

Computing Point-of-View: Modeling and Simulating Judgments of Taste

by

Xinyu Hugo Liu

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning,
on 1 May 2006, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Media Arts and Sciences

Abstract

People have rich points-of-view that afford them the ability to judge the aesthetics of people, things, and everyday happenstance; yet viewpoint has an ineffable quality that is hard to articulate in words, let alone capture in computer models. Inspired by cultural theories of taste and identity, this thesis explores end-to-end computational modeling of people's tastes—from model acquisition, to generalization, to application—under various realms.

Five aesthetical realms are considered—cultural taste, attitudes, ways of perceiving, taste for food, and sense-of-humor. A person's model is acquired by reading her personal texts, such as a weblog diary, a social network profile, or emails. To generalize a person model, methods such as spreading activation, analogy, and imprinter supplementation are applied to semantic resources and search spaces mined from cultural corpora. Once a generalized model is achieved, a person's tastes are brought to life through perspective-based applications, which afford the exploration of someone else's perspective through interactivity and play.

The thesis describes model acquisition systems implemented for each of the five aesthetical realms. The techniques of 'reading for affective themes' (RATE), and 'culture mining' are described, along with their enabling technologies, which are commonsense reasoning and textual affect analysis. Finally, six perspective-based applications were implemented to illuminate a range of real-world beneficiaries to person modeling—virtual mentoring, self-reflection, and deep customization.

Thesis Supervisor: Pattie Maes

Title: Associate Professor, Program in Media Arts and Sciences

**Computing Point-of-View: Modeling and
Simulating Judgments of Taste**

by

Xinyu Hugo Liu

Thesis Committee:

Advisor _____

Professor Pattie Maes
Associate Professor of Media Arts and Sciences
Massachusetts Institute of Technology

Thesis Reader _____

Professor William J. Mitchell
Alexander W. Dreyfoos, Jr. (1954) Professor
Professor of Architecture and Media Arts and Sciences
Director, Design Laboratory
Massachusetts Institute of Technology

Thesis Reader _____

Professor Warren Sack
Assistant Professor of Film & Digital Media
University of California, Santa Cruz

Contents

Abstract	3
Aperitif	13
1 Introduction	14
1.1 Thesis summary	14
1.2 Roadmap and contributions	32
2 Approach	35
2.1 Related work in user modeling	35
2.2 Person modeling from everyday texts	38
2.3 Model acquisition	40
2.4 Model generalization	47
2.5 Model application	52
3 Techniques & Technologies	58
3.1 Reading for affective themes	58
3.2 Culture mining	69
3.5 Commonsense reasoning	74
3.6 Textual affect analysis	79
4 Acquisition systems: implementations & evaluations	85
4.1 Cultural taste realm: ‘taste fabric’	85
4.2 Food realm: ‘synesthetic cookbook’	93
4.3 Perception realm: ‘escada’	96
4.4 Attitudes realm: ‘what would they think?’	102
4.5 Humor realm: ‘catharses’	110
5 Perspective-based applications	113
5.1 Art that’s always tasteful	113
5.2 Facilitating social introductions	120
5.3 An identity mirror	121
5.4 Virtual mentors, and pundits too	125
5.5 Foraging for food with the family	131
5.6 A jocular companion	133
6 Philosophical underpinnings	135
6.1 Bourdieu’s framework from taste	136
6.1 Post-structural aesthetics	137
6.2 The aesthetic consistency hypothesis	140
6.3 The stability of viewpoint	143
6.4 Dialogics	146
7 Conclusion	151
References	153

List of Figures

Figure 1-1. Triptych summarizing the thesis's approach to modeling persons in aesthetical realms.	18
Figure 1-2. A walkthrough for cultural taste modeling	20
Figure 1-3. Topological features in the learned taste fabric—identity hubs (left) and taste cliques (right)	22
Figure 1-4. A walkthrough for generalizing attitudes	25
Figure 1-5. The space of possible ways of perceiving	27

Figure 1-6. Summary of built applications	31
Figure 2-1. A semantic sheet representation	43
Figure 2-2. A semantic diversity matrix	46
Figure 4-1. Mining algorithm for cultural taste space	86
Figure 4-2. Taste Fabric’s instance types and ontology sources	86
Figure 4-3. Evaluation of the generalized cultural taste model, and the effect of identity hubs and taste cliques, in a complete recommendation task	92
Figure 4-4. Results of ten-fold cross validation showing blog-level classification accuracies	100
Figure 4-5. Learned feature weightings for single-scale classification.	101
Figure 4-6. How reflexive memories get recorded from weblog excerpts	104
Figure 4-7. Political viewpoints of major newspapers	106
Figure 4-8. Evaluating model-based predictions of persons’ reactions to news stories.	109
Figure 5-1. A semiotic model of aesthetic transaction	114
Figure 5-2. The effects of perspective on artistic rendition – the words “sunset” (top-row) and “war” (bottom-row) are rendered with a Thinking-Seeing bias (left-column) versus with an Intuiting-Feeling bias (right-column)	118
Figure 5-3. Aesthetiscope’s aesthetic impressions of the four season keywords (columns) rendered through the monadic optics of each Reader taken alone (rows)	119
Figure 5-4, a-d. Reflections in the Identity Mirror – (clockwise from upper-left) as performer approaches the Identity Mirror, his reflection gains descriptive granularity, passing from subcultures (a), into genres and artists (b), into films and albums (c), into foodstuffs, activities, and songs (d)	122
Figure 5-5. A panel of virtual AI mentors react affectively to a passage of text that the user is reading	126
Figure 5-6. Clicking on virtual Roz Picard triggers an explanation dialog	127
Figure 5-7. Evaluation of WWTT in a person-learning task	129
Figure 5-8. Interactions with the Synesthetic Cookbook	132

List of Tables

Table 4.1. Excerpt from Synesthetic Cookbook's lists of ingredients and descriptive keywords	95
Table 4.2. Political culture: attitudes of the Democratic and Republican parties	105
Table 5.1. Signalling efficacy of Aesthetiscope's five dimensions	116

Aperitif

Wither the future battles of humankind? I believe they will increasingly be fought in the aesthetic plane. Media systems and commerce begin to unravel the use function of aesthetics, leading to more systematic productions of poetics. The willful construction of perspective, authenticity, and image up the ante in the worlds of ideology and marketing. Today's artificial intelligence personalizes search and recommends books, but tomorrow's will likely design our life-styles and proustian recommenders will select wines and spirits to release particular memories and desires submerged within each of us.

While the explicit topic of this dissertation is point-of-view, its underlying thematic is certainly aesthetics. A point-of-view, after all, may be recognized as a coherent and comprehensive system of aesthetics, efflorescing from the limitless ecology of aesthetics that is life. What renders point-of-view such a challenging study is precisely its complex etiology—just as each snowflake is constituted idiosyncratically by an unknown mixture of passing clouds, so too are our perspectives shaped by psychological predispositions, past experiences, and culture embeddedness.

Point-of-view and its aesthetics have long been studied poetically and rhetorically. Here, I embark upon yet another such philosophical investigation of the topic, this time informed by a computational perspective. Over the past four years, I have implemented a cadre of experimental computational systems, which automatically model and simulate particular persons' judgments of taste in various aesthetical realms. At some point in this process, the versatility of the approach that was being applied across these realms became aware of itself, and informed by semiotics and cognitive science, a methodology began to crystallize. So, these systems and their results are presented in this thesis—now unified under a common theoretical discourse and computational framework.

1 Introduction

Our capacity for aesthetics and affectedness is one of the most celebrated bastions of humanity. Underlying our explicit knowledge and rationality is a faculty for judgment – the impulsion to prefer, to view the world through our individual lenses of taste. An interesting intellectual question is: can a computer model a person’s tastes, attitudes, and aesthetics richly enough to predict their judgments? This thesis explores one answer to the question.

Our investigation flies under the banner of point-of-view for two reasons. Firstly, the term reflects an understanding that individual tastes are seated in, and articulated against a social and cultural fabric. Secondly, ‘point-of-view’ is developed to mean not isolated taste judgments, but rather, a coherent and systematic apparatus that engenders such judgments.

1.1 Thesis summary

Each person has their own tastes, attitudes, and ways of perceiving the world. I believe that these aesthetic dimensions of our selves are revealed in our everyday writings—weblog diaries, commentary-rich papers, social network profiles, instant messenger

conversations, personal emails, and so on. I focus on the genre of everyday texts, which offers first-person and self-expressive accounts of everyday happenstance, and I point out that in this Information Age, these texts are *par excellence* portraits of who we are. Unlike esoteric ‘user models’ of persons acting within the context of particular computer applications, everyday texts portray persons in the most general sense possible—their domain of discourse is social, pragmatic, everyday life. In this thesis, I show that it is possible to build models of persons’ tastes, attitudes, and ways of perceiving the world by reading their everyday texts. The accuracy of these models and their applications will also be addressed. The rest of this thesis summary 1) discusses related work; 2) introduces the approach taken; summarizes the premise, methods, and results of person modeling in the three primary realms of 3) cultural taste, 4) attitudes, and 5) ways of perceiving; and finally, 6) reviews six implemented applications for the produced models. Subsequent chapters will address all of these topics in detail.

§

Before describing the approach taken by this thesis, I will frame the modeling portion of this work within related work in user modeling, natural language processing, and semiotics.¹

In user modeling and its related literatures of user-adaptive systems and recommender systems, we find a great deal of prior work on predictive models of people. Two important paradigms are category-based models and behavior-based models. Category-based models—such as Elaine Rich’s (1979) groundbreaking book recommender system, GRUNDY—models users by first collecting a profile of attributes describing the user, and second generalizing the user model from these attributes, based on *a priori* stereotypes. For example, women 26-35 years old could stereotypically prefer “romance novels.” Whereas each attribute activates a different set of stereotypes with varying numerical strengths, the final recommendation is entailed by the strongest stereotype overall. The categorical approach is a sound one, but models built from *a priori* stereotypes tend to be overly generic descriptions of individuals. The other paradigm of behavior-based models is *a posteriori* and predominantly data-driven. Software sensors collect users’ behavioral traces through an application and then generalize each user’s traces via statistical inference. Key examples of the behavior-based approach include collaborative filtering (Shardanand & Maes 1995), and Bayesian goal inference models (Horvitz *et al.* 1998). This

¹ The literature review in the thesis summary is merely indicative. A more complete survey of related work is achieved in later chapters—for user modeling (Section 2.1), interactive user interfaces (Section 2.5), computational reading (Section 3.1), text mining and social network analysis (Section 3.2), commonsense reasoning (Section 3.3), and textual affect analysis (Section 3.4).

approach has seen a great deal of success, and has been applied to a great variety of recommender systems and intelligent tutoring systems. A caveat related to this approach is that behavioral traces are often specific to an application domain, and may not well describe a person *in general*. These two paradigms are complementary though, and some systems including Selker's (1994) COACH system and Orwant's (1995) DOPPELGANGER system demonstrates aspects of both. While many of the behavior-based systems are more statistical (e.g. collaborative filtering recommenders), systems such as the COACH intelligent tutoring system are also knowledge-based. DOPPELGANGER also exhibited aspects of both paradigms. The current behavior of users (e.g. 'hacking', 'writing', 'idle') was predicted with Markov models, thus adopting a behavior-based approach. Dynamic categorical models generated for communities that the user was a member of served to supplement each user's profile of interests; thus harkening to the category-based approach.

In the computational linguistics literature, we find related work on methods for computing the subjective and affective dimensions of text. Wiebe, Wilson, Cardie and others have taken a corpus-annotation approach to tracking subjectivity in third-person narratives (Wiebe 1994) and to characterizing the degree of subjectivity in textual passages (Cardie *et al.* 2003; Wiebe *et al.* 2004). A drawback of the corpus-annotation approach is that it is supervised. Another approach to computing textual affect is by exploiting dictionaries annotated with *semantic orientation*, or, *prior polarity*, of words and concepts (Turney & Littman 2003). With this method, numeric priors on words become building block for statistical estimation of larger pieces of text, and some research has nuanced the application of polarity dictionaries by considering negations and intensifiers (Grefenstette *et al.* 2004a; Polanyi & Zaenen 2004, Wilson, Wiebe & Hoffman 2005). Liu *et al.* (2003) suggested a complementary method for appraising text's event structure using common sense knowledge, in order to also account for event-level connotations, e.g. "be(person, fired)."

In the semiotics literature, we find helpful frameworks for structuring readings of text. Greimas's (1966) isotopy model of reading describes the reading process convergently, as *monosemization*. The meanings of words in a textual passage mutually disambiguate one another, converging upon a system of stable themes which represent the overall understanding had from the text. Paralleling Greimas's bottom-up movement from words to themes, Zholkovsky (1970b) describes a top-down movement from larger theme to specific *expression devices* which manifest bits of the theme. Finally, in literary stylistics, cf. Todorov's (1968) introduction, it is thought that the most important evidence for a writer's style and tastes can be found in the emotive dimensions of a text's themes – its *affective themes*, so to speak. Sack's (1994; 2001) SpinDoctor exemplifies how natural language processing techniques can be employed to operationalize semiotic frameworks. SpinDoctor was

able to detect the ideology implied by a news story by reading for ideologically motivated actor-role bindings (Greimas 1987)—e.g. “Oliver North” is an actor who may be bound to various roles, such as “criminal” or “patriot,” by authors with different ideologies.

§

My approach to modeling a person’s tastes, attitudes, and ways of perceiving from their everyday texts relates and contributes to each of these three bodies of literature. Existing textual affect analysis techniques (via dictionaries of word affects and via commonsense-based sensing) are leveraged as a building block technology to advance the greater goal of reading for stable affective themes emergent from a person’s everyday texts—it is hoped that these affective themes will accurately depict a writer’s tastes, attitudes, and ways of perceiving. However, it is recognized that what directly results from such readings will be sparse and disconnected. To generalize a more comprehensive model of a person from fragmentary textual evidence, the person’s evidence is situated in cultural patterns of tastes and attitudes that underlie our society. To acquire these cultural patterns, we engage in comparative readings of hundreds of thousands of people’s everyday texts, finally arriving at a ‘map’ of culture’s topological space. Each person’s textual evidence is located on the map, and can be generalized by activating the neighborhood that surrounds the person’s location. The idea of locating persons in the cultural space inherits from both the behavior-based paradigm and the category-based paradigm. The topology of the cultural space is acquired by statistical inference, but in essence, the knowledge embodied by the topology is still a stereotype, albeit a data-driven one.

This thesis considers models of persons within five aesthetical realms. The three primary realms which are narrated in this thesis summary have already been enounced—attitudes, cultural tastes, and ways of perceiving. Two other realms discussed only in later thesis chapters are taste for food, and sense of humor. These five aesthetical realms are not claimed to be canonical, nor are they claimed to be independent in scope. Rather, realms were chosen opportunistically based on available everyday texts, and my own sense of interesting directions to explore.

A brief description of each realm is now given. An attitude is defined as an affect *about* some topic, so the realm of attitudes considers a person’s feelings toward every topic under the sun. In the realm of cultural tastes, a person is seen as a pattern of consumption, over a field of consumer interests and identities (e.g. books, music, films, foods). In the realm of perception, a person’s psychological dispositions are modeled as a coordinate location in a four-dimensional Jungian space whose axes are think, feel, sense, and intuit—proposed by Jung (1921) as four basic psychological functions. A perception model addresses the question—is the

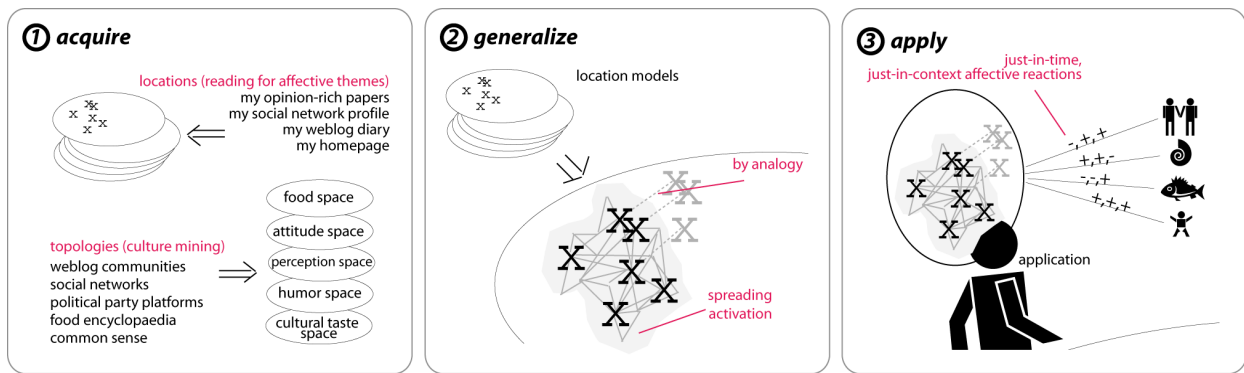


Figure 1-1. Triptych summarizing the thesis’s approach to modeling persons in aesthetic realms.

person a realist who senses and thinks or is she a romantic who intuits and feels? In the realm of food, a person’s taste is modeled as a pattern of liking and disliking over a densely connected semantic fabric of foodstuffs (e.g. flavors, sensations, ingredients, dishes). The realm of humor is motivated by Freud’s (1905) theory of *tendentious jokes* as outlets for psychic tensions. Thus, a person’s humor model resembles an attitude model—it considers the person’s psychic tension levels toward every topic under the sun.

The person modeling process which is applicable to each of the five realms can be summarized as three steps (Figure 1-1)—acquire, generalize, and apply. Each of the five realms has a topological space, and some of these spaces need to be modeled by employing culture mining to analyze cultural corpora—for example, the connectedness of nodes in the semantic fabric of cultural interests is acquired by analyzing pairwise affinities between interests, across 100,000 social network profiles. A person’s everyday texts are read for stable affective themes, and this textual evidence is regarded as *location* in the topological space. In a second phase, a generalized model of a person is produced. Her location is expanded into a more general surrounding neighborhood, which is implied by the topology of a particular realm. Spreading activation and analogy are two methods used to perform this generalization. In the third phase, the generalized model is employed by a range of applications to simulate the taste, attitude, or perceptual perspective of a person on arbitrary input.

To illustrate and concretize the described approach, the next three subsections narrate the person modeling process for the three realms of cultural taste, attitudes, and ways of perceiving.

§

To model a person’s cultural tastes, social network profiles were focused upon as an everyday text most suitable for acquiring person models. Today tens of millions of internet users are members of

social networking sites—such as Myspace², Friendster³, Orkut⁴, and Facebook⁵-- on which they maintain a text profile of their cultural interests (e.g. favorite books, music, films). In modeling cultural tastes, we assume that tastes are not arbitrary—that there is an unconscious gestalt that unifies each person’s selections of cultural interests. This coherence assumption is consistent with recent consumerist theory stating that people’s consumptive choices tend to form gestalts—McCracken (1988) termed these “Diderot unities” and Solomon & Assael (1987) called them “consumption constellations.” Proceeding from this premise, cultural taste modeling aims to analyze a corpus of 100,000 social network profiles, and extract from them a topology of cultural tastes, which can be used to generalize a single person’s profile.

An algorithm for ‘culture mining’ from 100,000 social network profiles is now described. Each social networking website, although having a somewhat different design, does observe certain conventions in how it elicits and displays users’ profiles. The convention is that a user’s interests are broken down into organizing categories, the most common being ‘books’, ‘music’, and ‘movies’. Within each category, interests are typically given as a token-delimited list. There is typically also an overarching category, called variously ‘passions’ and ‘general interests’—to emphasize their importance, our present modeling maps these descriptions into an ontology of identity descriptors (e.g. ‘fashionista’, ‘book lover’). Figure 1-2a shows a typical social network profile, taken from the Orkut social networking site—identity and interest categories are depicted. A first step of processing is to normalize as many of these natural language fragments as possible, using ontologies of interests and identities assembled from folksonomies such as DMOZ⁶ and Wikipedia⁷. Figure 1-2b supposes that some subset (the four nodes shown in black) of the profile’s descriptions have been mapped into recognized identity and interest descriptors. In actual processing of the 100,000 profiles, the rate of recognition was 68% of natural language fragments, including 8% false positives. Also shown in Figure 1-2b are red edges and red nodes—these are metadata associated with the black nodes. They are added to each user’s profile to improve the robustness of the profile, but at a discounted strength of 0.5 per link-hop traversed away from a black node.

² <http://myspace.com>

³ <http://friendster.com>

⁴ <http://orkut.com>

⁵ <http://facebook.com>

⁶ <http://www.dmoz.org>

⁷ <http://wikipedia.org>

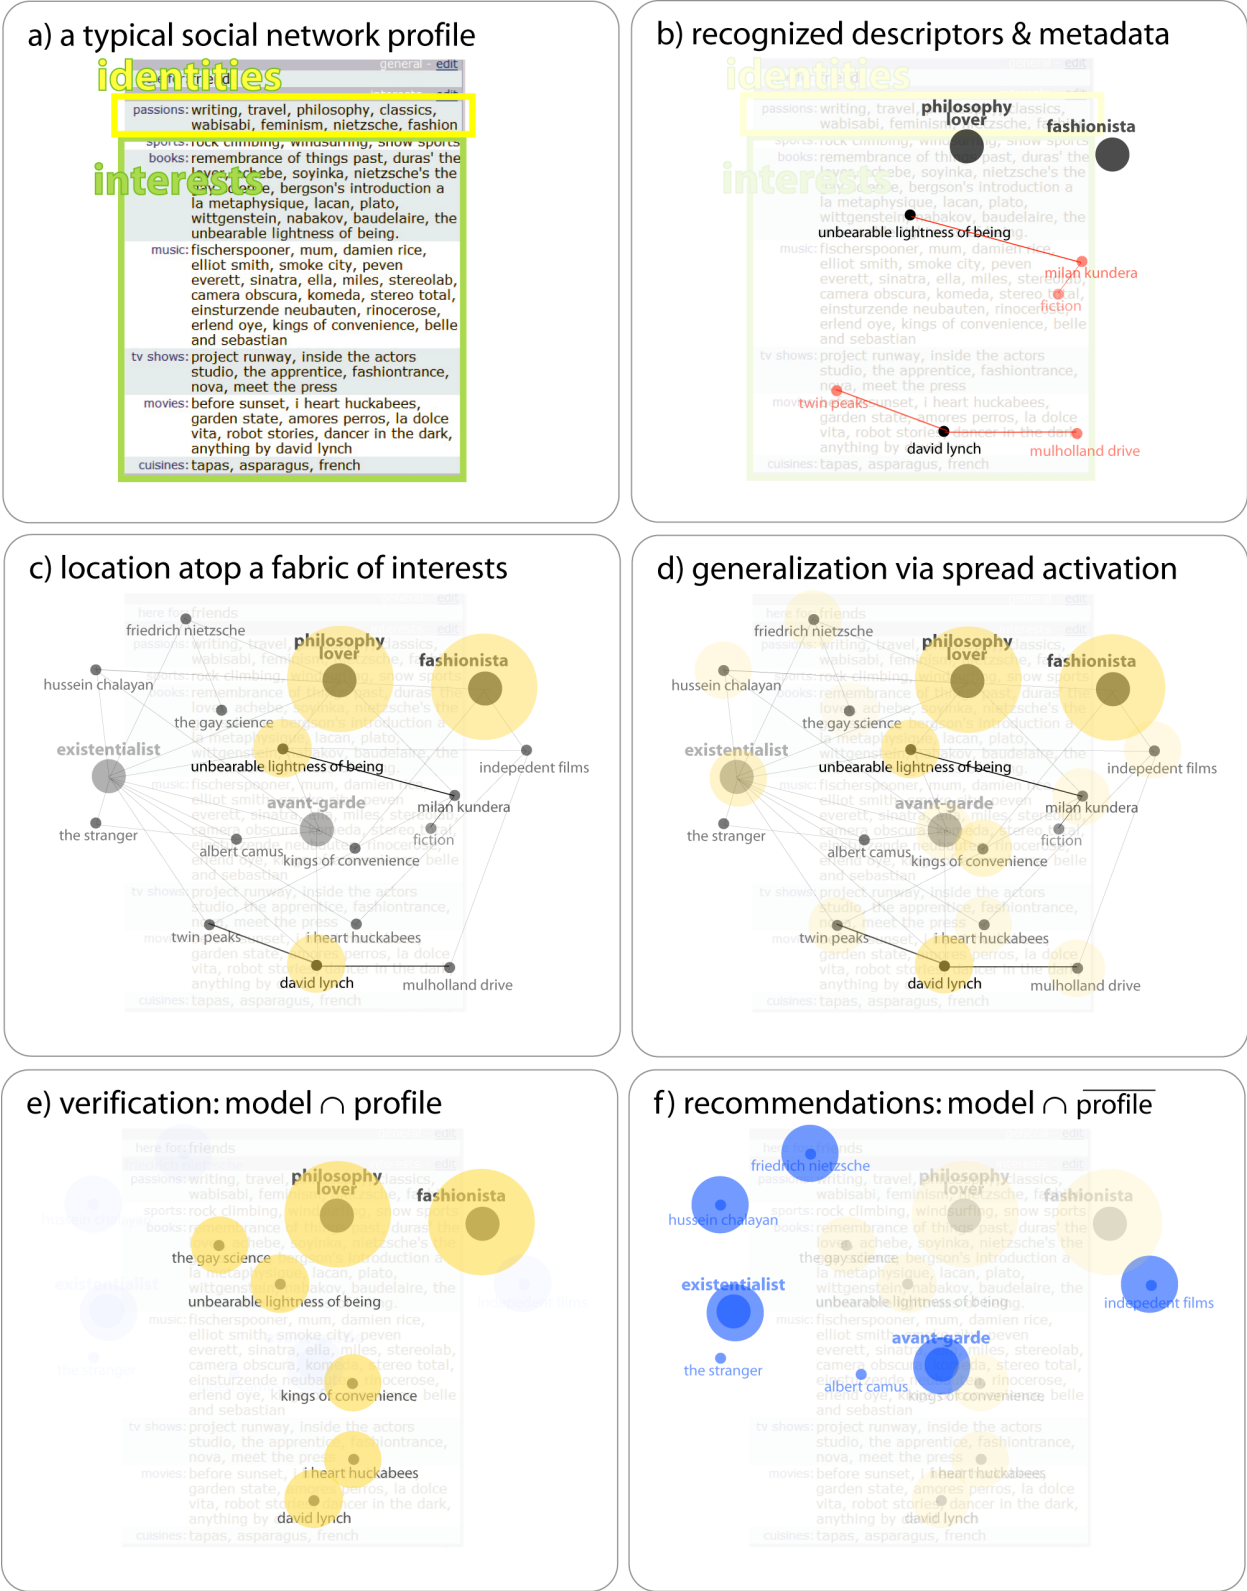


Figure 1-2. A walkthrough for cultural taste modeling

Next is a machine learning step which operationalizes the aforementioned assumption that each profile has taste coherence. The goal is to learn the numerical strength of affinity between every pairwise combination of interest and identity descriptors. In a parsed profile, each possible pairwise combination of descriptors is recorded as a co-occurrence. Using the pointwise mutual information (PMI) (Church & Hanks 1990) measure of semantic similarity, the aggregate co-occurrence data for the 100,000 profiles is analyzed, and a PMI score representing affinity is calculated for each pair of descriptors. After pruning, what results is a 12,000 by 12,000 correlation matrix of learned affinities—I will call this a *semantic fabric of cultural taste* (or *taste fabric* for short) to emphasize the graph’s density and to re-introduce the spatial metaphor.

Using this taste fabric which has captured cultural patterns of taste, a person’s profile can be generalized. An algorithm for producing a generalized model of a person’s cultural tastes is now described. Suppose the profile shown in Figure 1-2a is inputted to the model generalizer. First the profile is segmented into natural language fragments, those fragments are normalized via ontology recognition (Figure 1-2b) and metadata is added at a discount, and the normal and metadata nodes are located into the just-acquired taste fabric (Figure 1-2c). Via spreading activation outward from these nodes along the fabric’s edges (e.g. discount = 0.5), a generalized model of the person’s taste is produced (Figure 1-2d). It is well described as an *activation cloud*, or, to emphasize that it has captured the general interests of the person, we can call this a *taste ethos* formation. It is useful to illuminate two interpretations of this generalized model. Taking the set intersection between this model and the original profile (assuming that its contents were fully normalizable), Figure 1-2e shows that in addition to the recognized descriptors, the generalized model has rediscovered other descriptors that were present in the profile, but failed to be recognized; a soon to be described evaluation treats these rediscovered nodes as verification that the algorithm has worked properly. Finally, Figure 1-2f shows the complementary set of nodes which are not in the profile—we can interpret these as the identities and interests that are suitable to offer up as recommendations.

Some interesting results from taste fabric production are its emergent topological features such as cliques (i.e. a cluster of nodes with strong mutual connectedness) and hubs (i.e. nodes with strong connections to an unusually high number of other nodes, sometimes called ‘stars’). Figure 1-3 depicts two topological features—the *identity hub* ‘existentialist’ in the left pane, and two *taste cliques* straddled by ‘brian eno’ in the right pane. To produce the unexpectedly sparse semantic network shown in this visualization, most of the graph edges in the original 12,000 x 12,000 correlation matrix were thresholded away, until only thousands of the strongest edges remained; edge strength is not shown in the visualization. Identity hubs and taste cliques are interesting because they are disproportionately influential features in the spreading activation

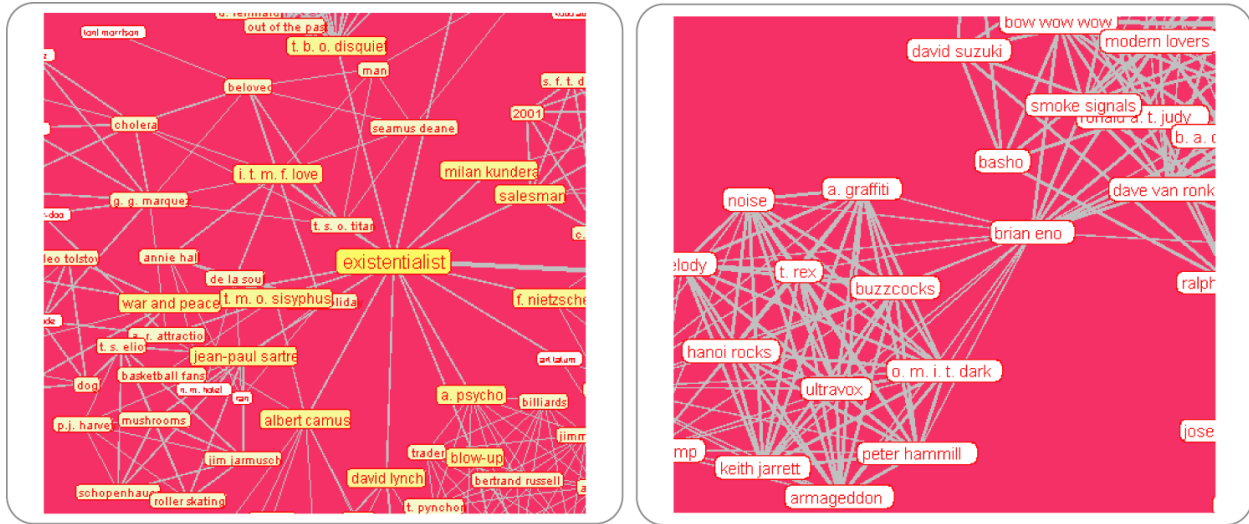


Figure 1-3. Topological features in the learned taste fabric—identity hubs (left) and taste cliques (right)

process. A telling statistic is that after profiles are normalized into ontology, each identity descriptor in the ontology occurred in the corpus of 100,000 profiles on average 18 times more frequently than the typical interest descriptor. We hypothesize that in the graph, identities behave as *indexicals*, serving as hubs around which interests organize—as shown in Figure 1-3’s left pane, ‘existentialist’ is an identity hub which can be seen to unify a variety of interests matching intuition, such as ‘albert camus’, ‘friedrich nietzsche’, ‘death of a salesman’, etc. Likewise, taste-cliques shown in the right pane may behave like identity hubs in spreading activation, because as the size of the clique increases, so does its influence—measured as the number nodes leaving the clique. We may think of taste cliques as unnamed *indexicals*. That taste cliques and identity hubs seemingly organize the taste fabric is consistent with consumer theorists’ observations that taste is shaped under ‘Diderot unities’ (McCracken 1988) and ‘consumption constellations’ (Solomon & Assael 1987).

The accuracy of the generalized model of cultural taste was evaluated against three control systems in a five-fold validation experiment over the corpus of 100,000 profiles. Given the task of producing a *complete recommendation*—total rank-ordering of interest and identity descriptors, a control system approximating item-item collaborative filtering (Sarwar *et al.* 2001) yielded an accuracy of 0.74, which was exceeded by our system’s accuracy of 0.86. It was also found that identity nodes improved recommendation accuracy by 0.05, and taste cliques were also beneficial, though less so. These results provide one measure of validation of the described method for generalizing a model of persons’ cultural tastes from their social network profile.

§

We now shift gears to consider modeling a person's attitudes from their everyday texts. We begin with a working definition for attitude—an *attitude* is a discourse topic, imbued with some affective tint; that is to say, attitude is how you feel *about* some thing, some one, some event. This definition is consistent with Ortony, Clore & Collins's (1988) position on emotion—that they always result from cognitive appraisal *about* some thing, some one, some event. How is the affective tint quantified? We chose Mehrabian's (1995b) PAD model of affect—which considers affect in a three dimensional Cartesian space whose axes are pleasure-displeasure, arousal-nonarousal, dominance-submissiveness. To model a person's attitudes, we regarded attitudes as affective themes to emerge from readings of weblog diaries and commentary rich papers. By strategically choosing everyday texts, which are defined as first-person and self-expressive, we hoped to avoid the difficult problem reported by Wiebe (1994) of tracking character viewpoints in third-person texts. With everyday texts, it is a reasonable assumption to attribute affective themes to the writer's attitudes.

The algorithm used for Reading for Affective Themes (RATE) is now described. The algorithm is consistent with Greimas's (1966) isotopy model of reading, which stipulates that in reading, meaning is convergent, and the end product of reading is an isotopy—a system of stable themes (Greimas called these *classemes*). In the beginning is personal everyday text, for example, a weblog diary. The diary is segmented, tokenized, and a surface syntactic parse of the text is made with the MontyLingua natural language processor (Liu 2002). The diary's original space of words are now reduced to keywords, key phrases, and parsed event structures, where event equals verb(subject, object, indirect objects*). Each of these textual entities are polysemous and taken alone, they have numerous mutually incompatible connotations. However, by intersecting all the entities' connotations, certain threads of consistent meaning emerge, and these include discourse topics. Accompanying each textual entity is also an affective characterization scored in the PAD model—for example, (P:+0.9,A:-0.2,D:+0.5) might correspond with the affect 'smug'. Just as textual entities are summed up into emergent topics, so each textual entity's PAD affect can be statistically averaged into the emergent topic's stable affect. The association of topic plus its PAD affect is called the stabilized attitude. Proceeding along these lines, a system of attitudes emerges from the weblog diary.

Reading for affective themes in this simple associative way was accomplished with lightweight natural language processing tools, embodying a knowledge-based approach. Natural language parsing tasks were all performed by MontyLingua (Liu 2002). Topic identification was performed using the `guess_topic()` function in the ConceptNet (Liu & Singh 2004) commonsense reasoning toolkit, augmented with topic hierarchies mined from DMOZ. Eagle *et al.* (2003) had made similar use of ConceptNet for topic spotting in spoken conversations. Affective analysis of text in terms of the PAD

model was achieved by hybridizing a superficial sensing approach with a deep sensing approach. Superficial sensing utilized lexical affect dictionaries which state prior polarities for words—one such dictionary employed was ANEW (Bradley & Lang 1999) and another custom-built dictionary was created out of sentiment headword classes in Roget's (1911) English Thesaurus. For deeper textual analysis via common sense inferences about event structures in the text, Emotus Ponens (Liu, Lieberman & Selker 2003) was used.

The explicit descriptors extracted from a social network profile were an incomplete model; likewise, the explicit system of attitudes read out of a weblog diary are also incomplete, and require generalization. An algorithm for producing a generalized model of a person's attitudes is now described. One way to visually represent a person's location in the space of possible attitudes is via *semantic sheets*. As shown in Figure 1-4a, the upper sheet is a grid enumerating all possible topics. A person's attitudes are captured in the lower sheet, as a grid of PAD affects associated with each possible topic. Walking through generalization, the process begins with the results of RATE being plotted onto these semantic sheets. Figure 1-4b supposes that two stable attitudes resulted from such a reading—general arousal (P:?, A:+, D:?) about 'feminism' and anger (P:-,A:+,D:+) about 'drugs'. A first generalization technique is to propagate these attitudes to topics closely related to 'feminism' and 'drugs'. Spreading activation along topic hierarchy lines (gotten from DMOZ and ConceptNet) achieves this (Figure 1-4c). Note that certainty of the spread attitudes is discounted. A second technique is structure-mapping analogy (Gentner 1983). Figure 1-4d shows that an attitude about 'aids' is inferred by analogy with 'cancer', based on the topics' shared attributes. ConceptNet is used to perform analogy. Certain cases of analogy tend to produce wrong inferences, such as inferring that attitudes about 'dog' will translate to the taxonomically and functionally similar 'cat'. These exceptions are discussed in Chapter 2. Next, Figure 1-4e depicts the idea that your attitudes can be supplemented by introjecting the attitudes of your *imprimers*—who Minsky (forthcoming) describes as those whose goals and values you mimic, such as mentors and parents. Currently, the identification of imprimers requires supervision. An algorithm identifies imprinter candidates by reading for the co-occurrence of self-conscious emotions (e.g. embarrassment, pride) with mention of persons, but as yet, a corpus of texts cannot be automatically assembled for imprimers. Figure 1-4e shows that an imprinter's explicit attitudes can align with and supplement the model of the one who is imprinted (these also occur at a discounted certainty). Finally, we can see that the result of generalization (Figure 1-4f) has spread the original explicit attitudes into its surrounding semantic neighborhood.

