it requires users maintain a maximum distance from a conversation to be able to participate in it. This system seems a more natural interface for social interaction than other graphical chat systems, perhaps because its conditions for conversation are more similar to those in real life.

Lastly, there is the issue of multiple presences a text based chat system. This is not as much of a problem in graphical chat systems since there usually can be only one visual instantiation of each person in these systems. In text chat systems, however, a person can be a member of more than one channel at once. Generally, IRC channels are viewed as disjoint spaces; though there is a single person logged into the IRC network, the conversation space is at the level of the channel, and users do not think of the network as a larger, all encompassing space. Is this a limitation for how we see and think about these spaces? Channels in IRC are a way for users to make sense of the space. They gather around subjects instead of objects.

With these differences in mind, we can take the framework proposed by Cattell, and the group characteristics suggested by Hemphill and Westie, and derive sensible measurements for online social interaction. We must also keep in mind the ease with which any of these dimensions can be numerically measured. In the real world, any statistic can be gathered by a human that can already understand and interpret the group that he is observing. If we are interested in tracking online social interaction automatically with a computer, we must be able to program the computer to measure each dimension. For dimensions that require an already present understanding of group interaction, such as determining the leader of a group-which would require a significant amount of linguistic processing on the computer's part-then the task of programming the measurement of that dimension becomes non-trivial or even impossible with current technology. Hopefully, at least some defining characteristics will also be easily measurable ones.

Two of the Hemphill and Westie's dimensions can immediately be rejected-intimacy and potency-since they would require information not available in IRC. Measurement of these qualities would have to be done through explicit questioning of the user. Though not necessarily so, the stratification of a group seems irrelevant to the types of groups we are studying, and could be rejected initially.

Application of the remaining eleven dimensions is less clear. Control and participation are more binary measures of how a group is initially defined, rather than functions of the ongoing syntality of the group. Looking deeper into these two dimensions would require qualitative analysis of the group, as well as an analysis of the social network within a group using linguistic techniques. Both of these processes are difficult to measure with a computer. Flexibility would require the recognition of turn taking procedures and group defined procedures, which could be done with difficulty using textual analysis of posts to a group.

Measurement of hedonic tone is a domain of study unto itself within affective computing. Detecting the emotional content through simple keyword lookup is possible, but would
not completely cover this dimension. Polarization might also be measured using keyword lookup, perhaps by trying to determine if the subjects that are being discussed fall into one or a number of subjects. This measurement may be useful in determining how many subgroups exist within a particular channel. The permeability of a group can be measured by tracking the response to a post within a group. Though more difficult in IRC than in Usenet due to lack of post threading in the former, this measurement may be important in determining the barrier to entry of channel or subgroup within a channel. Also, since a channel can be protected by a password, noting permeability on even a binary level may be important.
Cite projects that might help in the above measurements.

Through the comparison of users personal information (including where they are logging in from) as well as a language style analysis, the homogeneity of a group can be determined. Determining where users overlap in personal information may provide some insight into the shared psychological goals of the group, perhaps allowing us to determine more natural lines of group segmentation than that of channels (as discussed above as the second difference between real life and online group social interaction). \&interesting idea

The final four characteristics are the easiest to measure. Size is an obvious dimension, since it is a trivial matter to count the number of people within a channel, assuming a channel provides a meaningful segmentation of people into a group. (There is some evidence that the artificial groups imposed by the system are meaningful for Usenet [9]. IRC channels are similar in function to newsgroups on Usenet. ** get to these earlier ${ }^{* *}$ ) By tracking a user's membership and postings to multiple channels, we can measure a group's autonomy and viscidity. If a group has members that do not belong to other groups, then that group has a high level of autonomy. If a group has a low level of autonomy but its members tend to frequent the same set of channels, then that group has a high viscidity. These two measurements may help address the multiple presence issue discussed above. By looking at how users move from channel to channel (some of this
research has been done by Sherry Turkle in Life on the Screen), we can discover how they view the space in which they interact [10].

