Mapping the Mental Space of Game Genres

short paper **submission ID vgsp_0010**

J.P. Lewis*
Stanford Computer Graphics Lab

Pamela Fox
University of Southern California

Morgan McGuire
Williams College

Abstract

In this short paper we explore the application of manifold learning algorithms to build maps of the "space" of game genres. The gaming community has produced an informal classification of games that defines genres such as first-person shooters, real-time strategy games, etc. While these genres are generally accepted, producing a more formal taxonomy of game types is challenging, both because it involves subjective opinions and because the mental space of game types is clearly high-dimensional. Psychological research has in some cases successfully addressed the problem of creating visual models of subjective mental concepts using multidimensional scaling. Recent manifold learning approaches provide a way to make similar maps of potentially high-dimensional data, though they have been mainly targeted to physical rather than conceptual data. We illustrate applying manifold learning to pairwise similarity ratings between games. The resulting maps shows clustering of related games into spontaneously arising "genres". Maps such as these can help support discussion and analysis of games including exploration of age- and gender-related perceptions of games, provide a spatial framework for game recommender systems, and perhaps suggest opportunities for new types of games in the empty spaces in the map.

1 Introduction

People clearly have a subjective sense that The Sims is a different sort of game than Quake, and this is reflected in the classification of Quake as a "first-person shooter", which The Sims is not. As games become a subject of research, it may be valuable to explore game categorization more formally. This can be accomplished by direct analysis of game properties, and perhaps by opinion surveys.

In theory a "map of game space" could be an attractive alternative categorization, for several reasons:

- it would provide a visual framework for some types of analysis and discussion,
- it might suggest unexplored game opportunities corresponding to the empty spaces in the maps, and
- humans are sometimes adept at absorbing and remembering information presented in spatial formats such as maps (although this is a complex subject [Richardson et al. 1999; Satalich 1995]).

Producing such a map seems challenging, however, both because it involves subjective opinions and because the mental space of game genres is clearly high-dimensional – a game is distinguished in many independent attributes including single/multiplayer, first/third-person, and many aspects of theme, plot, and game play.

*jplewis@stanford.edu

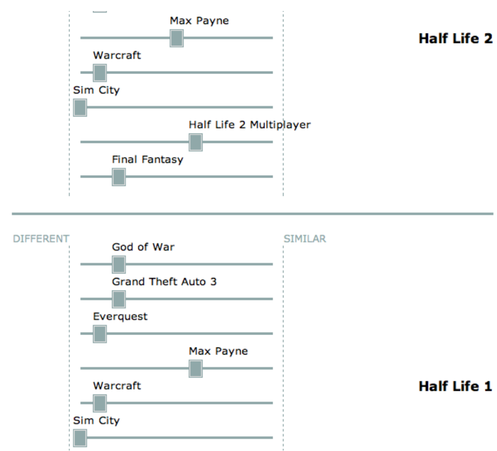

Figure 1: Survey participants were asked to rate the similarity of games using sliders, with no further instruction as to how "similarity" was to be interpreted. This screenshot shows `Half Life` titles being compared to other games selected by the user.

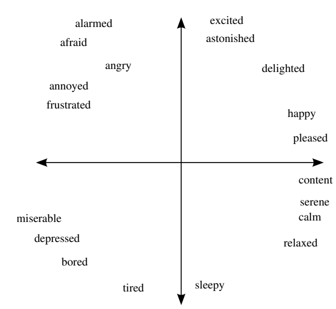

Figure 2: Schematic diagram of "emotion space" obtained from multidimensional scaling, simplified from [Russell 1980]. The horizontal and vertical axes encode displeasure/pleasure and low/high arousal dimensions respectively.

Each of these problems has been addressed, in separate communities. Psychological research has shown that valid visual maps of mental models can be produced in some cases [Russell 1980; Lokuge et al. 1996]. Recent research in the neural information processing and computer science communities has produced algorithms that successfully discover mappings from high-dimensional data to lower-dimensional maps.