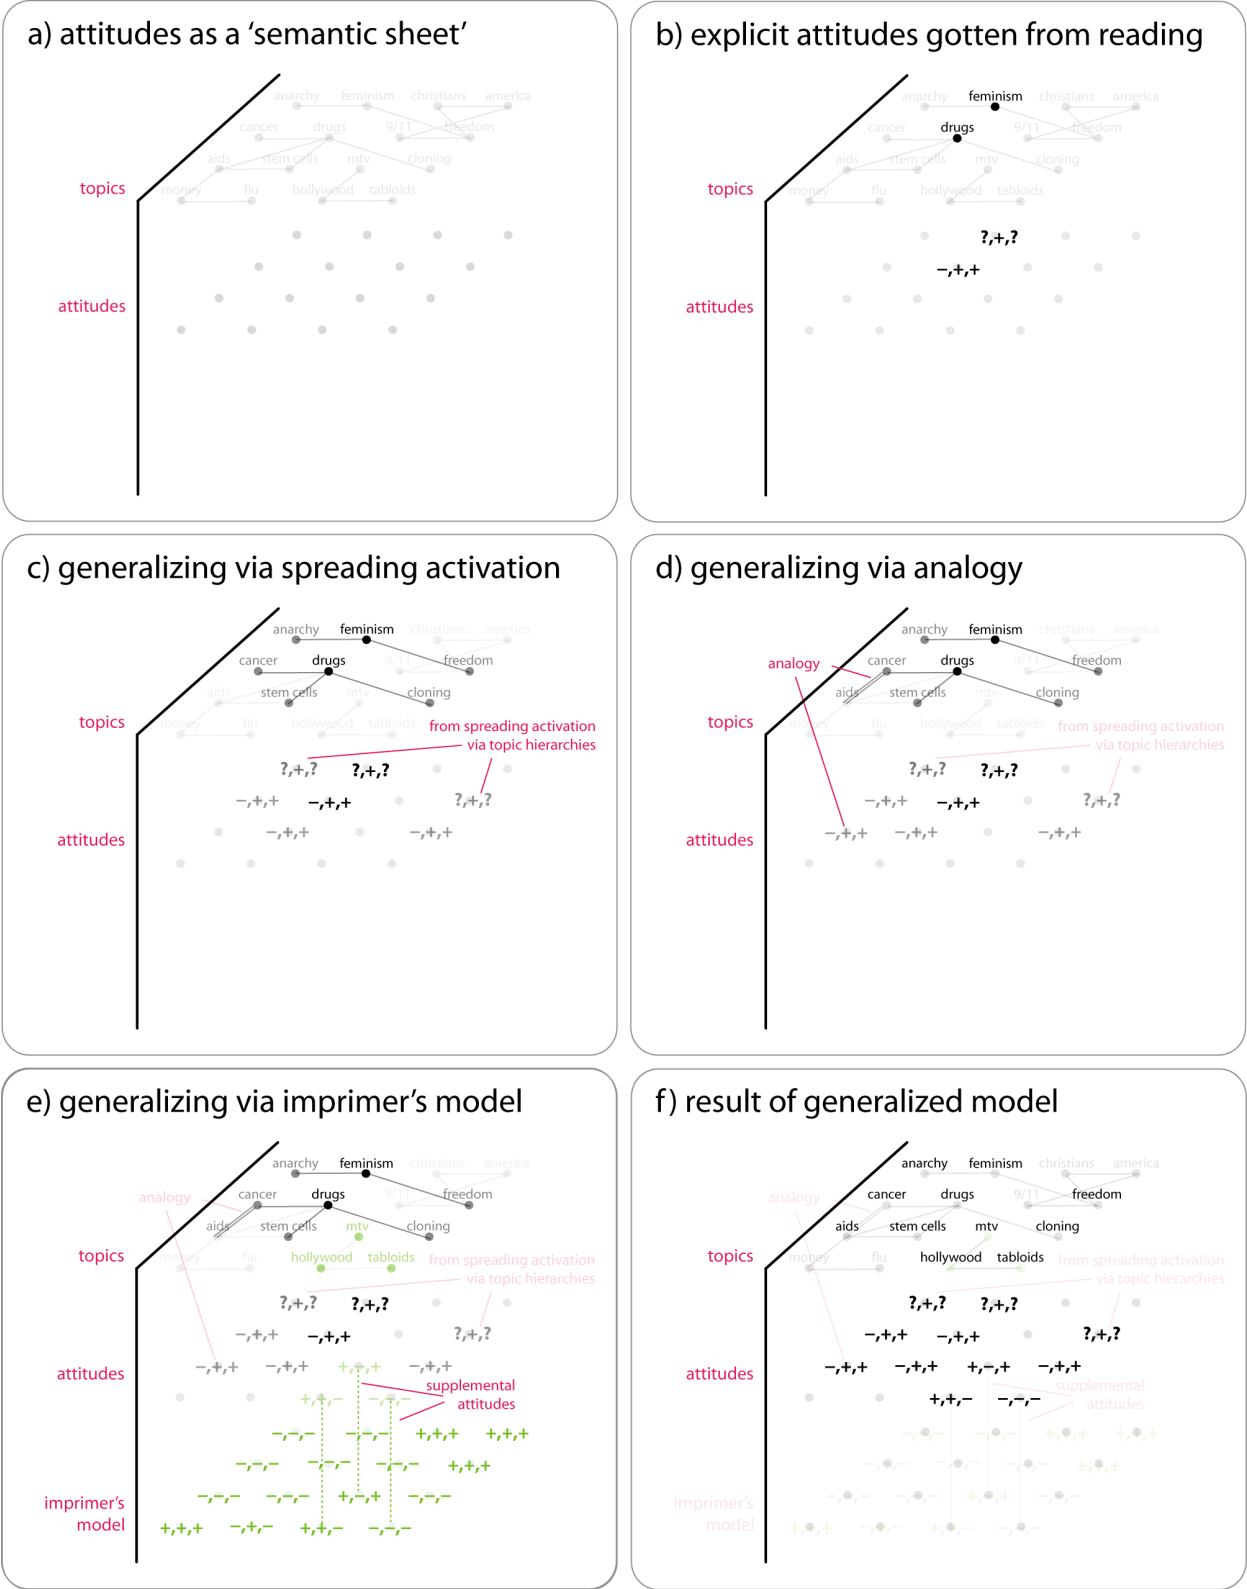


Figure 1-4. A walkthrough for generalizing attitudes

To evaluate the acquired model of attitudes, two experiments were conducted—their results are summarized below. A first experiment was conducted for political culture. Two political corpora were assembled from Democratic Party and Republican Party political speech transcripts and these were used to generate a “Virtual Democrat” and a “Virtual Republican.” Examination of their most positive and most negative attitudes were consistent with intuition, though some attitudes were seemingly incorrect. For example, Virtual Democrat held very negative attitudes about ‘god’ and ‘elderly’, and Virtual Republican held a very positive attitude about ‘poor’. These illustrate one drawback of the associative approach to reading for affective themes—Democrats were actually negative about the invocation of ‘god’ in civil affairs, but that nuanced attitude was incorrectly generalized to ‘god’. Using Virtual Democrat and Virtual Republican to define the poles of a political spectrum, the political bias of six major U.S. newspapers were calculated as a function of their degree of alignment with either pole’s attitudes. With some scale normalization, the results of this experiment were found consistent with a previous study which specified the media bias of these same newspapers (Groseclose & Milyo 2004). The only major discrepancy was that “Wall Street Journal” was found to be conservative leaning using the thesis system while the prior study found it to be liberal. With some investigation into political theory, the discrepancy was explainable—the thesis system looked at the editorial texts of ‘Wall Street Journal’, while Groseclose & Milyo looked at news articles.

A second experiment evaluated the accuracy of attitude prediction using the generalized model against human raters. Measuring the deviation of the model’s PAD reaction to news articles against the actual PAD reaction of the corresponding human raters, it was found that Arousal (A in PAD) prediction was most accurate, with an average deviation of 0.22 (out of 2.0 max). The prediction of Pleasure and Dominance was promising, but variance was large, and their 95% confidence intervals overlapped with one of two baselines. One interpretation of these results is the Arousal is easier to predict than Pleasure or Dominance because it is more amenable to additive calculation. More details of the studies are presented in Chapter 4.

§

We describe person modeling experiments for a third realm—the realm of perception. Consider that persons are disposed to perceive and engage the world in different ways. For example, realists and romantics seemingly interpret the world in antithetical ways. But what sort of theoretical framework can betray the differences between realists and romantics? With a salutary spirit, Carl Jung’s (1921) theory of psychological type was adapted as the framework to be experimented with. Jung proposed four basic psychological functions—sense, intuit, feel, think—and suggested that each person is disposed to engage these functions to various degrees. A most

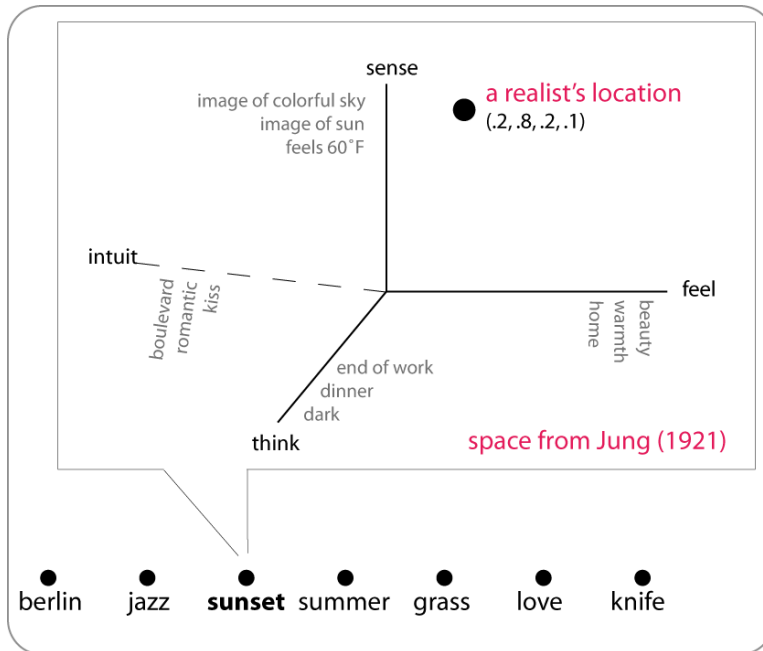


Figure 1-5. The space of possible ways of perceiving

simple model then, would be to take each function as an axis in a four-dimensional Cartesian space. Assuming each axis ranges from 0.0 (not engaged) to 1.0 (fully engaged), a four-tuple coordinate would indicate a person's location in perception space. As Figure 1-5 shows, the 4-d model of perception simulates a person's interpretive process by acting as a 'prism' that refracts meaning. For example, as shown the prism is passing over the topic of 'sunset'. Along each axis gather keywords forming one interpretation of sunset. For example, 'feel' interprets 'sunset' as 'beauty', 'warmth', 'home'. The system predicts that a realist, located at (0.2, 0.8, 0.2, 0.1), would roughly adopt 20% of keywords from 'feel', 80% of keywords from 'sense', and so on. A person's way of perceiving is thus taken to be a mixture of the Jungian interpretations.

How can such a model be acquired by reading a person's everyday texts? A first-pass experiment was completed, to see if affective patterns of communication in a weblog could implicate a rough location of the person in the Jungian space. In this proposed reading, the affective themes being read for differ slightly from the reading scheme proposed for attitude modeling. The following affective themes resulted from the proposed reading:

- ❖ EGO-PAD (writer's average PAD-level)
- ❖ ALTERS-PAD (other persons and things' PAD-level)
- ❖ INCOMING-PAD (PAD flowing from alters into ego)
- ❖ OUTGOING-PAD (PAD flowing out from ego into alters)
- ❖ MENTAL-ACTIVITY (frequency with which mental hypotheticals were invoked, e.g. "I thought that")

- ❖ INTROVERSION-EXTRAVERSION-RATIO (ratio of passive acts e.g. ‘resent’ to active acts e.g. ‘murder’)

These themes are further explained in Chapter 4. A blog corpus for 3800 bloggers with known MBTIs was analyzed for the above statistics. MBTI is the Myers-Briggs Type Indicator (Briggs & Myers 1976), a popular psychological inventory of personality, and is based on Jung’s theory of psychological types. MBTI performs binary classification of persons on four scales—Extravert-Introvert (E-I), Sense-Intuit (S-N), Feel-Think (F-T), Judge-Perceive (J-P). For our purposes, only S-N and F-T classification are relevant. A machine learning algorithm called BoosTexter (Schapire & Singer 2000) was fed the affective theme features and from the 3800 blogs, learned binary classifiers for S-N and F-T. Evaluating this classification with ten-fold cross validation, average classification accuracies were 0.58 for S-N, and 0.62 for F-T. These exceeded a 0.5 lower bound corresponding with guessing, but were well short of upper bounds of 0.85 for S-N and 0.73 for F-T, estimated from MBTI’s five-week test-retest reliability statistics (Myers & McCaulley 1985), which represent a fundamental limitation on the stability of the MBTI scales. Looking at a decomposition of the classifier’s learned features, for S-N classification, affective exchange (incoming and outgoing PAD) was more important than ego and alter’s affects; for T-F classification, ego’s affect was more important than alter’s affect or exchanged affect. This result was nicely consistent with intuition about these Jungian dimensions. The overall result is promising, but this computational reading approach is still a ways away from being viable for real-world application.

§

Up to now we have explored how models of a person’s tastes, attitudes, and ways of perceiving can be acquired and generalized. The results of evaluation proved the general promise of the approach, but also illuminated some weaknesses, which suggest an agenda for further work. Another more tantalizing ‘result’ is a slew of implemented applications driven by these person models (Figure 1-6)—these concretize and motivate the thus described modeling. What is a person model in the aesthetical realms good for? They enable new tools for self-reflection, person learning, and deep customization. The six implemented applications are now introduced. Subsequently, I distill some interaction design principles gotten from reflecting on building these artifacts.

What Would They Think? (WWTT) (Liu & Maes 2004) is a panel of virtual mentors who reside on the computer desktop, and offer users just-in-time affective feedback. Figure 1-6a shows WWTT configured to display a panel of mentors from Artificial Intelligence—(from left to right) Rodney Brooks, Seymour Papert, Rosalind Picard, Marvin Minsky, Douglas Lenat. WWTT is envisaged as a novel way to learn about a mentor’s points-of-view.

As the user browses web pages, and types emails and papers, the virtual mentors continuously observe the user's read-write textual activity. The user's present textual context becomes input for these virtual mentors, and according to each mentor's generalized model of attitudes, a simulated reaction is produced to the input. Reactions, in the form of a PAD value, are graphically rendered according to this visual metaphor—green for pleasure, red for displeasure; brightness for arousal; blurry if submissive, and sharp if dominant. To learn more about the justification for a reaction, reacting mentors can be double-clicked—bringing about an offering of quotes from that mentor's personal texts that best justify their reaction. WWTT's suitability as a tool for person-learning was evaluated in a 36-person user study. The task was to answer a multiple choice test about strangers' personalities, explicit attitudes whose evidence was located in their weblog diaries, and implicit attitudes whose evidence is not stated in their diaries, but whose correctness was verified by each diary author. Study participants formed three groups—one browsed strangers' weblogs, another used a text-search version of WWTT, and a third used the full WWTT system. Results of the study showed with statistical significance (95% confidence) that WWTT users outperformed text-search WWTT users on personality questions; WWTT users outperformed both text-search WWTT and weblog users on questions about explicit attitudes. All three groups have equally poor performance on implicit attitude questions, which tested knowledge of attitudes not stated in the weblog diary. These results are strong evidence for the utility of models of personal attitudes presented in the WWTT interface.

The *Identity Mirror* (Liu & Davenport 2005) shows you your cultural identity as a swarm-of-keywords (Figure 1-6b). Monitoring your social network profile and generalizing your taste ethos from that profile produces a visualization of the person's cultural identity in the form of an abstract mirror. Basic image recognition and image tracking was implemented to allow a viewer to interact with his reflection. Walking to and fro, the viewer traverses granularities of identity, from broad descriptions (e.g. identity descriptors, music genre, book genre) of a far-away viewer, to specific descriptions (e.g. song titles, book titles, film titles) when the viewer is up-close. Finally, a viewer's reflection is time-variant—representing the understanding that as culture's priorities and desires change, so does cultural identity, which is always articulated against cultural priorities. An advantage of the *taste ethos* representation is that it is can be biased by activating contextual nodes in the taste fabric to represent the present concerns of culture. A topic parser monitors live news feeds and automatically biases the taste fabric. For example, during Oscars season, nodes relating to film and entertainment are activated in the taste fabric, and as a result, the mirror's reflection reveals a more glamorous entertainment-oriented facet of your taste ethos.

The *Aesthetiscope* (Liu & Maes 2005b; Liu & Maes 2006) is a perspective-driven abstract art bot (Figure 1-6c). In the spirit of Ellsworth Kelly and early twentieth century abstract expressionist artwork that took the form of a semantic color grid, the *Aesthetiscope* renders inspirational texts (e.g. a word, a poem, song lyrics) into abstract color grid artwork. A model of the viewer's preferred ways of perceiving and interpreting the inspirational text controls the chosen combination of colors. The viewer's model is given as four 0.0-1.0 numbers corresponding to the viewer's disposition for the four scales—think, feel, intuit, and sense. Just as wine may be chosen to pair with foods, the *Aesthetiscope* generates perspective-specific artwork that pairs with a user's music playlist and choice of poems. A user study of the *Aesthetiscope*'s generated artwork validates the claim that artwork that is specific to an inspirational text has greater aesthetic efficacy than artwork with mismatched text. An interesting implication of the *Aesthetiscope* beyond its deep customization is the ability for viewers to communicate their differing perspectives to each other, and to explore the intersection of their perspectives.

The Synesthetic Cookbook (Liu, Hockenberry & Selker 2005) is an interactive recipe browsing interface backed by a mined semantic fabric of food consisting of 60,000 recipes, 5000 ingredient keywords, 1000 sensorial keywords, 400 cooking procedures, 400 nutritional terms. In the cookbook, *virtual tastebuds* (Figure 1-6d) simulate persons' reactions to recipe selections. A tastebud in food space is akin to *cultural taste ethos* in cultural taste space. A virtual tastebud is represented as an activation cloud, and is generated by spreading activation across the food fabric, from a starting profile of a user's likes and dislikes. To account for the importance of disliking in taste for food, the food fabric implements negative activations, or inhibitions, in addition to the more standard positive activation.

Ambient Semantics is a wearable information system that offers wearers of its tag-reading wristband form-factor just-in-time feedback on books that the wearer picks up, and on people that the wearer meets. Each wearer is backed by their social network profile and a model of their attitudes generated from their corpus of everyday texts. Figure 1-6e is a screenshot of the system's social introduction faculty. When two strangers meet, the system tells them what they have in common. More than intersecting their keywords, the system explains shared identities, and cultural interests that are in the shared context of their intersecting taste ethoi. By illuminating shared context rather than explicit keywords, there are more opportunities to seed ice-breaking conversations and make social bridges.

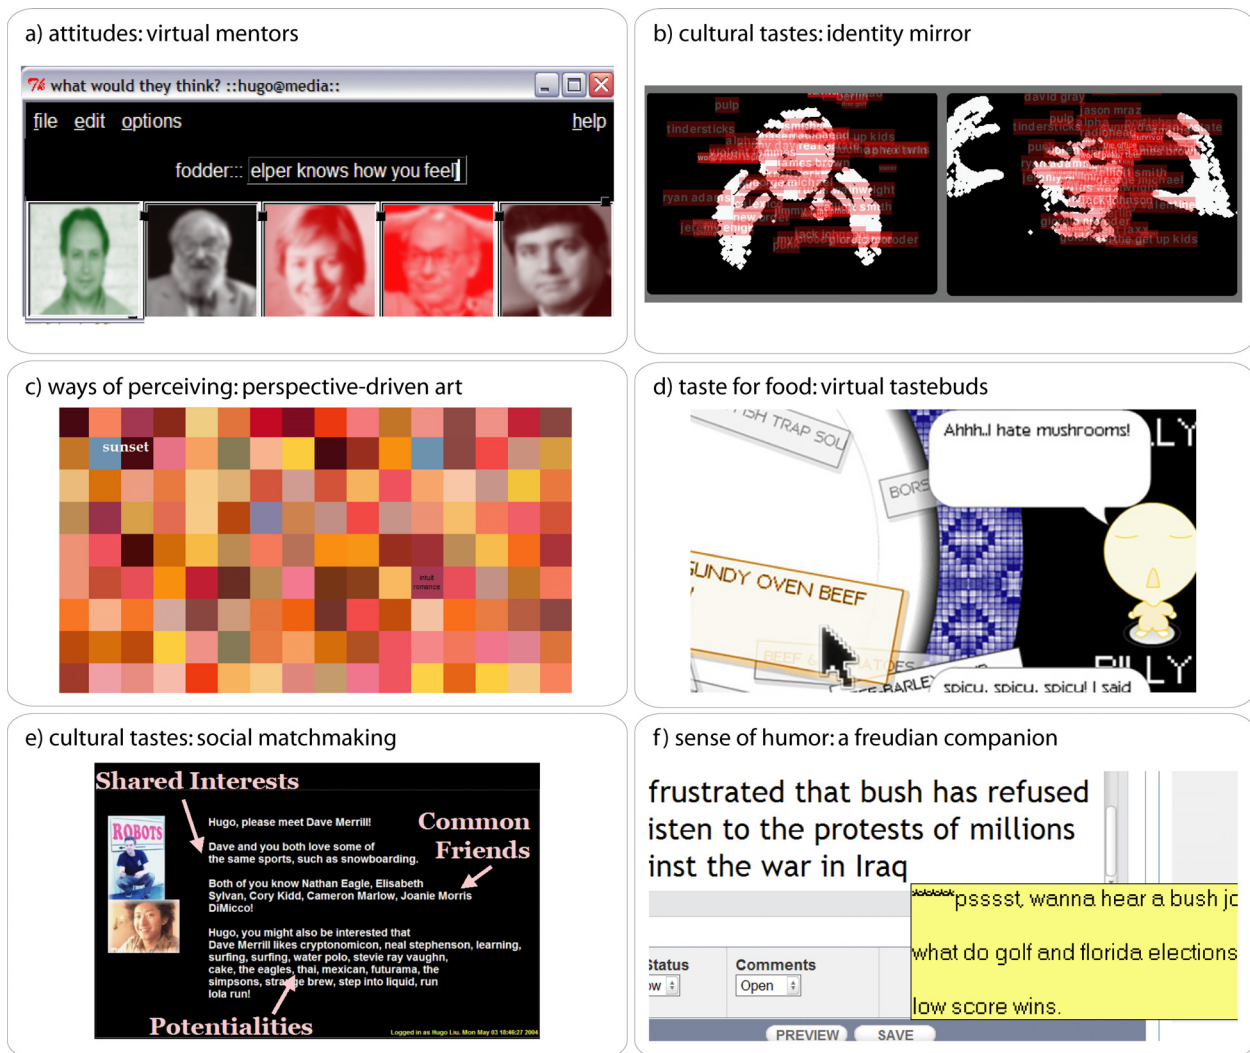


Figure 1-6. Summary of built applications

Cartharses (Figure 1-6f) is a Freudian joke-teller which observes a user's read-write activity, monitors his/her psychic tension levels about various topics, and delivers jokes by mapping the user's pattern of tensions into archetypal patterns of tension associated with each niche family of jokes. Rather than inventing jokes, *Cartharses* selects jokes to tell from a repository of 10,000 jokes. Freud's (1905) theory of tendentious jokes stated that the function of these jokes was to give catharses to bottled up psychic tensions. The reason why jokes are often grouped into ethnic families is because an *ethnos* has shared upbringing, and thus, shared patterns of psychic tension. A person's tension model is represented using semantic sheets similar to an attitude model, except that a measure of psychic tension supplants the PAD measure of attitude. In *Cartharses*' backend (pun intend), a corpus of 10,000 jokes was categorized into humor niches—such as blonde jokes, foreigner jokes, sexual jokes, Bush jokes, Clinton jokes—and archetypal tension patterns were

learned for each family of jokes, from a corpus of bloggers who appreciated each type of joke.

The process of conceiving, prototyping, implementing, and refining these artifacts was certainly any evolutionary one. These artifacts make use of rich generalized models of a person's tastes to simulate not just bits of a person's preferences, but a systematized and comprehensive account of their preferences—their *points-of-view*, if you will. Because the artifacts have novel capability, they necessitate novel interaction design. Certain ideas worked better than others, and eventually themes emerged. In the spirit of reflexive practice, three design lessons for this class of *perspective-based applications* were distilled from the application development process.

- ❖ #1—continuous observation of a user's textual activities and immediate feedback ensures an interesting 'walk' through a perspective, which only becomes apparent through animation.
- ❖ #2—feedback given *just-in-time* assures that a user benefits from the perspective when interest is piqued, and feedback given *just-in-context* provokes the user's imagination and critical abilities, since a reaction is synthesized for a new textual situation whose attitudes may never have been stated explicitly in the person's everyday texts.
- ❖ #3—the perspective should be tinkerable so that users can better grasp the capabilities and limitations of the artifact—e.g. tastebuds can be reprogrammed, perspective in the art bot can be shifted via sliders, and virtual mentors can be prompted to explain their reactions.

In this section, I summarized the thesis' experiments in person modeling and their results. The next section further distills these thesis results into contributions and a roadmap to the subsequent thesis chapters.

1.2 Roadmap and contributions

This thesis is fueled by a series of experimental systems that were built to model persons from their everyday texts, within five aesthetical realms—their cultural tastes, their attitudes, their ways of perceiving, their taste for food, and their sense of humor. Supported by evaluations and implemented applications, three broadly stated contributions of the thesis are:

- ❖ Through several parallel lines of investigation (via five aesthetic realms), it was demonstrated that rich 'person models' can be built by reading a person's everyday texts—her so-called 'textual traces' (cf. 'behavioral traces')—and can be generalized according to the underlying topology of the cultural space via spreading activation, analogy, and imprinter supplementation.

- ❖ Textual affect analysis was developed to implement an associative reading for affective themes. This technique complements previously used ‘topic spotting’ techniques for attitude mining. By focusing on first-person everyday texts like weblog diaries and social network profiles, the difficult task of viewpoint attribution was obviated, and it could be assumed that the text’s emergent affective themes were the writer’s own attitudes, tastes, etc. The strengths and limitations of this reading approach were made evident in three model evaluations.
- ❖ Six perspective-based applications were implemented, demonstrating a range of applications entailed by the thesis’s approach to person modeling—they constitute tools for self-reflection, person learning, and deep customization. The utility of at least one of these applications in a person-learning task was strongly supported by user study.

The rest of the thesis is structured as follows.

Chapter 2 presents the thesis’s approach to modeling a person’s tastes in the aesthetic realms, structuring discussion into three prongs of ‘acquire’, ‘generalize’, and ‘apply’. Related work in User Modeling prefaces introduction of the thesis’s approach. Related work in Intelligent User Interfaces accompanies discussion for ‘apply’.

Chapter 3 presents techniques and technologies for person modeling from everyday texts. Two techniques—‘reading for affective themes’, and ‘culture mining’ are discussed. ‘Reading for affective themes’ is situated in related work on computational reading. ‘Culture mining’ is situated in related work in musical similarity, social network analysis, and text mining. Two key technologies—‘commonsense reasoning’ and ‘textual affect analysis’—are presented and situated in their respective related works.

Chapter 4 narrates implementation-level details for each of the five built acquisition systems—cultural tastes, attitudes, ways of perceiving, taste for food, and sense of humor. Evaluations for three of these systems—cultural tastes, attitudes, and ways of perceiving—are presented in full.

Chapter 5 presents six perspective-based applications, which are enabled by person modeling. 1) an art bot which creates art suited to a person’s ways of perceiving the world; 2) a kiosk that facilitates serendipitous social introductions based on shared tastes; 3) a textual mirror to support self-reflection; 4) virtual mentors and pundits which give just-in-time feedback; 5) a synesthetic cookbook which simulates the tastebuds of family mentors; and 6) a jocular companion to parlay everyday woes into opportunities for humor. Evaluations for two of these applications are reported.

Chapter 6 reflects upon the philosophical underpinnings of the computational theory and methodology presented in this thesis, offering a theory-inclined digestion of the presented material, and betraying a direction for future work.

The thesis concludes in **Chapter 7**.

2 Approach

This chapter details an approach to modeling persons as they are in the social everyday, and parlaying these models into a range of perspective-based applications. Examining evidence of personal self-expression in everyday texts enables modeling in this most general of domains, and distinguishes the character of the present investigation from related works in the user modeling literature.

Section 2.1 discusses seminal and related works in the user modeling literature. Section 2.2 states this thesis's approach, finding situation for the present work in the user modeling literature. Sections 2.3-2.5 entertains separate discussions and related works for the three prongs of the approach—acquire, generalize, apply.⁸

2.1 Related work in user modeling

The field of user modeling is mature and over 25 years old. Because user modeling research is usually motivated by improvement of application performance, some closely related literatures which have themselves become developed are recommender systems, user-adaptive systems, and adaptive hypermedia. All the while, the user modeling field has also maintained a methodological wing focusing on generic techniques, and 'user modeling shells'—backbone systems which can be populated with domain-specific data and

⁸ It should be noted that discussions of particularly substantial techniques and technologies have been lifted out of this chapter and moved to Chapter 3 in order to improve the flow of discussion.

rules. Numerous proper surveys of user modeling can be found elsewhere—Kobsa (2001) reviews user modeling shell systems through 2000; Brusilovsky (2001) betrays the related field of adaptive hypermedia; a Communications of the ACM special issue (Resnick & Varian 1997) surveyed the related field of recommender systems. What follows is certainly an improper and incomplete survey—particular systems and themes are focused upon in order to invite comparison to specific themes in the literature.

From the viewpoint of the present work, user modeling can be divided into two approaches—category-based modeling, and behavior-based modeling. A third approach—discourse-based modeling—is also touched upon, though the area is still a relatively recent development resulting from cross-over research in corpus-based computational linguistics.

Category-based modeling is of interest because knowing the category of someone affords generalization of a user's model based on properties of the category. Though it was one of the first works in user modeling, Rich's (1979) book recommender system, GRUNDY, is still a *par excellence* example of category-based modeling. GRUNDY consisted of a user profile acquisition component, and a generalization and recommendation component. A profile of a user's demographic characteristics was acquired by a questionnaire-style interview. The generalization component consisted of *a priori* stereotype rules, mapping demographic categories into book preferences associated with those categories. The generalization algorithm consists of firing rules activated by the user's profile in order to produce a generalized user model. For example, suppose that the category of women 26-35 years old stereotypically prefers "romance novels" with strength S . If a user's profile implicates her into that category with uncertainty U , then a preference for "romance novels" is added to the user's generalized model, with strength S and uncertainty U . By iterating through other attributes in the user's profile, these preferences began to add up and overlap, eventually, the generalized user model should triangulate or converge onto several key recommendations.

GRUNDY's approach resembled and paralleled developments in AI expert systems, as GRUNDY was essentially a system of hand-crafted rules backed by an uncertainty model, and was purposed for recommendation. Some subsequent user modeling shell systems formalized the stereotype and rule idea to allow user models to be assembled for any domain, by populating a rule base. GUMS (Finin & Drager 1986) and UMT (Brajnik & Tasso 1994) were two shells that allowed for the definition of hierarchical stereotypes and mapping rules, thus circumscribing the scope of GRUNDY. As both took a logic-based approach, they incorporated truth maintenance capabilities. The category-based approach of GRUNDY and later related modeling shells is a sound one, but the quality of models built from *a priori* stereotypes depend entirely on the insightfulness and profundity of those hand-crafted stereotypes.

Behavior-based modeling is interesting as an approach because it states that user models should be driven by actual 'behavioral traces' produced by the user, and not the result of *a priori* assumptions. Hence, user preferences are not deduced from stereotypes but are rather induced from a *history of user actions*. Social information filtering (Shardanand & Maes 1995), or collaborative filtering, as is the most preferred term today, is a popular algorithm for predicting user preferences. In that algorithm, a user's history of actions (e.g. purchase history, browsing history) is represented as a high-dimensional vector of ratings over a field of items (e.g. products, webpages). By comparing one user's vector with the vectors of all other users, mathematical metrics such as *cosine similarity* are used to find similar users and potentially desirable yet undiscovered items. User-user collaborative filtering has also been varied as item-item collaborative filtering (Sarwar *et al.* 2001), which is more commonly used in e-commerce applications such as Amazon's⁹ product recommendation mechanism. Today, there are several recommender toolkits built on variations of collaborative filtering, such as GroupLens (Breese *et al.* 1998; Herlocker *et al.* 1999), which supports predictive modeling from various behavioral traces such as the user's profile elicited via online forms, navigation history, and transaction history.

Collaborative filtering is one popular technique for behavior-based modeling. Another technique is Bayesian modeling. One classic example, the Lumiere project (Horvitz *et al.* 1998) used Bayesian modeling over histories of users' actions (and contexts) in the Microsoft Excel spreadsheet application to infer likely user goals. A key difference between collaborative filtering and Bayesian learning is that the Bayesian approach excels in complex decision spaces where the presence of contextual features changes drastically the interpretation of user actions. Collaborative filtering presumes that features are more or less of a homogenous class, and are not conditioned on one another. Other than Bayesian learning, a variety of other machine learning and statistical techniques may also be applicable—many are surveyed in (Zukerman & Albrecht 2001). For instance, the DOPPELGANGER user modeling shell system (Orwant 1995) allowed for interchangeable machine learning algorithms such as Markov models, clustering, and linear prediction, to predict current user state (e.g. 'hacking', 'writing', 'idle') from hardware and software sensors. However, DOPPELGANGER, along with knowledge-based intelligent tutoring systems such as COACH (Selker 1994) also exhibit qualities of the earlier described approach of category-based generalization. Each user's profile of interests was generalized by the models of user communities that a user belonged to—such models were community-maintained. Significantly, DOPPELGANGER'S category-based models were not *a priori*, but were rather driven by community data.

⁹ <http://amazon.com>

A third approach for user modeling is assigned a working term here—*discourse-based modeling*. Discourse-based user modeling is likely the result of cross-over research from corpus-based computational linguistics and artificial intelligence, especially dialogue systems. One survey of research that lies between user modeling and natural language processing systems is presented in (Zukerman & Litman 2001). A *discourse model* is a model of the field of topics that are either covered by a conversation, or is a semantic field defining the vernacular of a domain. The former sense of the term is relevant to natural language dialog systems, in which users are modeled as their accrued topics and goal states from conversation with the system. The latter sense of the term has been used *system-centrally* in AI and HCI systems to describe a system's field of allowable syntax and semantics, e.g. (Wahlster 1991); in other words, the field of possibility. However, discourse models can also exist for entities larger than a user's conversation, and smaller than the total space of possibility; for example, communities and cultures have preferred discourses. Knowing the discourse models of several communities, by comparing the user's discourse against those community discourse models, it is possible to characterize the user as inheriting from certain communities. This then becomes a basis for generalization. Cassell & Bickmore (2003) described embodied conversational agents that modeled a user's discourse as breadth and depth of familiarity over a range of traversed topics (e.g. weather, baseball). This user discourse model then activated nodes in a spreading activation discourse planner, and a number of next conversation steps were calculated. One way to think about the spreading activation planner is that it has captured the common sense discourse of how conversation topics are related to one another. By plotting a user's discourse onto the planner's topology and spreading activation, a generalization is made.

2.2 Person modeling from everyday texts

The thesis approach to 'person modeling' is now described and situated in the user modeling literature. Collaborative filtering and other behavior-based approaches have the virtue of being driven by actual data, but behavioral traces taken to represent a user (e.g. transaction history, browsing history) are typically application-specific—thus these traces constitute context-dependent data portraits of a person. In fact, the name 'user modeling' itself reflects a common motivation for such systems—tailoring and personalizing application behavior around its users. GRUNDY (Rich 1979) actually considered persons more generally, since what demographic stereotypes capture are intuitions about cultural patterns of preference. However, such patterns were entered in *a priori* and not learned from actual cultural data.

This thesis extends the data-driven approach of user modeling to what I term *person modeling*, by considering evidence of persons in the social everyday rather than considering histories of user actions, which are inevitably specific to their application's domain. To

facilitate computation, this domain of the social everyday is further subdivided into realms of concern. The experimental systems described by this thesis explore models of persons within five aesthetical realms—cultural tastes, attitudes, ways of perceiving, taste for food, and sense of humor. As much as possible, the representations of persons in each of these realms is inspired by and consistent with theories of people and cultures known in the humanistic literatures such as psychology, cultural theory, and semiotics.

The approach to person modeling espoused here is data-driven, but whence evidence of persons in the social everyday? One plentiful source is the corpus of his/her everyday texts—weblog diaries, social network profiles, homepages, instant messenger conversations, etc. By applying natural language processing techniques to read this text, *textual traces* are extracted. The parsed and normalized traces are then amenable to machine-learning generalization just like behavioral traces.

Reading a person's texts yields explicit textual traces. But in order to create a predictive person model, these textual traces must be generalized. GRUNDY performed generalization via a priori stereotypes, while DOPPELGANGER generalized users' interest on the basis of community-maintained models, which underlied each user. Our approach to generalization is to locate a person's textual traces onto the nodes of a cultural topology that prescribes affinities and relationships between nodes. This cultural topology is not *a priori* like GRUNDY, but is data-driven like DOPPELGANGER. However, whereas DOPPELGANGER'S community models were neatly represented and directly computable, cultural topology sometimes needs to be mined from natural language text; this represents an additional challenge. The technique of *culture mining* is described in Chapter 3 as a way to extract cultural topology from cultural corpora, such as 100,000 social network profiles, and political texts of Democrats and Republicans. Once a cultural topology is acquired as a graph structure, generalization proceeds via *spreading activation* (Collins & Loftus 1975). In the realm of attitudes, *analogy* (Gentner 1983) and *imprimers* (Minsky forthcoming) are additional techniques for generalization. Spreading activation, analogy, and imprimers are admittedly heuristic techniques, requiring the manual setting of various parameters like discount for spreading, discount for fan-out, etc. Although spreading activation along cultural topology seems very different from collaborative filtering's vector-similarity method for generalization, they are not so dissimilar. Mining cultural topology is an information-theoretic process of calculating affinities between cultural items that are true across a population—this procedure can be compared with the clustering effect achieved with item-item collaborative filtering.

A final resonance in the user modeling literature is with discourse-based modeling. First, extracting textual traces from a person's everyday texts can be compared with modeling a user's discourse

history through a conversation with a natural language dialog system, as both tasks concern gisting evidence from text. Second, a person's textual trace resembles her personal discourse; a culture's topology resembles wirings in the cultural discourse. Cassell & Bickmore's (2003) embodied conversational agents located a user's discourse model onto a spreading activation discourse planner. While the motive of spreading activation there was to generate temporal moves, more generally the act of plotting a user's topics onto a network of pre-wired topics and spreading activation outward is an act of topological generalization. This same technique is used in the present approach to expand a person's discourse along the connections of the learned cultural topology.

Once a generalized model of a person's tastes, attitudes, ways of perceiving, etc. is produced, the model is used in applications to simulate a person's reactions to arbitrary textual input. This 'reactive' posture on simulation deviates from the more standard 'prescriptive' posture in user modeling, which is that a user model predicts preferences, and these predictions are simply offered up as recommendations. In some sense, asking a recommender system to react to some thing of your choosing could be as strange as bringing some medication up to a pharmacist, and asking her if she thinks you should take that medication. The motivation for using a general model of a person to simulate reactions is to enable a class of what I call *perspective-based* applications—such as virtual mentors and pundits that constantly react to a user's context. Strategies for generating reactions include spreading activation and memory-based reasoning.

Having situated the thesis's approach to person modeling in the user modeling literature, the next three sections will detail the three phases of the described approach—acquire, generalize, and apply—focusing on how phases are operationalized for each of the five aesthetic realms that were modeled.

2.3 Model acquisition

Three key tasks are identified in the acquisition phase of person modeling.¹⁰ First, knowledge representations are devised to best describe each of five aesthetic realms being addressed—cultural tastes, attitudes, ways of perceiving, taste for food, and sense of humor. Second, the strategy of reading persons' everyday texts for affective themes is concretized for each realm as a *reading schema*—a process that yields explicit textual traces. Third, the cultural topology is acquired through culture mining. Representational choice is the primary decision, while reading strategy and culture mining are entailments of that choice. Thus, the second and third tasks will be developed as techniques in Chapter 3. This section, then, focuses, on

¹⁰ A fourth, "zeroth" step is deciding on sources of personal texts that are a good source of the information for person models in each realm.

the first task—introducing the rationale for, and features of each realm’s chosen representation.

Devising knowledge representation for the aesthetic realms was treated as a series of design problems. Three considerations were kept in mind through the process—1) what information could be gotten out of source texts, 2) how representational features might be consistent with humanistic scholars’ assessments of the aesthetic realms, and 3) how the representation could facilitate later tasks of generalization and simulating reactions. Below is a discussion of knowledge representation issues for the five realms; to be followed by a comparative discussion analyzing the commonalities exhibited across realms.

§

Realm of cultural taste. For the realm of cultural taste, social network profiles were chosen because several cultural theorists had already framed the issue of cultural taste in terms of patterns of consumer choices (Bourdieu 1984; Haug 1986), which is precisely what is captured in each social network profile. To capture a person’s particular interests, a representation for the realm needs to enumerate all possible interest and identity descriptors, so a hierarchy of 21,000 interest descriptors and 1,000 identity descriptors was assembled from various folksonomies. These descriptors are interconnected by a web of metadata relationships (e.g. parentOf(book author, book title)). But the folksonomic tree fails as a knowledge representation, because many serendipitous taste affinities between descriptors are not captured by metadata relationships in tree. Thus, inspired by recent consumerist theories that view a person’s consumer choices as *gestalts* (McCracken 1988; Solomon & Assael 1987), a method was derived to learn the mutual information between each pair of descriptors. After pruning, this resulted in 12,000² pairwise numerical affinities. Overlaying these affinities with metadata relationships, what results is an almost fully connected graph.

I term this a *semantic fabric* representation to reintroduce the spatial metaphor. We should think of a person’s explicitly stated interests as a pattern of activation in this fabric. Furthermore, the fabric terminology is consistent with terms in cultural theory—for example, Geertz described culture as webs—“man is an animal suspended in webs of significance he himself has spun, I take culture to be those webs” (Geertz, 1973: 4-5).

The topology of the semantic fabric is lumpy. In particular identity nodes and taste cliques can be identified as *semantic mediators* in the fabric, acting as connector hubs and thus exerting disproportionate influence. When a person’s pattern of interests activates the fabric, activation tends to flow into and out of these attractors; thus these

attractors arguably introduce an aspect of prototype-based inference into the person model generalization process. The existence of taste cliques and identity hubs are consistent with Solomon & Assael's (1987) prediction of 'consumption constellations'.

Realm of food. In this realm, a person's generalized model might well correspond to the common sense of the word, *tastebuds*. The field of foodstuffs is as broad, and as richly connected as the field of cultural interests, so it was decided that the same representation would be adopted. The realm of taste for food is represented as a semantic fabric of interweaving recipes, ingredients, cooking procedures, sensorial keywords, and so on. Again hierarchies of food metadata were assembled from a variety of web sources, and from the Thought for Food corpus—thousands of sentences embodying cooking common sense. Also, from a corpus of 60,000 recipes—the taste coherency assumption was again made, in order to learn the implied affinities between foodstuffs. Combining metadata with mined affinities, what resulted was a richly connected semantic fabric of the foodstuffs, whose nodes include—60,000 recipes, 5000 ingredient keywords, 1000 sensorial keywords (e.g. 'spicy', 'chewy', 'silky', 'colorful'), 400 cooking procedures, and 400 nutritional keywords.