Lastly, stability can be measured by tracking the membership of a group over time. This measurement is a particularly interesting one in light of the second difference in online social interaction presented above. A channel is still a channel on IRC regardless of its membership. Yet, the particular personality of that channel, and the type of interaction within it will change with its membership. A stable group should not change its character significantly over time, while a group with a constantly changing membership could be different every time we view it.

Aside from these four final dimensions, the majority of the characteristics of group social interaction presented by Hemphill and Westie require a significant amount of processing to measure automatically. Are there other, more easily measured dimensions of online groups? Keep in mind that the evaluations above of these group characteristics were based on interaction over IRC. For other Internet communications media, such as Usenet, the facility of making the above measurements will be different. Also, since Usenet represents a different type of social interaction, the importance of each variable will also change.

Interestingly, these four final dimensions are all structural variables within Cattell's categorization of group and individual variables. Other variables might be the percentage of people within a channel that post, the rate of conversation within a channel, and the length of an average post in a channel. We must also look for characteristics that fall into Cattell's other two categories, population variables and syntality variables in order to expand our measurements.

For population variables, we might look at the average length of a person's posts, as well as the amount of time they spend logged into the system and the number of groups to which they belong. Also, we might look at the rate at which they join and leave groups, and the times that they usually log on. Syntalic variables are more difficult to determine,
since they are generally more qualitative in nature than either population or structural variables. One characteristic that falls is XXXXXXXXXXXXX. Some of these variables are ones that Whittaker et al. used when looking at how large groups of people interact on Usenet [9]. Flesh out.
[COBOT]

It is important to consider the timeliness of these measurements. Though not discussed in the literature on traditional social psychology and sociometry, it is implied that groups change over time. The changes in these measurements over time are even more important for social interaction online because of the artificial establishment of groups within IRC. Much of the traditional research in real life group interaction is short term and about selfconfiguring groups. With online social interaction, we have the opportunity to study groups over a long period of time, on the order of months or years. We expect that there will be several time scales for change within these groups. Especially since IRC is a real time communications system, changes in these variables can happen over the course of minutes or days. Longer term changes, over weeks, months, and years, are also likely (similar to the time scales studied by Whittaker et al.) so rates of change for any variable of group interaction online are going to be important.

Evaluating the importance or relevance of each of these variables in the domain of online interaction is difficult. Those dimensions that are fundamental to an IRC channel will not make themselves obvious until we begin looking at all of the possible variables to determine which ones change from group to group. Also, our goal is opposite from that of traditional social psychology. We are interested primarily in making the structure and dynamics of group social interaction online explicit and visible, whereas traditional study attempts to categorize and define the types of group interaction. We will need to take as given a particular model of social interaction, and visualize how the social interaction occurs. Thus the utility of particular measures of individual and group interaction will be necessarily different for us versus traditional social psychology.

## Crowd Perception

When we look at a group of people, what do we see? What lands upon the back of our eyes is merely color. There is a great deal of cognitive processing that goes on between the light entering our eyes and the percept of the group. Indeed, we see a group of people on two levels: the group as a whole, and the individuals within it. As we shift our attention to from individual to individual, or from group to individual, what allows us to pick out the interesting visual information from the mass of percepts that our eyes take in?


Figure 5.16
(a) Perceived as horizontal rows or vertical columns or both. (b) Perceived as vertical columns.

For our purposes, we will look at perceptual processing at a sufficiently high level that the visual paths of neural circuitry in our brain have already extracted recognizable form, color, and motion. Using these perceptual components, our brain begins the cognitive processing that allows us to understand what we see. In the 1920's, the Gestalt school of psychology provided the first basic principles of the process of perceptual organization. Max Wertheimer, one of the founders of Gestalt psychology, proposed the radical idea that the perceptual whole is not equal to the sum of its parts, meaning that our perception of objects is not based merely upon the appearance of the components that make them up [11]. Instead, there are fundamental cognitive processes, or rules, that allow us to infer information that is not present perceptually.

In his book Sensation \& Perception, E. Bruce Goldstein describes the Gestalt laws of visual perception:


Figure 5.42
Texture segmentation by texton differences. (a) Difference in orientation. (b) Difference in line crossings. In (c) there is no texton difference, so texture segregation does not accur. (From Nothdurft, 1990.)