In this short paper we illustrate the application of manifold learning algorithms to build maps of the "space" of game genres. Note that "genre" is informally and operationally defined; in particular, groupings of games arise directly from the survey community's similarity ratings, rather than from any a priori considerations.

2 Related Work

The notion of game "genres" is generally accepted, but there are differing theories as to how games might be classified into gen-

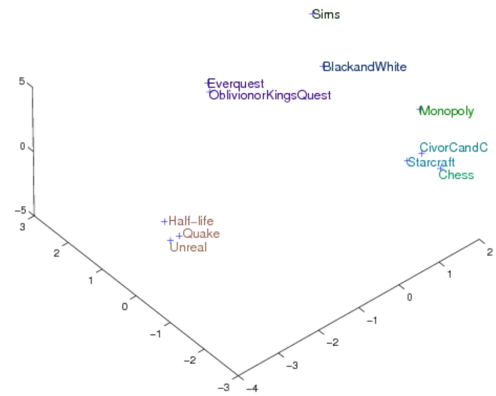

Figure 3: Three-dimensional map obtained from a pilot study of 11 games rated by 18 people. The three dimensional coordinates of each point are also encoded in the color of the game names. Note that the global orientation of this plot is arbitrary – there is no intrinsic "up" direction.

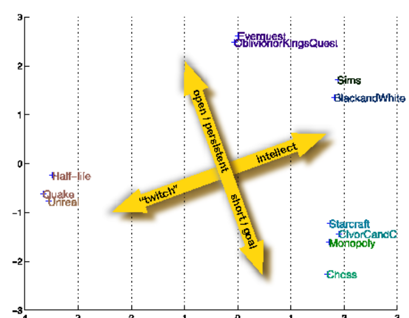

Figure 4: A two-dimensional projection of Fig. 3. The axes of the space may be speculatively interpreted; several gamers suggested that these two dimensions might be labeled as "twitch" (fast reaction time) versus intellect games, and short-term / goal-oriented versus open-ended and persistent games.

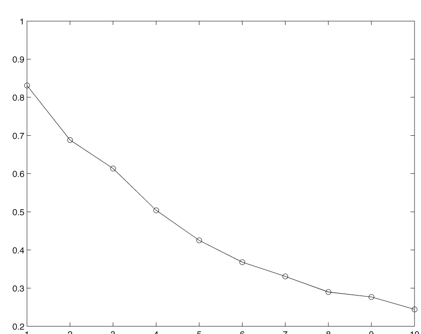

Figure 5: Embedding distance error (vertical axis) versus dimension for a set of 92 games rated by 124 people.

res. For example, the Entertainment Software Association identifies Sports, Action, Adventure, Strategy, Shooter, Role-Playing, and Family genres [ESA 2005]. A long line of video game critical work exemplified by Juul [Juul 2005] classifies games along theoretical axes such as progressive versus emergent. Cukier [Cukier 2006] exemplifies this work – he rejects the standard marketing genres and proposes five of his own criteria, including number of players, synchrony, and realism. Whereas these analyses are based on top-down reasoning and to some extent represent a single author's opinion, we provide a bottom-up analysis of how a population of gamers actually perceive game similarities.

Multidimensional scaling (MDS) techniques provide the basis for our approach. These techniques attempt to place points in a reconstructed Euclidean space of specified dimensionality (typically a plane) so that inter-point distances are similar to given distances between points in a (potentially unknown) space. When the given distances correspond to points of potentially high dimensionality but lying on a linear manifold of the same dimensionality as the desired reconstruction (such as inter-city distances measured from a planar map), MDS can reproduce the original configuration of points, subject to an arbitrary rigid transform.