Just as identity hubs and taste cliques acted as the semantic mediators for cultural taste, in the realm of food, cuisine nodes (e.g. 'Chinese', 'dessert') and basic flavor nodes (e.g. 'sour', 'spicy') act as semantic mediators for tastebuds. Tastebuds, like cultural tastes, are represented as activations of the food fabric. In the case of tastebuds, both positive (food likes) and negative (food dislikes, food allergies) activations are allowed.

Realm of attitudes. There are several possible interpretations of attitude. Jung, for example, introduced 'extraversion' and 'introversion' as two basic psychological attitudes. However, we define attitude in a simplistic and common sense way—an attitude is your feeling toward some topic, which may be a thing, or person, or event. Regarding attitude as topic+affect is consistent with the cognitive appraisal theory of emotions (Ortony, Clore & Collins 1988)—which states that emotions result from cognitive appraisal of some thing, person, or event, and is affect that is directed toward that cognitive target. Weblog diaries was chosen as a promising candidate for acquisition of personal attitudes. Just as cultural interests were backed by metadata hierarchies, topic hierarchies were mined from DMOZ (e.g. *subtopicOf(feminism, philosophy)*) and from ConceptNet (e.g. *isA(burger, food)*).

However, the application of 'culture mining' to acquire a semantic fabric was not effective over the corpus of weblog diaries. First, whereas 20-50 descriptors were available with high confidence in each of the 100,000 social network profiles, attitudes gotten from reading were fewer and more tenuous—due to the relative sparseness of cultural interests being mentioned in weblog diaries,

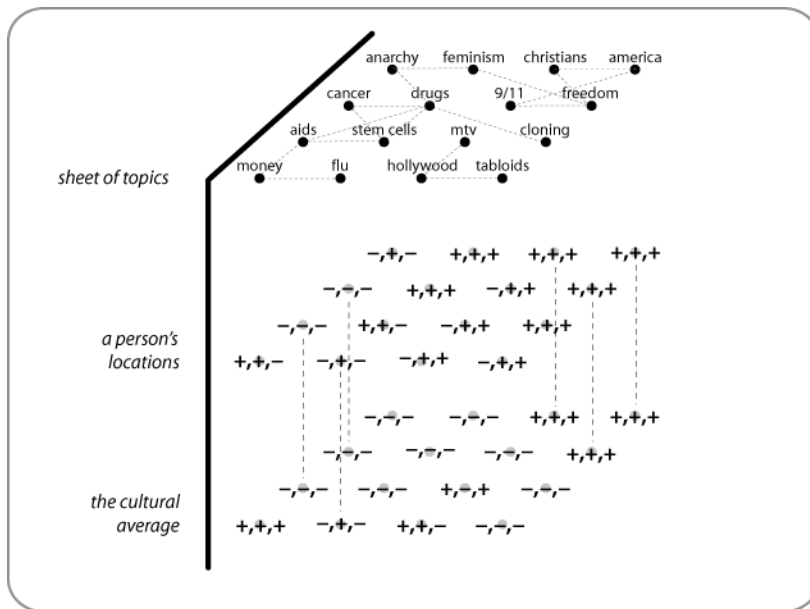


Figure 2-1. A semantic sheet representation for the realm of attitudes

and due to the difficulty in assessing with certitude that cultural descriptors mentioned in the weblog diary are intended as expressions of personal interest. Second, we hypothesize that the self-reflective context of a social network profile, which leads to a desired coherent enumeration of interests, is not as pronounced in weblog diaries. Third, whereas cultural interests and tastes for food may be described in everyday texts as clearly a preference or dispreference, an attitude's affect is not given as binary choices in weblog diaries.

As a result, a *semantic sheet* representation (Figure 2-1) was designed to facilitate visual understanding of the realm¹¹. In it, a sheet of topics supposes an enumeration of all possible topics, linked by their metadata gotten from topic hierarchies. Then, each person's attitudes are represented as a sheet of affects which corresponds to the sheet of topics. Affect is represented with Mehrabian's (1995b) three-dimensional PAD model. In this representation, both persons and cultures can be represented as sheets. As will be demonstrated in a political culture scenario in Chapter 4, a person's affinity toward a culture can be measured (and visualized) as the degree of alignment (shown as vertical dashed lines) between their two sheets of attitudes.

The semantic mediators conceived for this realm are Minskian *imprimers* – mentors, parents, or cultures whose attitudes prime ours.

¹¹ though admittedly, it is in essence a memory-based representation. Organizing memories into sheets here visually affords metaphors like 'attitude alignment'

As was shown in Figure 1-4, imprinter models are sheets, and in generalization, these sheets supplement a person's missing attitudes. For example, the model of a person who is imprinted by Warren Buffet in topics relating to business can be supplemented by attitudes on those topics from Warren Buffet's person model.

Realm of humor. A Freudian approach was adopted to representing sense of humor. To model a person's sense of humor, again, weblog diaries are a preferred source. The space of humor is likewise modeled using semantic sheets, except that instead of PAD values, there is a unary value standing for tension level. A person's sheet is her own pattern of psychic tensions about various topics, which are presumed to be dually rooted in conditions of her upbringing, as well as in her more recent life frustrations.

There are many accounts of humor in literatures computational and humanistic. Some postures are that jokes are 1) motivated by feelings of superiority (traced back to Burke) and thus, derisive; 2) motivated by need to relieve frustration (cf. Freud); 3) motivated by expectation violation (cf. Kant); 4) delightful and informative as a mechanistic caricature (cf. Bergson) and thus examples of how not to be (cf. Minsky); and 5) a defense mechanism (cf. Piddington). The experiment developed here proceeds from the relief premise. Freud's (1905) account of *tendentious* jokes reflects early psychoanalysis's conceptualization of the affective unconscious as an expanding and contracting hydraulic bag, which alternately stores and gives catharsis to psychic energies and tensions. Freud distinguished between innocent and tendentious jokes. Innocent jokes elicit just a smile or chuckle and are not emotionally heated, while tendentious jokes have a sexual or aggressive character, and can draw out aggressed, howling laughter. Tendentious jokes are a boon to the psychic economy of the listener, says Freud. They relieve psychic tension that has pent up around particular topics.

Freud theorized that people of similar cultural background, such as members of a lifestyle or ethnic group—due to shared upbringing and experiences with family, gender, sex, and money—also tend to share a pattern of social inhibition and psychic tension. Cultural humor, then, effects aggressive laughter as a way to give catharsis to tensions created by social inhibition. Sexual jokes are prevalent, in part because almost all societies inhibit sexuality to some degree. Inspired by Freud's account, we represent niche humors (e.g. 'bush jokes', 'jewish jokes', etc.) as pre-formed sheets of tension—in fact they are *archetypal tensions* because they are shared by a cultural grouping of persons. We may also think of the space of tension/topic humor as being semantically mediated by niche humors. Finally, generalization can be performed by calculating the degree of alignment between a person's sheet and sheets of the various niche humors. For example, a person with a sheet of tension on topics such as 'Iraq', 'freedom', 'fascism' would best align with the archetypal sheets for 'Bush jokes' and 'political jokes', while a person tense about topics such as 'virtue', 'scandal', and

'government' would best align with 'Clinton jokes' and 'political jokes'.

Realm of perception. We were inspired by the following premise—why do realists and romantics perceive the world so differently? Looking to existing theories, Jung's (1921) theory of psychological type directly addressed this question. In his theory of type, Jung proposed four fundamental psychological functions—sense, intuit, think, feel—to account for all the differences in human perception. In this vocabulary, the difference between realists and romantics can be re-viewed as a difference in disposition—when asked to interpret the word sunset, a romantic leans on feelings (e.g. 'romance', 'warmth') and intuitions (e.g. 'embrace'), while a realist leans on ratiocinations (e.g. 'off work', 'dinner') and sensations¹² (e.g. 'dark').

More formally, we decided to represent the space of possible perceptions as a dimensional space whose axes are Jung's four fundamental psychological functions—sense, intuit, think, feel. A person's particular disposition is formalized as a coordinate location in the space. For example, a romantic is located in high feeling, high intuition, low sensing, low thinking.

A first-pass experiment to model person's rough location in this space considers person's patterns of affective communication evident in their weblog diaries—patterns such as PAD affect associated with ego (e.g. 'I', 'me'), PAD affect associated with alters, and PAD that is pushed from ego to alters, and from alters to ego. The success of this approach was mixed. Though not yet implemented, the notion of archetypes as semantic mediators of perception can also be considered, as identifying mediators has worked well in the other realms. Archetypes such as 'realist', 'romantic', could be identified as locations in the space, each associated with some textual signature, and these could facilitate the location of persons into the four-dimensional space.

§

A comparative discussion now reviews some shared themes exposed in devising knowledge representation for the above realms.

One theme that crosses knowledge representations and realm is the notion of *semantic mediators*. In the cultural taste realm and food realm, supernodes like 'identities' and 'cuisines' helped to organize the cultural topology—more than other nodes, these tended to be vehicles of consistency and coherency in spaces that were otherwise non-hierarchical. In the realm of attitudes and the realm of humor, again 'imprimers' and 'niche humor' served as semantic mediators, helping to organize and prime the attitudes and tensions of

¹² Jung means his key words 'think', 'feel', 'intuit', and 'sensing' differently from their common sense interpretations.

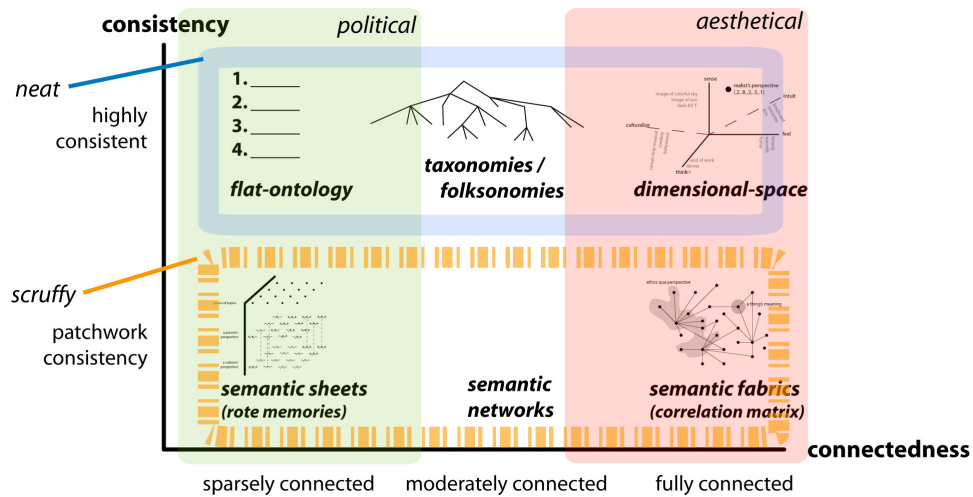


Figure 2-2. A semantic diversity matrix

individuals. Finally, in the realm of perception, we speculated that archetypes such as ‘realist’ and ‘romantic’, though unimplemented, could be a useful semantic mediator in that realm.

Another theme is a preference for the maximum amount of connectedness possible, as topological connectedness bolsters the power of model generalization, and affords perspective-based artifacts the comprehensiveness necessary to react to arbitrary inputs. One reason why rule-based user modelers tended to be prescriptive rather than reactive is that prescription is easier than reaction. To react to arbitrary input implies that the semantic distance between the user’s model and the input is known. A space wired with a few connections, such as stereotype rules, does not guarantee a sound distance calculation because its connections make too many assumptions. This is related to the *semantic distance fallacy* that is often talked about in regards to the poor decision of relying on semantic network link-hops as a measure of semantic distance (Collins & Quillian 1969). By contrast, a semantic fabric is a connectionist representation, and affords a much better distance calculation from a person’s model to any input. Hence, high degree of connectedness enables the reactive mode for perspective-based artifacts. Of course, the semantic sheet knowledge representation is far less connected and ideal than semantic fabrics, but we make do. Generalization over semantic sheet is opportunistic, and connections can be made via spreading activation along topic hierarchies, via analogy, and via imprimers.

To make sense of these two themes—that semantic mediators lend *consistency* and that there is a preference for *connectedness*, we attempt to plot some knowledge representations along these two dimensions, to see what insight might be developed. Inspired by Marvin Minsky’s (1992) “causal diversity matrix” in which Minsky plots the space of intelligence problems along the axes of number of

causes and number of effects, Figure 2-2 presents a “semantic diversity matrix,” which summarizes the representational tradeoffs.¹³ I ask you the reader to indulge me for a moment in the following speculation. Representations of the top row (i.e. ontology, folksonomy, dimensional-space) are symbolic and neat, while the bottom row of representations (i.e. sheets, networks, fabrics) are phenomenal and scruffy. Neat representations seem to be centralized and prescriptive, and are thus prevalent in symbolic AI; scruffy representations seem to be decentralized and descriptive, and are thus prevalent in corpus-based research. Now interpreting the columns—the leftmost column (i.e. ontology, semantic sheets) tends toward political by emphasizing disconnectedness, ideology, denotation, and control; whereas the rightmost column (i.e. dimensional spaces, semantic fabrics) tends toward the aesthetical, emphasizing both weak and strong connections, which afford connotation and evocation. If these trends are true trends, then perhaps the semantic diversity matrix can be of assistance to researchers in computational modeling, who would like to reflect upon the appropriate knowledge representation for their problem. Consider, for example, that Sack (1994) chose actor-role analysis for representing ideology. As actor-role bindings may be plotted somewhere between ontology and semantic sheets, we could, without knowing the problem being modeled, guess that the problem domain was political, and somewhat scruffy.

Having characterized the knowledge representation issues in model acquisition, the next section discusses the generalization phase of person modeling.

2.4 Model generalization

Once a person’s explicit textual traces are acquired from her everyday texts, those traces seed the inference of a generalized person model. The leap from textual traces to generalized model can be idealized as the leap from discrete data points to a comprehensive model—meaning that a reaction can be calculated for virtually any input conceivable in a realm. A goal for the generalized person model is to approach the sophistication and comprehensiveness of a person’s *perspective*. To achieve generalization, three heuristic approaches are tried—spreading activation (over cultural topology and metadata hierarchies), analogy, and imprinter model supplementation. This rest of this section 1) argues that the generalization process is convergent and exhibits truth maintenance properties; and presents scenarios for generalization via 2) spreading activation, 3) analogy, and 4) imprinting.

§

Convergence. Generalization from fragmentary textual traces should see convergence because a person's expression of those textual traces was a non-arbitrary emanation of that person's perspective. We define convergence as the observation that each new textual trace add to the agglomerating person model offers less and less information not already known to the model. Convergence would be important to observe when trying to create a comprehensive model of a person's taste because it signals that the model is nearing circumscription of the person's taste. We opine on some reasons for believing in the coherency and consistency of perspective. First, human memory and cognitive function are shaped by the principle of economy. Given that physical memory is resource constrained, mental structures optimize for maximum utility—e.g. dream-work can be regarded as a garbage collection process consolidating recent relevant experiences and rejecting extraneous aspects; creativity means re-appropriating knowledge tied to one context, into other unexpected contexts, thus the knowledge which tends to persist are those with overloaded utility, versatile enough to solve problems under a variety of contexts. Second, competency in the social world motivates a need to communicate oneself, and consistency of perspective facilitates one's social intelligibility. That people compose themselves around intelligible *personae* is consistent with Goffman's 1959 thesis that we interact with the world as a performer acts through various masks. This line of reasoning also supports a belief that person's coherencies should be captured in a mined cultural topology because *personae* are the product of social and cultural negotiation. Third, tastes are intrinsically organized to minimize dissonance. Individuals' consumptive choices tend to cohere around their design for a lifestyle—clusters of goods having common implications to lifestyle are known as 'Diderot Unities' (McCracken 1988) and 'consumption constellations' (Solomon & Assael 1987). Cognitive dissonance (Festinger 1957), the unpleasant experience of mental self-conflict, leads persons to re-organize their possessions and their identity to eliminate glaring dissonance.

Truth maintenance. Truth maintenance (Doyle 1980) is the detection and elimination of intra-system contradictions. We argue that in the context of model generalization, methods like spreading activation has a truth maintenance effect because nodes may be activated from multiple other nodes (evidence corroboration). Usually, spreading activation is thought of only as a mechanism for *semantic expansion*, but in a densely connected network, spreading activation also maintains truth—the strength of that property increasing proportionally with the density of connections in the activation network's topology. The following example compares spreading activation in a sparse network versus in a dense network. Suppose there is a semantic network whose connections are sparse, and suppose activation is spread outward from a seed set of nodes. In this case, it is unlikely that activations seeded from different nodes will cross paths one-hop away or two-hops away. This has the effect of semantic expansion, but does not have a truth

maintenance property. Suppose now that there is a semantic fabric with very dense connections. Suppose again that spreading activation proceeds from a seed set of nodes. In this case, what is different is that activations seeded from different nodes will cross paths almost immediately, providing checks and balances on the information embodied in each. Two-hops away, it is likely that all the seed nodes' activations will have crossed. This crossing action, we argue, has the effect of truth maintenance—if two crossing activations are both strong, their activations are additive, and the node at which the crossing takes place is dually corroborated; if in the case of the food fabric, a positive activation crosses a negative activation, their effect is to contradict each other, resulting in a neutral (non) activation. Finally, we note that semantic fabrics present much more ample opportunity for corroboration and contradiction than did the sparse logical truth maintenance and contradiction detection features in previous user modeling systems such as GUMS (Finin & Drager 1986) and UMT (Brajnik & Tasso 1994).

§

Spreading activation. Collins & Loftus (1975) originally posed spreading activation as a psychological theory of memory. The rationale between that associative reasoning spreads across nodes of semantic memories with some energy, which decays proportionally with the number of steps traversed. Salton & Buckley (1988) affirmed the usefulness of spreading activation as a heuristic technique for semantic expansion in information retrieval tasks. They also articulated further techniques such as a system for assigning discounts. Each time activation traversed a link, the activation a was discounted by some coefficient, resulting in e.g., $0.5a$ activation. Furthermore, nodes with a large number of outgoing links incurred a fan-out discount (sometimes called a branching factor discount). Our use of spreading activation over topic hierarchies and over cultural topology follows this standard discount scheme—applying discounts to spreading and to fan-out. Next, we walk through two scenario of use for spreading activation.

In the realms of attitudes and humor, activations seeded by textual traces are spread along the lines of implicit topic hierarchies. Consider that an individual's attitude toward the topic of 'water conservation' is pleasurable-arousing-dominant (+1.0,+1.0,+1.0), and consider that the same individual's attitudes toward the topics 'environmental protection', and 'recycling' are unknown. Generalizing from the known attitude for 'water conservation', activation is spread along the edge `supertopicOf('environmental protection', 'water conservation')`, which is gotten from folksonomies such as DMOZ. Because activation flows 'upstream' from topic to super-topic, an upstream discount of 0.75 is applied, resulting in 'environmental protection' (+0.75,+0.75,+0.75). For simplicity, separate tallies of PAD and uncertainty are not maintained and

uncertainty is folded into PAD. To make this simplifying assumption, PAD values must range (-1.0 to +1.0), and definitely not (0.0 to +1.0). Next, assuming that 'recycling' is the only sub-topic of 'water conversation', with a downstream discount of 0.5, the relaxation result would be 'recycling' (+0.5,+0.5,+0.5). Topics with multiple sub-topics impart their affection downstream with an additional 'fan-out' discount, which is heuristically set to be inversely proportional to the log of its number of sub-topics. The mechanism of spreading activation continues propagating activation from topic up to supertopic to super-supertopic, and from topic to subtopic to sub-subtopic, each time applying the appropriate discount, until the post-discount PAD value is negligibly small (neutral).

In the realm of cultural taste, a person's textual traces are relaxed over the cultural topology, resulting in a taste ethos. Textual traces are presumed to be a fragmentary expression of a more profound but ineffable *ethos* (character). Spreading activation over the semantic fabric of taste should converge upon the ethos. Suppose that A, B, and C are three nodes with strong mutual affinities, forming a clique. Supposing then that only A and B are known in the textual trace, the result of spreading activation is to implicate C also into the ethos of the individual, as both activations from A and B will corroborate C. Activation is also mediated by identity hubs, and taste cliques. If item A has a strong affinity to identity X, and X has strong affinity to items A, B, C, and D, then when activation spreads into X, all strong members of X are together pulled into the gathering ethos. An individual's ethos is more likely to stumble into a few focused identity hubs (thus, not threatened by fan-out) and pick up all their member spokes than to be constituted by nodes that are not interconnected themselves by semantic mediators. This observation implies that spreading activation in the presence of semantic mediators acquires a hierarchical quality.

§

Analogy. In attitude model generalization, analogies are automatically made, mapping topics in the model to new topics outside the model, and propagating PAD values to those new topics. While spreading activation along topic hierarchies is often a *context-preserving* act (e.g. 'environmental protection' and 'water conservation') analogy is more often a *cross-context* act (e.g. 'war' and 'pollution'). As such, analogy is definitely a heuristic technique for generalization. Analogy for the purposes of this thesis is defined as a relationship between concepts with similar attributes, though not usually sharing a direct taxonomic parent. Structure mapping (Gentner 1983; Fauconnier & Turner 2002) is a way to find analogy, and this method is operationalized in the ConceptNet commonsense reasoning system. For example, suppose an environmentalist's textual traces are asked to react to the concept of 'war'. Feeding 'war' into ConceptNet's `get_analogous_concepts()` function

results in a rank-ordered list of analogous concepts, among which are some topics known in the environmentalist's textual traces. Actual output from ConceptNet is shown below, edited for legibility.

```
[war is like storm] the concepts share:
  ==PropertyOf==> bad
  ==PropertyOf==> violent
  ==PropertyOf==> dangerous
  ==CapableOf==> destroy property

[war is like pollution] the concepts share:
  ==PropertyOf==> evil
  ==CapableOf==> kill
  ==CapableOfReceivingAction==> cause
  ==CapableOfReceivingAction==> stop
```

Based on these shared attributes, the topic 'war' inherits the environmentalist's attitudes about 'storm' and 'pollution', but each time with a 0.5 uncertainty discount, just as in spreading activation. ConceptNet is capable of resolving not only keyword topics, but also second-order topics such as 'eat burger', 'open door', etc. Its capabilities are surveyed in Chapter 3.

§

Impriming. Minsky (forthcoming) introduced the notion of 'imprimer' as a mentor or confidant that one forms an attachment to, and engages in mimesis of. Minsky's notion is supported by similar notions in the psychology literature, such as Freud's (1915) theory of 'introjection'—children unconsciously emulating their parents' values—and Ogden's (1979) theory of 'projective identification'. According to Minsky, parents imprime their children, an advisor imprimes a student, and even fictional characters and cults can imprime. Minsky's litmus test for an imprimer is one who can stir self-conscious emotions such as pride and embarrassment within oneself. In attitude modeling, imprimer relations are thus detected by looking for co-occurrences of self-conscious affect with potential imprinting entities such as persons and cultures. Once identified though, corpora must be assembled for imprimers in order to model their attitudes. This unfortunately is still a supervised process. In generalization, imprimer's models are taken to supplement the person's own model in parts of the topic space where the individual's model offers no insight. It cannot be emphasized enough that generalization via imprimers is experimental, heuristic, and is motivated by a desire to ultimately create perspective-based applications that can react to any input. Such applications are ultimately fail-soft because they are vehicles for provocation and have other means of verifiability. When heuristics like imprimers produce wrong inferences, which they will, the results are not catastrophic. Interestingly, the effect of imprimers on attitude prediction was measured in an evaluation of attitude modeling, and will be presented in Chapter 4.

This section characterized the generalization process as convergent and truth-maintaining, and provided details on three generalization mechanisms of spreading activation, analogy, and imprinting. Next, the final section focuses on issues in model application.

2.5 Model application

Having acquired and generalized a person model, two remaining issues are now foregrounded—1) how a generalized person model can be applied to simulate reactions to arbitrary input; and 2) how can these models be embedded in *perspective-based applications*, and what interactions should such applications afford? The rest of this section 1) presents a simplistic model of simulation, also considering its limitations; 2) builds a context of related work for perspective-based applications; and 3) reflects upon design lessons learned from the experience of building perspective-based applications.

§

A simplistic model of simulation. Since the generalized model has propagated the implications of a person’s textual traces as far as possible, simulation is reduced to a matter of mapping arbitrary input into the generalized model, and reading the reaction off the model. Using the computational reading approach that reduces a person’s everyday texts into a set of explicit textual traces (e.g. attitudes, interest & identity descriptors, etc.), each input likewise results in n textual traces. The interpretation of these traces is additive, meaning that reaction to the n textual traces is a linear combination of reactions to each textual trace.

In semantic fabric representations (i.e. realm of cultural taste, realm of food), simulation looks like distance measurement. A reaction to some input is positive if the input’s nodes are proximal to the individual’s *taste ethos*, and is neutral or negative if the input is distant, or if the input is proximal to a negative activation (in the food realm). The scalar value of this reaction, though, is not likely to be meaningful in isolation; the value *is* meaningful when reactions to two different inputs are compared—e.g. a person is predicted to like one thing *more* than something else. Hence, the design of applications like Identity Mirror and virtual tastebuds encourage constant reaction to a variety of input, in order to tease out the value of comparison for users of these applications. This notion of simulation via distance measurement is consistent with some work in experimental psychology. Montgomery’s (1994) study argued that the valence of reaction generated by someone’s perspective was proportional to the psychological distance between the experimental subject and object.

In semantic sheet representations (i.e. realm of attitudes and realm of humor), reactions are simulated via memory-based reasoning. An input is processed via computational reading and either a set of

topics is extracted (if the input is simple e.g. 'bush, war'), or a set of attitudes is extracted (if the input is opinionated e.g. 'war is bad'). Topics evoke a set of attitudes in the generalized model, and the overall reaction is a linear combination of the PAD values for these evoked attitudes. Inputs containing attitudes also evoke their corresponding attitudes in the generalized model, but the Pleasure scale in the overall reaction is instead replaced with the degree of agreement. For example, a pro-war model will react pleasurably to the input 'war', but will react displeasurably to the input 'war is bad'.

Limitations of simulation. At least two criticisms can be fairly leveled on this simple simulation strategy—at worst, generalization methods like 1) spreading activation and 2) analogy may in fact lead to incorrect inferences; and 3) even at best, a generalized model captures a stereotype of a person, but while persons defy their own stereotypes, straight-forward memory-based simulation ignores this level of human dynamism.

First, some doubt can be cast on the foundation of generalization. We have supposed that spreading activation can work because choosing from the field of cultural interests is additive. However, suppose that some cultural interests acted as contextual operators. An example is when a person lists in his profile a bunch of kitschy and over-popular interests, but lists one devastatingly insightful and hip interest. For a real person who is reading this profile, the presence of the latter acts as a contextual operator, forcing re-interpretation of the kitschy interests as a conscious act, and this transforms understanding about the tastes of the person represented by the profile. Bayesian modeling can handle contextual operators, but simple spreading activation cannot.

Second, generalization via topic hierarchies and analogy may also go awry. Ideology and other forces may violate the general assumption that attitudes toward a topic and its analogs are mutually consistent. 'Tree' and 'rock' share the super-topic of 'nature', so aesthetic consistency presumes that positive attitudes about 'tree' can induce symmetrically positive attitudes about the sister concept of 'rock'. However, aesthetic consistency is violated in the instance of 'dogs' and 'cats', which share the super-topic of 'pets'. It is far from clear that a sympathetic attitude toward 'dogs' can predict a sympathetic attitude toward 'cats'. Looking at empirical data handled in attitude modeling, 'dogs' and 'cats' tend actually to form an aesthetic opposition—dog lovers tend toward distaste for cats, and vice versa. Pet preference seems to be a politicized and ideological space exempt from aesthetic consistency—perhaps because pets are so often invoked in the present culture to signify the personality of their owners. Dog lovers are presumed to be social, whereas cat lovers are presumed to be asocial. A hack to minimize these effects is to avoid spreading activation directly across sister nodes (i.e. nodes sharing a parent) in topic hierarchies, and also to focus on attributes

(e.g. 'propertyOf') rather than taxonomic features (e.g. 'isA') when generalizing via analogies.

Third, even when a generalized model has successfully captured a stereotype of a person, the formation of reactions cannot be so simplistic; reactions may in real life deviate from expectation due to human dynamism and dialogical nature. A person's reaction is more sophisticated than simply reiterating their views. Most views are complex enough such that a person may adopt both positive and negative attitudes about a topic under different social or political contexts, and views also evolve over time. Future work should thus look to a Bayesian dimension to modeling. An input may also provoke a particularly clever, even self-contradictory reaction, as an individual hopes to achieve sarcasm, irony, or self-overcoming. In "Principles for a Sociology of Cultural Works", cultural theorist Bourdieu (1993) explained that theorists, for example, are keenly aware of their location in the space of criticism, and thus their reactions often deviate from what is expected of them because they are constantly playing games with their self-stereotype. The sort of reflexivity that causes reactions to play with expectation is related to the idea of dialogism (Bakhtin 1935). Unfortunately the sophistication of this dynamism makes it difficult to simulate. At present it is beyond the computational scope of this thesis work, though thoughts for future address of this aspect are presented in Chapter 6.

§

Having presented a simplistic model of simulation, we now discuss *perspective-based applications*, which are enabled by taste and attitude simulation. Six such applications were implemented, and they were introduced in Chapter 1—virtual mentors; virtual tastebuds; perspective-based art; an identity mirror; a system that facilitates social introductions; and a Freudian joke teller. These perspective mirrors and simulators employ the just-in-time information retrieval (JIT-IR) paradigm in their interface, as a way to create a journey for users through someone else's perspective. Related work on JIT-IR systems and perspectives in the interface is now surveyed.

Related work. Just-in-time information retrieval systems (JIT-IR) are software systems that monitor the user's context, build queries from this context in the background, and *pushes* content to the desktop that a user may find relevant at this particular moment. Most technology users have experienced basic examples of JIT-IR—such as auto-completion in operating systems and word completion in mobile phones. Drummond (1992) pioneered auto-fetching of relevant programming libraries based on observations of a user of a programming shell, terming the technique *active browsing*. LETIZIA (Lieberman 1995) embedded an observational agent in a web browser, and loaded other web pages of potential interest to the user in a side pane. The Remembrance Agent (Rhodes & Starner 1996)

observes users writing emails and papers, and displayed past ‘remembered’ documents based word frequency similarity.

Many other systems followed the JIT-IR paradigm (Joachims, Freitag & Mitchell 1997; Badue, Vaz & Albuquerque 1998; Budzik & Hammond 1999; Kulyukin 1999; Maglio *et al.* 2000), leading to recapitulation and formalization of the paradigm as autonomous interface agents (Lieberman 1997), and just-in-time information retrieval agents (Rhodes & Maes 2000). One criticism of JIT-IR systems is that pushing documents to the interface detracts from rather than enhances user’s task performance because they forces users out of the flow that is their task context. Another criticism is that JIT-IR systems incorrectly assume that document ‘similarity’ entails ‘relevance’ and ‘usefulness’, with Budzik *et al.* (2000) demonstrating experimentally some tasks in which similarity-based retrieval was not particularly useful.

At the confluence of interface design and storytelling, we also identify a body of work that explored the notion of perspectives in the interface. Apple Computer’s Guides project (Oren *et al.* 1990) was a multi-character interface that assisted users in browsing a hypermedia database. Each guide embodied a specific character (e.g. preacher, miner, settler) with a unique “life story.” Presented with the current document that a user is browsing, each guide suggested a recommended follow-up document, motivated by the guide’s own point-of-view. Don’s (1989) “We Make Memories” was an interactive multimedia installation featuring a great-grandmother who would tell stories in fragments. The particular trajectory of storytelling was triggered by a viewer’s sensed context. Guides and “We Make Memories” distinguish themselves from other character-based interfaces like Microsoft’s Bob (also now, the Microsoft Office paper clip) because they tried to impart the psychology and memories of specific persons unto a user’s present context, whereas Bob did impart any subjective substance but was merely an anthropomorphic device. Point-Counterpoint (Budzik & Hammond 2000) was a proposed JIT-IR system that would retrieve documents by both a *similar query* for documents according with the user’s textual context but also by an *opposite query* for documents by experts who dissented with the present context.

§

We reflect upon lessons learned in the development of the thesis’s six perspective-based applications. The aim of this class of applications is not so much to improve user performance in conventional tasks. Rather, they explore new capabilities for computers such as supporting self-reflection, person learning, and deep customization. Previous insight into JIT-IR systems is now parlayed into a working framework for thinking about how models of persons’ perspectives can be communicated through the interface effectively. This framework is articulated as three design

principles—1) continuous observation and feedback, 2) just-in-time and just-in-context, and 3) tinkerability. These lessons introduce and incorporate related work in intelligent user interfaces.

Continuous observation and feedback. It may be said that just as light has no resting mass, perspective is not intelligible in stasis. To fully appreciate and intuit someone’s simulated perspective as captured in a generalized model, it should be animated and allowed to react to a broad many things. In the interaction design literature, Bill Gaver (1991) has foregrounded the idea that an artifact is easy-to-use when its *affordances* are perceptible. According to Gaver, perceptual psychologist J.J. Gibson (1979) first coined the term ‘affordance’ to describe the ability of people to intuit an object’s potential for actions from perceptual cues and feedback. For example, thin vertical door handles afford pulling, while wide horizontal doorplates afford pushing. A person’s aesthetic perspective as embodied in a generalized person model, however, affords more complex actions. According to Gaver, systems of complex actions require more active perception, and many exploratory engagements with the system. Complex objects can often be dissected into nested affordances revealed over time. Thus, the affordances of a perspective can be conceptualized as consisting of *semantically nested affordances*—that is, aesthetic consistencies latent in the perspective can suggest divisions of affordance-space. A useful way to expose the affordances of a person’s tastes, attitudes, ways of perceiving, etc. could be for an application to continuously observe a user’s current textual context, and proactively offer feedback. Because information push treats all user browsing and writing activity as implicit input, it avoids the time cost of users manually formulating queries. Thus, exploration of affordances can be had with less effort, and animating a model resembles an exploratory walk through a perspective. What exploratory walks achieve is an increase in the *surface area* of a perspective’s lessons.

Just-in-time, just-in-context. Perspective-based applications continuously observe the user’s context and offer feedback. The nature of this feedback should be *just-in-time*—meaning that feedback given should pertain to the user’s present textual context, such as the text on a webpage that the user is scanning over with the mouse, or the sentence that the user just finished typing. JIT-IR interaction suits perspective-based applications because the reactions they simulate make the best sense when they can be readily bound to the user’s present context. But more than just-in-time, the nature of feedback should also be *just-in-context*—meaning that the reaction should pertain to the gestalt of the user’s present context, rather than considering only a subset of the context. For example, Lumiere (Horvitz *et al.* 1998) aimed to be *just-in-context* because its Bayesian model tried to make sense of all the pieces in the user’s present context to infer the user’s goal—achieving synesthetic reasoning. As a counter example, a Google Ad Words advertisement is embedded in a webpage and is dynamically populated based on the presence or absence of particular keywords on the webpage. But

these ads are not *just-in-context* because they appeal to only fragments of the text. Following Budzik *et al.*'s (2000) reasoning that similarity is incommensurate with relevance or usefulness, we argue that for a perspective to be communicated effectively, it should have to attempt reactions to the gestalt of a user's context. Such reactions are exciting—they play with the idea of how prior attitudes and tastes might combine and compose under novel contexts. Taste and attitudes in their general sense, are intertextual, so they demand to be explored as such. Finally, just-in-context reactions can provoke users and motivate them to further investigation.

Tinkerability. While proactive feedback gives users of perspective-based applications many perceptual entrees into a perspective, there needs also to be a way for users to tinker with the perspective itself. If just-in-context reactions is the provocation for critical thinking, then tinkering is the follow-through. For example, WWTT allows a perspective to be changed in a text editor, and allows a user to dig deeper into any reaction by asking a virtual mentor to justify a reaction with a corpus of quotes from memory. The Aesthetoscope allows the user to tweak the perceptual perspective by moving horizontal sliders to change location in the Jungian perception space, and then immediately visualizing the consequences of this change. Avatars in the Synesthetic Cookbook can be reprogrammed with food keywords on-the-fly by clicking on the mouth to reveal the tastebuds. Tinkerability also ensures that users learn about the limitations of computational person modeling—for example, seeing quotes from a virtual mentor that just reacted allows users to verify the accuracy of the generalized model.

In summary, this chapter opened with a survey of related work in user modeling. An approach to modeling persons from everyday texts was introduced and compared with the related work in the user modeling. The approach was structured into three phases—acquire, generalize, and apply. Apropos model acquisition, knowledge representations for five modeled realms were explained, and they were inter-related along the dimensions of connectedness and consistency. Apropos generalization, it was suggested that convergence and truth maintenance are characteristics of the generalization process. Details for spreading activation, analogy, and imprinting were then presented. Finally, apropos model application, a simplistic model for simulating reactions was presented, and framed by related work on just-in-time information retrieval systems, three design considerations for the class of perspective-based applications were articulated.

3 Techniques & Technologies

This chapter presents the techniques and technologies that are central to person modeling from everyday texts. Sections 3.1 and 3.2 present *reading for affective themes*, and *culture mining* – situating these techniques in their respective related works. Sections 3.3 and 3.4 articulate two technologies that are core to this thesis – *commonsense reasoning*, and *textual affect analysis* – they are also compared to related works.

3.1 Reading for affective themes

Reading for Affective Themes (RATE) is developed in this thesis as the means of acquiring models of – a person’s tastes, attitudes, ways of perceiving, taste for food, and sense of humor – from his/her everyday texts, such as weblog diaries, social network profiles, homepages, emails, etc. RATE is an associative reader that employs lightweight natural language processing techniques, a knowledge-based topic spotting skimmer, and a knowledge-based textual affect skimmer. Corresponding to each realm is a *reading schema* specific to that realm’s person model. A reading schema dictates how the outputs of skimmers should be translated into a realm’s model. For example, the reading schema for the realm of attitudes is most direct – a person’s everyday texts are skimmed for topics and affects, and each topic is associated with the affective contexts that underlied its instances in the text. Using statistical estimation over a person’s entire corpus, stable affect values can be bound to the topics¹⁴ – these constitute the affective themes gisted from the person’s text that then become input to the model generalization phase.

The rest of this section gives systematic presentation to the RATE technique – 1) related work in computational reading is discussed, and the RATE technique is compared to this literature; 2) the genre of first-person, everyday texts from which this modeling work proceeds, is characterized; 3) schemas specific to reading for each of the five aesthetical realms are presented; 4) an algorithm for RATE is narrated.

§

¹⁴ Since RATE readers need a large corpus of texts to gist from, we assumed that a person’s corpus of texts that often spanned months or years would still contain time-stable attitudes. The phenomenon of attitudes changing over time was not dealt with in this thesis.

Related work in computational reading. Reading for affective themes builds on related work in the interrelated literatures of computational reading, story understanding, text summarization, and structuralist models of reading. RATE's textual affect analysis aspect is expansive enough to warrant its own survey of related work – which will be given in Section 3.4 to accompany discussion of the textual affect analysis technology.

Formal and knowledge-based story understanding. Understanding stories is a classic modeling problem in artificial intelligence with a long history and vast literature. Some early story understanding systems (Charniak 1972; Charniak 1977; Cullingford 1978; Wilensky 1978; Carbonell 1979; Dyer 1983) modeled stories formally using elaborate hand-crafted knowledge in the form of predicate logic rules, frames, scripts, plans and goals, which could achieve in-depth understanding, but were too brittle to handle most stories found "in-the-wild." To improve robustness, a more recent incarnation of knowledge-based modeling employed a very large collection of commonsense knowledge to support inference making over stories (Mueller 1998).

Memory-based readers. Most story understanding systems assumed that understanding was objective, but some systems adopted a more cognitive posture. IPP (Lebowitz 1980) and CYRUS (Kolodner 1984) understood stories using episodic memories, and IPP was capable of generalizing episodes. These memory-based approaches were bolstered by Haberlandt *et al.*'s (1980) experimental finding that human readers naturally projected episodic structure onto stories. Shifting from objective understanding to subjective reading, Moorman and Ram's (1994) ISAAC was a creative reader that could focus, attend, and willfully suspend disbelief. AQUA (Ram 1994) could interleave and motivate reading with asking of questions. ISAAC and AQUA represent segues from story understanding into computational reading. Rather than regarding stories as containing truths to be recognized, computational readers skim text, constructed *situation models* (Zwaan & Radavansky 1998) to explain the text, and revise and consolidate these models to maintain consistency.