1. Pragnaz: every stimulus pattern is seen in such a way that the resulting structure is as simple as possible
2. Similarity: similar things appear to be grouped together.
3. Good Continuation: points that, when connected, result in straight or smoothly curving lines are seen as belonging together, and lines tend to be seen in such a way as to follow the smoothest path.
4. Proximity: things that are near to each other appear to be grouped together.
5. Common Fate: things that are moving in the same direction appear to be grouped together.
6. Familiarity: things are more likely to form groups if the groups appear familiar or meaningful.

Gestalt translates roughly to "organized structure," and forms a general system of predictions for how particular stimuli will be interpreted by the brain [12]. It is not a proper theory, since it does not provide an explanation of why these phenomena occur, but rather a framework through which we can describe and create visual stimuli that induce particular psychological interpretations.

More recently, work by Bela Julesz on textons and by Ann Treismann on feature integration theory (FIT), provide additional perceptual phenomena to guide our understanding of crowd perception. Textons are the properties that cause texture segregation-"the perceptual separation of fields with different textures" [11]. FIT proposes pop-out boundaries, those qualities of perception that facilitate the differentiation of particular objects from their surroundings. Both of these theories


Figure 5.44
Visual search stimulus. You can find the O among the V's almost instantaneously.

The target in this display is an R.


Figure 5.45
Visual search stimulus. Finding the R among the P's and Q's is not instantaneous unless, by chance, you happen to look at the R first.

propose that the perception of distinct objects in the visual field are caused by differences in color, form, size, and orientation of an object and its local environment. The greater any of these differences are, the more the object will pop out, and Treismann has shown in studies of visual search that this principle explains why certain objects are easier to pick out than others.

With these frameworks in hand, we can now ask why certain phenomena of group perception occur. When we look at a set of people, why do we sometimes see them as a group? Proximity says that if people are close together, then they will appear grouped. So when we see a room of people, and that collection has some subset of people who are closer to one another than the average distance between people, we might see that subset as a group. A social group of people is a very familiar segmentation for a set of people, and we understand that a social grouping has some meaning for
those people-they are interacting with one another-and thus the principle of familiarity applies here as well. Good Continuation is also at work: a group of people, such as the crowds in Milgram's photographs, form smoothly curving lines, so they are likely to be grouped. If we are to believe that the principle of pragnaz holds for people as well as things, then the simplest way to think of a set of people is as a group, and not as individuals.


How is it that we can pick out and attend to particular people within a group or crowd? Common Fate may be at work here if there is motion of the individual with respect to the group as a whole-this is why we can see the movement of groups within a larger whole. The pop-out effect is also particularly strong in the case of attending to individuals. We all have witnessed the phenomena of a person sticking out of a crowd "like a sore thumb." This is because their visual characteristics-the color of their hair or clothing, the shape of their dress and body, their directional orientation, or their physical size are all different from what we perceive visually as the norm for a group of people.

These cognitive principles for object perception are general; they explain how we will understand all types of visual stimuli, and not just the image of a set of people interacting socially. If we are going to create a display that will be easily understood, then we should adhere to these principles of perception. The case of perceiving a group of people visually is an example of how these principles work. Thus we should leverage the already innate abilities for seeing groups of people in the real world when we create a display intending to portray group social interaction online.

## Life-like representations

When we look at a billiard table and watch the cue ball strike another ball, the motions we see do not appear life-like, though they follow the principles of object perception and
group perception outlined above. There is a particular quality to the motion of people and other live things that is missing from the motion of billiard balls and other inanimate objects.

The embodiment of life in a moving image comes from perceived goal-directed action or intentionality of movement [13]. When a simple circle moves as if it has a goal, as opposed to merely being affected by forces within the environment, then it is seen as being alive. Disney discovered eleven related visual qualities that aid in this perception [14]. While the whole set of these qualities is important for creating an object with personality, only a handful are important for creating the illusion of life:

1. Squash and Stretch: defining the rigidity and mass of an object by distorting its shape during an action.
2. Anticipation: the preparation for an action.
3. Slow In and Out: the spacing of the in-between frames to achieve subtlety of timing and movement.
4. Arcs: the visual path of action for natural movement.