MDS has been used in psychology research to produce visual representations of mental maps and concept relations. Although the true "distance" between concepts is not known, an approximate distance can be obtained from a population average of rated pairwise similarity (or dissimilarity) between concepts. Russell [Russell 1980] applied MDS to similarity ratings of human emotions, resulting in the well known circumplex model in which emotions are radially arranged in a two-dimensional space having axes that have been labeled as 'pleasure/displeasure' and 'high/low arousal' (Fig. 2). In this map, an emotional state such as "relaxed" is on the 'pleasure' side of the first axis, but has a low value on the second (arousal) axis for example, whereas "bored" also has low arousal but is located on the displeasure portion of the first axis. Although this model is not universally accepted, it has been independently derived from rated pairwise similarities of emotion words in a number of languages, as well as from rated similarities of facial expressions [Plutchik and Conte 1996]. More recently, [Lokuge et al. 1996] applied MDS (as well as an alternate trajectory mapping technique) to reconstruct a "tourist content map" of Boston. The reconstructed locations provide a 'map' suitable for overlaying tourist site information such that similar types of destinations are clustered.

Despite the success of these applications, classic MDS can be expected to fail when the data are not arranged on or close to a linear manifold (hyperplane). This limitation is addressed by recent *manifold learning* techniques that attempt to deduce an approximately distance-preserving mapping from a potentially curved submanifold embedded in a higher dimensional space to a visualizable low-dimensional representation. Note that this sort of mapping is familiar in computer gaming technology in the form of the mapping from a point on the curving surface of a model in 3D space to the s/t coordinates in a planar texture map, and manifold learning-related algorithms are now being employed to automatically define texture charts with minimal stretching [Zhou et al. 2004]. While we do not have room to survey the existing algorithms (more than five distinct approaches exist currently), [Sudderth 2002] provides a survey that identifies some underlying relationships among several approaches.

Although manifold learning is generally considered as a mapping from given data points in a higher dimensional space, in several of the existing algorithms the original data is only used during calculation of pairwise distances. This is important, because it allows us to apply manifold learning to conceptual domains (such as this one involving the "space" of games) where the original points and their dimensionality is unknown, and only pairwise distances are

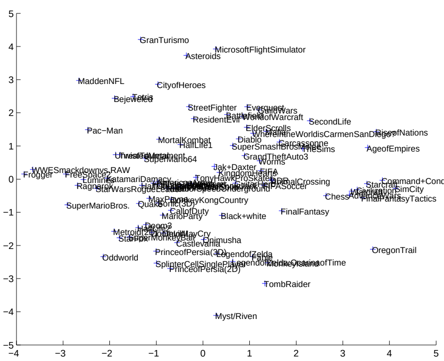

Figure 6: Projection of the eight-dimensional manifold on axes 2,4. Please enlarge to see details. See **Discussion** section for description.

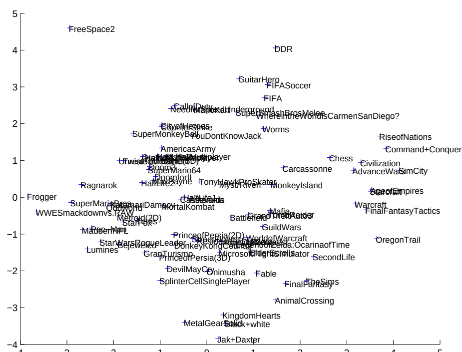

Figure 7: Projection of the eight-dimensional manifold on axes 2,3. Please enlarge to see details. See **Discussion** section for description.

obtainable.

In our experiments we use the *Isomap* algorithm [Tenenbaum et al. 2000]. This algorithm discovers data lying on a potentially curving lower-dimensional manifold embedded in a high dimensional space. In our application the high dimensional space consists of the many dimensions along which games can differ, and the lower dimensional manifold arises when the variation among a group of similar games can be effectively approximated with a smaller number of dimensions. In general one expects that the shortest path between widely separated points may depart from the manifold, whereas the distance between nearby points approximates the distance measured on the manifold (for example, the three-dimensional distance between cities on opposite sides of the globe is much shorter than their geodesic distance, but the distance between nearby cities is similar to their geodesic distance). Isomap thus identifies the (K) nearest neighbors of each point and assumes that the distances to these points are approximately the correct geodesic distances. Geodesic (on-manifold) distances between distant points are obtained by accumulating the shortest path between neighbors, e.g. using the Dijkstra algorithm. Once all pairwise distances on the manifold are known, MDS is employed to reconstruct the lower dimensional manifold. We used the closed-form spectral MDS described in [Kimmel 2003] in the first experiment reported next.