Information-extraction readers. Some other systems were also backed by frames and scripts, but used them to instead skim texts, leading to a more opportunistic, information retrieval approach to text understanding. FRUMP (DeJong 1979), for example, used semantic frames and case patterns to skim summaries from new stories. Plot units (Lehnert 1982) was another plan for narrative summarization. More recently, FERRET (Mauldin 1991) improved upon FRUMP's pattern-based approach. SpinDoctor (Sack 1994) used patterns sorted by ideology (e.g. (?criminal murdered ?victim)) and adopted actor-role schema from Greimasian to summarize the ideology implied by news stories.

Probabilistic and statistical readers. Another class of story understanding systems incorporated connectionist networks to support story comprehension (Charniak 1986; Dolan 1989; Lange & Dyer 1989; Miikkulainen & Dyer 1991; St. John 1992; Langston *et al.* 1995). The probabilistic aspect of these systems softened symbolic constraints and this in turn improved the handling of contextual issues in story processing. Recent work (Halpin 2003) makes use of both fully symbolic and also fully statistical approaches such as Latent Semantic Analysis to analyze story plot.

The role of affect and perspective in stories. The importance of affect as a parameter and organizer of cognition in humans and machines has been argued for (Simon 1967; Sloman & Croucher 1981). Tan (1994) conceived a model of story processing based on the premise of stories as super-episodes containing many emotion episodes. In story understanding systems, formal symbolic accounts of affect in relation to modeling characters and plot were presented in (Lehnert 1982; Dyer 1983; Elliott & Ortony 1992). Formal models of the related notion of narrative perspective were presented in (Wiebe & Rapaport 1988; Wiebe 1994).

Reading for affective themes (RATE). Reading for affective themes mixes the approach of information-extraction readers, with the approach of knowledge-based story understanding. At the large scale of reading, RATE takes an information extraction and statistical approach to estimating statistically salient topics and affects, but at the small scale of reading, it uses knowledge-based component technologies.

Rather than attempting in-depth understanding achieved by formal story models, RATE privileges recall and robustness over depth. Thus, it gathers textual evidence for affective themes at the sentence-level, stores the evidence in a memory database, and uses statistical methods to estimate average values and salience for affective themes using the memory base. RATE is also associative. For example, in the realm of attitudes, an attitude is a topic plus a statistically summarized affect value. A topic's stable affect value is measured by statistical estimation over all the affect values that co-occurred with instances of the topic in the text. This associative method for attributing affect to topics can be compared to work on learning the meaning of words from their surrounding context (Berwick 1989; Cardie 1993), and Yarowsky's (1992) statistical method for word-sense disambiguation. A difference is that RATE applies statistical disambiguation not to words but instead to thematic abstractions (topic and affect), which were themselves inferred from the text. To adapt RATE so that it can populate person models for all five realms, five reading schemas are used, which prescribe how the outputs of the text skimmers are to populate the various realms' models. The idea of reading based on different schemas can be compared to situation models for reading (Zwaan & Radavansky 1998), and to actor-role analysis (Sack 1994).

At the small scale of reading, a RATE reader skims text dually with a topic-spotting skimmer, and with a textual affect skimmer. These skimmers take a knowledge-based approach that may be compared with (Mueller 1998). The topic spotter belongs to ConceptNet (Liu & Singh 2004), which makes use of a large knowledge base of common sense relations to annotate passages of text with inferences, and then triangulates on topics that emerge in the space of inferences. The textual affect analyzer also uses a component, Emotus Ponens (Liu, Lieberman & Selker 2003) that makes commonsense inferences about a text's affective qualities based on its event structures.

Having compared RATE with related approaches in computational reading, we should keep in mind that RATE's approach is not being considered for arbitrary story texts. RATE is considered only for the genre of first-person, everyday texts from which the taste and attitude modeling work proceeds. Next, we characterize the properties, assumptions, and limitations of this genre of text.

§

The genre of everyday texts. Consider weblog diaries that allude to the same repertoire of friends and situations day after day, entry after entry—each time giving a slightly different take, but always seeming to recapitulate a theme. *Everyday texts* are the target of our person modeling. Examples of suitable everyday texts are weblog diaries, commentary-rich papers, personal emails, instant messenger conversations, personal homepages, and social network profiles. We define everyday texts as the genre of first-person texts that we author informally and as part of day-to-day life. As such, these texts are semi-private (they are certainly not formal) and they are *self-expressive*—they should have a strong editorial quality because they were written with free expression of the author's opinions about all the things that he may have encountered in everyday happenstance. Everyday texts are also *self-expressive* because we hypothesize that by modeling tastes and attitudes expressed in them, we have actually modeled their author. Some arguments are made for the suitability of everyday texts as a target corpus for person modeling.

First, everyday texts facilitate the ease of natural language processing tasks necessary for computational reading by being first-person. Third-person texts are constituted by potentially multiple narrative voices, characters, and perspectives which alternate, sometimes even without explicit segmentation. Segmenting and disambiguating between narrative voices is a difficult problem in narrative comprehension—the complexities of which are exposed in (Wiebe 1994). Early attempts at story understanding systems such as Charniak's (1972) were tenuous as scenarios, in part because the children's stories that were chosen as the textual corpus were riddled with narrative shifts between multiple characters. Everyday texts, in contrast, are first-person, and can be more successfully assumed to contain, for the most part, the attitudes, dispositions, and tastes

constituting the writer's own perspective; hence avoiding a hairy segmentation and disambiguation task.

Second, redundant expression is a boon to reading for affective themes. Everyday texts are fraught with redundant evidence of the writer's judgments about the world. Biographical texts would seem in some respects to be a superior source for person modeling because they contain crisp and summary propositions about a person's life experiences and attitudes. Ironically, while precision is appreciated by fatigue-sensitive human readers, shallow machine readers such as the RATE readers implemented for this thesis require redundancies and lengthy textual corpora in order to triangulate upon a correct model of a person. Everyday texts such as a weblog diary or a commentary-rich paper do not explicitly state the writer's attitude about a particular topic just once, but rather, the attitude toward that topic tints the subtext of several statements, thus creating multiple and redundant evidence. For example, consider weblog diaries that allude to the same repertoire of friends and situations day after day, entry after entry—each time giving a slightly different take, but always seeming to recapitulate a theme. In addition, whereas biographical texts often omit detailed instances of a person's judgment in favor of more general statements, everyday texts are not subject to this editorial pressure. That everyday texts find the writer opining over a very broad set of subject matters is a boon to the acquisition of rich person models.

In contrast to its advantages, everyday texts are herewith subject to two limitations—the publicity bias and the performance bias—as explained below.

Everyday texts are limited by a publicity bias because complete candor about one's views is not always possible due to anxiety of audience and anticipatory self-censoring. The sources of everyday texts enumerated in the above all anticipate an audience of more than just the writer herself. Authors of weblog diaries and social network profiles may anticipate that friends and co-workers could stumble upon the text. Boyd (2004), in her ethnography of social network profiles, points out that a social network profile that is at once subject to the audience of people from many different life contexts—such as friends, family, lovers, ex-lovers, and co-workers—is necessarily steered away from complete candor. In fact, it is more likely that a person composing a social network profile will only reveal aspects of self that are universally palatable to friends, lovers, and the boss; thus person models acquired from texts affected by these biases will capture who a person portrays herself as, not necessarily the “real” person that they really are. Notwithstanding authorial restraint instilled by anticipation of publicity, a supposition of RATE is that so long as attitudes, tastes, and ways of perceiving are insinuated by a text or present in the intratextual unconscious, there is hope for their indirect excavation.

Performance bias is the idea that everyday texts may contain not a single underlying perspective, but rather, may manifest perspectives associated with different personae, according to the social context represented by different audiences. Goffman's (1959) thesis in *The Presentation of Self in Everyday Life* was that social interaction could be likened to a dramatic performance—we are all capable of wearing different social masks, and the mask we choose to wear is negotiated by how we fit into varying social contexts such as—the workplace, with friends, with a lover, with family. This limitation may nonetheless be found acceptable because while an everyday text composed for one's research and a text composed for one's friends reveal different aspects of oneself, the two texts tend to express judgments over different topics—one set related to research, the other set related to social life. In such cases where contradictory affects about the same topic are detected across one individual's everyday texts, RATE's statistical estimation process would simply cause the contradictory affects to cancel out.

§

Reading schemas. Reading for affective themes employs two skimmers. A topic-spotting skimmer returns a scored list of topics for textual passages of arbitrary length. A textual affect skimmer returns a PAD affect value for textual passages of arbitrary length. However, the raw outputs of these two skimmers still need to be mapped into the five realm models, and their inter-relationship is not yet specified. To operationalize the mapping, five reading schemas are now presented.

But first, a uniform vocabulary for describing schemas is desired, and a vernacular that is appropriate to build on is found in Greimas's (1966) *isotopy* model of textual interpretation. The isotopy model gives a standard account of how coherent themes can arise out of a collection of textual fragments, which are themselves ambiguous. The model is also consistent with schemes for word sense disambiguation (*cf.* Yarowsky 1992). A brief review of the isotopy model follows. A simple vocabulary consists of the terms—*lexeme*, *seme*, *classeme*, and *isotopy*. A *lexeme* is a fundamental unit of a lexicon—e.g. "Frank Sinatra" is a lexeme in our lexicon of cultural taste. For the reader, the meaning of a lexeme is unconfirmed or ambiguous—that is, each lexeme could have multiple senses of meaning, or *semes*. But not so for the writer, who had a *seme* in mind when uttering a lexeme. To discover the *seme* that was intended, Greimas assumed that a text's meaning should converge; thus by co-occurring in a text, lexemes mutually disambiguate one another. Greimas termed the strategy of reading for convergence, *monosemization*. Adding an intermediate structure to the analysis, each *seme* is capable of a number of higher-level contextual entailments, known as *classemes*. For example the *seme* of 'barks' in the sentence "the dog barks" can be associated with *classemes* such as [+caninity] [+animalhood], etc. (NB the bracket and plus sign

convention for expressing classemes). During monosemization, lexemes are disambiguated into semes such that a system of complementary and mutually consistent classemes are selected. One whole system of classemes is called an isotopy. In RATE, *lexemes* correspond to unconfirmed and ambiguous concepts; *semes* correspond to significant (confirmed) and disambiguated concepts; *classemes* correspond to higher-level abstractions such as discourse-level topics. Because the basic isotopy model gives us a useful way to describe how themes are built up, we describe reading schemas by sub-typing its vocabulary. Schemas for our five realms are now presented.

Attitudes. The goal of reading for attitudes is to summarize a weblog diary into a list of its topics, each tinted with a stable affect. When the weblog is first parsed, many word-level ('party'), phrase-level ('bachelor party'), and event-level ('throw party') *topic-lexemes* will be recognized in the topic spotter's lexicon (ConceptNet plus topic folksonomy). The topic-spotting skimmer intersects the ambiguous meanings of these topic-lexemes, converging upon a list of *topic-classemes*. At the same time, each topic-classeme points back to the topic-lexemes that were its supporting evidence. The textual proximity of each topic-lexeme is associated with some unconfirmed *affect-lexeme* as scored by the affect skimmer, and given as a PAD value. Using statistical estimation, the affect-lexemes that underlied each topic-classeme's supporting evidence is summarized into an *affect-classeme*. Finally, the topic-classeme and its affect-classeme are bound into a duple, and this is called an *attitude-classeme*. The scored list of a text's attitude-classemes constitutes an isotopy which populates the attitudes person model.

Humor. The humor schema is explained as a transformation of the attitudes schema because both have semantic sheets as their target representation. *Topic-lexemes*, and *topic-classemes* are retained from attitudes schema. Instead of measuring affect in terms of PAD values, each 3-tuple *affect-lexeme* value outputted by the affect skimmer is transformed into a unary *tension-lexeme* value. Tension is scored between 0.0 and 1.0, proportional to the degree to which the incoming PAD resembles a prototype for psychic tension—displeasure, high arousal, and dominance. High arousal plus dominance represents Freud's notion of *tendentiousness*. Tension-lexemes are then summarized statistically into *tension-classemes* and bound to the *topic-classemes* to produce *humor-classemes*. The scored list of a text's humor-classemes constitutes an isotopy which populates the humor person model.

Cultural taste. Although the thesis summary only described how cultural tastes were extracted from social network profiles, this reading schema allows cultural tastes to also be extracted from any everyday text such as a weblog diary or homepage. Recall that since the model of cultural taste is represented on a semantic fabric, it does not attach a PAD value for each interest or identity descriptor; hence, PAD will again be transformed into a unary *significance-lexeme* value

that screens for significance of interest—based on PAD’s degree of similarity to the prototype: pleasure, arousal, dominance. A text is parsed and filtered through the cultural taste lexicon, resulting in *interest-lexemes* (e.g. books, films, authors, music genres) and *identity-lexemes* (e.g. ‘extreme sports lover’, ‘intellectual’, ‘goth’). Then, the significance-lexemes underlying all identical *interest-lexemes* are statistically summarized and bound to the interest descriptor to constitute an *interest-seme*. For example, we find that the television show ‘Fresh Prince of Bel Air’ occurred four times, and all had high significance, thus, ‘French Prince of Bel Air’ becomes an interest-seme, for we have confirmed its significance to the writer. Likewise, the significance-lexemes underlying identical identity-lexemes are statistically summarized, and we call what results is an *identity-classeme*.¹⁵ Finally, the scored list of interest-emes and identity-classemes constitutes the isotopy and populates the cultural taste person model. The stable significance scores associated with each interest and identity determine the initial activation energy of nodes on the taste fabric. Acquiring the cultural taste model from social network profiles did not benefit from this level of nuance.

Food. The food schema is explained as a transformation of the cultural taste schema because both have semantic fabrics as their target representation. *Significance-lexemes* carry over from the cultural taste schema. Instead of interest-lexemes and identity-lexemes, the raw text is parsed into *recipe-lexemes*, *ingredient-lexemes*, *cooking-procedure-lexemes*, *sensation-lexemes* according to the lexicon of the food fabric—and these are associated with their significances and summarized into their corresponding semes. Furthermore, because an extensive metadata system annotated recipes with their cuisine type and basic flavor types, the topic-spotting skimmer then builds up confirmed recipe-emes, ingredient-emes, etc. into *cuisine-classemes* and *basic-flavor-classemes*, for an additional level of unification. Finally, all the resulting semes and classemes constitute the isotopy.

Perception. The perception schema is the most unlike the other four schemas. Whereas cultural taste schema tracked a very large lexicon of interests and identities, perception schema only tracks—*ego-lexemes* (occurrences of ‘I’, ‘me’, or ‘my *’), *alters-lexemes* (any syntactic subject or object that is not an ego-lexeme), *action-lexemes* (verbs), *incoming-lexemes* (any sentence frame whose agent is an alters-lexeme and whose patient is an ego-lexeme), *outgoing-lexemes* (any sentence frame whose agent is an ego-lexeme and patient is an alters-lexeme), and *mental-lexemes* (invocations of thought states, e.g. ‘I thought that’). The affect skimmer output PAD affects for each of these lexemes. The PAD affect of an incoming or outgoing lexeme is

¹⁵ Significant identities are called classemes, whereas significant interests are called semes. This was a deliberate semantical choice. A classeme is a particularly dominant seme that is also more abstract and contextual than a seme. Identities are extracted inexactly from the text, hence classemes; whereas interests are explicit in the text, hence semes.

based on the PAD of the agent, action, and patient—this calculation is complex and is covered in Chapter 4. Finally, *classemes* are produced from these lexemes, as the statistical summary of their associated PAD values. Two exceptions are that action-lexemes produce the *ratio-claseme*, which gives the statistical ratio of introverted to extroverted actions; and mental-lexemes produce the *mental-claseme*, which estimates the density of thought state propositions in the text. These classemes constitute a perception isotopy, and are later fed into a learned classifier that yields their location in the Jungian dimensional space.

Five reading schemas were presented to prescribe how the output of text skimmers should be used to populate each realm’s model. Using these schemas, we achieve a generality for acquisition—for example, we note that cultural taste can be acquired not only from social network profiles, but using its schema it can be acquired from weblog diaries or homepages. Next, we narrate a four-step algorithm for RATE that uses these schemas.

§

An algorithm for RATE is now presented. The goal of implemented RATE readers is to digest a corpus of everyday texts into the schemas of the five realms. This algorithm is applicable to weblog diaries, homepages, and other free-form everyday texts; this discussion will make less sense for social network profiles, which are already semi-structured, though it is not incompatible with those texts. The algorithm has four phases—1) natural language normalization; 2) topic skimming; 3) affect skimming; and 4) statistical summarization.

Step 1—natural language normalization. The first step of processing is to prepare texts for computation by natural language normalization, which includes tasks like tokenization, semantic recognition, part-of-speech tagging, lemmatization, anaphora resolution, phrase chunking, phrase linking, and syntactic frame extraction. The MontyLingua natural language processor (Liu 2002) is used as the subsystem to perform these normalization tasks—so only a very brief review is given here.

- ❖ **Tokenization.** An open text is split into word and punctuation tokens. Punctuation marks are assumed distinct tokens, modulo marks used in abbreviations. Contractions such as “can’t”, “shouldn’t” are resolved into “cannot” and “should not”.
- ❖ **Semantic recognition.** Using the lexicons specific to each realm, textual fragments are normalized into those lexicons. In addition, named-entities are recognized in preprocessing because they confuse the tagger and chunker. Examples of entities and their normalized forms (e.g. “John W. Stewart” ==> \$NAME_JOHN_W_STEWART\$), temporal expressions

("last Tuesday" ==> \$DATETIME_1142022711\$), and ontological interest items (e.g. "Nietzsche" ==> \$BOOKAUTHOR_FRIEDRICH_NIETZSCHE\$).

- ❖ **Part-of-speech tagging.** Tokens are assigned part-of-speech tags such as VB (verb, root form), NN (common noun singular), and JJR (adjective, comparative form) using the Penn Treebank tagset, based on Eric Brill's (1992) transformation-based learning tagger for English.
- ❖ **Lemmatization.** The lemma, or normal form, for nouns and verbs are generated. Lemmas are important supplemental information added as annotations to tokens. Lexical features such as number are stripped from nouns (e.g. "robots" ==> "robot"), and tense is stripped from verbs (e.g. "went" ==> "go").
- ❖ **Anaphora resolution.** An anaphor is a referring expression such as a pronoun (e.g. "he", "they") whose referent usually lies in the immediately antecedent sentences. As the reader scans the textual tokens sequentially, a deixis stack of possible referents such as noun phrases, are maintained. When an anaphor is encountered, it is resolved with the aid of the deixis stack, according to the knowledge-poor resolution strategy outlined in (Mitkov 1998).
- ❖ **Phrase chunking.** From a flat sequence of tagged tokens, phrases will emerge as the boundaries of phrases are identified, e.g.:

"John/NNP likes/VBZ to/TO play/VB board/NN games/NNS"
 ==>
 (NX John NX) (VX likes to play VX) (NX board games NX)

Here, NX denotes noun chunks, and VX denotes verb chunks. Moving from words to the level of chunks allows text to be regarded on the conceptual level. Phrase chunking is accomplished by a set of regular expression patterns operating over the stream of words and tags.

- ❖ **Phrase linking.** To inch toward a syntactic parse tree, verb chunks need to be linked to their noun phrase and prepositional phrase arguments. Accomplishing this requires some heuristics about verb-argument structure, as well as semantic selectional preferences, gotten from common sense knowledge in ConceptNet. The following example illustrates one successful resolution in light of ambiguity:

(NX John NX) (VX robbed VX) (NX the bank NX) with (NX the money)
 ==> (John
 (robbed
 (the bank
 (with (the money))

Note that "the money" was linked to "the bank" and was not instead implicated as the second argument to the verb "robbed". This mechanism makes the common sense

assumption that “the money” was not the instrument used by John to perform the robbery, though such a scenario is certainly possible. More challenging cases of linking arise in the encounter of surface movement phenomena such as subject-verb inversion (e.g. “have you the money?”), topicalization (e.g. “to the bank I go”), and passive voice (e.g. the subject is nested in an agentive “by” phrase in the utterance “The bank was robbed by John”).

- ❖ **Syntactic frame extraction.** Finally, the event structure of each sentence can be captured concisely as a syntactic frame, e.g.:

```
SENTENCE FRAME
[verb] “robbed”
[subject] “John”
[direct object] “the bank” “with the money”
```

A sentence frame may contain any number of direct and indirect objects. A sentence frame is constructed for each clause in the text. A dependent clause has a frame which is linked to the frame of the clause upon which it depends.

Step 2—topic skimming. After open texts are normalized into phrases, events, and lexical items, the topic-spotting skimmer is run, according to the five specific reading schemas. Topic extraction is performed using ConceptNet common sense reasoning toolkit, but augmented with various folksonomies in order to gist topics for realms with custom lexicons like cultural taste realm. The ConceptNet topic extraction mechanism is presented in Section 3.3.

Step 3—*affect skimming.* A mechanism for textual affect analysis scans over text and annotates it—at the various granularities of phrase, sentence, paragraph, and document—with its affect valence score, given in the Pleasure-Arousal-Dominance format of Mehrabian (1995b). The reading schema prescribes the particular granularity that is used. Some schemas also prescribe a transformation of PAD into measures of tension or significance. The affect skimmer is constituted from knowledge of both—1) *lexical affect*—a database of sentiment words (e.g. ‘cry’, ‘sad’) annotated with affective valence plus a database of non-sentiment words (e.g. ‘homework’, ‘recess’) annotated with their typical affective connotation as measured in psychological focus groups; and 2) *event-centric affect*—a database of commonsense knowledge affording typical affects for everyday concepts and events (e.g. ‘be(person,fired)’). By combining lexical and eventual treatments, affect can be sensed as a more robust combination of both surface language and deep meaning. Section 3.4 presents details for textual affect analysis technology.

Step 4—*statistical summarization.* Having transformed a corpus of raw everyday texts into normalized concepts, topics, and affects; affects are grouped by the concepts and topics they underlie, and are statistically summarized, according to the needs of each reading

schema, using heuristically determined statistics. For example, to gist stable attitudes for the attitudes realm, a reinforcement statistic (details in Chapter 4) is used in order to simulate a reflexive memory. The reinforcement statistic has rather low tolerance for topics having contradictory PAD values—in which case a neutral PAD is learned. A more forgiving statistic, the *first-order moment*, is used to stabilize affects in the schemas of cultural taste, humor, food, and perception. Recall that the goal of monosemization is to converge upon stable themes. These statistics, by definition, converge.

One shortcoming of statistical summarization is that it cannot handle certain texts, such as jokes and sarcastic passages, whose action is to provoke *ruptures of isotopy*, as Greimas called it. For example, a joke narrative leads a reader to converge upon a false isotopy, only to have that isotopy overturned by a punch-line, which forces systematic re-semization of the text, resulting in the actual isotopy. Statistical convergence certainly would draw wrong conclusions from such texts, as a statistical summarizer is unable to recognize the rupture, and unable to selectively overturn things that it has already read. So, it is assumed that sarcasm and jokes within the chosen everyday texts are limited to isolated episodes within the textual corpus. The statistical estimation mechanism can tolerate a certain amount of these episodes as ‘noise’ in the text.

In this section, the ‘reading for affective themes’ (RATE) technique was presented. RATE was first articulated against a backdrop of previous work in computational reading—finding resonance with previous approaches in statistical reading, knowledge-based reading, and memory-based reading. Second, the genre of everyday texts to which RATE applies was characterized as first-person and self-expressive. Third, we described five reading schemas which specify how the outputs of RATE’S topic skimmer and affect skimmer should be mapped into person models for the five realms. Finally, an algorithm was given.

Whereas RATE can acquire a person’s textual traces, creating a generalized person model demands that we interpret a person’s traces as articulations against a backdrop of cultural patterns. The next section introduces a technique that models the backdrop of cultural patterns.

3.2 Culture mining

Culture mining is a technique for acquiring the connectedness of a cultural topology by text processing a cultural corpus. The technique was developed to support person modeling in the realms of cultural taste and taste for food. It depends on, and is compatible with work in text mining and information-theoretic measures of similarity. What is added to the literature is an end-to-end macroscopic account of mining—from corpus selection, to

normalization, to statistical analysis, to post-processing, resulting in cultural topology as a semantic fabric. The rest of this section 1) reviews related work in text mining and situates our approach; and 2) narrates an “algorithm” for culture mining.

§

Related work. I would like to implicate several thematics in related work—musical similarity, social network analysis, and text mining.

Musical artist similarity. Whitman and Lawrence (2002) inferred a similarity matrix model of musical artists using a text mining approach and a peer-to-peer similarity approach. For their text mining approach, musical artists were annotated with “community metadata”—adjectives and noun phrases scraped by typing the artist into a web search engine—and artist similarity was assessed as the degree to which their community metadata overlapped. For their peer-to-peer similarity approach, they harvested user’s file folders on the peer-to-peer network OpenNap, and learned artist-to-artist similarity using the TF-IDF metric (Salton & McGill 1983) by treating same-folder artists as co-occurrences. Related to (Whitman & Lawrence 2002) is Whitman and Smaragdis’s (2002), which reported that cultural signatures for music genres could be used, in conjunction with the auditory signal, to classify unknown artists based on style similarity. Ellis *et al.* (2002) infers a similarity matrix of similarity between 400 artists, evaluating the performance of several metrics including community metadata, and variations of the Erdős measure (cf. six degrees of Kevin Bacon). Baumann & Hummel (2005) continues measurement of artist similarity based on shared text in web search results, and essays a characterization of this method as ‘cultural metadata’.

Social network analysis. Much research has examined the explicit structure of social networks, and studied their topologies via graph theory. Newman (2001) mined scientific coauthorship networks and found that collaborations ‘funneled’ through gatekeeper scientists. In our taste fabric identity hubs constitute a similar topological feature. Jensen and Neville (2002) mined structured metadata relations from the Internet Movie Database (imdb.com) and learned a Bayesian network model to represent and predict item distances probabilistically. They also model the relational semantics of social network relations implied between movie actors from the Internet Movie Database and the Hollywood Stock Exchange (www.hsx.com). Finin *et al.* (2005) examine how the FOAF (“friend-of-a-friend”) ontology applies Semantic Web concepts to enable efficient exchange of and search over social information, illustrating how social networks could develop with its semantics already explicit. Finally one work which considers the semantic content entailments of social network users is McCallum, Corrada-Emmanuel, and Wang’s (2005) modeling of Author-Recipient-Topic correlations in a social network messaging system. Given the topic

distributions of email conversations, the ART model could predict the role-relationships of author and recipient. The work considered group clusters and dyadic relationship dynamics but does not consider cultural aggregates.

Text mining. Work in text mining has employed information-theoretic metrics of similarity (Lin 1998) to perform unsupervised feats such as characterizing the meaning of their words by their surrounding context. Latent semantic analysis (LSA) (Landauer & Dumais 1997) is a popular method for comparing the meanings of texts using the Singular Value Decomposition (SVD). Pointwise mutual information (PMI) (Church & Hanks 1990) is the similarity metric currently used for culture mining. Turney (2001) adapted PMI into PMI-IR (similar to Bayes) to support the mining of synonyms based on PMI interpretations of search engine results. The machine learning step of ‘culture mining’ can be most directly compared with the induction of term-term similarity by assuming document closure (Brin *et al.* 1997); that work used a metric called *interest*—a measure of variable dependence, i.e. PMI without the logarithm.

Culture mining. The technique introduced here prescribes a method for modeling cultural topology from a cultural corpus of semi-structured texts, i.e. 100,000 social network profiles. The premise that a large-scale social network community can support taste modeling assumes an emergent semantics is possible in this domain (Staab *et al.* 2002; Aberer *et al.* 2004). Our systematic finding of pairwise similarity between a large number of items in a space can be compared with (Brin *et al.* 1997; Whitman & Lawrence 2002; Ellis *et al.* 2002). Our method advocates recall-maximizing heuristic normalization of text fragments into ontological descriptors and surrounding metadata. The addition of metadata annotations to descriptors can be compared with Whitman & Lawrence’s (2002) community metadata annotation approach. This work’s item-item rather than user-user approach to modeling social networks distinguishes it from user-centric analyses (Newman 2001; Jensen & Neville 2002; McCallum, Corrada-Emmanuel, & Wang 2005). The result of culture mining is a semantic fabric that can be used for recommendation without further preserving original user vectors or item vectors, making the technique an alternative to user-user collaborative filtering (Shardanand & Maes 1995) and item-item collaborative filtering (Sarwar *et al.* 2001).

§

Algorithm. An “algorithm” for an end-to-end modeling of cultural topology from cultural corpora is presented. Unlike a real algorithm, the one presented here prescribes not only the machine learning metric but also a walk through the various decision spaces

encountered in the modeling process. The algorithm generalizes from experience in mining taste fabric from social network profiles. The steps are—1) considering the fitness of a cultural corpus for mining; 2) assembling knowledge resources and disambiguation heuristics for maximizing extraction of descriptors from the cultural corpus; and 3) machine learning and post-processing.

Cultural corpus selection. To judge the fitness of a cultural corpus, we need to define culture and decide what of its aspects are most important to capture. In semiotics, Barthes (1964) proposed culture to be the set of symbols and interpretive strategies salient to the unconscious of a population; plus the system of connections that organize these symbols, and the valence of each symbol—its degree of *privilege*. Similarly, Geertz (1973) described cultures as ‘webs of significance’ spun by people. What sort of textual corpus can fairly capture the symbolism, connotations, and values of a culture? Two promising candidates can be considered—1) a very large sampling of the everyday texts of participants in a culture; or 2) a few prototypical texts that epitomize the breadth and depth of the culture. The latter kind of corpus is *authoritative*—it assumes that a culture’s symbols, connotations, and values can be idealized, and that there are essential texts that also fairly represent the culture. Authoritative corpora have the advantage of being clean and having a high signal-to-noise ratio; disadvantages are that it can be hard to find enough quantity of this kind of text for mining, and that intuitions leading to selections of these texts may encapsulate a bias. In contrast, the former kind of corpus is *populist*—it does not presuppose that the culture has been essentialized in textual form. Such a corpus is appropriate for modeling motivated by ethnography because it is more likely to yield new insight; disadvantages are that the sample embodied by the corpus may represent a skewed distribution, and it may be harder to assess the systematic biases of a large sampling than the bias of a few authoritative texts. Other considerations for corpus gathering include noting the time-span of texts collected into the corpus, and noting how amenable the text is to normalization via semantic recognition.

Heuristic normalization. A large and multi-layered lexicon of cultural symbols needs to be assembled—preferably one that exhaustively enumerates all the symbols that are possible and mentioned in the cultural corpus. For example, a hierarchy of 22,000 interest and identity descriptors was assembled from numerous folksonomies such as Wikipedia, DMOZ, Allmusic, and Allrecipes, to normalize social network profiles. Identity descriptors were distinguished from interest descriptors because they are possible semantic mediators, which belong at a higher level of abstraction; this harkens to Barthes’s idea that cultural symbols form heterogeneous rather than homogeneous systems. Using these 22,000 descriptors and considering error associated with natural language processing, 68% of segmented profile tokens were recognized into this ontology, with 8% false positives, which we suggest is a high number considering that a person’s enumerated

interests are not constrained by any vocabulary in the profile elicitation form. To boost machine-learning performance, recognition rates should be maximized heuristically. One heuristic is to make pragmatic assumptions about ambiguous terms in the corpus—for example, popularity data associated with the hierarchy of 22,000 descriptors was used to disambiguate abridgements into their most likely referent, e.g. “bach -- > J.S. Bach.” Another heuristic is to consider gathering initial statistics on frequencies of words and phrases in the raw corpus to guide ontology gathering—for example, a percentage of the 1,000 assembled identity descriptors were templates of the form “___ lovers,” filled with what corpus statistics said were the most prevalent interest descriptors. The allowed handling of identities which were premised from prominent interests, e.g., “Star Trek lovers,” “chocolate lovers.” A third heuristic to improve the coverage of machine learning is to add metadata annotations to recognized descriptors, plus some uncertainty discount. So instead of submitting ‘The Unbearable Lightness of Being’ to machine learning, its uncertain metadata (‘Milan Kundera’, 0.5), (‘fiction’, 0.25), (‘philosophy’, 0.25) should also be submitted; thus, robustness is improved and more general regularities can be uncovered.

Machine learning. Similarities between descriptors can be calculated using a variety of metrics such as TF-IDF (Salton & McGill 1983), Pointwise Mutual Information (PMI) (Church & Hanks), Bayes, or any number of other metrics and variations (Lin 1998). One decision is—will similarity be learned between two descriptors, or between multiple descriptors? PMI is a symmetric measure of similarity, and we assumed that descriptors were symmetrically similar; however Tversky (1977) pointed out cases in which similarity is asymmetric (e.g. when two descriptors are not peers). Finally, in order to learn the connectedness between descriptors, examples of connectedness are needed, and this requires the identification of set identities in the corpus. In the case of social network profiles, we assumed that a person’s whole profile had a set identity, namely, its descriptors were bound by *aesthetic closure*. This assumption was consistent with cultural theories suggesting personal taste and style coherence (Bourdieu 1984; Solomon & Assael 1987; McCracken 1988). Some other assumptions could have been to assume closure within interest categories, or closure within tightly knit groups of people in the social network. Once set identity was chosen, all pairwise combinations resulting from that identity were collected as examples for machine learning via the PMI statistic. After an n by n similarity matrix is learned, a post-processing step should threshold away links whose affinities fall below the minimum score that could be considered significant. What is left is a semantic fabric that captures the cultural topology. To produce visualizable subsets of this dense network, further thresholding needs to be performed; the remaining links may be used to produce self-organizing maps and spring-loaded graphs, e.g. Figure 1-3.

To summarize this section, an approach to ‘culture mining’ was compared to related work in musical similarity, social network analysis, and text mining. Then, an “algorithm” for ‘culture mining’ was narrated—the decision space of assembling a cultural corpus was described, a strategy for robust normalization was given, and machine learning and its assumptions were discussed. In the next two sections, two technologies that support person modeling are discussed.

3.3 Commonsense reasoning

Commonsense reasoning via ConceptNet (Liu & Singh 2004b) supports multiple aspects of person modeling—its topic-gisting faculty was augmented to constitute RATE’S topic skimmer; its analogy-making faculty was used in model generalization; and it provides one of the technical foundations for another key technology: textual affect analysis. This section—1) introduces ConceptNet’s approach with respect to related work in commonsense reasoning systems; 2) examines the structure of commonsense knowledge in ConceptNet; and 3) explains ConceptNet’s mechanisms for topic-gisting and analogy-making.

§

Approach and related work. ConceptNet is a freely available commonsense knowledge base and natural-language-processing toolkit which supports many practical textual-reasoning tasks over real-world documents including topic-gisting, analogy-making, and other context-oriented inferences. The knowledge base is a semantic network consisting of over 1.6 million assertions of commonsense knowledge encompassing the spatial, physical, social, temporal, and psychological aspects of everyday life. ConceptNet is generated automatically from the 700,000 sentences¹⁶ of the Open Mind Common Sense Project (Singh *et al.* 2002) — a World Wide Web based collaboration with over 14,000 authors.

What is commonsense knowledge? Of the different sorts of semantic knowledge that are researched, arguably the most general and widely applicable kind is knowledge about the everyday world that is possessed by all people — what is widely called ‘commonsense knowledge’. While to the average person the term ‘common sense’ is regarded as synonymous with ‘good judgment’, to the AI community it is used in a technical sense to refer to the millions of basic facts and understandings possessed by most people. A lemon is sour. To open a door, you must usually first turn the doorknob. If you forget someone’s birthday, they may be unhappy with you.

¹⁶ Numbers given here were accurate as of summer, 2004

Commonsense knowledge, thus defined, spans a huge portion of human experience, encompassing knowledge about the spatial, physical, social, temporal, and psychological aspects of typical everyday life. Because it is assumed that every person possesses commonsense, such knowledge is typically omitted from social communications, such as text. A full understanding of any text then, requires a surprising amount of commonsense, which currently only people possess. The ConceptNet common sense reasoning toolkit was created as a way to afford machine readers some basic contextual reasoning faculties needed to achieve a humanesque understanding of text and narrative.

The size and scope of ConceptNet make it comparable to the two other large-scale semantic knowledge bases in the literature—Cyc (Lenat 1995) and WordNet (Fellbaum 1998). However, there are key differences. While WordNet is optimized for lexical categorization and word-similarity determination, and Cyc is optimized for formalized logical reasoning, ConceptNet is optimized for making practical context-based inferences over real-world texts. That it reasons simply and gracefully over text is perhaps owed to the fact that its knowledge representation is itself semi-structured English (a further discussion of reasoning in natural language can be found in (Liu & Singh 2004a)). ConceptNet is also unique from Cyc and WordNet for its dedication to contextual reasoning. Of the 1.6 million assertions in its knowledge base, approximately 1.25 million are dedicated to different sorts of generic conceptual connections called k-lines (a term introduced by Minsky (1980)). Contextual commonsense reasoning is highly applicable to textual information management because it allows a computer to broadly characterize texts along interesting dimensions such as topic and affect; it also allows a computer to understand novel or unknown concepts by employing structural analogies to situate them within what is already known. The former affordance is invoked to enable RATE processing's topic skimmer. The latter affordance is invoked in model generalization and simulation as a way to anticipate taste judgment responses to previously unseen input.

§

Structure of knowledge. ConceptNet has a natural language knowledge representation, which is seen as a boon to contextual reasoning. Unlike logical symbols, which have no *a priori* meaning, words are always situated in connotations and possible meanings. That words carry prior meanings, however, is not a bad thing at all, especially in the context game. By posing ConceptNet's nodes as semi-structured English phrases, it is possible to exploit lexical hierarchies like WordNet to make node-meanings flexible. For example, the ConceptNet nodes 'buy food' and 'purchase groceries' can be reconciled by recognizing that 'buy' and 'purchase' are in some sense synonymous, and that 'groceries' are an instance of 'food'.

The ConceptNet knowledge base is formed by the linking together of 1.6 million assertions (1.25 million of which are k-lines) into a semantic network of over 300,000 nodes. The present relational ontology consists of twenty relation-types, distributable into eight genres—k-lines (1.25 million assertions); things (52,000 assertions) e.g. is-a, property-of, part-of, made-of, defined-as; agents (104 assertions) e.g. capable-of; events (38,000 assertions) e.g. prerequisite-event-of, first-subevent-of, subevent-of, last-subevent-of; spatial (36,000 assertions) e.g. location-of; causal (17,000 assertions) e.g. effect-of, desirous-effect-of; functional (115,000 assertions) e.g. used-for, capable-of-receiving-action; and affective (34,000 assertions) e.g. motivation-of, desire-of.

Some examples of ConceptNet's nodes and relations:

```
(ConceptuallyRelatedTo 'bad breath' 'mint' 'f=4;i=0;')
(ThematicKLine 'wedding dress' 'veil' 'f=9;i=0;')
(IsA 'horse' 'mammal' 'f=17;i=3;')
(PropertyOf 'fire' 'dangerous' 'f=17;i=1;')
(PartOf 'butterfly' 'wing' 'f=5;i=1;')
(MadeOf 'bacon' 'pig' 'f=3;i=0;')
(CapableOf 'dentist' 'pull tooth' 'f=4;i=0;')
(FirstSubeventOf 'start fire' 'light match' 'f=2;i=3;')
(SubeventOf 'play sport' 'score goal' 'f=2;i=0;')
(LocationOf 'army' 'in war' 'f=3;i=0;')
(EffectOf 'view video' 'entertainment' 'f=2;i=0;')
(DesirousEffectOf 'sweat' 'take shower' 'f=3;i=1;')
(UsedFor 'fireplace' 'burn wood' 'f=1;i=2;')
(CapableOfReceivingAction 'drink' 'serve' 'f=0;i=14;')
(MotivationOf 'play game' 'compete' 'f=3;i=0;')
(DesireOf 'person' 'not be depressed' 'f=2;i=0;')
```

Through graph-based inference techniques including structure-mapping (Gentner 1983) and spreading activation (Collins & Loftus 1975), ConceptNet can perform these semantic tasks—finding contextual neighborhoods, making structural analogies, extracting topics from text, and estimating the emotive valence of a text.

§

Mechanisms for person modeling. ConceptNet's topic-gisting and analogy-making faculties are applicable to person modeling, so their mechanisms are now reviewed.

Topic-gisting. ConceptNet's topic-gisting mechanism is built upon a more basic premise. With all of the complexities associated with the term 'context', we can begin at one very simple notion. Given a concept and no other biases, what other concepts are most relevant? The ConceptNet API provides a basic function for making this computation, called `get_context()`. For example when fed the input 'go to bed', `get_context()` returns the following contextual neighborhood of nodes, with relevance scores omitted:

```
sleep
rest
take off clothes
close eye
dream
go to sleep
brush tooth
have nightmare
be tire
have sex
take nap
snore
relax
insomnia
```

Technically speaking, the contextual neighborhood around a node is found by performing spreading activation radiating outward from that source node. The relatedness of any particular node is not simply a function of its link distance from the source, but also considers the number and strengths of all paths that connect the two nodes.