Using these visual techniques, we can create an animation that seems goal-directed. Squash and Stretch makes a shape organic. Anticipation of the upcoming movement through a "wind-up" or pre-squash, coupled with slow in and out makes the motion seem intentional, and not just a result of the environment. Arcs ensure the motion seems believable.

## Viewing online social interaction

Considering the visual characteristics described above, how do current graphical chat systems portray group interaction? There are two classes of graphical chat programs currently in use today: avatar systems and abstract representational systems. For our purposes, however, both of these systems portray groups of people in the same way, by letting users position their representations explicitly. A good example of such a system is Chat Circles, which represents users as colored circles within a large, rectangular room. (Viégas, et. al.1999) Avatar systems are similar, but portray users with more complex shapes, often resembling humans or anthropomorphized creatures.

ComicChat, an avatar based chat system built upon IRC, is worth noting because it generates views of conversation groups automatically, without user intervention. (ComicChat) This particular process is very useful for making clear the conversations that take place within the text-based chat system. However, the images it generates fall short of presenting a realistic view of the group as a whole. The frames that are generated only show the active participants in the conversation at the time of the snapshot, ignoring all other members of the channel. Also, the frames are static, so understanding of the dynamics of the group based on motion information is unavailable.

For graphical chat systems like Chat Circles, the users of the system explicitly create the groups that are graphically portrayed. Through their movement within the virtual space, the individuals forming a social group will do so visually as well as conversationally. In systems that force users to move close to the people with which they are having a conversation, such as Chat Circles, the formation of groups visually is a necessity. For other avatar-based chat systems, however, while users generally move into the same room as the other members of their conversation, they don't necessarily form cohesive conversation groups visually, instead relying on the fact that the other members of the group can read their posts merely by being present within the same room.

Once these visual groups are formed, however, most users, both in abstract representation chat systems and avatar chat systems, tend to remain in one place. Motion in this type of interface represents merely the movement from one group to another. Since this is not necessarily the most important information that can be displayed about a social group, it may be a waste of a visual dimension to portray it using such a powerful display as that of movement. More interesting is to highlight the active conversants within each social group. Further, body language and movement within a social group are two important information channels portrayed through visual motion. In these graphical chat systems, both of these channels are lost.

A significant limitation of most graphical chat systems is that they discard all but the most recent posts within a group. Almost all graphical chat systems present a user's
posting in a bubble above or within that user's representation on screen. These posts fade quickly from view, thus if we have not been paying attention to the screen, we loose all information about the recent chat activity. Textual chat systems do not have this problem since all posts are logged in the chat window, however all chat systems lack a method for keeping track of this information when a user isn't logged on: we can only see information about a chat for which we were present. While this restriction holds in the real world as well, we should attempt to provide added value to using online chat, and include features and information that we cannot have in the real world.

A third abstract graphical chat system, The Fog, addresses some of these issues of history of postings in a chat. The Fog presents users as colored squares labeled with the users name. Much like Chat Circles, users can move their representations around within a room. This system, however keeps a history of posts to the chat room, and integrates this information into the interface. Older posts fade into the background as new posts are

added. This history feature allows participants to keep track of a conversation's recent past, and allows new members of the conversation or users who left the conversation for a short while to bring themselves up to date. While this feature is very useful, and goes beyond other systems in integrating historical information into the chat interface, it provides only a short-term memory for the conversations. Still, we can get a sense for recently active participants by noticing which colors are present in the conversation history. Also, since a post maintains its location even when a user moves, the location of posts can give a map of the movement of the conversation throughout the space.

## Measuring a social group online

Using techniques from sociology and visualization, as well as both the good and bad features of current graphical chat systems, we can now attempt to build a visualization system for online social interaction. $\leftarrow$ word better Coterie is such a system: It provides a display of the group social interaction within an IRC channel based on our perception of social groups in real life, uses knowledge of visual cognitive processes to create a richly informative display that is clear and intuitive for the viewer, and includes both long- and short-term historical and persistence information about user activity in the chat space to contextualize each user's interactions.