3 Method

As described above, in order to apply manifold learning to map game genres, we need a matrix of pairwise similarities for various games. Because the goal is to let relationships arise as much as possible from the data, we obtained ratings on actual game titles rather than on pre-defined genres (such as first-person shooter, etc.) that might unnecessarily influence the subsequent ratings. Defining similarity between even a large subset of the many existing games is clearly prohibitive however due to the $O(N^2)$ number of pairwise similarities, however.

To approach this issue we proceeded in two steps. First we undertook a pilot study using only a representative subset of well-known games. In consultation with several gamers we identified the following set of nine games, Quake, Oblivion, The Sims, World of Warcraft, Black and White, Starcraft, Half-Life, Unreal Tournament, and Civilization. To increase the probability that survey participants would be familiar with most of the games, two alternate titles were identified: Oblivion or King's Quest, and Civilization or Command and Conquer. This choice was based on the judgment that these titles are sufficiently similar to the originals (relative to the granularity of our sampling of titles) that they would serve as an equivalent choice. We included Chess and Monopoly to provide two prominent non-computer games as landmarks.

The resulting set of 11 games requires $N(N-1)/2 = 55$ questions to enumerate the pairwise similarities, a reasonable number. To obtain a sampling of these pairwise ratings, members of an on-campus gaming organization were invited to complete a web-based survey in which each pair of games was rated on a 1..5 scale, 5 being most similar. The survey was implemented as a web browser form, and eighteen people accessed the survey. Pairwise distances between games were obtained by negating the averaged similarity rankings (specifically, distance $D = 1 + \max(S) - S$ with S being the matrix of averaged similarity rankings, and $D_{i,i} = 0$).

Figs. 3-4 show results from this pilot study. The resulting three-dimensional map is shown in Fig. 3. This figure shows a plausible grouping and partial clustering of games: the first person shooters (Half-life, Quake, and Unreal) are tightly grouped and separated from the "wizards&dragons"-themed games (Everquest, Oblivion/King's Quest), which are in turn separated from games with a strategy emphasis (Civilization/Command and Conquer, Starcraft, Chess), and from "God games" (The Sims, Black and White).

As in Russell's mental map of emotions (Fig. 2), it is tempting to try to assign labels to the axes. With the reminder that the given orientation is arbitrary, several people familiar with games were able to assign interpretable axes to two-dimensional projections of this data (see Fig. 4). These early interpretations should be viewed with some hesitation, however.

The pilot study suggested that Isomap might give reasonable results when applied in this fashion. On this basis we undertook a larger scale survey. People were recruited from undergraduate classes,

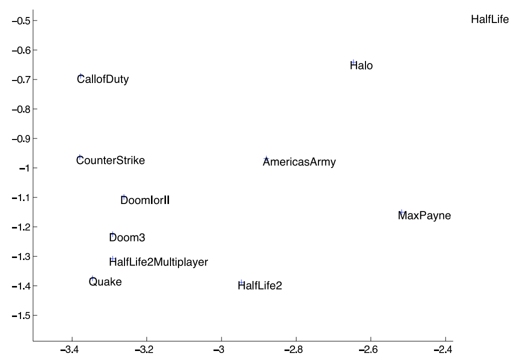

Figure 8: Clustering of FPS games appears clearly in this detail from the projecton on axes 1,2. Please enlarge to see details.

and using a posting to an on-line gaming site. The survey asks people to select 10 games from larger (and expanding) list currently containing 110 games. Participants then rated the pairwise similarity between these games (Fig. 1). After removing problematic responses, 124 participants rated 92 games. (Due to a bug, two versions of the MortalKombat series had the same data tag, and participants who selected both had to be omitted. For other participants, both titles were treated as equal representatives of the series).