Topic-gisting via `guess_topic()` is a straightforward extension of the `get_context()` feature to accept the input of real-world documents. Its value to information retrieval and data mining is immediately evident. Using MontyLingua, a document is gisted into a sequence of verb-subject-object-object (VSOO) frames. Minor transformations are applied to each VSOO frame to massage concepts into a ConceptNet-compatible format. These concepts are heuristically assigned saliency weights based on lightweight syntactic cues, and their weighted contextual intersection is computed by `get_context()`.

The `get_context()` function used in this way serves as a naïve topic spotter. To improve performance it may be desirable to designate a subset of nodes to be more suitable as topics than others. For example, we might designate ‘wedding’ as a better topic than ‘buy food’ since ConceptNet has more knowledge about its subevents (e.g. ‘walk down aisle’, ‘kiss bride’), and its parts (e.g. ‘bride’, ‘cake’, ‘reception’). Previous to the addition of this feature to ConceptNet, Eagle *et al.* (2003) used `get_context()` in similar fashion to extract topics from overheard conversations. Previous research in text summarization has used the link structure of WordNet in a similar manner to perform topic detection via “lexical chains” (Barzilay & Elhadad 1997; Silber & McCoy 2002). Hovy and Lin (1997) recognized the need for symbolic general world knowledge in topic detection, which is a key component of summarization. SUMMARIST (Hovy & Lin 1997) give the example that the presence of the words ‘gun’, ‘mask’, ‘money’, ‘caught’, and ‘stole’ together would indicate the topic of ‘robbery’. However, they reported that WordNet and dictionary resources were relationally too sparse for robust topic detection. ConceptNet excels at this type of natural language contextual task because it is relationally richer and contains practical rather than dictionary-like knowledge.

To build the topic skimmer for RATE, ConceptNet's `guess_topic()` function is augmented with realm-specific knowledge, such as cultural taste realm's hierarchy of 22,000 descriptors, and the DMOZ folksonomy that supports the attitudes realm. Augmentation adds domain-specific knowledge to topic-gisting, but preserves the overall mechanism—ontological nodes are mapped into ConceptNet's nodes, and metadata relations (e.g. `x subtopicOf y`) into graph edges.

Analogy-making. Like context manipulation, analogy-making is another fundamental cognitive task. For people, making analogies is critical to learning and creativity. It is a process of decomposing an idea into its constituent aspects and parts, and then seeking out the idea or situation in the target domain that shares a salient subset of those aspects and parts. Because AI is often in the business of dissecting ideas into representations like schemas and frames, analogy-making is quite prevalently used. It goes by pseudonyms like fuzzy matching, case-based reasoning, structure-mapping theory, and high-level perception. While in principle, a basic form of analogy is easy to compute, a large-scale, domain-general repository of concepts and their structural features, such as what ConceptNet contains, is required to produce commonsensical analogy-making to some approximation.

Gentner's structure-mapping theory of analogy emphasizes formal, shared syntactic relations between concepts. In contrast, Hofstadter and Mitchell's (1995) 'slipnets' project emphasizes semantic similarities and employs connectionist notions of conceptual distance and activation to make analogy more dynamic and cognitively plausible. Analogy in ConceptNet can be coaxed to resemble either structure-mapping or slipnets depending on whether weakly semantic relations (e.g. 'LocationOf', 'IsA') or strongly semantic relations (e.g. 'PropertyOf', 'MotivationOf') are emphasized in the analogy. Analogy in ConceptNet also has a slipnet-like connectionist property in that connections between nodes are heuristically weighted by the strength or certainty of a particular assertion.

Stated concisely, two ConceptNet nodes are analogous if their sets of back-edges (incoming edges) overlap. For example, since 'apple' and 'cherry' share the back-edges, [(PropertyOf `x` 'red'); (PropertyOf `x` 'sweet'); (IsA `x` 'fruit')], they are in a sense, analogous concepts. Of course, it may not be aesthetically satisfying to consider such closely related things analogous (perhaps their shared membership in the set, fruit, disqualifies them aesthetically), our simple discussion will not indulge such considerations here. As with the `get_context()` feature, it may also be useful to apply realm-filtering to dimensionally bias the `get_analogous_concepts()` feature. For example, preferring to variously emphasize functional similarity versus affective similarity versus attribute similarity, certain relation-types can be weighed more heavily than others.

Analogy-making supports model generalization in the realm of attitudes. For example, suppose an individual's attitude about 'macrome' is known, but her attitude about 'crafts' is not known. Structural analogy allows the precedented attitude about 'macrome' to be relaxed unto the unprecedented attitude about 'crafts'. In other words, the meanings of novel symbols are resolved within the known semantic space via structural analogy. The effect of structural analogy is to expand the semantic reach of a person's attitudes model. As discussed earlier, one pitfall to this approach is that attitudes inferred by structural analogy are sometimes invalid when they take place in semantic areas which lack aesthetic consistency—for example, although 'dogs' and 'cats' are structural analogs along many taxonomic dimensions, such as both being 'pets' and 'animals', they are aesthetically quite inconsistent. According to empirical data in WWTT, 'dogs' and 'cats' tended to form an aesthetic opposition—dog lovers tend toward a distaste for cats, and vice versa. Pet preference seemed particularly to be a heated and ideological space. One "hack" to avoid these problems was to modify `get_analogous_concepts()` to weight the aesthetic and affective dimensions of commonsense knowledge (e.g. motivation-of, desire-of, property-of relations) more heavily than the formal, taxonomic dimensions (e.g. defined-as, is-a, part-of relations) of knowledge. The hypothesis is that affective dimensions of knowledge tend to be more *aesthetic-preserving* than taxonomic dimensions.

In summary, commonsense reasoning via the ConceptNet toolkit and API is a key technology that constitutes a topic skimmer for RATE, and supports model generalization with analogy-making. A final contribution of commonsense reasoning to person modeling is that it provides one of the technical foundations for RATE's textual affect skimmer. The technology of textual affect analysis is addressed in the next section.

3.4 Textual affect analysis

Textual affect analysis is a key technology that constitutes the affect skimmer in reading for affective themes. This section 1) presents related work in sentiment and affect analysis of text, and situates our approach; 2) presents the PAD representation of affect that is used in our approach; and 3) describes a comprehensive mechanism for textual affect analysis that considers surface sentiment, lexical sentiment, and "deep" sentiment.

§

Related work. Related work in textual sentiment and affect analysis is presented along various themes, below.

Hand-crafted lexicons. Elliott's (1992) Affective Reasoner was simulation-world program that used mood keyword detection and hand-crafted heuristics for modeling textual affect. Subasic & Huettner (2001) tagged words with multiple affective attributes such as 'intensity' and 'centrality' in order to create fuzzy thesaurus for characterizing documents. Das & Chen (2001) used a manual lexicon, along with other methods, to analyze postings on investment bulletin boards for correlations between positivity and stock price. Numerous databases of lexical affect are in existence. Roget's (1911) Thesaurus contains sentiment classes that were annotated and used in our approach. General Inquirer (Stone *et al.* 1966) pooled several previous sources such as the Lasswell dictionary into a large database of lexical annotations. Three further sources—Whissell's (1989) dictionary of affect, Pennebaker, Francis, & Booth's (2001) Linguistic Inquiry and Word Count (LIWC) database, and the ANEW database (Bradley & Lang 1999) are databases of lexical affect derived from psychological data such as focus group responses.

Semantic orientation. Hatzivassiloglou & Mckeown (1997) classified the polarity or semantic orientation of adjectives using corpus examples of conjunctions (adj and adj) and disjunctions (adj but adj) as similarity and dissimilarity sets. Turney & Littman (2003) measured the semantic orientation of words and phrases as their PMI affinity to positive and negative paradigm words. Altavista search engine's 'NEAR' operator was used to define the co-occurrence sets. Turney (2002) classified the sentiment of movie reviews using the average semantic orientation of adjectival noun phrases. More nuanced applications of semantic orientation dictionaries for characterizing the affect of larger textual passages (Grefenstette *et al.* 2004a; Polanyi & Zaenen 2004; Wilson, Wiebe & Hoffman 2005) modify a word's valence by looking for *contextual valence shifters* in its surrounding words, thus accounting for intensifiers and negations.

Subjectivity classification. Employing corpus annotation and probabilistic classification, Wiebe *et al.* (1999) found that the presence of adjectives was the significant indicator of subjectivity in text. Hatzivassiloglou & Wiebe (2000) reported that dynamic adjectives (e.g. 'kind' and 'careful') and gradable adjectives (i.e. accepts intensifiers) were better indicators of subjectivity than typical adjectives. Moving beyond adjectival features, Riloff & Wiebe (2003) bootstrapped subjectivity classification by learning idiomatic patterns such as "x drives y up the wall." More recent work examined the application of subjectivity classification for question answering (Cardie *et al.* 2003). A survey and summary of the corpus-annotation approach to subjectivity classification was given in (Wiebe *et al.* 2004).

Movie and product reviews. Pang, Lee & Vaithyanathan (2002) explored various machine learning techniques for binary classification of UseNet movie reviews, and found a best

performance of 82.9% using SVM with unigram features. Pang & Lee (2004) reported that minimum-cut information further improved movie review classification, suggesting that such an approach does a better job of capturing cross-sentence context. Dave, Lawrence & Penncock (2003) classified product reviews using SVM over unigram and bigram features.

Narrative affect. Liu, Lieberman & Selker's (2003) Emotus Ponens system sensed sentence-level affect of user's everyday text via six basic emotion categories (Ekman 1993) by measuring events' connotations against commonsense knowledge extraction from the Open Mind corpus. Liu, Selker & Lieberman (2003) combined that approach with lexical affect and smoothing to sequence the affect of Project Gutenberg stories. Alm, Roth & Sproat (2005) used supervised machine learning to classify sentences from fairy tales into a canon of emotions based on Ekman's. Mishne (2005) compared semantic orientation and SVM approaches for binary classification of mood in blog entries and blogs. Mihalcea & Liu (2006) used top terms learned from happy and sad blog posts to bootstrap a linguistic ethnography of thematic differences between bloggers' happy and sad states—time of day, day of week, socialness, activity, and human-centeredness.

Our approach. Reading for affective themes demands textual affect analysis on several scales—lexical affect, sentence-level affect, and affect of larger textual passages. RATE also requires more granular descriptions of textual affect than is offered by binary classification, e.g. positive-negative, subjective-objective. Thus we used a comprehensive textual affect analysis approach that integrates surface, lexical, and event-level analysis. Surface sentiment is captured via mood keywords; this can be compared with (Elliott 1992). Lexical sentiment is captured via ANEW (Bradley & Lang 1999); incorporation of lexical affect is comparable to some hand-crafted approaches (Subasic & Huettner 2001; Das & Chen 2001) and to classifications based on average semantic orientation (Turney 2002). Finally, “deep” sentiment is sensed using the commonsense connotation method (Liu, Lieberman & Selker 2003).

§

PAD model of affect. In choosing a representation of affect, a desirable model would be generic enough to be attributable to both persons and objects, and sensitive enough so that even nuanced affects, could be captured elegantly. Discrete ontological models such as Paul Ekman's (1993) basic emotion ontology of Happy, Sad, Angry, Fearful, Disgusted, Surprised, derived from the study of universal facial expressions, can be problematic for computation because what they capture are discrete and classifiable emotional states. To unify them, Liu, Selker & Lieberman (2003) tried learning Markov models to measure their fungibility. A computational desideratum is to avoid the discreteness problem by choosing a

representation that provides a *continuous* account of affect, including trans-states that are subtle or warrant no linguistic label.

The dimensional PAD model of affect proposed by Mehrabian (1995), specifies three nearly orthogonal dimensions of Pleasure (vs. Displeasure), Arousal (vs. Nonarousal), and Dominance (vs. Submissiveness). In this thesis, the 3-tuple notation is used, e.g. (-, +, +), or more granularly, e.g. (P0.25 A0.5 D0.2). P0.0 is displeasure, P1.0 is pleasure; A0.0 is nonarousal, A1.0 is arousal; and D0.0 is submissiveness, D1.0 is dominance. Because PAD is dimensional, the distances between affects is easily computable as a Cartesian distance. This affords the simplicity of finding prevailing moods from a cloud of data points by simply calculating the first-order moment of the points in the cloud. Rubin, Stanton & Liddy (2004) verified high inter-rater agreement on classification using a similar unification model based on valence-arousal-engagement (Watson & Tellegen 1985), and suggested its suitability for NLP.

The PAD representation is a unification model into which most of the other models of affect can be mapped. For example, (P0.25 A0.5 D0.25) might correspond to sadness, while (P0.25 A0.7 D0.8) would correspond to anger. One exception to PAD's representational capability is that it does not represent directionality of affect. For example, "resentment" is an inwardly kept affect whose corresponding outwardly directed affect is "anger." In cases outside of the perception realm, directionality is not necessary to compute. In the perception realm, textual affect analysis accounts for directionality by assessing affect with respect to the *transaction* between ego (i.e. the writer) and alters (i.e. all other persons and textual entities). In each transaction, affect is either incoming (from alters into ego), or outgoing (from ego to alters), or collects about ego or alters. Herein lies an adaptation of the PAD representation to allow for the directionality of affect.

§

Mechanism. A comprehensive textual affect analysis system was built as a combination of a surface sentiment analyzer to recognize explicit mood keywords, a lexical sentiment analyzer to recognize affect information in non-mood keywords, and a deep sentiment analyzer to recognize the affective connotations of everyday events, such as 'getting fired'. Each analyzer skims the same textual passage and outputs a PAD assessment. The outputs are merged into an overall PAD assessment by taking the linear combination of the scores, each term weighted by manually determined coefficients standing for the efficacy of each analyzer. The textual affect analysis system scores a textual passage at different granularities: concept-granularity, sentence-granularity, paragraph-granularity, and document-granularity, to accommodate the varying needs of different RATE readers. Each analyzer is described below.

Mechanism for surface sentiment. Classifying text by spotting overtly emotional mood keywords like ‘distressed’, ‘enraged’, and ‘sad’ is an accessible naïve solution to affect sensing. Clark Elliott’s (1992) Affective Reasoner invoked mood keyword detection and hand-crafted heuristics for assessing textual affect. Considering that affect emanates from both surface language and deep meaning, sensing surface sentiment, though a naïve solution, is nonetheless an important part of a complete solution. Two dictionaries of sentiment words and their emotive valences were created for the surface sentiment mechanism. First, Peter Roget’s (1911) lexical sentiment classification system, taken from his 1911 English Thesaurus, features a 10,000 word affective lexicon, grouping words under 180 affective headwords, which can be thought of as very fine-grained and well nuanced affect classes. Each of the 180 affective headwords were manually assessed with their PAD implication, plus an uncertainty score (as some headwords were less informative than others). In addition, an affective lexical inventory produced by Ortony, Clore and Foss (1987) was manually assigned PAD scores. Based on the combination of Roget’s sentiment classes and Ortony, Clore, and Foss’s affective lexicon, a simple scoring mechanism skims a text’s normalized tokens, and assigns the whole textual passage a PAD score as the linear combination of its token’s PAD values, weighted by each token’s certainty score.

Mechanism for lexical sentiment. There is affective information contained even within words and patterns of words that are not mood keywords. For example, ‘homework’, ‘recess’, ‘lover’, ‘crime’ all imply some distinct affect when a commonsense interpretation is applied, but these words are likely missed by surface sentiment dictionaries such as Ortony *et al.*’s and Roget’s. At the very least, it should be acknowledged that, if considering the hypothetical set of all invocations of all words as lexicographers often do, most words would have a probabilistic leaning toward some affect or other. For example, the word ‘accident’ might be assigned a 0.75 leaning toward fear. One approach to gathering lexical affect statistical estimation—‘semantic orientation’ and ‘contextual polarity’—of adjectives, words in general, phrases, and of documents (Turney & Littman 2003; Wilson, Wiebe & Hoffman 2005). Another approach to gathering lexical affect has been via psychological experiments. Whissell’s (1989) dictionary of affect, Pennebaker, Francis, & Booth’s (2001) Linguistic Inquiry and Word Count (LIWC) computer program, and a corpus of psychologically normalized affect words called ANEW (Bradley & Lang 1999) are some examples. The third resource, ANEW, is used here for assessing lexical sentiment because it operates in the PAD dimensions needed for our mechanism. Each word in the ANEW list is assigned a PAD score. The lexical sentiment mechanism thus assesses the affect of a textual passage as the linear combination of affects of all ANEW words present in the passage, with each word weighted by the logarithm of its frequency in the passage, so as to prevent single errant words from skewing the sensed affect. To handle directionality of affect in the perception realm, the ANEW lexical sentiment dictionary is also applied to

appraise sentences by viewing each sentence as a transaction between the sentence's subject and direct object. For example, the sentence "John hit the baby," can be abstracted to this transaction—"John does something-negative to something-positive." This mechanism attempts to capture the affective dimension of how certain persons relate to other persons and entities.

Mechanism for "deep" sentiment. While the surface sentiment analyzer captures surface affect, such as the negative affect in the utterance, "I had a terrible day," it does not consider the deep semantics being communicated, i.e. the event and its common sense entailments; thusly, the complex affect in the utterance "I got fired today," whose affect is more subtextual than explicit, would be missed. The lexical sentiment analyzer can sense affect that is partially on the surface, partially submerged, but even it would miss more subtle entailments, for example, of the utterance, "my wife left me, she took the kids, and the dog." A commonsense knowledge based approach is needed to interpret events against the backdrops of their everyday connotations. Emotus Ponens (Liu, Lieberman & Selker 2003), a textual affect sensing system related to ConceptNet, is delegated the deep sentiment analysis task. Emotus Ponens parses a text's sentences into events and then evaluates the affective connotations of those events using a ConceptNet-like semantic network. For example, "I got fired today" connotes fear, anger and sadness, because getting fired often happens in a recession (-), getting fired is often the consequence of incompetence (-), and the effect of getting fired is not being able to buy things (-). By examining the semantic entailments of an event like getting fired, the valence of its deep affect is estimated.

In summary, this chapter presented the key techniques and technologies that enable person modeling from everyday texts. Reading for affective themes (RATE) was presented as an information-extraction approach to reading—constituted by a topic skimmer and a textual affect skimmer, and driven by reading schemas for each of the five realm. Culture mining was presented as an approach to inferring cultural topology from text corpora via natural language normalization and item-item similarity measurement. Commonsense reasoning via ConceptNet enabled topic skimming in RATE, partially supported affect skimming in RATE, and enabled analogy-making in model generalization. A comprehensive mechanism for textual affect analysis was presented, leveraging both word-level and event-level analyses. In the next chapter, implementations and evaluations of model acquisition systems for the five aesthetical realms are presented.

4 Acquisition systems: implementations & evaluations

This chapter presents implementations details and evaluation results for five model acquisition systems, covering the realms of – cultural taste (4.1), taste for food (4.2), ways of perceiving (4.3), attitudes (4.4), and humor (4.5).

4.1 Cultural taste realm: ‘taste fabric’

A person’s cultural taste is modeled and acquired by the Taste Fabric (Liu & Maes 2005a; Liu, Maes & Davenport 2006) system. The space of cultural taste is modeled using the semantic fabric representation, as an interweaving of nodes representing consumerist interests (e.g. music, books, sports, foods, films), and subcultural identities (e.g. “Book Lovers,” “Fashionistas”). The space is semantically mediated by topological features including subcultural identity hubs, taste cliques, and taste neighborhoods. An individual’s situation on the fabric is modeled by text processing his social network profile or RATE processing his personal homepage, followed by semantic relaxation into an *ethos* formation. The formation of *ethos* is greatly influenced by proximal topological features. A corpus-driven cross-validation evaluation of the taste fabric and ethos formation confirms the positive impact of topological features on taste prediction, and yielded accuracies comparable to competing approaches such as collaborative filtering. The rest of this section 1) discusses how cultural taste space is mined; 2) describes an individual ethos and the effects of topological features; and 3) presents an evaluation of the acquisition system.

§

Mining social network profiles for cultural taste space. The acquisition of cultural taste space is implemented in approximately 3,000 lines of Python code. As depicted in Figure 4-1, the architecture for mining and weaving the taste fabric from social network profiles can be broken down into five steps: 1) acquiring the profiles from social networking sites, 2) segmenting them to produce a bag of descriptors, 3) mapping of natural language fragment descriptors into formal ontology, 4) learning the correlation matrix, and 5) discovering taste neighborhoods via morphological opening, and

TASTE FABRIC'S INSTANCE TYPES








category	types	ontology sources
 identities	subculture, __ lover, taste echelon	wikipedia's "list of subcultures", dmoz
 films	auteur, film title, film genre	imdb, dmoz
 books	author, title, genre	amazon, wikipedia, dmoz
 music	artist, album, song, genre/decade	allmusic, amazon, dmoz
 foods	dish name, ethnicity, ingredient, course	allrecipes, foodsubs
 sports	name, genre	dmoz, amazon
 television	show name, genre	tvguide's "showguide", dmoz

Figure 4-2. Taste Fabric's instance types and ontology sources
 Figure 4-1. Mining algorithm for cultural taste space

labeling the network topology. Discussion of the last step is located elsewhere (Liu, Maes & Davenport 2006).

Step #1—scraping. The present implementation of the taste fabric sources from a one-time crawl of two web-based social network sites, which took place over the course of six months in 2004, yielding 100,000 social network profiles. Approximately 80% of these profiles contained substantive content because about 20% of users elected to not make their profile details publicly visible to the webpage crawler. The anonymity of social network users is protected because the normalization process wipes away all traces of individual users, as well as their idiosyncratic speech. From the 100,000 seed profiles, only the text of the categorical descriptors (e.g. "music", "books," "passions/general interests") is kept. Two social networks were chosen rather than one, to compensate for the demographic and usability biases of each. One social network has its membership primarily in the United States, while the other has a fairly international membership. Both however, had nearly identical descriptor categories, and both sites elicited users to specify punctuation-delimited descriptors rather than sentence-based descriptors. One drawback is that there is by our estimates, 15% membership overlap between the two sources so these twice-profiled members may have disproportionately greater influence on the produced fabric.

Step #2—segmenting profiles. Profile texts are easily segmented based on their interest categories. The format of profiles is for texts

to be distributed across templated categories, e.g., passions/general interests, books, music, television shows, movies, sports, foods, “about myself.” Typically “about myself” is populated with free-form natural language text while natural language fragments populate the specific interest categories. For the passions/general interests category, text is likely to be less structured than for specific interest categories, but still more structured than “about myself.” For each profile and category, its particular style of delimitation is heuristically recognized, and then applied. Common delimitation strategies were: comma-separated, semicolon-separated, stylized character sequence-separated (e.g. “item 1 \../ item 2 \../ ...”), new-line separated, commas with trailing ‘and’, and so on. Considering a successful delimitation as a category broken down into three or more segments, approximately 90% of specific categories were successfully delimited, versus about 75% of general categories. “About myself” and unsegmentable categories were discarded.

Step #3—ontology-driven normalization. After segmentation, descriptors are normalized by mapping them into a formal ontology of identity and interest descriptors (Figure 4-2). Newly segmented profiles are represented as lists containing casually-stated natural language fragments referring to a variety of things. They refer variously to authorships like a book author, a musical artist, or a film auteur; to genres like “romance novels,” “hip-hop,” “comedies,” “French cuisine”; to titles like a book’s name, an album or song, a television show, the name of a sport, a type of food; or to any combination thereof, e.g. “Lynch’s Twin Peaks,” or “Romance like Danielle Steele.” To further complicate matters, sometimes only part of an author’s name or a title is given, e.g. “Bach,” “James,” “Miles,” “LOTR,” “The Matrix trilogy.” Then of course, the items appearing under the general interests categories can be quite literally anything.

Figure 4-2 presents the ontology of descriptor instance types for the present taste fabric. At the top-level of the ontology are six specific interest categories plus one general interest category (i.e., “identities”). Also, as shown, there are roughly 25 second-level ontological types. There are a total of 21,000 recognizable interest descriptors, and 1,000 recognizable identity descriptors, sourcing from ontologies either scraped or XML-inputted from The Open Directory Project (dmoz), the Internet Movie Database (imdb), TV Tome, TV Guide, Wikipedia, All Music Guide, AllRecipes, and The Cook’s Thesaurus. Figure 4-2 only lists the primary sources, and lists them in order of descending saliency. The diversity and specificity of types ensures the maximal recognition capability over the free-form natural language in the profiles.

The ontology of 1,000 identity descriptors required the most intensive effort to assemble together, as we wanted them to reflect the types of general interests talked about in our corpus of profiles; this ontology was taken from Wikipedia’s extensive list of subcultures, from The Open Directory Project’s hierarchy of subcultures and hobbies, and finished off with some hand editing.

Identity descriptors in the form “(blank) lovers” were generated, where blank was replaced with major genres in the rest of our ontology, e.g. “book lovers,” “country music lovers,” etc. Some profiles simply repeat a select subset of interest descriptors in the identity descriptors category, so having (blank) lovers would facilitate the system recognizing these examples. The mapping from the general interests category into the identity descriptors ontology is far more indirect a task than recognizing specific interests because the general interests category does not insinuate a particular ontology in its phrasing. Thus, to facilitate indirect mapping, each identity descriptor is annotated with a bag of keywords which were also mined out from Wikipedia and The Open Directory Project, so for example, the “Book Lover” identity descriptor is associated with, *inter alia*, “books,” “reading,” “novels,” and “literature.” A consequence of employing two parallel mechanisms for identity descriptors, i.e. cultures versus (blank) lovers, there is overlap in a few cases, such as “Book Lovers” and “Intellectuals” or “Indie Rock Music Lovers” (genre of music) and “Indie” (subculture). Most cases of overlap, however, are much more justified, as the cultural lexicon, just as natural language, cannot be flattened to a canon. Perhaps the most controversial design choice for the sake of bolstering recognition rates, was up-casting subordinate identity descriptors into their subordinate descriptors. For example, while “Rolling Stones” is not in the ontology of identity descriptors, we automatically generalize, or up-cast, it until it is recognized, or all generalizations are exhausted; so the case of “Rolling Stones” is up-cast into “Classic Rock Music Lovers.”

To assist in the normalization of interest descriptors, aliases for each interest descriptor were gathered, along with statistics on the popularity of certain items (most readily available in The Open Directory Project) which the system uses for disambiguation. For example, if the natural language fragment says simply “Bach,” the system can prefer the more popular interpretation of “JS Bach” over “CPE Bach.”

Once a profile has been normalized into the vocabulary of descriptors, they are relaxed semantically using a spreading activation strategy from the formal ontology, because more than simply being flat wordlists, the ontological instances are cross-annotated with each other to constitute a fabric of metadata. For example, a musical genre is associated with its list of artists, which in turn is associated with lists of albums, then of songs. A book implies its author, and a band implies its musical genre. Descriptors generated through metadata-association are included in the profile, but at a spreading discount of 0.5 (read: they only count half as much). This ensures that when an instance is recognized from free-form natural language, the recognition is situated in a larger semantic context, thus increasing the chances that the correlation algorithm will discover latent semantic connections.

In addition to popularity-driven disambiguation of, e.g. “Bach” into “JS Bach,” several other disambiguation strategies were applied. Levenshtein (1965/1966) edit distance is used to handle close misspellings such as letter deletions, consecutive key inversions, and qwerty keyboard near-miss dislocations, e.g. “Bahc” into “Bach.” Semantically empty words such as articles are allowed to be inserted or deleted for fuzzy matching, e.g. “Cardigans” into “The Cardigans” (band).

Using this crafted ontology of 21,000 interest descriptors and 1,000 identity descriptors, the heuristic normalization process successfully recognized 68% of all tokens across the 100,000 personal profiles, committing 8% false positives across a random checked sample of 1,000 mappings. This is a good result considering the difficulties of working with free text input, and enormous space of potential interests and identities.

Step #4—correlation. From the normalized profiles now each constituted by normalized identity and interest descriptors, correlation analysis using classic machine learning techniques reveals the latent semantic fabric of interests, which, operationally, means that the system should learn the overall numeric strength of the semantic relatedness of every pair of descriptors, across all profiles. Choosing to focus on the similarities between descriptors rather than user profiles reflects an item-based recommendation approach such as that taken by Sarwar *et al.* (2001). Technique-wise, the idea of analyzing a corpus of profiles to discover a stable network topology for the interrelatedness of interests is similar to how *latent semantic analysis* (Deerwester *et al.* 1990) is used to discover the interrelationships between words in the document classification problem. For the present task, the information-theoretic machine learning technique called *pointwise mutual information* (Church & Hanks 1990), or PMI, was chosen. For any two descriptors f_1 and f_2 , their PMI is given in equation (4.1).

$$PMI(f_1, f_2) = \log_2 \left(\frac{\Pr(f_1 f_2)}{\Pr(f_1) \Pr(f_2)} \right) \quad (4.1)$$

Looking at each normalized profile, the learning program judges each possible pair of descriptors in the profile as having a correlation, and updates that pair’s PMI. What results is a 22,000 x 22,000 matrix of PMIs, because there are 21,000 interest descriptors and 1,000 identity descriptors in the ontology. After filtering out descriptors which have a completely zeroed column of PMIs, and applying thresholds for minimum connection strength, we arrive at a 12,000 x 12,000 matrix (of the 12,000 descriptors, 600 are identity descriptors), and this is the raw interest fabric. This is too dense to be visualized as a semantic network, but less dense semantic networks can be created by applying higher thresholds for minimum connection strength, and this is the reason why clustering seem to

appear in the InterestMap (Liu & Maes 2005a) taste fabric visualization.

§

Ethos formation. An individual can be located atop the built taste fabric by text processing his social network profile according to the above described normalization procedure, or by processing his personal homepage according to the RATE technique presented in Chapter 3. Such a reading returns explicit textual traces—a scored list of interest descriptors and identity descriptors. A further refinement step—ethos formation—is necessary to transform textual traces into the final generalized model. To form the individual’s ethos, semantic relaxation is performed. The raw n by n correlation matrix is re-viewed as a classic spreading activation network (Collins & Loftus 1975). That is to say, activation spreads outward from all the nodes from the person’s explicit textual traces to all the connected nodes, then from all connected nodes to each of their connected nodes. The strength of the spread activation is proportional to the strength of the PMI along any edge in the graph. The energy of the spreading is also inhibited as the number of hops away from the origin grows, according to a per hop discount rate (i.e. 0.5). The resultant contextual neighborhood of nodes and their associated activation weights, constitutes the individual’s taste ethos.

The formation of the individual’s ethos is greatly influenced by semantically mediating features intrinsic to the topology of the taste fabric. Two of these features—subcultural identity hubs and taste cliques—are described here, while a third feature—taste neighborhoods—is presented elsewhere.

Subcultural identities as hubs. Far from being uniform, the raw fabric is lumpy. One reason is that identity hubs “pinch” the network. Identity hubs are *identity descriptor nodes*, which behave as “hubs” in the network, being more strongly related to more nodes than the typical *interest descriptor node*. They exist because the ontology of identity descriptors is smaller and less sparse than the ontology of interest descriptors; each identity descriptor occurs in the corpus on the average of 18 times more frequently than the typical interest descriptor. Because of this ratio, identity hubs serve an *indexical* function. They give organization to the forest of interests, allow interests to cluster around identities. The existence of identity hubs allows us to generalize the granular location of what we are in the fabric, to *where in general we are* and what identity hubs we are closest to. For example, it can be asked, what kinds of interests do “Dog Lovers” have? This type of information is represented explicitly by identity hubs.

Taste cliques as agents of cohesion. More than lumpy, the raw fabric is denser in some places than in others. This is due to the presence of *taste cliques*—a n -clique formation with strong internal

cohesion. While the identity descriptors are easy to articulate and can be expected to be given in the special interests category of the profile, taste cohesiveness is a fuzzier matter. For example, a person of a Western European aesthetic may fancy the band “Stereolab” and the philosopher “Jacques Derrida,” yet there is no convenient keyword articulation to express the affinity between these pairs. However, when the taste fabric is woven, cliques of interests seemingly governed by nothing other than taste clearly emerge on the network. One clique for example, seems to demonstrate a Latin aesthetic: “Manu Chao,” “Jorge Luis Borges,” “Tapas,” “Soccer,” “Bebel Gilberto,” “Samba Music.” Because the cohesion of a clique is strong, *taste cliques* tend to behave much like a singular identity hub, in its impact on network flow. Hence its influence on ethos formation is commensurate to that of an identity hub.

§

Evaluation. The accuracy of the generalized model of cultural taste was evaluated in a corpus-driven validation, and results are summarized here. The hypothesis of this evaluation was that if the taste fabric accurately reflects the space of cultural taste, then the interests and identities possessed by each individual atop the fabric will tend toward coherency, i.e. those items will be more proximal to each other in the taste fabric. The evaluation is posed as a cross-validation interest recommendation task. Three controls are introduced to assess the impact that identity hubs and taste cliques have on the quality of recommendations; and to measure the performance of the fully enabled system (FULL-SYSTEM) against a conventional recommendation approach.

In the first control (IDENTITY-OFF), identity descriptor nodes are simply removed from the network, and taste ethos is used. In the second control (IDENTITY-OFF-CLIQUE-WEAK), identity descriptor nodes are removed, and n -cliques¹⁷ where $n > 3$ are weakened¹⁸. The third control (PMI-TALLY) produces recommendations without spreading activation, rather it simply scores each descriptor based on their PMIs. PMI-TALLY emulates an item-item collaborative filtering approach, such as the one given in (Sarwar *et al.* 2001).

Five-fold cross validation was performed to determine the accuracy of the taste fabric in recommending interests, versus each of the three control systems. The corpus of 100,000 normalized and metadata-expanded profiles was randomly divided into five segments. One-by-one, each segment was held out as a test corpus and the other four used to train a taste fabric using PMI correlation analysis.

¹⁷ a qualifying clique edge is defined here as an edge whose strength is in the 80th percentile, or greater, of all edges

¹⁸ by discounting a random 50% subset of the clique’s edges by a Gaussian factor (0.5 μ , 0.2 σ).

	PMI-tally (avg $\pm \sigma$)	identity OFF, cliques WEAK	identity OFF	full system		
S1	0.72 \pm 0.16	0.77 \pm 0.14	0.80 \pm 0.15	0.84 \pm 0.10	1	PMI-tally (control)
S2	0.73 \pm 0.13	0.80 \pm 0.16	0.81 \pm 0.18	0.87 \pm 0.12	2	identity OFF, cliques WEAK
S3	0.73 \pm 0.15	0.77 \pm 0.15	0.80 \pm 0.13	0.88 \pm 0.10	3	identity OFF
S4	0.74 \pm 0.16	0.81 \pm 0.13	0.83 \pm 0.13	0.86 \pm 0.11	4	full system
S5	0.75 \pm 0.17	0.80 \pm 0.15	0.83 \pm 0.16	0.87 \pm 0.11		
avg	0.73	0.79	0.81	0.86		

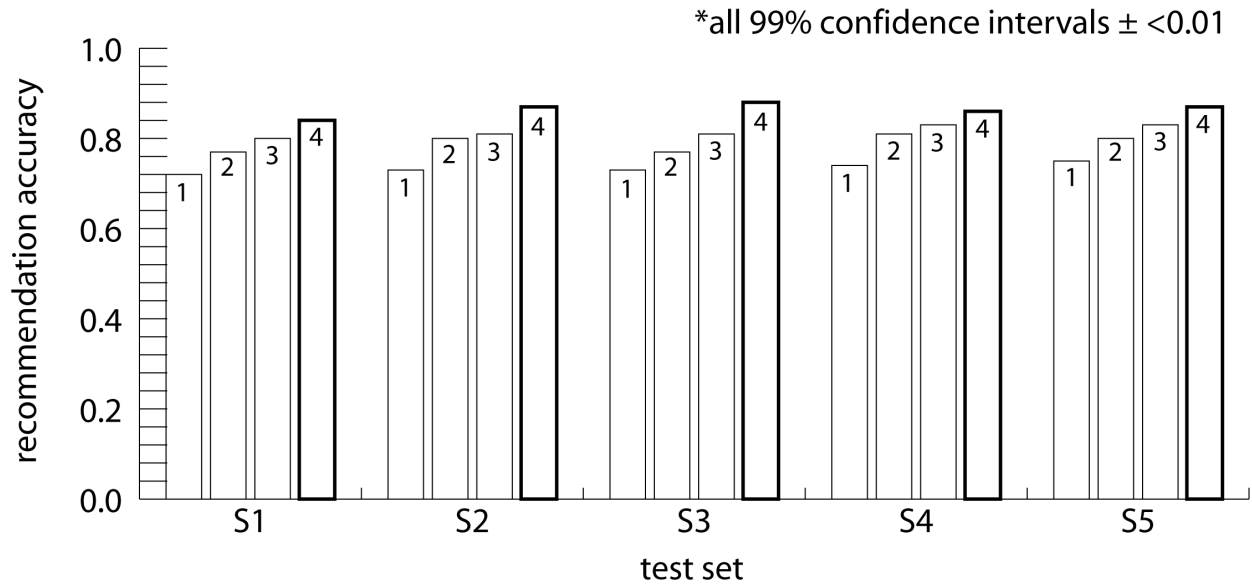


Figure 4-3. Evaluation of the generalized cultural taste model, and the effect of identity hubs and taste cliques, in a complete recommendation task

Within each normalized profile in the test corpus, a random half of the descriptors were used as the “situation set” and the remaining half as the “target set.” Each of the four test systems uses the situation set to compute a *complete recommendation* – a rank-ordered list of all interest descriptors; to test the success of this recommendation, we calculate, for each interest descriptor in the target set, its percentile ranking within the complete recommendation list. As shown in (4.2), the overall accuracy of a complete recommendation, $a(CR)$, is the arithmetic mean of the percentile ranks generated for each of the k interest descriptors of the target set, t_i .

$$a(CR) = \frac{1}{k} \sum_{i=1}^k \text{percentile}(t_i, CR) \quad (4.2)$$

The accuracy of a recommendation is scored on a sliding scale, rather than requiring that descriptors of the target set be guessed exactly within n tries, because the size of the target set is so small with respect to the space of possible guesses that accuracies will be too low and standard errors too high for a good performance

assessment. For FULL-SYSTEM, IDENTITY-OFF, and IDENTITY-OFF-CLIQUE-WEAK, the spreading activation discount was set to 0.75. The results of five-fold cross validation are reported in Figure 4-3.

The results demonstrate that on average, FULL-SYSTEM recommended with an accuracy of 0.86. In IDENTITY-OFF, removing identity descriptors from the network not only reduced the accuracy to 0.81, but also increased standard deviation and standard error of mean by 38%. In IDENTITY-OFF-CLIQUE-WEAK, removing identity descriptors and weakening cliques further deteriorated accuracy slightly, though less notably, to 0.79. All three systems performed favorably against the PMI-TALLY baseline accuracy of 0.73. Because activation flows more easily and frequently through identity hubs and taste cliques than through the typical interest descriptor node, the organizational properties of identity and taste yield proportionally greater influence on the recommendation process.

This analysis of cultural taste in terms of consumer interests quite resembles the data that Bourdieu (1984) had worked with in his statistical analysis of cultural taste questionnaires in 1960s French society. The notable difference between the approaches is that while Bourdieu sought to implicate capitals as the fundamental dimensions of taste-space, no basic dimensions are assumed or discovered by Taste Fabric. Instead, taste distance calculation relies on the tightly interwoven nature of consumerist interests. The density of interconnections is made possible by the scale of mining 100,000 profiles in which 20-50 keywords are extracted from each. However, the topology of the fabric is far from smooth and uniform—local structuration is supplied by semantic mediators—such as special nodes representing subcultural identities (*cf.* Figure 1-3); cliques of tightly knit interests; and larger regions of interests called ‘taste neighborhoods’.

4.2 food realm: ‘synesthetic cookbook’

The Synesthetic Cookbook implements an acquisition system for persons’ tastes for food. The food realm is represented as a semantic fabric of food— interweaving recipes, ingredients, flavors, moods, and cooking procedures. Cuisines and basic flavors act as semantic mediators on the ‘food fabric’—serving a connector/gateway function. Food fabric’s connectedness is acquired by mining a cultural corpus of texts about foods for semantic correlations—including a database of 60,000 recipes, encyclopedias about foods, and the Thought for Food corpus of food common sense. An individual’s tastebud can be acquired by a RATE processing of any of their everyday texts, or by a history of their preferences for dishes in the Synesthetic Cookbook application. Similar to a cultural taste model, a tastebud is represented here as an ethos formation; different from a cultural taste model, a tastebud’s textual traces

distinguishes between food preferences and food necessities, and accepts both positive (e.g. “spicy”) and negative (e.g. “not spicy”) leanings. The rest of this section 1) gives a brief history of the Synesthetic Cookbook; 2) describes how the food fabric is mined; and 3) describes how a tastebud is represented and acquired.