Coterie is built using the client/server model. The server provids an IRC bot to gather posts from selected IRC channels, a database in which these posts are stored and within which statistics about interaction are computed, a user and group modeling system that takes the information from the database and evolves a group model for each channel in real time, and a server to provide update information to a Coterie client. The Coterie client provides a connection to the server for update events to the user and group models as well as a dynamically created visualization of this model. All components are written in Java, and the system as a whole provides a real time visualization for group social activity within a set of IRC channels.

A distinction similar to Catell's regarding measurement dimensions is made between information gathered about each member of the group and the group as a whole. The Coterie database computes the following statistics for group activity:

1. lastLogon: the time of the last logon to this channel
2. lastLogoff: the time of the last logoff to this channel
3. lastPostLength: the length of the last post to this channel
4. logonRate: the rate of logons per minute over the past 60 minutes to this channel
5. logoffRate: the rate of logoffs per minute over the past 60 minutes to this channel
6. averagePostLength: the average length of a post to the channel
7. postRate: the rate of posts per minute over the last 60 minutes to this channel
8. populationGrowthRate: the increase or decrease of members in this channel per minute over the past 60 minutes
9. numberOfUsers: total size of this channel
10. percentageOfTalkers: percentage of members of this channel who have posted within the last 60 minutes

The last logon and last logoff numbers provide a measure of when a channel has last grown or shrunk in size. This is analogous in the real world to noting when someone has joined or left a group of people, and perhaps indicates that someone has gained or lost interest in interacting with the group. Coupled with last post length, these measurements provide an up to date snapshot of the group's most recent activity. The logon rate, logoff rate, and post rate provide a short-term history of the internal behavior of the group as a whole. Since they are computed over the past 60 minutes, they measure recent trends in this behavior. Note that population growth rate is computed using by adding logon rate and logoff rate, and is maintained as a convenient magnitude-invariant measure of group growth. If a group's population growth rate is zero, then it is maintaining its size, regardless of how many members are joining and leaving.

The average post length is a measure of the average magnitude of a group's interaction. Coupled with post rate, we can determine how a group interacts over time. If a group has a large average post length, but a small post rate, then that group has a small amount of interaction between its members, but each interaction is significant. Likewise, if post rate is high, but average post length is small, then the group supports quick, short messages among its members, perhaps indicating a more commentative conversation style rather than the perhaps deeper, more thought-out comments of the previous example.

A channel's number of users is merely a measure of its size. This is a very basic measurement, and the only one that also appears in Hemphill and Westie's list of real world group dimensions. However, this simple measurement may be extremely important. Milgram suggests that sheer magnitude of a crowd of people will indicate whether that group has enough mass for certain crowd phenomena to occur.

The percentage of talkers within a group provides a measure of the magnitude of interaction within the channel. A channel with a large percentage of talkers is one that is very cohesive; its members tend to all interact within the group, perhaps indicating a significant familiarity between those members. Since IRC channels are artificial groups, we cannot say for sure that a cohesive group is necessarily uniformly cohesive. It may be that there are a number of conversational sub-groups within the larger group of channel members, and each of those groups contains members that know each other well, but who don't know well the members of other groups. Still, since they share a common space, they may be at least aware of each other's presence, and since all activity within the channel takes place in public, all members who interact are visible to all other members. Alternatively, a channel may have a small percentage of talkers, which indicates that a channel has a large number of people who are logged in only to watch the interaction of others. This may be an example of the real world phenomenon of "people watching."

For these statistics, time is important: since a group is constantly changing, both in how it interacts with itself and who its members are, such measures of group behavior are only valid over a short term. An hour was chosen as the time window for group statistics based
on informal observation of general IRC behavior. Longer or shorter time windows for averaging group measures may be appropriate, depending on the group. Placing emphasis on current and recent social activity events is similar to Mark Ackerman's work on
 social activity indicators. (Ackerman 1995)

The Coterie database computes these measures for individual activity: last logon, last logoff, logon rate, logoff rate, last post length, average post length, and post rate. These measures are analogous to those computed for the group as a whole, except that the time window for recent events is five minutes instead of 60 minutes. The change in time window was made because it is assumed that individuals change their social behaviors over a short time period, whereas groups as a whole tend to change slowly, since each individual has only a small effect on the behavior of the group.