These 92 games yield 4186 possible comparisons, whereas the 124 participants each giving 45 participants yielded 5580 comparisons. However, the coverage is uneven, and some game pairs are not compared. We thus altered the Isomap algorithm to use an MDS method that provides for missing data (Matlab's `mdscale`). An additional issue is that, since participants are selecting different subsets of games, the meaning of "similarity" varies across participants. That is, one person might select a subset that contains only FPS titles, whereas a second person could select a more varied sample of games. A rating of "very different" from the first person reflects a smaller actual "distance" than a similar rating from the second person. While this issue increases the variance of the resulting visualizations, we treat the similarity rating for a particular pair, averaged across all people who rated that pair, as an (unbiased) estimate of the desired distance. This practice of averaging across people is commonly required with psychological data.

Isomap also requires the manifold dimension as an algorithm parameter. As with other dimensionality reduction techniques, the true dimensionality of the data can be estimated by examining the decrease in representation error on the reconstructed manifold as its dimensionality increases (Fig. 5). In Fig. 5 we see that the embedding error for the full data set decays slowly with the embedding dimension with no obvious "knee", whereas the pilot study had a knee at three dimensions. We choose an 8-dimensional embedding for the figures of the full survey.

4 Discussion

Figures 6-8 show several interesting projections of the eight-dimensional reconstructed manifold. Games are traditionally grouped in stores and marketing data by platform (e.g., PS3, Windows, Wii, Xbox360) and genre (e.g., FPS, puzzle, adventure, racing). Surprisingly, this decades old scheme does not appear to actually represent the mental categories into which game players consciously or subconsciously group games. Instead, players appear to group games first by setting or aesthetic and second by the primary mechanic of the game. In other words, players appear to categorize by assets and primary mechanics, independent of how a game is marketed. This can be seen in Figures 6,7, as described next. On the other hand, clear clusters respecting existing genres are evident at small scales (Fig. 8).

In Fig. 6 the vertical axis roughly corresponds to aesthetic and the horizontal axis to the primary mechanic. First examine the vertical axis. At the top of the figure are the relatively realistic simulation games (Asteroids was fairly realistic compared to contempraries like PacMan). Below are a series of idealized, almost, comic-book like characters, including anime heros, fantasy warriors, and superheros (as well as football players). Below that are the goofy animated characters: frogs, gorillas, cows, and cartoon plumbers. Towards the bottom are the more family-friendly fantasy worlds of Riven, Oddworld, Price of Persia, and Zelda. Tomb Raider is an outlier on the "family-friendly" aspect, however it does exhibit a strong adventure-fantasy aesthetic that matches Myst and Zelda. Along the horizontal axis we see the primary mechanic of jumping and spatial puzzles (e.g., Frogger, Katamari, Super Mario Bros) give way to strategic puzzles (Tetris, Prince of Persia) and combat (Zelda, Diablo, GTA 3, Everquest) and finally to classic minimax planning with inhomogensous units (Age of Empires, Chess, Starcraft, Command and Conquer). This intuitive axis interpretation is imperfect because these notions of aesthetic and mechanic don't really form linear spectra.

The general classification argument is more strongly supported when directly examining the clusters. In Fig. 6, the center right contains the traditional "Real-Time Strategy" games; supporting the idea of genre, yet Chess, Oregon Trail, and Final Fantasy come very close to the cluster because the primary mechanic is minimax over inhomogeneous units, regardless of the "real-time", "combat", or "tech tree" aspects that characterize the RTS genre. In the center bottom, Splinter Cell, Prince of Persia, Tomb Raider, Zelda, and Castlevania cover different genres from modern spy combat to adventure, but all involve similar combat mechanics. In Fig. 7, we see the rhythm and apttern games DDR, Guitar Hero, and FIFA soccer clustered near the top. The 3D run and shoot games just above and right of the center–notice how Super Mario 64 becomes grouped with the traditional "FPS" games, as do third player perspectives like Max Payne because the mechanic is the same. Near the center bottom we see the anime-influenced aesthetic of Devil May Cry, Spinter Cell, Onimusha, Final Fantasy, Animal Crossing, and Kingdom Hearts.