§

Background. The Synesthetic Cookbook is rooted in three years of research on computer representations of food, done in collaboration with Barbara Wheaton, food historian and honorary curator of the food collection at Harvard’s Schlesinger Library. Its earlier version, HyperRecipes, was exhibited at several open houses, where hundreds of people have played with the system; many technologists, designers, mothers, and food critics have experimented with the system and contributed to its evolving design. HyperRecipes sought to expose the cultural context and taste-practices which underlied recipes, and to challenge the notion of authenticity by using statistical mining approaches to characterize authenticity as formed by a predictable set of practices. The current version of the Synesthetic Cookbook is based on a food fabric created from the mining of numerous cultural corpora.

§

Cultural corpora. Several cultural corpora are spliced together to create the food fabric. A database of 60,000 recipes was collected from the web, in the standardized Meal Master format. This corpus has a great variety of recipes from different sources, but there is nonetheless a clear bias toward American cooking and baked goods. Barbara Wheaton’s database of food terms, USDA’s online nutrition encyclopedia, and an online website encyclopedia of ingredients and cuisines were collected. The Thought for Food corpus was also included—containing 21,000 sensorial facts about 4,200 ingredients (e.g. “lemons are sour”), and 1,300 facts about 400 procedures (e.g. “whipping something makes it fluffy”).

Mining. Culture mining, as described in Chapter 3, was applied to the total cultural corpus in order to mine out interrelations between food items and food ideas. An ontology of foods, cuisines, flavors, and moods was first compiled from these corpora by identifying the popular and oft-used adjectives and noun phrases. This ontology supports the natural language normalization phase of culture mining. The textual passages in the recipe, and in the online encyclopedia of ingredients and cuisines, were treated to topic extraction and textual affect sensing in order to transform those texts into lists of significant food items. Recipes were treated both as a single food item (i.e., the ‘dish’), and as a bag of co-occurring ingredients, cooking procedures. Finally, the same pointwise mutual information learning algorithm already applied in mining the

Table 4.1. Excerpt from Synesthetic Cookbook’s lists of ingredients and descriptive keywords

some top ingredients			some descriptive keywords		
worcestershire	oregano	cumin	acidic	healthy	pineapple
cinnamon	cayenne pepper	bay leaves	aha	heartly	pink
paprika	nutmeg	heavy cream	alcohol	heavy	piquant
cornstarch	mustard	wine	alcoholic	herb	popeye
soy sauce	ground beef	mayonnaise	american	herbal	pork
chili powder	tomato paste	sesame oil	appealing	holiday	potassium
vinegar	chicken broth	tomato sauce	apple	holidays	potato
vanilla extract	ketchup	tomatoes	apples	homecooked	pregnant
honey	margarine	lemon	arabia	homey	primitive
sour cream	thyme	tabasco sauce	arid	hot	purple
bay leaf	parsley	yeast	aromatic	indian	raspberry
buttermilk	carrots	garlic salt	art	iron	red

cultural taste fabric was used to learn affinities between food items, and this constitutes the food fabric.

Food items and mediators. The food fabric weaves together 5,000 ingredient keywords (e.g. “chicken”, “Tabasco sauce”), 1,000 descriptive keywords including flavors, cuisines, and moods (e.g. “spicy,” “chewy,” “silky”, “colorful”), and 400 nutrient keywords (e.g. “vitamin a”, “manganese”), as well as all the negations (e.g. “no chicken,” “not spicy”). Table 4.1 excerpts from the list of ingredients and the list of descriptive keywords. Cuisines (e.g. chinese, indian) and basic flavors (e.g. sweet, spicy, salty) acts are semantic mediators because while they are a limited set of terms, they have a wide coverage over the corpus of recipes and foods, and thus, they are correlated to other food items in greater numbers and proportion than the typical food item. In addition to cuisines proper, e.g. ‘indian’, there are also cuisines improper, e.g. ‘indianish’. Proper cuisines suffixed with ‘-ish’ or ‘-y’ denote cuisine resemblances. For example, during machine learning, a recipe for ‘Jambalaya’ –which calls for many spices such as ‘bay leaves’, and involves the cooking procedure of ‘reduction’ –is labeled as similar to an indian curry dish. Thus, it is labeled with the improper cuisine keywords ‘indianish’ and ‘indiany’, and those words then become correlated with the jambalaya dish and all its ingredients and procedures. Cuisine resemblance embodies a post-structuralist aesthetic, and challenges the tradition notion of ‘authenticity’ with respect to food stuffs.

§

Tastebud. An individual’s situation on the food fabric can be gotten by applying RATE processing to a personal homepage, or weblog diary. In this case, a synonym dictionary to bolster coverage first expands the ontology of food items, and the RATE reader is set to recognize words and topics belonging to this semantically expanded set of terms. Since much of the food vocabulary is also applicable to

discourses not concerning food (e.g. “a spicy dish” versus “a spicy night out”), an interesting and indirect set of textual traces can be discovered by the RATE reader. For more focused textual traces, a tastebud acquisition interface allows tastebuds to be programmed with textual traces directly, or a history of a user’s interactions with the Synesthetic Cookbook application can be recorded, and that used can seed a tastebud model. This history stores a list of all descriptive keywords the user has searched for, and all the recipes the user has chosen to view.

To transform what the textual traces captured by RATE into a generalized tastebud model, an ethos is formed from textual traces by spreading activation outward to define a contextual neighborhood in the food fabric. A design choice warranted by observations about the nature of people’s food cravings led to the segregation of items in the tastebud acquisition interface into two classes—food preferences and food necessities. Preferences are weak demands, while food necessities are requirements—for example, ingredients which must be utilized, ingredients which are not available, and allergies. In addition, both positive (e.g. “spicy”) and negative (e.g. “not spicy”) leanings are accepted. The ‘no’ and ‘not’ operators can negate any food item or descriptive keyword. Food preferences are used to form the ethos, while food necessities acts as a filter which disqualifies nodes from the food ethos. Negative preferences are spread as negative activations.

4.3 Perception realm: ‘escada’

Experimental System for Character Affect Dynamics Analysis (ESCADA) (Liu 2005) performs a RATE processing of an individual’s everyday texts, such as a weblog diary, in order to produce a perception model. Of all the realms’ models, this is the most experimental, and its results are the most tenuous. Recall that the space of perception is framed by Jung’s four fundamental psychological functions—sense, intuit, think, feel—taken as orthogonal axes of perception space. Also recall that in the reading schema for perception realm, there are perception-lexemes and perception-classemes. Perception-lexemes are instances of affective communication between the writer, called ‘ego’, and other textual entities, called ‘alters’. For example, the utterance “I laughed at John so hard” is abstracted into an affective transaction—a passing of the valence (-,+ ,+) associated with the phrasal verb “laugh at” from ego into ‘alters’. This instance is called a perception-lexeme, since affective transaction is hypothesized as a basic unit of perceptual disposition. ESCADA transforms a weblog diary into a bag of perception-lexemes using textual affect sensing, especially invoking the lexical sentiment analyzer. The machine learner, Boostexter (Schapire & Singer 2000), was run over a large corpus of annotated blogs, and a mapping was learned between perception-lexemes and perception-classemes. The rest of this section 1) overviews the Character Affect Dynamics theory; 2) discusses how lexeme-to-

classeme mappings were learned from a corpus of weblog diaries; and 3) presents an evaluation of the RATE reader's performance.

§

CAD theory. Character Affect Dynamics is a theory which posits that latent patterns of affective communication in a narrative betray the time-stable perceptual dispositions of the characters of the narrative. In the case of a first-person narrative such as an everyday text, it is the writer's affective engagement with other persons and things that is of interest. CAD theory has a cognitive linguistic basis. Talmy's (1998a) force dynamics theory models linguistic utterances as forces exchanged between agents and objects (e.g. 'the door could not be opened'). Force dynamics theory in fact proposes applications to social interactions and to internal psychodynamics. Following force dynamics, CAD theory examines the affective forces present between characters in a narrative. CAD suggests that much of this analysis can be modeled as the passing of affective tokens between textual entities—since textual entities approximate agents and objects.

For example, the utterance "I stole Mary's ice cream," can be interpreted as affective token pushing: "I [negative-act] [positive-object]. In this transaction, the writer does something bad to something valued by Mary. As a result, it could be concluded that the writer has negative affect, that he is aggressive, that Mary's 'ice cream' henceforth bears the traumatic connotation of something negative (due to emotion's contagion). Furthermore, if the next utterance is "Mary resented me," there is confirmation that the previous act was negative, and the fact that Mary's retaliatory nature is disclosed. "Resent" is known in ESCADA's lexical knowledge base as a passive-act, thusly, Mary is passive-aggressive.

The above scenario suggests some advanced capabilities of CAD-enabled story understanding. The ESCADA (Experimental System for Character Affect Dynamics Analysis) system implements an extend version of this scenario. Deep story understanding, though, is understandably brittle, but CAD theory's claim does not rely on deep understanding, only on shallow reading—i.e., emergent patterns of affect token passing between characters can predict their perceptual dispositions—tendency toward thinking or feeling, sensing or intuiting. To test this claim, an initial set of these patterns were implemented in ESCADA:

- ❖ EGO-PAD (main character's PAD-level)
- ❖ ALTERS-PAD (other characters' PAD-level)
- ❖ INCOMING-PAD (PAD flowing from alters into ego)
- ❖ OUTGOING-PAD (PAD flowing out from ego into alters)
- ❖ MENTAL-ACTIVITY (quantity of invocations of mental hypotheticals, e.g. "I thought that")

- ❖ INTROVERSION-EXTRAVERSION-RATIO (ratio of passive acts e.g. ‘resent’ to active acts e.g. ‘murder’)

§

Learning lexeme-to-classeme mappings. RATE processing of an individual’s weblog diary results in the identification of many instances of affective communication from the text. These instances are compiled into these fourteen affective statistics by averaging over each blog entry:

- ❖ EGO-PLEASURE (scale: -1.0 to 1.0)
- ❖ EGO-AROUSAL (scale: -1.0 to 1.0)
- ❖ EGO-DOMINANCE (scale: -1.0 to 1.0)
- ❖ ALTERS-PLEASURE (scale: -1.0 to 1.0)
- ❖ ALTERS-AROUSAL (scale: -1.0 to 1.0)
- ❖ ALTERS-DOMINANCE (scale: -1.0 to 1.0)
- ❖ INCOMING-PLEASURE (scale: -1.0 to 1.0)
- ❖ INCOMING-AROUSAL (scale: -1.0 to 1.0)
- ❖ INCOMING-DOMINANCE (scale: -1.0 to 1.0)
- ❖ OUTGOING-PLEASURE (scale: -1.0 to 1.0)
- ❖ OUTGOING-AROUSAL (scale: -1.0 to 1.0)
- ❖ OUTGOING-DOMINANCE (scale: -1.0 to 1.0)
- ❖ MENTAL-ACTIVITY (scale: nonnegative integer)
- ❖ INTROVERSION-EXTRAVERSION-RATIO (scale: 0.0+)

A mapping must be learned from these statistics into these four perception-classes—thinking, feeling, intuiting, sensing. More accurately, thinking-feeling and intuiting-sensing are binary oppositions, so either pole from each opposition must be selected, but not both. To learn this mapping, machine learning is fed a corpus of weblog diaries already annotated with the correct classes. Whence such a corpus? Conveniently, the desired classes can be found in the Myers-Briggs Type Indicator (MBTI) (Briggs & Myers 1976) inventory of personality. Thinking-feeling and intuiting-sensing make up two of the four MBTI scales. MBTI is derived from Jungian’s psychological functions, and is widely used in pop cultural psychology tests such as Bloginality¹⁹; in fact, because it is possible to search for all blogs which feature their author’s Bloginality test result, hence our annotated corpus.

About the MBTI. MBTI has four scales: Extraversion-Introversion, Sensing-iNtuition, Thinking-Feeling, and Judging-Perceiving. The first three scales were found to be independent, while the fourth was found slightly co-dependent on SN with S predicting J (Myers & McCaulley 1985). By combining the four scales, MBTI allows for sixteen Jungian types, e.g. ENFP, ISTJ, ISFP, etc. This evaluation examines the performance of the four individual scales. While in the actual MBTI assessment these scales are continuously-valued, for simplicity this evaluation treats each scale as a dichotomy.

¹⁹ <http://bloginality.love-productions.com/>

Corpus. A corpus of roughly 3800 blogs was assembled, for which the MBTI of the blogger is known. “Known MBTI” is accepted as having met at least one of the following two conditions:

- ❖ Blogger has listed MBTI type in their profile, and not listed any other competing/conflicting MBTI types there as well)
- ❖ Blogger has featured in their blog a cut-and-paste entry stating the results of an online MBTI test they took, such as Bloginality (MBTI-clone), and not listed any other competing/conflicting MBTI test results as searchable in their blog.

From the 3800 blogs, 85,000 combined blog entries were mined, averaging 22 entries per blog. The average time spanned by the blog entries from each blog is 8 weeks.

Sanitizing. To further prepare the blog entries, noisy entries had to be identified and discarded. A common practice in blogging is the use of occasional canned entries or favorites lists. For example, a blogger may cut-and-paste the results of various online temperament tests and create a blog entry from that. Or, a blogger could fill in her responses to a ‘20-questions’ type of personality inventory and make a blog entry from that. Canned entries were identified using clone detection (similar language, similar graphics) across all blog entries. Entries with long numbered lists were also discarded. Finally, null entries and entries without the presence of at least the pronouns “I” or “me” were discarded, as these texts are not likely to be egocentric. Finally, the corpus was pruned such that equal numbers of blog entries were available for each of the sixteen MBTI personality types (as this would create equal proportions of E-I, S-N, F-T, J-P, as a necessary testing condition).

Generating MBTI classifier. After RATE processing proceeds and the fourteen affective statistics are computed for each blog, the statistics along with known MBTI-labels, are fed into a machine learning algorithm to learn optimal numerical weights on each of the 14 profile features. Not only is this an unbiased way to learn a heuristic MBTI classifier for blogs, it is also a way to uncover the relative importance and efficacies of our ESCADA statistics. Boostexter is the machine learning system used, configured for 200 rounds of boosting, and n-grams up to two. Using the produced classifier weblog diaries can be used to roughly locate their authors in the Jungian perception space, though not with excellent granularity. Two of the MBTI scales learned by the MBTI classifier, are not used to create the person’s perception model.

§

Evaluation method. This evaluation challenges ESCADA to read blogs and classify bloggers into their Jungian personality type, as given by the Myers-Briggs Type Indicator (MBTI). The subset of results which are interesting to perception modeling are those which pertain to just the thinking-feeling and intuiting-sensing scales of




Figure 4-4. Results of ten-fold cross validation showing log-level classification accuracies

MBTI. Notwithstanding, all four scales are presented here, for completion. To fairly simulate the efficacy of the ESCADA-derived classifier on unseen data, hold-one-out ten-fold cross-validation was used over the corpus of 3800 MBTI-annotated weblogs. The whole corpus was randomly divided into ten sections. Taking each section in turn as the testing set, the other nine sections served as the training set. Boostexter was again configured for 200 rounds of boosting, and n-grams up to two.

Bounds on performance. Performance on each of MBTI’s scales is bounded below by fair chance guessing (50%), and bounded from above by MBTI test-retest reliability statistics. Because there were equal numbers of blog entries of each MBTI in the corpus, a lower bound on classifier performance is 50%, achieved by a classifier which tosses a fair coin to decide on the value for each of the four-scales. To note, the distribution of the sixteen MBTI types in the overall population is quite uneven, and in our experience gathering the online corpus of MBTI blogs, typing was also very uneven.

A loose upper bound on performance is the MBTI four-to-five-week test-retest reliability statistics. This bound hints at the underlying (in)stability of the MBTI personality inventory, notwithstanding still the de facto popular psychology assessment of personality. Myers and McCaulley (1985) survey continuous score correlations from ten studies for the four-to-five-week test-retest interval. They found reliability coefficients of .77 to .93 for EI, .78 to .92 for SN, .56 to .91 for TF, and .63 to .89 for JP. Assuming roughly binomial distribution for these scores, we estimate cursory median reliabilities of EI .84, SN .85, TF .73, JP .78. Given that the average blog found in our corpus has entries covering a time-span of 8 weeks, we regard the four-to-

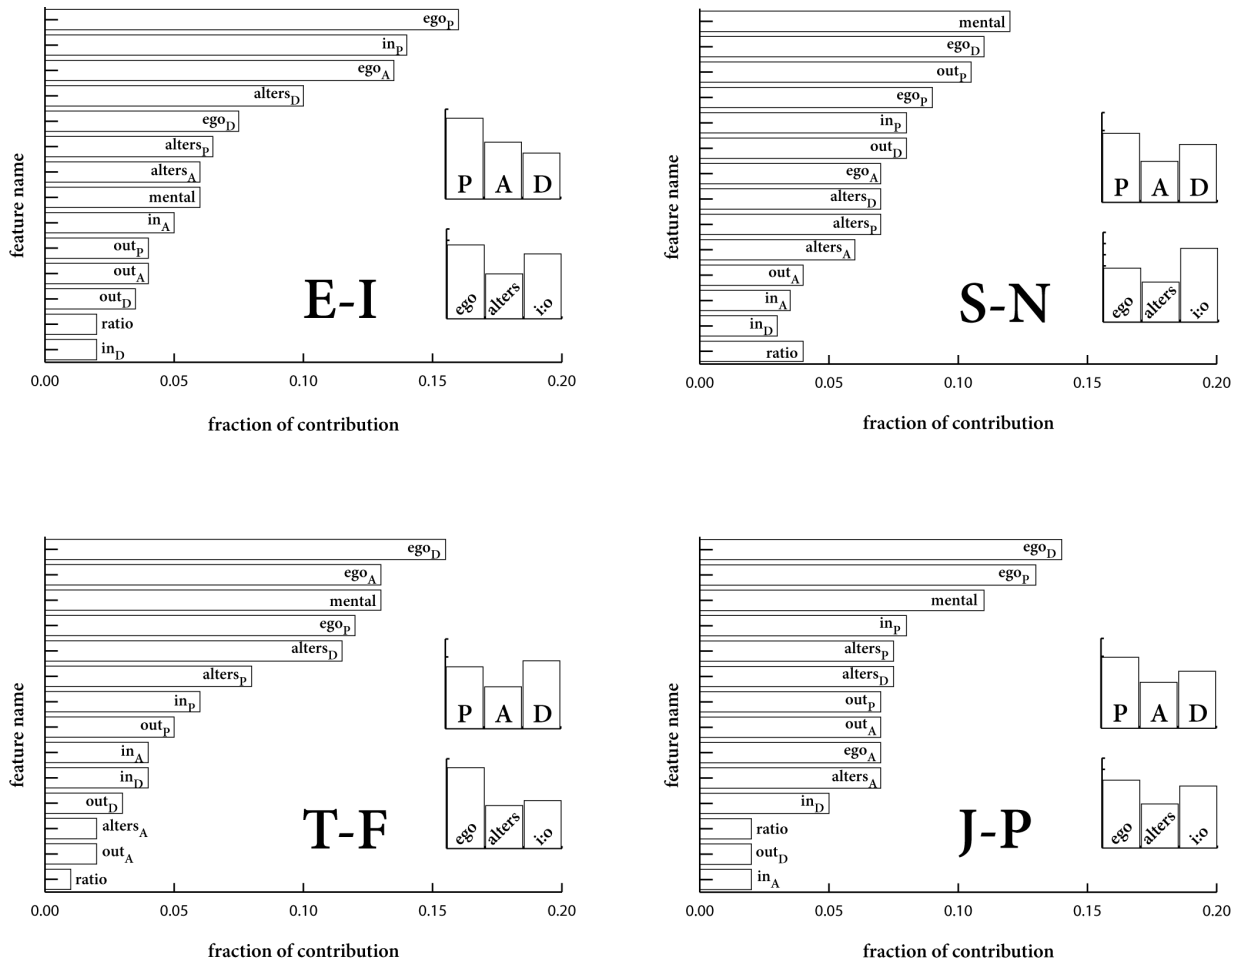


Figure 4-5. Learned feature weightings for single-scale classification.

five-week test-rest reliabilities as loose upper bounds on performance for each respective scale.

Results. Following hold-one-out ten-fold cross-validation, MBTI classifiers were trained for each of the ten training sets. The classification accuracies of these classifiers applied to their corresponding validation test sets are given in Figure 4-4. Average accuracies ranged between 0.58 and 0.67 for the four independent scales. Classification of E-I was most successful, at 0.67, while S-N was least successful, at 0.58. The scale classifiers demonstrated that they contain information by outperforming the lower bound of 0.50. On average, the classifiers underperformed their corresponding upper-bounds by margins of E-I 0.17, S-N 0.27, T-F 0.11, J-P 0.18. Under this context, T-F most closely approached optimal prediction, while classification of S-N was most ineffective using the ESCADA statistics over the blog corpus.

To ascertain the usefulness of the individual ESCADA statistics to each of the four MBTI scales, an analysis of Boostexter's outputted

.SHYP (strong hypothesis) files was undertaken. The .SHYP files contain the rules which constitute each classifier. For each scale, there were ten classifiers learned from the ten validation sets. The .SHYP file corresponding to each classifier was parsed, and the numerical weights and feature-names implicated in each of the rules were extracted. Based on the combined weights for each feature, and averaged over ten classifiers for each scale, the relative contribution of each feature was calculated. Results are given in Figure 4-5.

According to Figure 4-5, ego's affect was most important, followed by the mental activity index, then by alters' affect. Incoming and outgoing affects were more tenuous, while the introversion-extraversion statistic was not reliable for MBTI classification. For the E-I scale, pleasure and arousal of the ego, as well as pleasure flowing into the ego, were the more useful features. For the T-F scale, the ego-centric features, and in particular, the ego's dominance dimension were most useful. The S-N and J-P scales appraised usefulness in similar fashion and shared common top features, suggesting some mutual information between those scales. The aggregate of incoming-outgoing features was more useful than the ego features and alters features for the S-N scale, suggesting that Sensing bloggers and iNtuiting bloggers can be distinguished by their different affective postures toward alters. By contrast, the greatest utility of ego features in the T-F scale accords with the intuition that T-F can be appraised more solipsistically than the other three scales. The mental activity index—which measures the quantity of vocalizations of mental hypotheticals, e.g. "I thought that"—was a top-three useful feature in S-N, F-T, and J-P, but not in E-I. One could take this result to suggest, counter the intuition of some, that extraverted and introverted bloggers can hardly be distinguished by how they vocalize their thoughts and opinions. Or, this result could be owed to the nature and culture of blogging, which is arguably a revealing activity, and a venue for dramatic performance (Boyd 2004).

Of pertinence to perception modeling, results were mixed. Location within thinking-feeling most closely approached its upper bound, with an accuracy of 0.62 +/- 0.05. Location within sensing-intuiting fared more poorly with an accuracy of 0.58 +/- 0.05, far from the upper bound of 0.85, and only slightly better than guessing.

4.4 Attitudes realm: 'what would they think?'

What Would They Think? (WWTT) (Liu & Maes 2004) is an acquisition system for attitude modeling. To model a person's attitudes, a corpus of everyday texts is compiled—from commentary-rich research papers, instant message conversations, personal emails, and

weblog diaries. The corpus is fed to WWTT's RATE reader, which infers an attitude isotopy from the texts. An attitude isotopy consists of a set of topics extracted from the texts, each associated with a statistically average emotive valence, given as a PAD score. To infer the average attitude model of a culture, either an aggregate of individual textual corpora could be fed into the RATE reader, as was done for Xanga weblog communities, or a culturally representative text corpus could be fed as input, as was done in order to model political culture. Finally, to bolster the coverage of individuals' models, Minskian *imprimer* relations were identified between individuals, such that each person's model is bolstered by the models of their *imprimers*. The rest of this section 1) presents a detailed example of how the RATE reader constructs an isotopy; 2) discusses a use-example describing how WWTT was used to plot the attitudes of periodicals in the space of political culture; 3) discusses how Minskian *imprimers* enrich an individual's model; and 4) presents an evaluation of the quality of attitude capture with WWTT.

§

Attitude isotopy is captured by an affective reflexive memory.

WWTT implements attitude isotopy as an affective reflexive memory which stores lexemes, or *exposures*, as they are encountered during a skim of the text. For clarity, discussion of isotopy now briefly shifts into the vocabulary of memory. The idea that the reader has types of memories is owed to psychological theories of memory. Tulving (1983) describes both a 'reflexive memory' and a 'long-term episodic memory' (LTEM). He equates LTEM with "remembering" and reflexive memory with "knowing" and describes their functions as complementary. While long-term episodic memory deals in salient, one-time events and must generally be consciously recalled, reflexive memory is full of automatic, instant, almost instinctive associations. Both memories were implemented in WWTT, but long-term episodic memory was found to be far less useful in user studies of the system. Their discussion is thus left elsewhere (Liu & Maes 2004). Reflexive memories are formed through the conditioning of repeated *exposures* rather than one-time events. The conditioning process also acts as a noise filter against any incorrect textual affect classifications.

The affective reflexive memory is represented by a lookup-table. The lookup-keys are concepts which can be semantically recognized as a topic—such as a person, action, object, or activity. Associated with each key is a list of *exposures*, where each exposure describes a distinct instance of the concept appearing in the inputted everyday texts. An exposure, *E*, is represented by the triple: (date, affect valence score *V*, saliency *S*). At runtime, the affect valence score associated with a given conceptual cue can be computed using the formula given in Eq. (4.3)

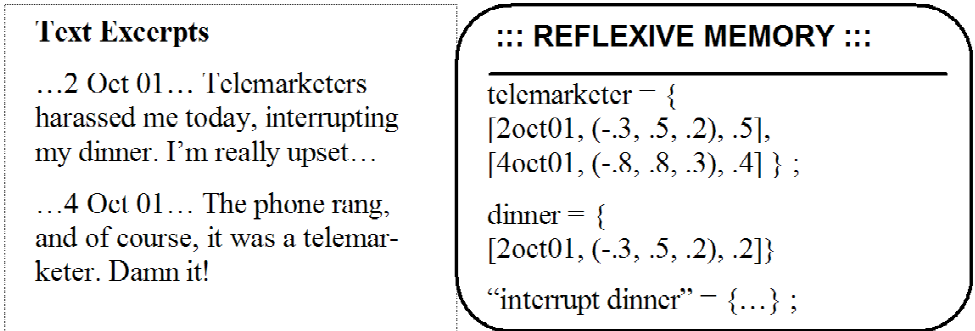


Figure 4-6. How reflexive memories get recorded from weblog excerpts

$$\frac{1}{n} \left[\log_b (\max(n, b)) \right] \left[\sum_{t=startdate}^{enddate} S(E_t) V(E_t) \right] \tag{4.3}$$

where n = the number of exposures of the concept; $b = 2$

This method of calculating the stable affect associated with a topic further nuances the reading schema for perception modeling that was given in Chapter 3. Instead of taking a naive average of the PAD valences, reinforcement is used to determine the stable PAD. Eq. (4.3) gives the valence of a conceptual cue averaged over a particular time period. The term, $[\log_b(\max(n, b))]$, rewards frequency of exposures, while the term, $S(E_t)$, rewards the saliency of an exposure. In this simple model of an affective reflexive memory, we do not consider phenomena such as belief revision, reflexes conditioned over contexts, or forgetting. To give an example of how affective reflexive memories are acquired from personal texts, consider Figure 4-6, which shows two excerpts of text from a weblog and a snapshot sketch of a portion of the resulting reflexive memory.

In the above example, two text excerpts are processed with textual affect sensing, and topics both simple (e.g. 'telemarketer,' 'dinner,' 'phone') and compound (e.g. 'interrupt dinner') are extracted. The saliency of each exposure is determined by heuristics such as the degree to which a particular concept is topicalized in a paragraph. The resulting reflexive memory can be queried using Eq. (4.3). Note that while a query on 3 Oct 01 for "telemarketer" returns an affect valence score of (-.15, .25, .1), a query on 5 Oct 01 for the same concept returns a score of (-.24, .29, .11). Recalling that this valence triple corresponds to (pleasure, arousal, dominance), we can interpret the second annoying intrusion of a telemarketer's call as having conditioned a further displeasure and a further arousal to the word "telemarketer".

How does conditioning help the system cope with noise? In Figure 4-6, "phone" also inadvertently inherits some negative affect. However, unless "phone" consistently appears in a negative

Table 4.2. Political culture: attitudes of the Democratic and Republican parties

Democrats		Republicans	
pleasing topics	displeasing topics	pleasing topics	displeasing topics
recommitment	religious	jobs	withdrawal
public	legislation	success	terrorist act
american jobs	religious leaders	productivity	terrorists
values	congress	tax system	iran
nation	religious tradition	taxes	hillary clinton
jobs	criminal charges	skills	enemy
dean	republican congressman	poor	fail to stop
democratic party	money laundering	insurance	significant
american energy	bribery	mississippi	support
science	former majority	growth	embassy
literacy	leader	economic security	nuclear
high standards	god	health care	brutal dictator
mankind	housing	workers	transitional government
parents	democrats	president reagan	new attacks
children	elderly	free trade	higher taxes

affective context in the long run, Eq. (4.3) will tend to cancel out inconsistent affect valence scores, resulting in a more neutral valence.

§

Measuring media viewpoints in political culture. The attitudes of a culture can be modeled in the same way as the attitudes of an individual, simply by regarding the culture’s texts as an individual’s texts. By creating “person” models for extreme viewpoints, polemic cultural spaces can be defined, and individuals can be located relatively as lying somewhere in the continuum between the poles. To demonstrate this, the space of political culture was modeled using WWTT. The goal of the modeling task was for WWTT to assess the political bias of some major media outlets, and to prepare a head-to-head comparison of results with a study of media bias recently conducted by Groseclose and Milyo (2004). In that study, the authors statistically analyzed the patterns in which major media outlets cited political think tanks and policy groups with known political leanings. Observations were made mostly between 1995 and 2004. From their analysis, they estimated ADA (Americans for Democratic Action) scores for twenty top U.S. media outlets. A full ADA score of 100 indicates an ideal Democrat position, e.g. Congresswoman Maxine Waters (D-CA) was scored 99.6. A lowest ADA score of 100 indicates an ideal Republican position, e.g. Congressman Tom Delay (R-TX) was scored 4.7. An ADA score of 50 was considered neutral, e.g. NewsHour with Jim Lehrer was scored 55.8. Of the 20 media outlets analyzed, 10 were primarily television programs (e.g. CBS Evening News), 1 was a radio program (i.e. NPR Morning Edition), 3 were magazines (e.g. Time), and 6 were daily newspapers.

(Liu 2006) (Groseclose & Milyo 2004)

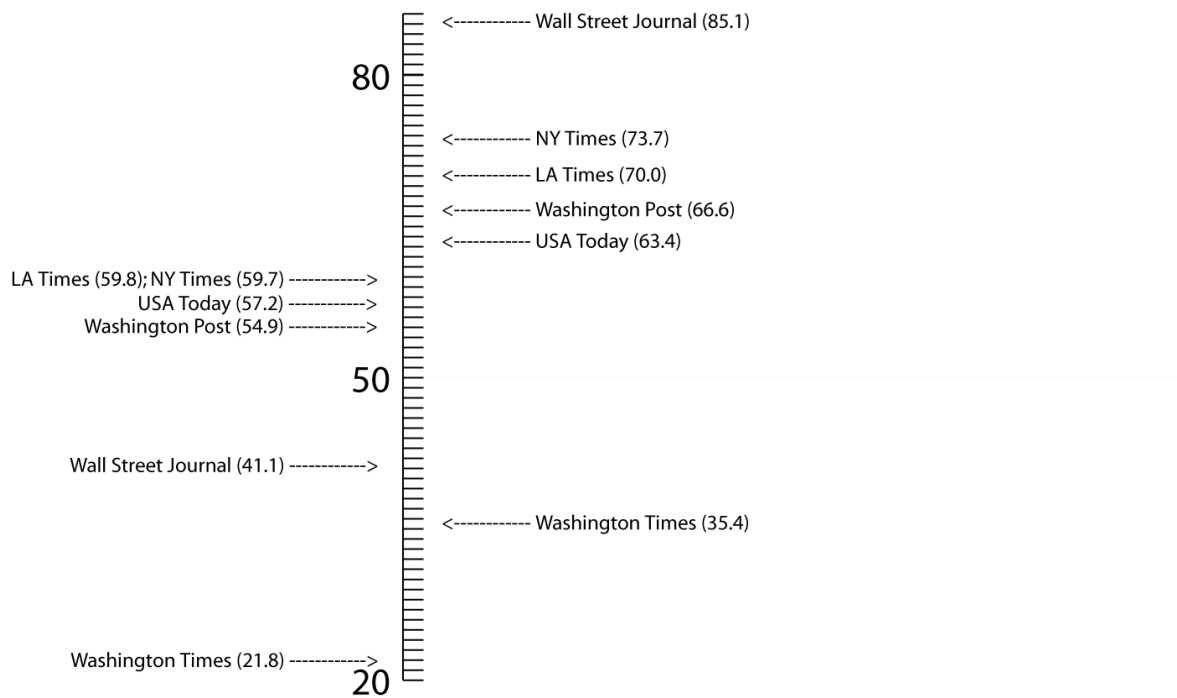


Figure 4-7. Political viewpoints of major newspapers

WWTT's model acquisition system was used to produce a head-to-head comparison with (Groseclose & Milyo 2004). Firstly, the viewpoints embodied by the Democratic and Republican parties were modeled, to generate the two poles of political culture. From the official websites of the Democratic and Republican parties, corpora of everyday texts were compiled from transcripts of each party's repository of political speeches, e.g. President Bush's speeches, and Democratic weekly radio address transcripts. A random subset of the available political speeches made from August 2000 to December 2005 were culled, resulting in 2 Megabyte plaintext corpora for both parties. From each corpus, WWTT compiled a person model. Table 4.2 lists the most pleasing (i.e. P in PAD) and most displeasing topics for the acquired models of Democratic and Republican attitudes, given in rank-order. Overall, the gist of these top lists appeal to common intuition, though the oft imprecision of RATE processing is apparent. For example, 'elderly' appears in the Democrat's most displeasing topics, though it is probably the case the Democrats are displeased by the neglect of elderly, rather than by the elderly. Similarly, Republicans are probably pleased that the 'poor' are being taken care of, rather than being pleased that folks are 'poor'. Imprecision may result in added noise when simulating judgments.

After the poles of political culture were modeled, the viewpoints expressed by major media outlets could be mapped into the continuum between ideal Democratic attitudes and ideal Republican attitudes. Because WWTT needs everyday texts (first-person and self-expressive), editorial texts were sought out as a close match because they are second-person and self-expressive. From Groseclose & Milyo's list of 20 media outlets, only the six major newspapers had editorials available on their websites. Some others, like NPR Morning Edition and the magazines, had featured columnists, but those were not representative of the outlet as a whole, so they were not used. For the six major newspapers, corpora of texts were formed from 1Mbyte of each's editorials compiled from each's website, bearing publication dates in the range of January 2004 to March 2006. Next, WWTT created attitude models for each of the six newspapers. An algorithm was run to align each newspaper's attitudes with each party's attitudes. The algorithm looked at each topic at the intersection of newspaper and party, and calculated the difference between the P-component of their PAD values. The average of all differences (directionality maintained) was recorded. For each newspaper, its alignment score with Republicans was subtracted from its alignment score with Democrats. Then, scores were normalized with a multiplier to the range 50 +/- 30, so that the results could be compared head-to-head with Groseclose & Milyo's results on the ADA scale, as shown in Figure 4-7. NB, WWTT's results, shown on the left-hand side of the ADA axis, do not imply that the absolute numbers are very meaningful, since then were re-normalized; however, the relative distances between newspapers is meaningful.

WWTT's placement of Los Angeles Times, New York Times, USA Today, and Washington Post to the "left" of center were consistent with the compared study. The far "right" of center outlier Washington Times was correctly identified. While most relative positions were preserved, the new placement of Wall Street Journal at the right of center disagreed with the compared study. The difference might be illuminated by the fact that the compared study examined the news articles in the newspapers, while WWTT was fed editorial articles. The difference seems consistent with observations also made by several political analysts, including Howard Kurtz of the Washington Post.

§

Minskian imprimers augment an individual's attitudes. While the basic attitudes model is sufficient to produce reactions to text for which there exists some relevant passages in the personal texts, a person's space of known attitudes are still often quite sparse in what they can react to. The addition of Minskian imprimers supplements the known attitudes in an individual's model with the models of imprimers. Marvin Minsky (forthcoming) describes an imprinter as

someone to which one becomes attached. He introduced the concept in the context of attachment-learning of goals, and suggests that imprimers help to shape a child's values. Imprimers can be a parent, mentor, cartoon character, a cult, or a person-type. The two most important criteria for an imprinter are that 1) the imprinter embodies some image, filled with goals, ideas, or intentions, and that 2) one feels attachment to the imprinter. Minsky theorizes that the images of imprimers can be internalized and their effects still realized *in absentia*. Internalized imprimers, or "mental critics," can do more than to critique our goals; enduring attachment can lead to willful emulation of a portion of their values and attitudes. Keeping a collection of these internal imprimers, they help to support our identity. From the supposition that we conform to many of the attitudes of our internal imprimers, when an individual is put into a novel circumstance for which he has not formed firm judgments, he may instead imitate an imprinter's goals and attitudes. Of course, a person's personality will affect the degree to which others influence their attitudes. This hypothesis is supported by much of the work in psychoanalysis. Freud (1915) wrote of the psychic process of 'introjection'—children unconsciously emulating aspects of their parents, such as the assumption of their parent's personalities and values. Introjection relates closely to other concepts in psychology, such as projective identification (Ogden 1979), internalization, and incorporation.

Minsky suggests that imprimers can be identified as those persons—fictive and real—who can evoke self-conscious emotions like pride and embarrassment in an individual. From this suggestion, WWTT implements imprinter identification by searching everyday texts for persons and subcultures (e.g. 'dog'-->'dog lovers') who elicit high arousal and high submissiveness, and collocate with self-conscious emotion keywords like 'proud', 'embarrassed', and 'ashamed'. The topic-context under which an imprinter exerts influence is also recorded, as the cluster of topics collocating with mention of the imprinter in the texts. One might like Warren Buffett's ideas about business but probably not about cooking. Once imprimers are identified, the imprinter's attitude model is linked to the present individual's model, and becomes invoked whenever the present individual's model cannot return any judgment and the imprinter is authorized to introject judgments for that topic-context. Next, the accuracy of model acquisition in WWTT is evaluated, and the contribution of imprimers is considered.

§

Model evaluation. The quality of model acquisition with WWTT was evaluated in a study with four subjects. Subjects were between the ages of 18 and 28, and had kept weblog diaries for at least 2 years, on average writing new every 3-4 days. A person model was generated for each subject from their weblog. Each person model had three components—the reflexive memory (i.e. the semantic sheet

		P_{deviation}	A_{deviation}	D_{deviation}		P_{deviation}	A_{deviation}	D_{deviation}	
subject₁	avg	0.39	0.27	0.44	subject_{avg}	avg	0.35	0.22	0.38
	σ	0.38	0.24	0.35		σ	0.35	0.20	0.34
	95%-ci \pm	0.22	0.13	0.20		95%-ci \pm	0.20	0.11	0.19
subject₂	avg	0.42	0.21	0.48	subject_{avg}	avg	0.40	0.28	0.44
	σ	0.47	0.23	0.31	(imprimer off)	σ	0.40	0.28	0.44
	95%-ci \pm	0.27	0.13	0.18					
subject₃	avg	0.22	0.16	0.38	baseline_{static}	avg	0.50	0.50	0.50
	σ	0.21	0.14	0.38					
	95%-ci \pm	0.12	0.08	0.22					
subject₄	avg	0.38	0.22	0.20	baseline_{random}	avg	0.67	0.67	0.67
	σ	0.33	0.20	0.32					
	95%-ci \pm	0.19	0.11	0.18					

..... 95% confidence interval




Figure 4-8. Evaluating model-based predictions of persons' reactions to news stories.

of attitudes), a long-term episodic memory (discussed elsewhere), and the person models of any identified imprimers.

In the interview, subjects and their corresponding models evaluated 12 short news snippets taken from Yahoo! News. The snippets are averaged 150 words long, and 4 snippets were selected from each of three genres: social, business, and domestic. The same set of texts was presented to each subject and the examiner chose texts that were generally evocative. The subjects were asked to summarize their reaction by rating three factors on the Likert-5 scale:

- Feel negative about it (1)... Feel positive about it (5)
- Feel indifferent about it (1) ... Feel intensely about it (5)
- Don't feel control over it (1)... Feel control over it (5)

These factors were then mapped onto the PAD format, assuming the following correspondences: 1→-1.0, 2→ -0.5, 3→0.0, 4→ +0.5, and 5→ +1.0. Subjects' responses were not normalized. To assess the quality of attitude prediction, the spread between the human-assessed and computer-assessed valences was recorded:

$$X_{deviation} = |X_{human} - X_{computer}| \quad (4.4)$$

The mean spread and standard deviation were computed, across all episodes along each PAD dimension. On the -1.0 to +1.0 valence scale, the maximum spread is 2.0. Figure 4-8 summarizes the results. Note that smaller spreads correspond to higher accuracy, and smaller standard deviation correspond to higher precision. Two baselines were considered. $Baseline_{static}$ presumed always neutral reactions, so its mean spread was 0.50. $Baseline_{uniform}$ generated random reactions from -1.0 to +1.0 assuming a uniform distribution, so its mean spread was 0.67. A more sensible random baseline might follow a Gaussian distribution rather than a uniform one – implying a mean spread between 0.50 and 0.67.