The dimensions used to measure group and individual activity within an IRC channel are ones for creating a visualization of the group activity. They certainly are not exhaustive, however they seem, in use, to describe most of the important features of a group. Each statistic is used to a greater or lesser extent, and only experimentation will validate the choice of statistics. Also, these statistics fit well with the visualization we will now describe. It is possible that different sets of statistics will be useful for different visualizations.

## Making an online social group visible

Once these statistics are computed, the user and group model is updated on the server, and the visualization on the client end is modified accordingly. Though they exist on different machines, the user and group model and the visualization are heavily coupled,
so we will describe them simultaneously. All of the data for the visualization is kept in the user and group model, so any reference to a visual characteristic implies that the actual data for that characteristic is maintained in the model and expressed in the visualization.

Each group that is being viewed is placed in its own section of the client window and has the name of that group displayed in that area. Multiple groups can be viewed using Tufte's technique of "small multiples". (19XX) The display also follows Tufte's principles of design, providing a clean and aesthetic interface, only showing important information. The visualization represents each user within a channel as a colored oval. The ovals are arranged at the bottom of the channel. Surrounding the center of the channel is the conversation area. This area is where all users are placed who have recently posted. Beyond the conversation are, on either side of the channel display, is where the edge crowd of the channel exists. If a user is a lurker, he will be placed here. This formation is modeled after images of crowds in Milgram's The Individual in a Social World. We can think of the display as a vertical slice through a circular crowd.

We chose to use ovals to represent people in Coterie because they are simple forms, and easy to manipulate. As in Chat Circles, we have attempted to provide a form that is representative of a human, but one that does not support the reading of false personality attributes. Ovals were chosen over circles since they are more indicative of the human form as viewed from the front or the side. The motion of the ovals utilizes technique of slow in and out and arcs for all actions. Ovals always move in a continuous and nonabrupt manner: when a user logs on or logs off, he will fade in or out visually; when he becomes a lurker, he will move continuously into the edge crowd, and not instantly reappear there. Since our perceptive abilities in the real world expect that a person's motions are bound by the laws of physics, we have tried to maintain that expectation in Coterie's visualization.

Group statistics drive the overall patterns of a channel's representation. The width of the conversation is determined by the percentage of talkers: the higher the percentage of talkers, the smaller the width of the conversation. This mapping is based on the idea that a group with a large percentage of active people is more cohesive, and should appear more tightly knit. The width of the edge crowds is based on the logon and logoff rates. In Milgram's discussion on circular crowds, he notes that a group that has a large number of people who come and go will also tend to have a diffuse edge. He calls this particular characteristic of a crowd its permeability. For online crowds, permeability refers more to the psychological permeability of a group-the ease with which a person can join or leave. If a group has a large logon or logoff rate, then it may be easy for people to join or

leave the group, thus the edge crowd is made more diffuse.

A person's horizontal location in the channel changes over time. When a user posts a message, he jumps up from his current location and bounces into the conversation area (unless he is already there, in which case he just jumps in place). Jumping was chosen as the primary display for a user's posting action since a post is generally intended to draw the attention of the group, and jumping is one way in which people signal they want attention (or rather a jumping person will often draw attention).

The height of the jump is directly related to the last post length. Thus the amount that a person has to say is reinforced visually. If a person is agreeing with another member of the group by posting "yes," then that person should only jump a little, indicating that something was said, but not drawing the viewer's attention like a longer post should. Since a user's post could be very large, we mediate the effect that post length will have on the height of the jump. An inverse exponential curve is used to create a smoothly changing function:

$$
\frac{1}{e^{1 / 2}}+0.1 \text { postlength }
$$

Using this formula, the height of the jump changes smoothly with a change in post length, but the larger the post, the smaller the increase in the height of the jump.