Lastly, while a game can be characterized from any number of viewpoints including the underlying technology, assets (art, sound, text, story), and mechanics (the core game rules), and others, these are mostly independent of the publisher/producer's view of games as targeting a demographic or genre. In these figures we find that the player's mental space of games more closely matches the developer's categorization, although most players are probably unaware of the developer's process and choices.

5 Limitations and Conclusion

A major issue in this work (as in any manifold learning or dimensionality reduction technique) is that of whether the discovered dimensions are interpretable and represent useful aspects of the overall variation between entities. Psychological applications of dimensionality reduction have the additional challenge that the possible meaning of the resulting "dimensions" (if any) must be independently interpreted. A map of game genres as derived here admits both of these issues. Certainly games differ in an enormous number of possible "dimensions", and it may be the case that the variation expressible in a low-dimensional map is too coarse to serve as a useful summarization.

On the other hand, some applications of "mental maps" (such as that in [Lokuge et al. 1996]) do not require complete success from the point of view of these dimensionality reduction criteria in order to be useful; rather, it is only important that distance on the map roughly correspond to psychological distance. A recommender system in the form of a spatial map would be one such application. More generally it is known that people can absorb information more efficiently when it is coherently structured than when it is arbitrarily scattered, and we speculate that maps should be an effective form of information visualization for the gaming community, given the strong spatial skills required in many forms of game play.

The work shown here might be regarded as a suggestive experiment and illustration of an approach to deriving visualizations of conceptual maps. The maps shown in Figs. 6,7 should not be regarded as a definitive representation, however, both because of the limited size of the survey and number of games surveyed, and because alternate choices of various design parameters (e.g. the specific algorithms used to convert similarity to distance, and to select the nearest neighbors) will lead to maps that differ at least slightly. Nevertheless we hope that this work is a small step toward applying research methods to the domain of games and might lead to further work on formalizing game concepts.

References

CUKIER, J. 2006. For a new classification of game genres. Gamethink column, Oct 12, `gamethink.net/For-a-new-classification-of-game.html`.

ESA. 2005. Entertainment Software Association / NPD Group 2005 Sales and Genre Data. `http://www.theesa.com/facts/sales_genre_data.php`, viewed March 20, 2007.

JUUL, J. 2005. *Half-Real: Video Games between Real Rules and Fictional Worlds*. MIT Press.

KIMMEL, R. 2003. *Numerical Geometry of Images: Theory, Algorithms, and Applications*. SpringerVerlag.

LOKUGE, I., GILBERT, S. A., AND RICHARDS, W. 1996. Structuring information with mental models: a tour of boston. In *CHI '96: Proceedings of the SIGCHI conference on Human factors in computing systems*, ACM Press, New York, NY, USA, 413–419.

PLUTCHIK, R., AND CONTE, H., Eds. 1996. *Circumplex Models of Personality and Emotions*. American Psychological Association.

RICHARDSON, A., MONTELLO, D., AND HEGARTY, M. 1999. Spatial knowledge acquisition from maps and from navigation in real and virtual environments. *Memory and Cognition 27*, 4, 741–750.

RUSSELL, J. A. 1980. A circumplex model of affect. *J. Personality and Social Psychology 39*, 1161–1178.

SATALICH, G. 1995. *Navigation And Wayfinding In Virtual Reality*. Master's thesis, U. Washington HIT Lab.

SUDDERTH, E. 2002. Nonlinear manifold learning part II. `web.mit.edu/6.454/www/www_fall_2003/esuddert/manifold2_summary.pdf`.

TENENBAUM, J. B., DE SILVA, V., AND LANGFORD, J. C. 2000. A global geometric framework for nonlinear dimensionality reduction. *Science 290*, 2319–2323.

ZHOU, K., SYNDER, J., GUO, B., AND SHUM, H.-Y. 2004. Iso-charts: stretch-driven mesh parameterization using spectral analysis. In *SGP '04: Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, 45–54.