The acquired person models outperformed the baselines, excelling particularly in predicting arousal, and having the most difficulty predicting dominance. Standard deviations were very high, reflecting the observation that predictions were often either very close to the actual valence, or very far. The results along the arousal dimension recorded a mean spread of 0.22, and mean standard deviation of 0.20. This suggests that our attitude prediction models confidently outperform baselines in predicting arousal.

For each news snippet, reflexive memory was triggered an average of 21.5 times, episodic memory 0.8 times (hence its discussion was left elsewhere), and imprimers' reflexive memories were triggered 4.2 times. The experiment was re-run to measure the effectiveness of each type of memory (for details, see (Liu, 2003b)). We found that episodic memory did not contribute much to attitude prediction because of its low rates of triggering (it was hard to map personal episodes to news story episodes). A pleasant surprise was that imprimers seemed to measurably improve performance, which is a promising result.

4.5 Humor realm: 'catharses'

A *tendentious* joke, according to Freud (1905), elicits howling laughter because it gives catharsis to one's pent up psychic tensions and inhibitions. Premised on this observation, *Catharses* was implemented as a system for acquiring one's humor model from a RATE processing of everyday texts, such as a weblog diary. *Catharses* compiles a semantic sheet of one's tensions toward various topics, and also compiles archetypal tension sheets for various niche cultures. Tension here is calculated as a derivative of a PAD score—as displeasure, high arousal, and dominance. High arousal and dominance together signals aggression—a key characteristic of tendentious jokes.

A genre of humor – such as blonde jokes, political jokes, sexual jokes, foreigner jokes, catholic jokes, Jewish jokes – is justified as a vehicle

of catharsis for members of corresponding niche cultures. For example, those brought up in Jewish families will share experiences, inhibitions, frustrations, and embarrassments, which lead to the formation of a common pattern of psychic tension. Jewish jokes are effective and therapeutic to the niche culture of those brought up in Jewish families because they give catharsis to the archetypal psychic tension of that group. For Catharses, a corpus of 10,000 jokes was compiled, decomposable into twenty niche humors. For each niche culture, ten or so persons and their blogs were identified manually as *exemplars* of folks who would most appreciate that niche's jokes. By producing tension sheets for each exemplar, and by intersecting the exemplars of a single niche culture, archetypal tension sheets are produced. An individual's sheet is located in the space of niche humors by identifying the niche whose sheet of tensions best matches the individual's sheet. In the rest of the section, 1) Freud's hydraulic model of humor is presented, 2) the mining of archetypal tension from niche humor culture is described, and 3) personal humor is viewed as location in the space of niche humors.

§

Freud's hydraulic model of humor. Though many before and after Freud have articulated different understandings of the reason for and mechanism of humor, Freud's (1905) hydraulic model of tendentious jokes remains an authoritative account. It is also fully compatible with the present person modeling approach. Freud betrayed the economics of the psyche. The unconscious is conceptualized as an expanding or contracting hydraulic bag of emotion. There, psychic energies can be stored—saved willingly or pent up unwillingly—and released—saved energy can be spent, or pent up energy can find release. According to Freud, jokes are of three types which correspond with three life phases—first, a child first takes delight in verbal play when they discover that each word is invested with psychic energy, dammed up inside of it; second, as the child's intellect matures, mere play is replaced with jest, or the innocent joke, in which the joke does not yet perform a function other than to delight and pleasure; third and finally, jokes become tendentious, and serve the purpose of releasing psychic tensions produced by prior observance of social inhibitions. As vehicles of catharsis, tendentious jokes may be aggressive—offering an outlet for pent up hostility—or they may be obscene—allowing repressed desires to expose. Thus tendentious jokes have a dual function—they release psychic tension, and the release of the tension is itself a playful and pleasurable act. Freud marveled at the tendentious joke's 'economy of psychic expenditure'.

§

Archetypal tension. By Freud's model a joke is only cathartic if it manages to address the listener's tension. The fact that there are many different genres of jokes—each apropos to a different sort of

person—suggests that listeners’ tensions have some cultural regularity. The appreciators of each niche humor constitute a culture of persons who share a sense-of-humor. Connoisseurs of Bush jokes have all pent up tension about Bush; likewise, for Clinton jokes. What each ‘humor culture’ has in common is termed here as an *archetypal* pattern of psychic tension. The archetypal pattern associated with each of the twenty genres of jokes in a 10,000-joke corpus can be distilled from a set of exemplars. An exemplar is a person with a weblog diary who would appreciate a particular genre of joke. Ten exemplars were assembled for each of the twenty genres, and Catharses produced a semantic sheet of tension-topic pairs for each exemplar. Then, their sheets were distilled into a single sheet representing the shared model of the niche culture by adding the sheets together and renormalizing. Archetypal tension sheets are normalized such that the sum of all tension values is equal across all the culture’s sheets.

§

Locating personal humor in the space of niche cultures. Figure 4-7 depicted the location of newspaper’s political attitudes in the bipolar space of political culture. Through a similar approach, Catharses locates individuals in the multi-polar space of humor culture. Poles correspond to the twenty niche cultures. An individual’s distance toward a pole represents the degree to which any archetypal tension sheet releases the individual’s tension. Release is calculated by subtracting the tensions of the archetypal sheets away from the corresponding tensions in the individual’s sheet, and tracking just how much tension was relieved. The niche producing the greatest release can be regarded as the optimal genre of joke for that individual.

Of course, tensions vary from day to day, so a more just-in-time model of an individual can be produced by creating a tension sheet for just the everyday texts of the current day—recent weblog entries, instant messaging conversations, email writing, etc. In Chapter 5, Catharses qua artifact is revisited as a cathartic jocular companion.

This chapter presented a detailed discussion about the five implemented model acquisition systems. Quantitative evaluations of acquisition accuracy were also presented for the three primary systems—TasteFabric, ESCADA, and WWTT. Next, Chapter 5 will present six implemented perspectival artifacts, which leverage the produced person models. By simulating a person’s taste judgments to support various tasks, perspectival artifacts constitute novel tools for learning, self-reflection, and deep recommendation.

5 Perspective-based applications

What are acquired and generalized person models good for? The prospect that individuals' perspectives could be captured begs the possibility that those perspectives might also be played back in ways interesting to application users. This chapter presents a series of six implemented perspective-based applications, which accomplish this feat. Each application encapsulates a generalized person model and simulates a person's reactions to achieve various ends. Section 5.1 describes the *Aesthetiscope*—an art bot capable of creating art customized to a perceptual perspective—and presents an evaluation. Section 5.2 presents *Ambient Semantics*—a taste-based social introduction facilitator. Section 5.3 presents an *Identity Mirror*—reflecting a person's cultural taste identity back unto her. Section 5.4 presents *What Would They Think?*—a panel of virtual mentors and pundits who offer just-in-time feedback about a user's present context—and describes an evaluation of WWTT in a person-learning task. Section 5.5 presents the *Synesthetic Cookbook*—a taste-model-driven interface for food foraging. Finally, 5.6 presents *Catharses*—a Freudian joke-telling companion which foils life's frustrations with humor. Collectively, these applications demonstrate the potential for person models in support of person learning, self-reflection, and deep customization.

5.1 Art that's always tasteful

What is an artwork and how could a machine become artist? If an art bot were to assume the aesthetic perspective of the viewer, could artistic creation proceed such that the resulting artwork would suit especially the viewer's tastes? To investigate these questions, *Aesthetiscope* (Liu & Maes 2006) was implemented as a perspectival art bot which creates mythic color grid artwork reminiscent of early twentieth century abstract expression pieces of Ellsworth-Kelly, Albers, and Rothko. The art bot takes as input a word or a poem, and interprets the semantics of the input into the evocative world of colors through the aesthetic perception perspective of a particular

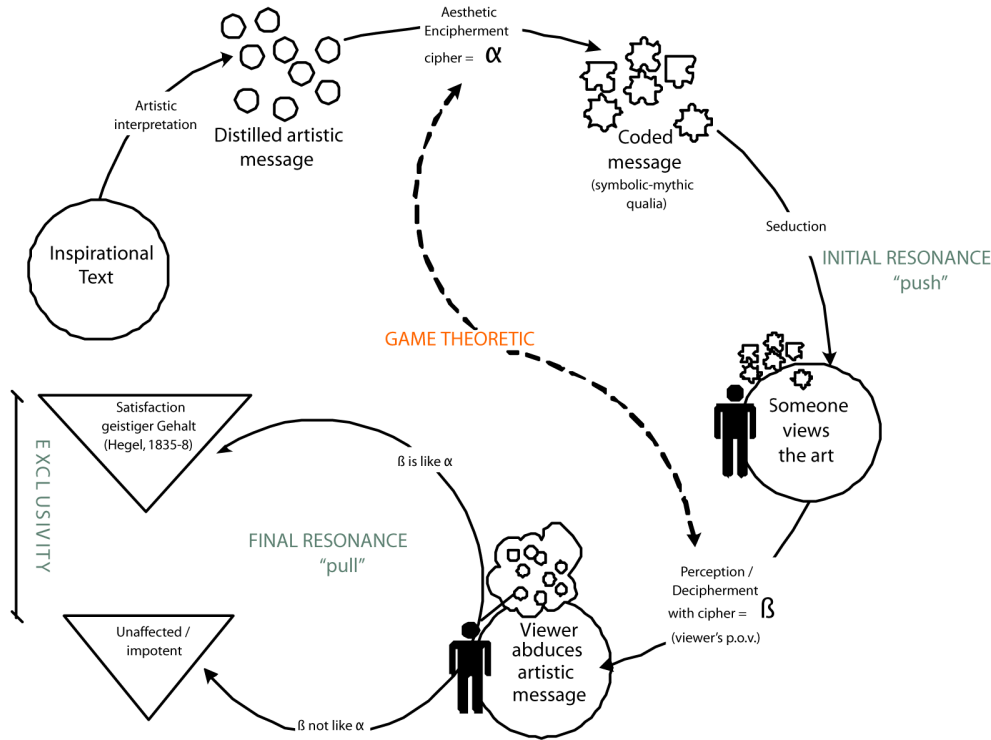


Figure 5-1. A semiotic model of aesthetic transaction

person, as modeled by the ESCADA system. The idea is that artwork can be custom-tailored to each person's preferred way of perceiving imagery and symbolism from a narrative. The rest of this section 1) overviews Aesthetiscope's semiotic framework and artistic mechanism, 2) evaluates the aesthetic efficacy of the Aesthetiscope, and 3) outlines some scenarios for use.

§

Semiotic framework. Aesthetiscope is premised on a semiotic theory of aesthetic transaction (Figure 5-1), based in Dewey's (1934) account of aesthetic experience, and narrated as follows. An artist works with some inspirational 'text', which can be an idea, some imagery, or whatever. Exercising interpretation and judgment, the artist distills the text into a more concise artistic message—via a *perspectival reading*. The artist chooses some artistic medium, e.g. a color grid. The expression of the message into the medium can be thought of as an aesthetic encipherment—the cipher being the artist's intent for and semantic appropriation of the medium. Next, the seduction phase involves motivating viewers to view and perceive the artwork. Viewers may be enticed toward an artwork by any number of factors—by its superficial beauty (e.g. shiny interesting things), by cultural endorsement (e.g. da Vinci's Mona Lisa), by intrigue or mythic qualities (e.g. sunset over a lake, the woods at dusk, a glimmering color grid), or by virtue of a thing

being located in a place-for-perception (e.g. a museum). Marcel Duchamp, for example, impelled viewers toward his Fountain, a urinal placed on its back, by virtue of the surrealism of seeing such a crass and lowly object within a museum. When a viewer perceives an artwork, he deciphers it through an understanding of the medium's semantics and by adopting various perceptual strategies—the most natural of which is his own perceptual perspective. If the viewer deciphers the artwork as it was intended, he will unlock the artistic message and feel satisfaction. If deciphering is unsuccessful or incorrect, it will leave the viewer unaffected. Decipherment should be viewed along a graded scale of efficacy, especially for abstract artwork. Finally, two principles govern the aesthetic efficacy of this transaction. Aesthetic experience is strongest 1) when the viewer finds that the message of the art is one that he, given his perspective, is more competent to receive than an arbitrary person is (exclusivity principle); and 2) when the meaning of art is not worn on the surface but is something that must be excavated from the artwork by the viewer's own wits (final resonance principle).

Perspectival reading. The viewer's perception person model, acquired by ESCADA, is applied toward a perspectival reading of the inspiration text. Five evocative readers—one for each of the perception model's four Jungian dimensions, plus another dimension for cultural perception—read the same inspiration text, and each produces a set of keywords signifying a particular understanding of the text. For example, Robert Frost's poem, "Fire and Ice," produces at least these keywords:

ThoughtReader → "earth", "cold", "hot"
 SightReader → "photos of fire", "photos of world", "photos of ice"
 IntuitionReader → "hot", "engine", "red", "freezing", "summer"
 SentimentReader → "arousing", "pleasurable", "passionate"
 CultureReader → "crazy", "fashion", "racism"

A perception model specifies the percent engagement of each psychological function: Think10%-See40%-Intuit50%-Feel70%--Culturalize10%. ESCADA does not give the culturalize dimension, but it may be manually set or set to 0%. The keywords produced by each of the five evocative readers are blended into a final set of keywords, according to the proportions specified by the person model.

Rendering perspectival reading in color space. The final keywords are mapped into color space according to three color logics: naturalistically sampled colors as culled from a photo corpus depicting thousands of objects in the world (e.g. colors of a tree taken from a photo), mood colors (e.g. colors for love and fear) as prescribed by the color psychology of the Bauhaus (Itten 1961), and symbolic colors (e.g. apples are red, the sky is blue), gotten from ConceptNet (e.g. has-property(stop sign, red)). Details of color rendition are given elsewhere (Liu & Maes 2006). The goal of conveying the text's singular, complex aesthetic character to the

Table 5.1. Signalling efficacy of Aesthetiscope’s five dimensions

	<i>Plausibility - 100 Poems/Songs</i>					<i>Plausibility - 100 Evocative Words</i>				
	Think	Culture	See	Intuit	Feel	Think	Culture	See	Intuit	Feel
Judge1	2.3	2.2	3.6	3.6	3.8	3.0	2.6	3.1	4.0	3.5
Judge2	2.0	2.3	3.3	3.3	3.8	2.5	1.8	2.9	3.5	3.6
Judge3	1.8	1.9	3.1	2.6	3.5	1.9	2.0	2.3	3.6	4.0
Judge4	2.5	2.3	3.7	3.4	4.3	2.6	2.5	2.6	3.5	4.5
Avg Score	2.2	2.2	3.4	3.2	3.8	2.5	2.2	2.7	3.6	3.9
Avg StdDev	±0.9	±0.7	±0.6	±0.8	±0.7	±1.1	±1.2	±1.6	±1.0	±0.8
Kappa' (avg)	0.31	0.33	0.51	0.40	0.56	0.48	0.42	0.68	0.70	0.75

perceiver is facilitated by the expectation that the viewer will blend these colors together in the mind’s eye, and attend to their undeconstructed gestalt rather than to each square individually. In this manner, the aesthetic character is not a simple sum of individual color squares, but rather, it becomes that Spirit which lives in-between the color squares.

§

Evaluations. Two evaluations were performed to assess the efficacy of the Aesthetiscope. The first evaluation measured the signalling efficacy of each of the five reading dimensions. The second evaluation measured capability of Aesthetiscope for producing aesthetic artworks.

Signalling efficacy of single reading dimensions. In the first evaluation, four human judges, all graduate students in science, art, or architecture, scored Aesthetiscope renditions of 100 commonly known assorted poems and songs (e.g. Browning’s “How Do I Love Thee?”, first passage of “The Raven,” “I Know Why the Caged Bird Sings,” “I Can’t Get No Satisfaction”, Lenin’s “Imagine”, “Good Vibrations”), most in the range of 150-400 words, and 100 evocative common words (e.g. “God,” “money,” “power,” “success,” “crime”) chosen dispassionately by the examiner but with care to maintain diversity. Because some words were potentially unknown to Aesthetiscope, the examiner discarded unknown words and replaced them until 100 known words were arrived at. Image sets of the text laid over the color grid rendition (so judges could refamiliarize themselves with the text) were pre-computed for these 200 renditions. Each set contained five images, each image visualizing one of the reading dimensions. Judges were asked to score each of the 1000 total images on the following instruction: “How plausibly does this artwork communicate the thoughts|cultural notion|imagery|free intuition|feelings you had of this text?” Scores were recorded on a standard Likert 1-5 scale (1=not plausibly, 5=very plausibly). Kappa coefficients, a commonly used measure of inter-rater agreement in classification tasks, were calculated between every pair of judges, and the average scores

computed. We relaxed the definition of agreement as two judges giving Likert scores with difference 0 *or* 1. Results are shown in Table 5.1.

Results suggest that renditions from Think and Culturalize were fairly poor insofar as they fell short of employing colors to manifest the judges' Think and Culturalize readings of the text. Renditions from See were fairly plausible in the poems/songs task, but very inconsistent on the word task; its very high average standard deviation of 1.6 on words suggests that it completely failed to visualize some abstract words, e.g. "power," while succeeding perfectly on words corresponding to concrete things. Intuit and Feel performed the best, and were consistently plausible in their renditions. Standard deviations trended higher on the word task, while the average scores were on par with the poems/songs task – this indicates that each reading was more brittle on just the one word input; however, when a rendition was successful, it was more intensely successful on the one-word input than for poems/songs. The average Kappa statistics (0=pure chance, 1=perfect agreement) indicate a fair to good agreement amongst the judges, with the greatest convergence of opinion around Feel, and demonstrating greater agreement in the word task than in the poems/songs task. These results are promising, but reveal that Think, and Culturalize lead to weak renditions; however, because these categories also saw the lowest inter-rater agreement scores, we could conclude that either 1) these are difficult dimensions to computationalize for a general public, and we should try to personalize these models; or 2) these are dimensions not generally amenable to expression in color space, and perhaps colors are not strong enough stand alone signals for these dimensions, perhaps form is also required.

Aesthetic efficacy. Can Aesthetiscope produce a satisfying color impression of a text in a non-arbitrary manner? A manual perspective setting of Think10%-Culturalize10%-See40%-Intuit50%-Feel70% was chosen, as it is a most widely appealing and suitable perspective. Taking the text from the 100 poems/songs, and 100 words, Each text was laid over its own rendition, and also over the rendition of another random text within the same category (poems/songs and words are separate categories). This randomized rendition controls for interface and form and helps to isolate measurement to just the ability of the chosen perspective to judiciously and aesthetically express the gestalt of the text. Fifty-one undergraduate students from MIT and Harvard University were each asked to make twenty at-a-glance binary judgments on randomly selected items in each of the two task categories: poems/songs, and words. The instruction was: "this text inspired which of these two artworks?" The results were as follows: in the poems/songs category, the correct text was paired with the artwork with an accuracy of 75.2% across all judges; in the words category, accuracy was 80.7% across all judges. Kappa statistics were not calculated because each volunteer judged a randomly selected subset




Figure 5-2. The effects of perspective on artistic rendition—the words “sunset” (top-row) and “war” (bottom-row) are rendered with a Thinking-Seeing bias (left-column) versus with an Intuiting-Feeling bias (right-column)

of the available renditions. With these results, a measure of confidence is gained that Aesthetiscope’s color renditions produce an aesthetic in the vein of art, and its aesthetic is demonstrably and non-arbitrarily tied to, and inspired by a perspectival reading of a text.

§

Possible appropriations. One might imagine that with the Aesthetiscope, the art on the wall might be generated or selected to always suit the tastes of the viewer. Because the Aesthetiscope’s artistic creation mechanism is seeded from some inspirational text, such as a word, a poem, or song lyrics, perhaps the art on the wall will serve as an aesthetic bridge between a room’s occupant and the room’s current event. The Aesthetiscope was, for a time, installed in a living room of the future at the MIT Media Laboratory, projected unto one of the room’s walls, and constantly changing itself in reaction to a book of poetry being interacted with, or the song being played over the room’s sound system. Just as wine is chosen to complement a meal, the Aesthetiscope can complement an activity with custom artwork. The room visualizes a sunset with swaths of deep purple when a Realist occupies the room (top-left panel in Figure 5-2). Then, when the Realist leaves and a Romantic enters, the same sunset is portrayed with brilliant and warm hues of orange (top-right panel in Figure 5-2). When they stand together, the




Figure 5-3. Aesthetiscope’s aesthetic impressions of the four season keywords (columns) rendered through the monadic optics of each Reader taken alone (rows)

artwork renders an average of their perspectives, illuminating their communion.

The marvel of affording artwork customized to a viewer’s perspective is that aesthetic perception, usually held in deep privacy by a person, half-submerged in the unconscious, can now be worn on one’s sleeve, so to speak. The viewer sees how she impacts an artwork depicting ‘war’ (bottom-left panel in Figure 5-2), sees how another person impacts the same artwork (bottom-right panel in Figure 5-2), and is provoked into critical re-evaluation on self, other, and the sense-of-beauty. Instead, what if instead of setting an artwork to one’s own perspective, it is set to the perspective of a loved one on distant shores? Might having the company of that dynamic artwork, reacting to everyday happenstance, offer comfort, solace, and facilitate emotional intimacy with the remote other?

By instilling the perception model of friends, mentors, celebrities, or cultural archetypes into the Aesthetiscope, it might be possible to learn about and intimate with someone or some culture in an exciting new way. Consider that five unique prisms for perceiving

the four seasons refracts the aesthetic of the seasons into twenty unique visions (Figure 5-3)—a plethora of aesthetic possibility.

5.2 Facilitating social introductions

Ambient Semantics is a wearable contextual information system that supports users in discovering objects and meeting people through pithy *just-in-time feedback* given in the crucial first moments of an encounter. For example, driven by the cultural taste model of its wearer, the system detects when its wearer has picked up a copy of Marvin Minsky’s “Society of Mind.” Through a display projected near the wearer, the system can indicate to the wearer that he would find the book tasteful because Minsky qua author and the genre of Artificial Intelligence books resonate strongly within the wearer’s cultural taste ethos.

By intersecting the cultural taste ethoi of two strangers, Ambient Semantics can help break the ice—exposing interests and subcultural identities that the two are likely to share.

Calculating taste-similarity: quantitatively vs. qualitatively. There is more than one good way to use TasteFabric to calculate the taste-similarity of two people. The more direct way is to measure the intersection of two spread activations. Taking each person’s seed profile of interests and mapping it into the taste fabric, we arrive at an initial configuration. Spreading activation outward from this configuration defines a semantic neighborhood, which earlier in the paper we referred to as a person’s taste ethos. Taking the semantic intersection of two or more persons’ ethoi, we arrive at the quantitative calculation of taste-similarity.

However, another intriguing possibility is to make a qualitative calculation about taste-similarity. Although the intersection of two taste ethoi is mathematically satisfying, it is not easily explainable and articulated. In other words, having the system explain that “the two of you share taste because you both have interests x, y, and z in your spreading activation clouds” is inappropriate. More articulate would be to cite a shared habitation of taste neighborhoods, for example, this explanation—“the two of you share taste because both of you are adventurers and lovers of wine.” Here, the mechanism of the recommendation feels more transparent. To calculate qualitative similarity, each person’s taste ethos would be used to score the degree of a person’s habitation across the various taste neighborhoods, which as you recall, are centered around identity nodes. Like the classic k-nearest neighbors classification scheme, here we classify persons by their k-nearest taste neighborhoods. Having completed this mapping, the subset of neighborhoods shared among the two or more persons become those persons’ shared situation. To communicate shared neighborhoods to the persons, the neighborhoods could be effectively visualized on a

screen, or, neighborhoods are safely summarized by stating the identity nodes which live within that neighborhood.

5.3 An identity mirror

Identity Mirror (Liu & Davenport 2005) is a perspectival artifact that shows a viewer her dynamic textual reflection against the changing cultural fabric. Instilled with the cultural taste model of the viewer, acquired by the TasteFabric system, Identity Mirror visualizes the viewer's cultural taste ethos as a cloud of keywords swarming over a live artistic image of the viewer, captured via webcam. The mirror uses image recognition and object tracking to afford self-reflexive performance—the viewer qua performer uses distance from the mirror, body movements, and gesticulations to negotiate her textual identity. The viewer's cultural reflection is dynamic—it changes from day to day as culture herself travels through new priorities and foregrounds new aesthetics. Culture's priorities are measured by continuous reevaluation of the cultural topics that are foregrounded via live on-line news feeds. The rest of this section 1) motivates the this mirror in the context of self-reflection; and 2) discusses the performative and critical affordances of the Identity Mirror.

§

Self-reflection as a performance. The practice of self-reflection in everyday life is the venue for grasping one's situation in culture, rightfully synonymous with 'identity.' Self-reflection, motivated by a desire to understand, reveal, and identify, is a play atop the stage of the Cartesian Theatre of mind—there, an individual is all of performer, audience, and *meta-audience*. On the one hand, he replays his performance up to the present. On the other, he imagines in his mind's eye what the judgment of culture-as-audience might be. And then on the third hand, he is individual-as-meta-audience, watching the back-and-forth of performer and audience, and feeling self-conscious emotions like shame or pride. Self-reflection needs to recur because we continue to perform throughout life, and therefore need as often to reflect upon the previously unassessed. Also, culture's eye is capricious, its attention and emphasis over matters ebb and flow, thus we must continually re-anticipate culture's judgment over not only our nonce, but also over the whole history of our being in the world. In self-reflection, we re-perform ourselves in order to understand through culture's eyes what our performance means—this might aptly be termed a *self-reflexive performance*.

Self-initiated and focused self-reflection is hard—it would seem to require an enormous effort of the imagination, and an advanced intuition for the cultural *zeitgeist* to reflect deeply and successfully. The exteriority of our performance is hardly obvious. In our present society, reflexive provocation comes under the aegis of cinema and literature, where techniques such as 'suture' raise the critical consciousness of audiences by projecting the protagonist into viewer.

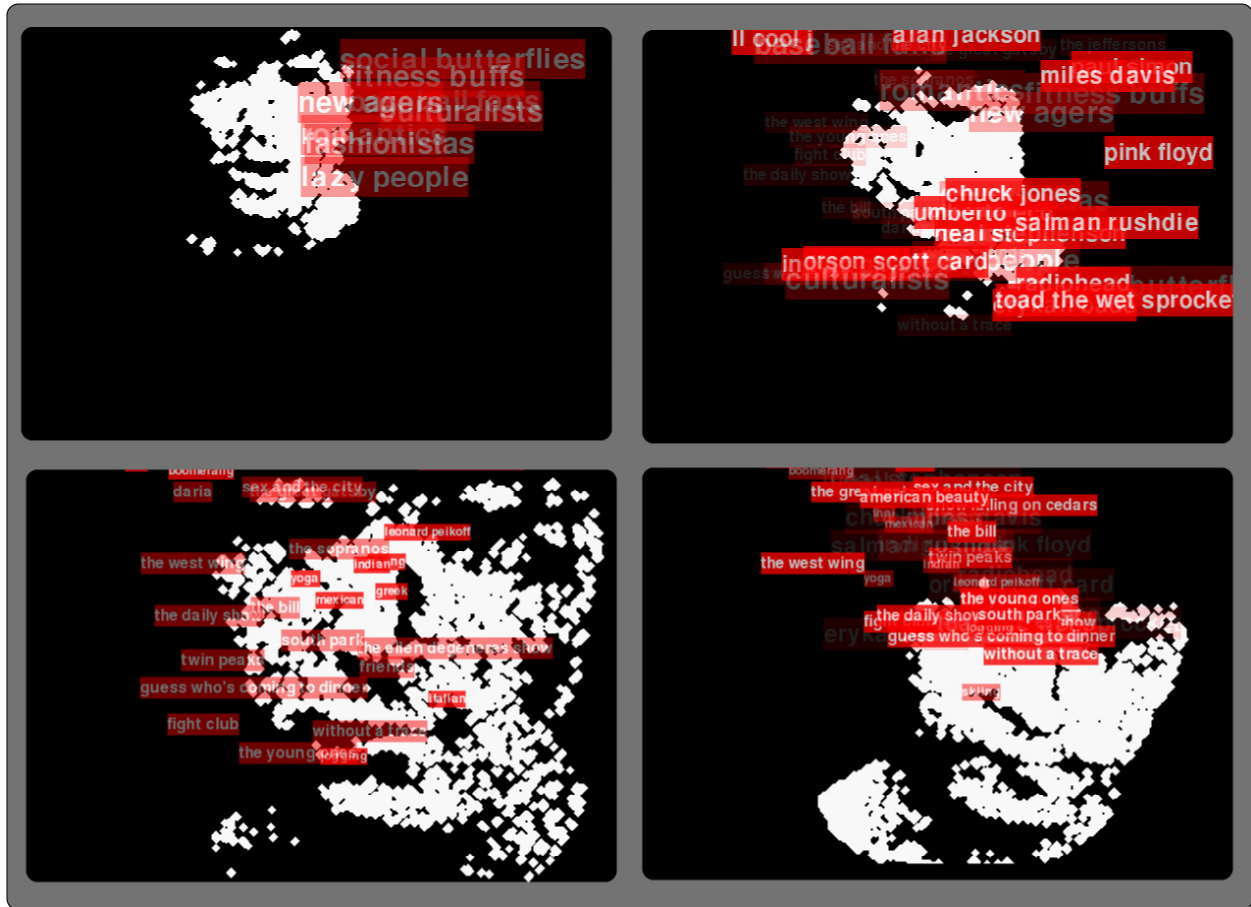


Figure 5-4, a-d. Reflections in the Identity Mirror – (clockwise from upper-left) as performer approaches the Identity Mirror, his reflection gains descriptive granularity, passing from subcultures (a), into genres and artists (b), into films and albums (c), into foodstuffs, activities, and songs (d)

But might there be a technological device to support vivid self-reflexive performance? The device should vitally support the performer’s critical imagination of culture-as-audience, as does film and narrative, yet the subject of the self-reflexive performance should be the performer herself.

The computation of the cultural fabric of taste in Taste Fabric, and the location of an individual’s cultural identity as an ethos atop that fabric presents an opportunity to create a self-reflexive technology. The Identity Mirror is a mirror instilled with a model of the viewer’s cultural identity. In self-reflexive texts like a social network profile, an individual will acknowledge certain interests and subcultures. But the Identity Mirror does not constitute the reflection with these self-acknowledged interests, rather, interests and subcultures which are entailed by the self-acknowledged interests are displayed – these represent the judgment of culture-as-audience upon the viewer. Because the viewer’s reflection is not self-acknowledged, herein lies reflexive provocation as viewer tries to see herself in that foreign reflection.

§

Performative affordances. Mirrors engage because they move us to self-awareness of body, movement, and intention. Building upon the success of mirror-based performance interfaces like Krueger's (1983) Video Place and Rokeby's (1995) Very Nervous System, Identity Mirror asks, what if instead of sound or imagery, an interface could immerse an individual in the abstract cultural flows of identity and symbolism?

The performer stands before a large screen (the mirror) displaying a swarm of keywords (cultural identity) hovering over a silhouette of the performer. A computed audience representing cultural judgment generates the keywords based on the implicit component of the performer's cultural taste ethos. The mirror, however, reflects not the performer's 'expressions given' (to echo Goffman), but rather culture's reception of the performer's 'expressions given off'—the mirror shows you your ethos and situation within the space of culture. Even if you insist that you are intellectual, if you love American Football and the movie *Top Gun*, then your cultural reflection in Identity Mirror will deny your intellect and instead brand your identity with "Republican Party." The judgment of culture-as-audience is not always kind or easy to swallow.

Figures 5-4, a-d depict the swarm-of-keywords format of reflections employed by IdentityMirror. This mode of describing could be variously characterized as hypertextual or intratextual. Keywords constitute a hypertext because of their nonlinearity and because subcultures expand into genres, genres into albums, albums into songs, and so on. A swarm-of-keywords serves also as an impressionistic device, with the sum of its descriptors insinuating a central intratextual thematic—one's identity, or sense-of-situation. Identity's mythic richness can be well preserved through intratextual portraiture—or "thick description," as Geertz called it (1973). Identity Mirror affords several performative interplays, described below.

First interplay—dancing with culture. A most immediate interplay between performer and culture-as-audience is body movement and physical gesture. Using real-time video tracking, the location and distance of the performer from the display can be sensed. With the performer standing far away, the reflection is comprised of culturally general descriptions (Figure 5-4a), labeling the performer's silhouette with subcultures, musical genres, book genres, and so on. As the performer approaches the mirror (Figures 5-4, b-d), general keywords fade out, supplanted by increasingly detailed descriptions like musical artists, authors, auteurs, cuisines. When the performer is closest to the mirror, he sees book titles, songs, sports, food dishes, which compose culture-as-audience's assessment of his ethos. Thus, movement toward and away, evoking the back-and-forth footwork

to many ballroom dances, allows the performer to physically negotiate the granularity at which culture judges him.

Of the keywords comprising the characterization, some are deeply rooted—deemed by culture as central to the performer’s identity, whereas other keywords more tenuously describe the performer’s ethos. If the performer is slow and deliberate in her movements, keywords will swim viscously in the interior region of her silhouette, and the tenuous keywords will be visible. However, if the performer should jerk or move too quickly, keywords will bounce around vigorously in the silhouette’s interior, and tenuous keywords will not be visible.

Walking to and fro the mirror affects the granularity of the keywords being shown, describing a far away performer with subcultural keywords, and an up-close performer with descriptors like song names, books, food dishes, etc. Reflecting the dynamicity of a person’s generalized cultural taste model, the reflection updates itself continuously as it monitors live news feeds which represent the pulse of the cultural *zeitgeist*. For example, when Oscars season strikes, the reflection portrays the more glamorous side of the performer.

Second interplay—a restless audience. As the mood of culture’s collective consciousness shifts from one day to the next, a corresponding shift can be felt in the attention faculty of culture-as-audience. Culture’s network of interconnected symbols is always in flux. As new connections emerge, other connections atrophy. These cultural shifts take place on a longer time scale. Day-to-day changes are shifts in mood. They are reflected in the energy levels of each symbol on the cultural fabric. When a symbol is highly energetic, it tends to contextually bias how a performer’s profile will be interpreted. Aspects of the performer close to the biasing symbol will be more prominent in the performer’s cultural identity. To mirror procession of the cultural *zeitgeist* and mood, Identity Mirror applies a machine reader to each day’s online news feed, extracting hot topics *du jour*, and using those topics to selectively energize and enervate the network of symbols. For example, immediately after September the 11th of 2001, the performer’s cultural identity would have appeared much more austere than immediately before that date. As cultural emphasis shifts, so unwittingly does the performer’s displayed identity, for identity is always articulated against culture.

Third interplay—off-stage performance. Viewing everyday life as a continuous performance, it would make sense that performance extends beyond time in front of the mirror. While only active reflection is self-reflexive, unaware actions in the world can still be judged by culture-as-audience. Off-stage, an individual builds a history of choices and behavior, to the extent that those aspects can be monitored and characterized. The individual listens to music,

buys books, plans a night out on the town. The next time that the individual is before the Identity Mirror, the complete history of off-stage choices and behavior is remembered and incorporated into culture-as-audience's perspective on the individual. Off-stage choices and behaviors performed by an individual within a particular context often suggests a facet of their persona. For example, the individual preparing for a Saturday night on the town listens to disco music and browses the Web for social events. Based on her performed acts within this context, culture-as-audience sees her not as her usual self but as a disco queen that night, so her reflection at that point in time is constituted by keywords belonging to her fun and entertaining facet. Facets and cultural mood shifts, *in toto*, demonstrates how the computation of culture can account for ephemera such as the passage of time, and the shifting spotlight of attention.

Fourth interplay—shadows. Performance casts many shadows, on the stage and over the audience, visibly and affectively. To disintegrate culture for a moment into its innumerable constituent realms, e.g. the world of fashion, the world of literature—each dimension behaves as a surface of sorts. Identity Mirror reifies the metaphor of surfaces and shadows, affording dancing with shadows as a further interplay. Whereas cultural reflection aims at a complete account of an individual in culture, dancing with shadows is phantasmagoric. Shadows of the performer against the surface of fashion, of food, or of literature alone are pale distortions of the whole self, but in multiplicity, shadows foment a dramatic nimbus of potentialities about the performer. Identity Mirror displays various shadows against fashion, literature, food, etc. When the performer stands to the left edge of the mirror, a shadow is cast to the right, as a dark swarm of keywords.

5.4 Virtual mentors, and pundits too

If we could get frequent and timely feedback from people whose opinions we value (*e.g.* family, friends, mentors, experts), then perhaps their perspectives would enhance our ability to interpret situations and make decisions. However, we often lack access to the people whose feedback we value, so we are forced to learn about their perspectives in other ways, *e.g.* by inferring attitudes and opinions from prior conversations or from books and papers. Forming a deep understanding of a person in this manner requires immense effort, and there is no guarantee that we can recall a person's opinion on a specific topic at the time we need that feedback the most. However, if a mentor's attitudes could be captured, shared with students, and played back in reaction to students' own readings and writings, it would represent a new mode of learning. *What Would They Think?* was described in Chapter 4 as an attitude model acquisition system. Here, it is reviewed as a perspectival artifact—a panel of virtual mentors and pundits who sit on a user's computer desktop, providing just-in-time affective feedback to the user's textual context. The rest of this section 1)



Figure 5-5. A panel of virtual AI mentors react affectively to a passage of text that the user is reading

describes some interactions with WWTT; 2) describes the interface's design; and 3) evaluates WWTT as a tool for learning about other persons.

§

Interactions. The application consists of a panel of advisors who sit on the desktop—be they mentors, pundits, friends, or self. A user can install anyone into the panel, by accessing the configuration menu, naming a new advisor, choosing a photo for the advisor, and pointing WWTT to a corpus of the advisor's everyday texts—such as a weblog diary or a bunch of commentary-rich papers. The system acquires the attitudes model of the advisor, and installs a new icon into the advisory panel, which is instilled with the model.

The advisors observe the user's textual activities—as she browses a webpage, writes an essay, or replies to an email—and the advisors constantly react to the current text being read or written. At any given moment, the user's textual context is passed in as input to the panel, and each advisor continuously emotes reactions. For

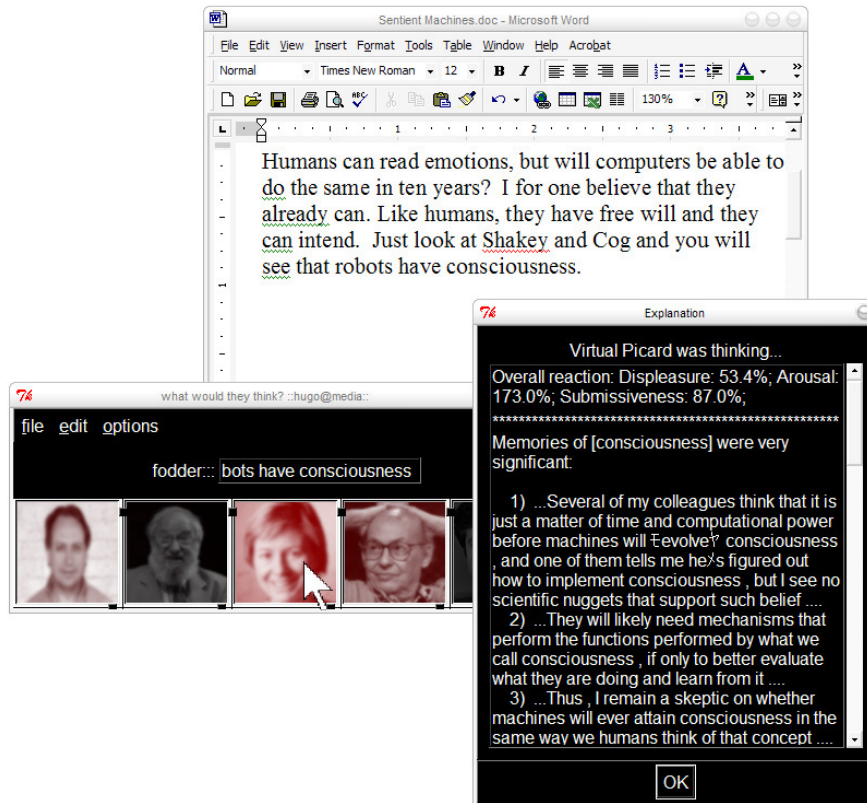


Figure 5-6. Clicking on virtual Roz Picard triggers an explanation dialog

example, Figure 5-5 depicts a panel of virtual AI mentors – (from left to right) Rod Brooks, Seymour Papert, Roz Picard, Marvin Minsky, and Doug Lenat—instilled with corpora of their commentary-rich papers. The user’s immediate textual context is her browsing of a project describing “Sociable humanoid robots.” As she highlights a paragraph of text, the AI mentors react affectively. A green-tinted face indicates an approving response, while a red-tint indicates disapproval. Intensity corresponds to affective arousal over a topic. Sharpness betrays dominance/confidence, blurriness corresponds to submissiveness/uncertainty.