After a short amount of time, a user will drift back to the edge of the conversation area and eventually into the edge crowd. In the real world, when a person is not an active participant in a conversation, they inhabit the sideline of the conversation, and look on at the action of the other members of the conversation. This principle drives the movement back to the edge of the crowd in Coterie. If a user isn't posting regularly, thereby participating actively in the conversation, then he is listening, and should be at the edge of that conversation. Is this true? What about people who are important but say little or talk infrequently?

When a user posts, along with bouncing into the center of the conversation, he will increase in size and become more opaque. The visual change serves to perceptually reinforce the grabbing of the viewers attention by the post. Both of these features also allow two recent posters to be easily differentiated. When comparing the brightness and size of two identical shapes side by side, small changes are noticeable, and will pop-out. Thus, the relative times of participants' recent posts are easily seen even when a member is partially overlapped. This same principle is used in the edge crowds to show the amount of time a user has been a lurker. Once the system has labelled a person a lurker, they will slowly shrink by $10 \%$ of their size. In this way, members who have recently become lurkers will appear bigger than those who have been lurking for a long time.

Historical statistics are used to visually differentiate users based on their recent activity in the channel. If a lurker posts a message, he will fade quickly back into the edge crowd. If a person is an active participant, he will maintain his presence in the center of the conversation for a longer period of time. The more active a person, the longer they will visually maintain presence. This allows the viewer to contextualize a message and categorize a user easily based on his motion. See 2 paragraphs before...

Since the statistics used to drive both the visual properties of both users and the group as a whole are computed using a sliding time window, trends in these statistics will cause changes over time of the visual representations of a channel and its members. For example, as a group becomes more active at the end of the day, its conversation space will shrink and more people will inhabit that area. One very powerful consequence of this sliding time window is that changes to the group structure, and even in the interaction levels of individual users, will be drawn out in time. In traditional graphical chat systems, changes in the activity level of users (which affects the group's statistics as well) go unnoticed. If users in a group become very active after an extended period of inactivity, they will, for a time, slip quickly into the background. For a group that has maintained this activity level for a while, the users will tend to maintain their presence in the center of the group. Thus it becomes easy to visually distinguish the changes in social interaction within a group.

## Evaluating the visualization

Coterie addresses two of the four differences between online and real life group social interaction, providing a bridge between our abilities to view and understand crowds in the real world and the invisible crowds in the virtual one. Space is created where before none existed. With Coterie's visualization, we can now watch group social interaction play itself out in real time, without the need for users to explicitly place themselves within that space. Also, since active members of a conversation are constantly moving, our attention is drawn to them as opposed to the ones moving from one group to another (as it is in Chat Circles).

One dimension of the visual display not used in the
 visualization is hue. While each channel is given a separate hue, this is not the best use for such a powerful visual differentiator, since there are already physical barriers between each channel. We could assign each user a hue to allow easier differentiation of specific people. This option, however, presents a problem in the edge crowd. Since there are typically many people in the edge crowd, compositing so many different translucent colors will make a mess of brown and black in those areas, completely preventing people from seeing how many people are actually in those areas, as well as destroying the aesthetics of the display. More work on this problem will have to be done, as the answer of how to integrate hue into the system is not obvious.

Coterie also ignores many of the dimensions described by Hemphill and Westie. Unfortunately, most of these dimensions require significant amounts of processing of user's posts. In this first implementation, we have attempted to draw out as much information from the user's activity as possible without implementing a language parser. While adding this functionality to the system would make Coterie more robust, there
have been other systems that provide social statistics based on language processing, such as Cobot, and we were interested in pursuing a different path of research.

Though we have attempted to create forms that evoke an impression of the human form and human motion, the visual stiffness of the oval used in the visualization detracts from its success as such an object. By using the animation techniques of stretch and squash and anticipation, we can create an object that seems more life-like. We can squash the oval when it's landing, and

a 3D crowds visualization
stretch it when it starts and ends its jump. We can also squash the oval just before the beginning of its jump to give the impression that the oval is intending to jump. Thus we can create impression of goal-directed action on the part of the oval, and make a more convincing visual display.

While the system is successful using only short-term historical information, it should take long-term persistence information into account as well. Users who are members of a channel for a long time should be visually different from new users, much as long time members of a real life group are differentiated by their knowledge of jargon and common jokes and phrases, as well as exhibiting a greater level of comfort with their surroundings. By keeping track of how long a user has been a member of a group, we can visually distinguish experienced participants from "newbies." Using long-term persistence information, we can also point out unusual appearances by people. In the same way that we might notice a particular person because they are deviating from their regular schedule, we can highlight when a user in a channel logs on at an unusual hour.

Representing multiple presences within a chat system is difficult. In the real world we expect that a single person is going to have a single presence. If we want to utilize our real world cognitive abilities for understanding group behavior in an online world, we should maintain this assumption. Since people can be in more than one channel at once, we might try overlapping the set of channels in one big window, and changing a user's color to move between the colors of the groups to which he belongs. This path may lead to more screen clutter, and create an interface that is more difficult to understand. We might have a person move from window to window, depending on his activity in all of the channels, but this leads to the false impression that he is leaving one channel to move to another, whereas he is actually just shifting his attention from one group to another. This is a difficult problem to solve, and no clear solution has yet been found.

Since most social interaction online takes place in non-graphical ways, we loose our ability to use well-honed crowd watching techniques. By creating a view of groups of people online interacting socially, we tap into those abilities. Designing a visualization so that important information is always portrayed in an easily understandable manner, driven by our cognitive and perceptual abilities, we enhance our ability to read and understand that display. Coterie was designed in such a way, and owes much of its success as a tool for visualizing group social interaction online to carefully following the
guidelines and ideas set forth by the domains of social psychology, perceptual science, and visualization.

## Bibliography

1. Hemphill, J. K., and Westie, C. M. "The Measurements of Group Dimensions," Journal of Psychology, 1950, 29, pp. 325-342.
2. Borgatta, E. F., Cottrell, L. S., Jr., and Meyer, H. J. "On the Dimensions of Group Behavior," Sociometry, Vol. 19, No. 4. December 1956, pp. 223-240.
3. Cattell, R. B., "New Concepts for Measuring Leadership, in Terms of Group Syntality," Human Relations, 1951, 4, pp. 48-63.
4. Newcomb, T. M., Turner, R. H., and Converse, P. E. Social Psychology. New York. Holt Rinehart and Winston, Inc., 1965.
5. Milgram, S. The Individual in a Social World: Essays and Experiments. Reading, Mass. Addison-Wesley Publishing Company, 1977.
6. Cherny, Lynn. Conversation and Community. Stanford, CA. CSLI Publications. 1999.
7. Reid, E.
8. Viégas, Fernanda B. and Donath, Judith S. Chat Circles. In Proceedings of CHI '99, ACM/SIGCHI, N.Y., 1999, pp. 9-16.
9. Whittaker, S., Terveen, L., Hill, W., and Cherny, L. The Dynamics of Mass Interaction. In Proceedings of CSCW '98. ACM, N.Y., 1998, pp. 257-264.
10. Turkle, S. Life on the Screen. New York, N.Y. Simon \& Schuster, 1995.
11. Goldstein, E. B., Sensation \& Perception. $4^{\text {th }}$ Edition. Pacific Grove, CA. Brooks/Cole Publishing Co. 1996.
12. Kanizsa, G. "Gestalt Psychology: What it is not." In Organization in Vision: Essays on Gestalt Perception. New York, N.Y. Praeger Publishers. 1979.
13. Premack, D. and Premack, A. J. "Infants Attribute Value to the Goal-Directed Actions of Self-Propelled Objects," Journal of Cognitive Neuroscience 9:6, pp. 848-856.
14. Thomas, F. and Johnston, O. The Illusion of Life.
15. ComicChat
16. Hannes Wallnöfer. The Fog. http://www.t0.or.at/~hwallnoefer/fog2.html. Accessed May 8, 2000.
17. Ackerman, M. S. and Starr, B. "Social Activity Indicators: Interface Components for CSCW Systems," In Proceedings of UIST '95. ACM, N.Y., 1995, pp. 159168.
18. WebAware
19. From Individual Human Agents to Crowds
20. Cobot
21. Tufte