Figure 5-6 depicts another example. This time virtual AI mentors constantly emote reactions as the user composes an essay on robots. As the user types “... that robots have consciousness,” panelists Rod Brooks, Roz Picard, and Marvin Minsky emote disapproval. The user wonders whence virtual Roz Picard’s intense disapproval, so he double-clicks virtual Picard and up pops an explanation dialog. Explanations are composed of salient quotes from a mentor’s everyday texts which support what is emoted. Here, virtual Picard’s reaction is defended by her quote— “Thus, I remain a skeptic on whether machines will ever attain consciousness...”

Polling community opinions. Generalizing from the idea of a panel of installed mentors, WWTT also supports the visualization of an entire online community. WWTT can be passed in a list of hundreds of weblogs, such as those that constitute a special interest group on the LiveJournal or Xanga websites. The blogs are automatically scraped, and miniature icons are created from each blog's profile photo. To poll the opinions of the community, the user can type a proposition such as "the Iraq war is justified" in WWTT's input box. The reactions of the community are visualized as patterns of red and green, intense and dim; these patterns are intelligible even when the icons themselves must be resized into miniature squares. Possible applications include computer-supported ethnography, virtual focus groups, and expert finding within organizations.

§

Interface design. Virtual mentors are represented visually with pictures of faces, which occupy a panel (or $n \times n$ matrix, to accommodate more personas). Given some input, each mentor expresses the affective reaction by modulating the graphical elements of its icon. Each mentor is also capable of explaining what motivated its reaction by displaying salient quotes from its repository of personal texts.

The iconographic face. A virtual representation of a person is given as a normalized, gray-scaled image of that person's face. Affective reactions are conveyed through modulations in the color, intensity, and sharpness of the face image. Early experimentation taught that faces were a more convincing visual metaphor something textual or abstract. Faces are quickly recognized and work well as a cognitive container for mentors' personalities. Affect is conveyed by modulating the image rather than by manipulating facial expression and gaze. The face is fraught with cues, so it is important to not portray more detail in the face than the person model has captured. Scott McCloud (1993) has explored extensively the representational-vs.-realistic tradeoff of face drawing in comics.

Spotlight metaphor. A straightforward scheme maps the three PAD dimensions of a mentor's simulated judgment (pleasure, arousal, dominance) onto the three graphical dimensions of color, intensity, and sharpness, respectively. The baseline image being modulated is gray-scaled and its brightness and contrast are equalized to be uniform across all images. Using a traffic light metaphor, a pleasurable or approving reaction tints a face green, while an unpleasurable or disapproving reaction tints a face red. An affectively aroused reaction results in a brightly lit icon, while a non-aroused reaction results in a dimly lit icon. A dominant (confident) reaction maps to a sharp, crisp image, while a submissive (unconfident) reaction maps to a blurry image.

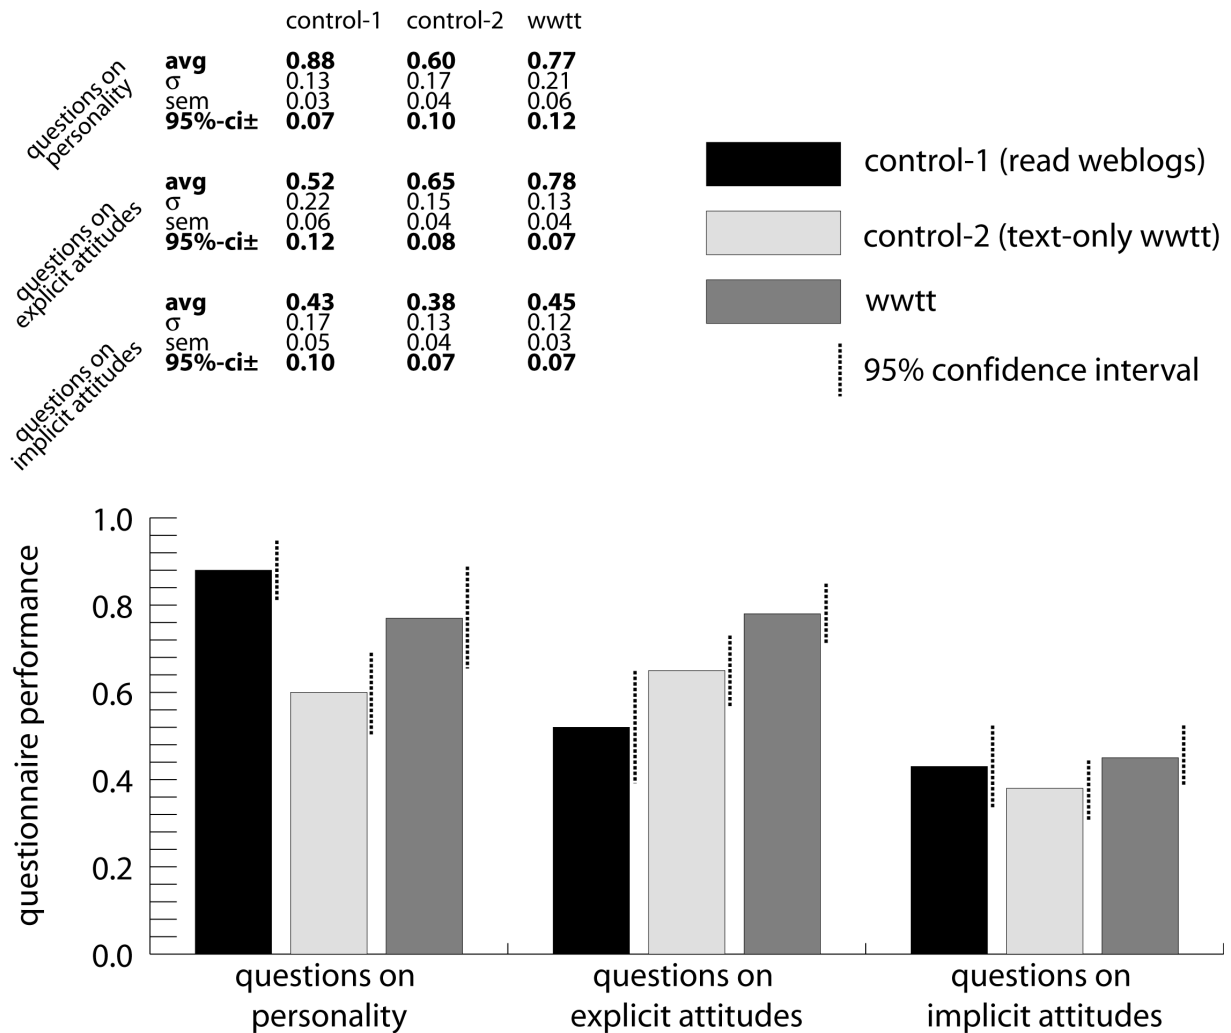


Figure 5-7. Evaluation of WWTT in a person-learning task

Explanation. A virtual mentor is capable of limited explanation. Clicking on a mentor’s reaction displays a collection of salient quotes from that persona’s text. These quotes are generated by back-pointers to the text stored by the person model. The presentation of the quotes is rank-ordered by saliency and relevance. Quotes which make the largest contribution or best exemplify the resulting affective reaction are promoted to the top of the explanation page. Triggered quotes offer indirect and partial justification for a persona’s reaction because the context of the quotes will not perfectly match the context of the input. It is up to users to verify from the indirect explanation whether or not an affective reaction is indeed justified. This lends the interface some fail-softness, as a user can recover if the system erroneously represents a person’s reaction.

§

Task-based evaluation. A user study was conducted to test the hypothesis that WWTT can help someone grasp the personalities and opinions of a panel of strangers more quickly and deeply than with baseline methods. The subjects were 36 college students, formed into three comparable test groups. The examiner used WWTT to acquire attitude models for four individuals from their extensive weblogs (at least two years of regular blogging). The study was posed as a “game” and the objective is for each subject to do their best to answer questions regarding the general personalities and specific attitudes of the panel of four strangers, who are previously unbeknownst to the subjects. They are allotted 20 minutes to answer 15 questions. Subjects in Group 1 were allowed to read through the weblogs of the four individuals with a basic text browser as their information interface. Group 2 used a textual version of WWTT. Group 3 used the real WWTT interface. The Group 1 baseline represents how a user typically learns about the attitudes of people without the assistance of technology. The Group 2 baseline provides a keyword-retrievable textual memory and uses the WWTT interface, controlling for all the non-affective elements of the application.

The textual version of WWTT warrants some further description. Textual WWTT does not use the color dimension (for expression of approval/disapproval) or the focus dimension (for expression of dominance/submission). Only arousal is expressed (through brightness). Concepts are extracted using identical natural language processing mechanisms as in the real WWTT. Whereas arousal is affective in the real WWTT, arousal in the textual version is set proportional to the number of textual memories triggered by the new textual episode. For example, suppose a new textual episode contained the concepts X and Y. X occurs 10 times in the personal texts, and Y occurs 19 times. Thus the total number of textual instances is 29, and the extent of the arousal reaction is proportional to this score. In the explanation mechanism of Textual WWTT, the quotes are rank-ordered to promote quotes which have the greatest number of concepts in common with the new textual episode being evaluated.

The 15 test questions, given in random orders, fall into three categories of knowledge: general personality traits, specific attitudes explicitly contained in the weblogs, and specific attitudes not contained in, but implied by the weblogs. The examiner is careful to ensure one clear answer for each multiple choice question. Answers to questions on implied attitudes not explicit in the weblogs were verified with the relevant individual. Questions on personality traits are of the vein, “who is the most shy?” Questions on explicit attitudes (e.g. “how does Sally feel about religion?”) are designed to test the information retrieval capabilities of each tested interface. Questions on implicit attitudes tell of how well each tested interface enables its user to project how a panelist might react to something

novel, e.g. “what would Sally think of Jim, given the bio on his web page?” The test results are summarized in Figure 5-7.

The results were promising. Group 3 consistently outperformed Group 2, came close to Group 1 in “personality traits,” and clearly outperformed Group 1 in “explicit attitudes.” All three groups struggled with “implicit attitudes” and performed comparably. It was observed that on average, Group 3 subjects spent less time answering each question than Group 1 and Group 2 subjects, and also had to make fewer last-minute guesses on unanswered questions than subjects in the other two groups.

Subjects in Group 1 reported that it felt easy to build an overall picture of a person by skimming an extended sample of their writing as in a weblog. However, searching for “explicit attitudes” handicapped subjects in Group 1, who had to use the search feature in the text editor, but could not query all four panelists in parallel as the WWTT interface enables. At first, subjects in Groups 2 and 3 struggled to come up with text to pose to the application. Many people in both groups came up with a surprisingly efficient strategy of passing in a string of keywords which would define the linguistic context probable to contain the information they wanted. For instance, to answer the question, “who loves to party the most?” a subject in Group 3 typed something like, “party clubbing booze drinking drinks threw up” and then clicked on each face, reading salient quotes in the explanation to verify the attitude. When faced with a choice of which face to click first, subjects in Group 3 usually clicked the one showing the highest arousal. Subjects in Group 3 spent less time sifting through explanation quotes to find a satisfactory answer than subjects in Group 2, suggesting that affective saliency is a useful way to order quotes.

The results of this study support the idea that WWTT allows a user to more quickly and deeply grasp the personalities and specific attitudes of a panel of strangers than either of two baseline approaches. The results suggest that an affective memory can in many cases be a more useful way of organizing and presenting information than a purely textual memory. Despite the poor precision of attitude prediction as suggested in the evaluation of the underlying attitudes model, WWTT’s fail-soft explanation mechanism bolstered the usefulness of the attitudes prediction. This study does not directly examine the usefulness of real-time feedback to a user engaged in some task, which would require a study of longer-term interactions.

5.5 Foraging for food with the family

What’s for dinner? *Synesthetic Cookbook* answers this question by a RATE processing of a user’s weblog or of their stream-of-

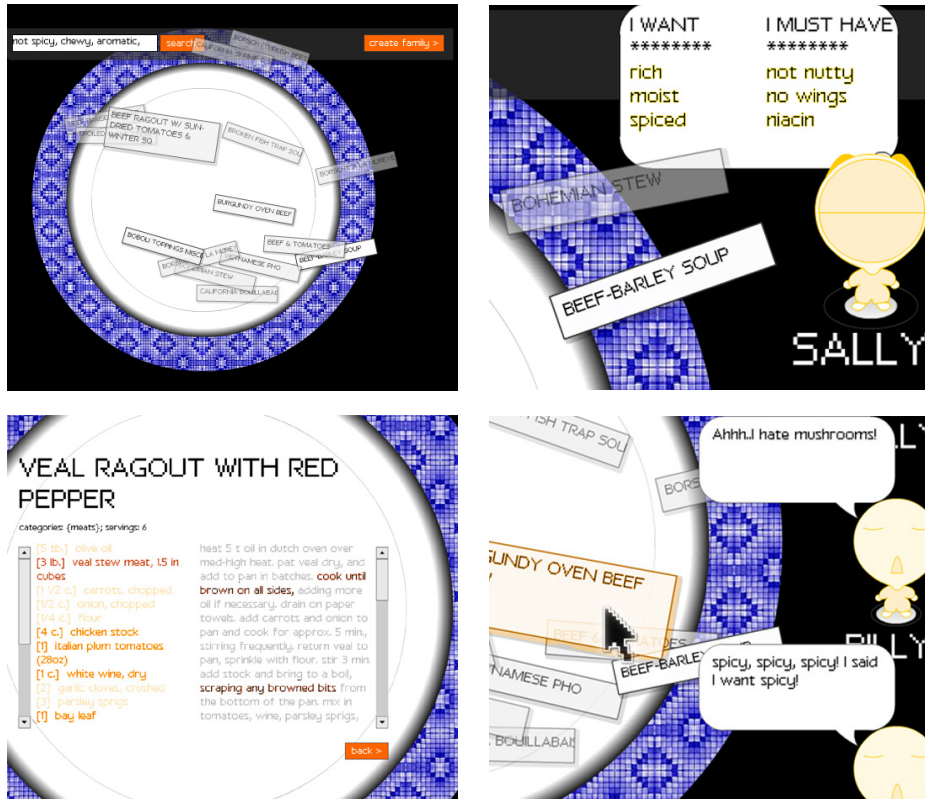


Figure 5-8. Interactions with the Synesthetic Cookbook

consciousness search terms. The cookbook models the tastebuds of family members so that as the user browses for dinner recipes, on-screen avatars anticipate their reactions, thus enriching the brainstorm with just-in-time family's feedback. Chapter 4 described the constitution of the food fabric and how a tastebud model could be acquired from a user's weblog or from a history of search terms. The rest of this section describes the interface of and interactions afforded by the visual recipe browser.

§

Interface. Synesthetic Cookbook is a flash-based application whose interface consists of a stream-of-consciousness search box, and to the right, a panel of three avatars, each instilled with the tastebuds of a loved one. As the food 'forager' types keywords (Figure 5-8 upper-left panel), evoked recipe suggestions from the underlying database of 60,000 recipes continuously fade in and out on little scraps of paper laid over a plate; the opacity of each suggestion scrap signifies how well it resonates with the search terms. The recipe suggestions fade in and out automatically, as the stream-of-consciousness search progress; hence, the information 'pull' of a traditional search is turned into a soft 'push' of suggestions; and the whole process is

more amenable to interactive refinement—for example, a forager sees some beef recipes, decides against beef, and so simply adds “no beef” or “not hearty” as the next term; all done without switching contexts. A wide variety of search terms are accepted—5000 ingredient keywords, 1000 sensorial keywords (e.g. ‘spicy’, ‘chewy’, ‘silky’, ‘colorful’), 400 cooking procedures, and 400 nutritional keywords—plus all negations of these terms. The goal is to handle all the possible descriptors under the sun—letting users articulate their cravings in terms of however it is that they imagine a dish to taste, e.g. ‘rich’, ‘spicy’, ‘homey’, or feel, e.g. ‘primal’, ‘misunderstood’, etc. Hence, a *synesthetic* cookbook. Finally, when the forager clicks on a recipe suggestion (Figure 5-8 lower-left panel), the recipe text is rendered with semantic highlighting such that the *essence* of their query is intelligible at-a-glance – e.g. forager searched for spicy, and now, all the spicy ingredients are highlighted in this recipe view.

Reactive tastebuds. By clicking the “create family” button in Figure 5-8’s upper-left panel, a panel of three avatars are shown. They can be named, and each avatar’s tastebuds can either be loaded with a family member’s stored food model (discussed in Chapter 4), or double-clicking on an avatar will allow tastebuds to be programmed directly. Figure 5-8’s upper-right panel shows the contents of Sally’s tastebuds. The “I want” column represents weak preferences and dislikes. The “I must have” column represents necessities and allergies, which are always strictly enforced. In the lower-right panel, the avatars can be seen emoting reactions to a recipe suggestion that the forager is mousing over. If an avatar is unhappy, she will speak her qualms. Each avatar’s tastebuds can also be populated by family members remotely. Additionally, a library of tastebuds can be stored for family members, celebrity chefs, whatnot, and by entering a name into the search bar, the forager will receive recipe suggestions selected on the basis of that person’s tastebuds.

5.6 A jocular companion

Everyday life is full of wrenches—frustrations at work, tensions with family, and anger with national political developments. Catharses is a desktop companion that seizes upon everyday gripes as opportunities for humor. Catharses tracks the humor model of an individual—her patterns of psychic tension—looking both at a prior model and a model with respect to the context of each day’s moods. Observing a computer user’s textual activities, such as writing emails, instant messaging conversations, and composing weblog entries, catharses delivers jokes to the desktop, apropos to the user’s recent frustrations. It attempts to nip tensions at their bud before they are repressed and turned into psychic baggage.

Catharses knows about 10,000 jokes, classified under twenty genres. Unlike JAPE (Binsted & Ritchie 1994), Catharses does not create new jokes, it simply identifies the user’s humor model with one of the twenty niche humors, and delivers jokes from that niche.

Catharses observes the user's global textual activities on the computer desktop—including in instant messaging software, word processing, email applications, and in web forms. It performs RATE processing on a window of the user's most recently typed text, and compares this just-in-context model to a prior model of the user's humor model. The prior model is updated periodically on the basis of a user's blog, but not everyday. If the just-in-context model captures tensions about topics which are already known to be tense under the prior model, then Catharses will deliver a joke to poke fun at this subject matter. Jokes are delivered as standard windows notification balloons which pop-up from Catharses's system tray icon. They may be dismissed, or they will auto-dismiss after a set period of time.

6 Philosophical underpinnings

Given that the thesis work claims to model cultural tastes, attitudes, and projections of persons into other aesthetical realms, it would be reckless and irresponsible to not qualify this feat with discussion of underlying philosophical and design assumptions. The sections in this chapter make a first attempt to articulate a philosophical rationale to support major suppositions that were made by person modeling. The reader is offered the caveat that the following ideas are not fully matured or well-defended—rather, they could more accurately be regarded as stubs and margin notes.

Section 6.1 explains how the thesis's approach to person modeling is consistent with Pierre Bourdieu's (1984; 1993) study of taste and culture. Section 6.2 explains that modeling persons as projections onto aesthetical realms rests on a *post-structural* understanding of aesthetics. Section 6.3 admits that generalizing a person from fragmentary impressions of everyday texts represents an *aesthetic consistency hypothesis*. Person modeling assumes that personal judgment is stable rather than ephemeral or fast-changing—stability is addressed in Section 6.4. Finally, generalized person models are admittedly stereotypical, and simplistic simulation does not do justice to how persons actually evaluate new situations; Section 6.4 describes how persons exceed their generalized models via dialogic dynamism, and suggests this as an avenue for future pursuit.

6.1 Bourdieu's framework for taste

Bourdieu's field-habitus-doxa model of taste suggests a theoretical basis for representing and computing taste. The layman's understanding of 'point-of-view' as a spatial metaphor is not only simple, but also powerful enough to entail a computational framework. In *Distinction* (1984)—a monograph on the social and cultural basis of taste—cultural sociologist Pierre Bourdieu theorized an elaboration of the 'point-of-view' metaphor, and appropriated it to explain differences in taste between social classes in French society. Although the self-proclaimed significance of *Distinction* is to critique the role of capitals and taste in reproducing France's social hierarchies, the work's semiotically suggestive aspects are most interesting and relevant to our present investigation.

Among the work's key words are three most relevant ones—*habitus*, *field*, and *doxa*. Bourdieu implicates an individual's faculty of taste judgment as being structured by set of personal dispositions forming a *habitus* (from habit). A habitus is constituted as instinctive patterns of consumption over social and cultural *fields* of goods. The acknowledged intersection of the personal habitus and cultural field is in turn the *doxa*—doxa is the site of an individual's self-assessed identity and location within the social and cultural field. Additionally, there may be other unconscious ways in which habitus aligns with field, such as via the *class habitus*—the predispositions inherited from one's social class, which forms a backdrop to one's own idiosyncratic habitus.

Habitus, field, and doxa, I suggest, quite resembles the generalized person model, the cultural topology, and a person's textual traces, respectively. Like the field, the cultural topology defines the limits of what is possible—be it a space of cultural goods, a space of possible perceptions, or a space of all possible attitudes toward political topics—and captures usual cultural behaviors. Fields considered in *Distinction* include the artistic field, the field of political opinions, and the field of life-styles. Like doxa, a person's textual traces constitute a self-report representing how the person admits to fitting into culture. Like habitus, a person's generalized model captures her unconscious and ineffable potentialities. We idealized both the habitus and the generalized model as being generative and capable of reacting to unlimited stimuli.

One of the most practical consequences of Bourdieu's framework for taste computation is the finding that taste judgments and social distance can be measured as distance between locations in the abstract, semantic space of the field. The analyses presented in *Distinction* were based on statistical data (Bourdieu 1984, 525-545) that Bourdieu had compiled from a lifestyle survey of 1200 or so French residents, conducted in the 1960s. Plotting that data as a cloud of points in n-dimension feature-space, Bourdieu used first-order statistics to determine *axes of inertia*. Normalizing the clouds of points along these axes, Bourdieu produced two-dimensional maps

of taste-space—such as the variants of petit-bourgeois taste (*ibid.*, 340), the space of life-styles (*ibid.*, 129), the space of food relative to economic and cultural capitals (*ibid.*, 186), and the political space (*ibid.*, 452). In Bourdieu’s diagrams, cultural and economic capital constituted primary organizing dimensions of taste. A person could be located in the various taste spaces by the amount of capitals he possessed, and based on this location, several predictions could be made relevant to our task by reading distances in the chart—1) what goods he is likely to consume, i.e. the goods at his location; 2) what goods he is likely to fancy or disdain, i.e. goods located just upstream or downstream of his location in economic capital; and 3) which other folks is he likely to share taste with, i.e. other folks overlapping his position in cultural and economic capitals.

Consistent with Bourdieu (1984; 1993) and Montgomery’s (1994) spatial conceptions of culture, this thesis used various distance metrics to simulate persons’ reactions. For example, a person is predicted to like cultural interests that are proximal to his person model in the taste fabric, with distance being measured by spreading activation. In the realm of attitudes, a person’s attitudes toward unknown topics are predicted based on their nearness to a person’s explicitly stated attitudes, with distanced measured by spreading activation over topic hierarchies and analogy.

6.2 Post-structural aesthetics

The thesis’s approach to point-of-view and aesthetic subjectivity presupposes a *social constructionist* stance upon the ideas of identity and aesthetics. The notion that a person’s aesthetic perspective is greatly shaped by their social and cultural milieus was not always a foregone conclusion in the history of Western intellectualism. From the Enlightenment up through post-structuralism, general understanding of aesthetics has shifted from pure and objective, to socially constructed and relative. A more recent event—the hyper-globalization and commodification of culture—again shifted the fostering of personal aesthetic perspective from few subjective experiences, to countless consumptive choices. The idea that a person’s tastes, attitudes, etc. can be composed out of selections from a space of cultural and aesthetic possibilities mirrors the post-structural shift away from paradigm toward increasing reliance on syntagm. Notwithstanding, describing individuals by cultural location is just a stereotype, albeit a specific one. An individual, located in cultural space, is still able to dictate how she wishes to internalize and appropriate the cultural goods available to her location. The rest of this section first narrates the two post-structural shifts of aesthetics—relativization and syntagmization; second, it portrays the interplay between the exteriority of location in culture, and the interiority of aesthetic perspective.

§

Shift #1 – aesthetic relativization. Immanuel Kant (1790), last of the Enlightenment thinkers, insisted that pure, normative, and objective aesthetic judgment was not only possible, but that all deviations from that aesthetic was a result of corruption and poverty. Kant's aesthetic formalism since dominated nineteenth century Europe, and had caused the idea of aesthetics to be invoked solely in reference to high art—as the “pure” practice *par excellence*. The Enlightenment's idealization of aesthetics and rational thinking indulged in a neo-Platonist metaphysics of presence and created a logocentric illusion that was nonetheless propped up by a rigidly structured society—the condemnation of popular aesthetics as vulgar became one of the bourgeoisie's most favored instruments of cultural hegemony. The thinking that aesthetic judgment could be pure is a fallacy pertaining to manifolds. Just as persons very localized on the Earth perceive it to be flat, so occupying a very narrow location in the realm of aesthetic possibilities affords the illusion that aesthetics is objective. Both the Earth, and cultural fields of possibility, are manifolds. In a sense, it was quite natural for persons of pre-globalised hierarchical societies to commit the objective fallacy because they were vested into a stable but quite narrow portion of the aesthetic spectrum. To quote the unknown poet, through Seneca, “the part of life we really live is small.”

The post-structural shift of aesthetics away from objective and formal, toward subjective and relative actually saw its seeds sown during Kant's time. Alexander Baumgarten coined the word ‘aesthetics’ in his *Reflections on Poetry* (1735), to mean aesthetics quite differently than Kant's more dominant interpretation. A philosophical descendent of Leibniz, Baumgarten was interested in the psychological value of art—he posed aesthetics as the sensorial and imaginative experience of poetry, and he conceived of the ideas in poetry as “clear and confused” as opposed to rational ideas which were “clear and distinct.” Psychological aesthetics, seeded in Baumgarten, and later advocated by Freud and Dewey, was not concerned with purity or objectivity—a thing was aesthetic only insofar as it could affect a person. For Freud (1919), the aesthetic was intimate and narcissistic-- we are affected most by things that mask and foil aspects of ourselves, our memories, and our desires—in other words, Freudian aesthetics is concerned with artwork-as-mirror. Dewey (1934) regarded aesthetic as a subjective experience between self and artwork—an artwork has the character ‘aesthetic’ if it can seduce viewers into a state of sensitive perception and vulnerability. While Freud and Dewey relativized aesthetics by illuminating subjective experience, others such as Adorno (1970) and Bourdieu (1984) de-objectivized aesthetics by point out its etiology in social and cultural conditions.

Shift #2—the syntagmization of aesthetics. In Europe of the Enlightenment, before hyper-globalization, media culture, and consumer culture, individuals cathected their identity around a few key events that were experienced deeply. Simmel (1971) described this as ‘subjective culture’—identities were constructed around

major events such as one's job, church membership, and social class. Relatively few aesthetic materials were known to the common sense of the entire cultural population. In this time, aesthetics was very strongly paradigmatic—each material, such as holding the job of a policeman, held tremendous implications for an individual's identity. Since the onset of media culture and consumer culture, a much greater range of cultural commodities have been available for individuals to construct identities out of. It is as if one's identity could be painted with increasingly finer strokes. Consequently, the descriptive strength of any given paradigm, such as holding the job of policeman, has been greatly diminished. Simmel, a Romantic, noted this fragmentation of identity across many small shards of description, and believed that it would cause the 'form' of an identity to more closely approach the authenticity of an individual. He regarded favorably the fracturing of culture into highly specific life-style niches, and termed this new mode of identification, 'collective individuality' (*ibid.*).

Erstwhile's few strong paradigms were in consumer culture replaced by countless micro-paradigms—consumerist interests such as books, music, film, foods, and subscription to various niche cultures such as the indie rock or goth movements. With market forces able to seek out, fulfill, and even create desire, the present society's space of cultural goods is vast and diverse, approaching Plato's vision of plenitude (McCracken 2006). Whereas a strong paradigm was rather an unambiguous signal of identity, micro-paradigms could be appropriated in many different ways, so the context in which goods are invoked has become increasingly important. In other words, the syntagmatic combination of materials has replaced paradigm as the dominant mode for individual aesthetic identification. In the literature, Bourdieu (1984) presupposed a 'consumer aesthetics', and Haug (1986), a 'commodity aesthetics'. Insofar as aesthetics is now described by phenomenal consumption, rather than prescribed by transcendental conference, Sack terms these the 'network aesthetics'. These new aesthetics now refer more to individual taste and life-style than to art—'the aesthetic' has shifted from artwork's intension to individual's intention. By afford individuals fine-grained and syntagmatic specification of their aesthetic perspectives, post-structural aesthetics harkens to Baumgarten's pronouncement of poetry as 'clear and confused', in opposition to paradigm's 'clear and distinct' nature.²⁰

²⁰ It is suggested that syntagmatic expression is the aesthetic mode of our present social and cultural reality. The validity of the present viewpoint computation is seated in this presupposition. However, this post-structural aesthetics is not more 'final' or closer to 'truth' than aesthetics of antiquity or of the Enlightenment. Nietzsche called our times the 'age of comparison', yet he presaged "a posterity that knows it has transcended both the completed original folk cultures, as well as the culture of comparison, but that looks back on both kinds of culture as on venerable antiquities, with gratitude." [HumanAllTooHuman, p.30, aph 23, transl: marion faber w/ Stephen lehmann]

§

Interiority, exteriority. The mere availability of the plenitude of cultural goods out of which a self can be constructed does not necessarily entail that all persons can afford to seize upon these materials. In Bourdieu's (1984) Marxian framework, the accessibility of a field of cultural possibilities to a person is dictated by that person's possession of various capitals—economic capital, cultural capital, education capital, etc. Due to the constraint of capital, a person's exterior location as a pattern of consumption over the field of cultural goods cannot be regarded as a necessarily accurate reflection of a person's true taste-dispositions or *habitus*. Hence, in Bourdieu's equation of social physics, “,” it is habitus and capital together which expresses as location.

Another discrepancy between an individual's aesthetic interiority and exterior behavior arises from the ambiguous appropriation of cultural goods within more specific semantic fields. Each cultural good from the overall social field can hold different and specific meanings when placed into the context of more specific fields to which an individual may belong to, such as the artistic field. Some specific fields may correspond tightly to the overall social field, while others may be rather autonomous; Bourdieu conceived of the autonomy of a specific field in terms of 'refraction', while Luhmann (1984), in his autopoietic systems theory, conceived of a specific field's relationship to the overall social world as 'resonant'.

For all its computational complexity, this thesis had to make some idealizing assumptions about the relationship between interior self and exterior description. It was assumed that an individual's exterior preferences, as those captured in a weblog diary or social network profile, were more or less reflections of the individual's interior and authentic aesthetic perspective. The constraints of capital were not considered. Distortion resulting from ambiguity in appropriation of cultural goods was assumed to diminish as the number of cultural goods involved in description of the individual increased. That is to say, the greater the number of data points in a pointillistic portrait of the individual, the better the fault tolerance for errant points.

6.3 The aesthetic consistency hypothesis

The aesthetic consistency hypothesis can be stated as making two complementary claims—1) semantic mediators effect aesthetic consistencies in culture; and 2) the principle of economy and Diderot unity effects aesthetic consistency within an individual habitus. Aesthetic consistency has been invoked across the various stages of person modeling as a central hypothesis upholding the work. This section first reflects upon semantic mediators of aesthetic consistency in social and cultural milieus. Second, aesthetic consistency within

the individual's habitus, and the application of semantic relaxation to create habitus out of location, is reviewed.

§

Aesthetic consistency in cultural milieu. As shown in Figure 2-2, a major organizational difference between different knowledge representations for person modeling is the varying level of consistency supported by each representation. Semantic sheets and semantic fabrics support patchwork consistency, while dimensional spaces imply full consistency. The full consistency of dimensional spaces such as Jungian space of perception is theoretically constructed. The patchwork consistency of semantic sheets and semantic fabrics is however more interesting because there are phenomenal semantic mediators which seemingly effect local consistencies between items in the space.

The space of cultural taste, constituted as a semantic fabric of consumerist interest nodes, is aesthetically consistent around *taste-cliques*, and *subcultural identity hubs*²¹. A subcultural identity acts as a point of convergence for taste, because most subcultures entail their own systems of life-style. Just some of the subcultures, excerpted from Wikipedia's list of subcultures²²: Anarchists, Aristasia, Bohemianism, Hacker, New Age, Vegans, Swingers, Emo, Hip-hop. In addition, subcultural identities and movements may arise around a single consumerist interest, e.g. Trekkie fandom, book lovers, etc. Once a movement has become identified in the media as a subculture, it gains a new cultural prominence, and becomes self-perpetuating. Cycles of promotion for the subculture's life-style entailments are begun in the media, creating demand for other consumer goods subsumed into this new aesthetic cathexis, and spurring the cultural production of new consumer goods consciously fashioned according to the prescribed life-style. In the taste fabric that was mined from the texts of 100,000 social network profiles, subcultural identity descriptors occurred 18 times more frequently than other interest descriptors. For example, a person would list 2 identity descriptors and 36 interest descriptors in a profile. Consequently, identity descriptor nodes behave as hubs in the resultant fabric, bringing together many spoke interests. Similarly, taste cliques are n-cliques of strongly interconnected nodes, which together as a cluster, function as unnamed identity hubs. Furthermore, Bourdieu (1984, ch. 2) identified certain key socio-economic variables and capitals as central, from which other variables and capitals derived. For example, that 'scholastic age' is a transformation of 'inherited cultural capital' implies that the latter has imposed a consistency across capitals.

²¹ Hubs are known as 'attractors' in the dynamical sociological systems literature.

²² http://en.wikipedia.org/wiki/List_of_subcultures

Aesthetic consistency in social milieu. In the space of attitudes, attitude mentors, which Minsky termed *imprimers*, act as semantic mediators. An individual's attitudes are often structured around the attitudes of exemplars in their life, whom they imitate in some domains. Minsky called these exemplars 'imprimers', and proposed that they could be identified on the characteristic that imprimers can make an individual feel self-conscious emotions such as pride and shame. An imprinter can be a real person, a fictive person, or even a cult idea. These imprimers—strong personality archetypes such as a Donald Trump for finance, one's grandfather for values, or the ideal Republican archetype, shape individual's attitudes, thus effecting consistencies and alignments between individuals' attitudes, and do for social milieu what subcultural identities do for cultural milieu. For example, all individuals imprinted by the ideal Democrat archetype will mutually align within the political subspace of topics. In the What Would They Think? system, it was shown in an evaluation that the inclusion of imprimers in attitude simulation improved prediction accuracy.

§

Viewpoint as intertextuality between archetypes. In the psychoanalytic semiotics of Kristeva (1986), echoed in Silverman (1983), echoed with respect to cultural diaspora in Bhabha (1994) and with respect to literary influence Bloom (1982), every particular textual entity—be it a novel or a personal viewpoint—is shaped through the unconscious influence of many texts—be they a predecessor's novel, or a cultural archetype such as subcultural identity and imprimers from the aforementioned. Thus an individual's aesthetic agency lies in the syntagma—the intertextuality between archetypes. For example, Silverman (1983, 217) explains intertextuality in Flaubert's *Madam Bovary*: "Thus Emma Bovary reconfirms her masochistic subjectivity with each repetition of the romantic scenario... They do so transparently, without any consciousness that the images and narratives with which they identify are historically and culturally specific." Even the possession of cultural goods evokes unconscious cultural thematics—"the hierarchies and classifications inscribed in objects (especially cultural products)" (Bourdieu 1984, 470-1).

Consequent to viewpoint's intertextuality, the aesthetic complexity of the individual habitus should be proportional to the number of possessed archetypes, which is a number far smaller than the number of possessed cultural descriptors. Rather, cultural descriptors should be aesthetic consistent along the lines of possessed archetypes.

Economy and aesthetic closure. A more idealized assessment would suppose that not only is a viewpoint structured around archetypes, but that these archetypes are melded together to shed dissonance such that the total aesthetical system of an individual is

self-consistent. Two reasons to believe in total coherence are the principle of economy and the Diderot effect. Aesthetics, ultimately pertains to feelings. The accommodation of dissonant feelings is ultimately harder than the accommodation of inconsistent facts, because affect's physiological seat insists on coherency of mood, and economy of affect. Thusly, tensions between dissonant aesthetics should quite urgently demand relaxation and resolution. Had one's aesthetical system more resembled a rote memory than an intuitive apparatus, taste judgments would not be so immediate or freely given. The tendency of an individual consumer toward aesthetic consistency has also been theorized by McCracken (1988) as the Diderot effect. A consumer, upon acquiring a cool new possession, will tend to re-evaluate other possessions constituting his life-style, considering their replacement in order to restore aesthetic consonance with the cool new possession.

Semantic relaxation. Aesthetic consistency's implication for person modeling is that it affords generalizations that magnify the predictive power of models in judgment simulation. On the supposition that an individual's generalized model will tend toward aesthetic consistency, this thesis performed a procedure termed *semantic relaxation*--to transform the rather jagged outlines of an individual's cultural location, into the more fluid generalized model. In the realm of cultural tastes, semantic relaxation proceeds from a cultural fabric--cultural materials, interrelated by numerical aesthetic affinities. Supposing that A and B were strongly connected, and knowing that an individual had expressed preference for A in his everyday texts, then semantic relaxation assumes, with some discount, that B should also be preferred, as that is the general tendency of cultural participants. Via semantic relaxation, an individual's dissonances are relaxed into a more fluid cloud or neighborhood of points, called an *ethos*.

6.4 The stability of viewpoint

Viewpoint is admittedly a useful abstraction for pondering the systematicity of a person's judgments, and it is hoped that what a generalized person model captures is viewpoint, but how stable can a viewpoint be? Consider the four-to-five week test-retest reliability statistics for the Meyers-Briggs Type Indicator inventory of personality, recapitulated here. Recall that the model of perception computed for this thesis shares two of MBTI's four scales--SN and TF. Myers and McCaulley (1985) surveyed continuous score correlations from ten studies for the four-to-five-week test-retest interval. They found reliability coefficients of .77 to .93 for EI, .78 to .92 for SN, .56 to .91 for TF, and .63 to .89 for JP. Assuming a roughly binomial distribution for these scores, cursory median reliabilities of EI .84, SN .85, TF .73, JP .78. could be estimated. The implication of this result is that in the course of four to five weeks, a Sensing type person crosses the threshold into being an Intuiting type person 15% of the time, and a Thinking type person crosses the threshold into being a Feeling type person 27% of the time. Retest reliabilities

could have been more informative if the percent spread between a person's location along the four continuous MBTI scales were assessed instead, as it would shed light as to what fraction of persons who retested differently experienced major shifts, versus minor adjustments which happen to cross the midpoint. Nonetheless, the results do illuminate a time-granularity—the stability of perceptual dispositions needs to be considered on the order of weeks to months. Methodologically, this knowledge would prescribe that perception models be recomputed at least once a month, if possible. The time-stability of person models within other aesthetic realms—attitudes, cultural taste, taste for food, and sense of humor—are other matters entirely. The rest of this section will raise three points about the stability of viewpoint. Firstly, the stability of viewpoint has largely to do with the investment traps of each realm. Secondly, the monosemic action of statistical reading will naturally produce stabilizations and expose instabilities. Finally, it is argued that, with respect to utility of the produced models, time-stability of viewpoint is always secondary to viewpoint's cross-context stability.

§

Investment traps stabilize viewpoint. A viewpoint will be slow to change if much has been invested in its formation, and if its revision is costly. Each of the aesthetical realms is associated with unique investment traps and costs of revision. Perception models are considered to be least stable by these criteria, because how we are disposed to gather (sense or intuit) and process (feel or think) information could vary quite freely depending on one's present life context. An individual may in fact possess many personas—one around a lover, one at work, one with friends—and may exercise different perceptual dispositions for each persona. If perceptual disposition can be tied to persona, then its stability should vary according to number and turnover rate of personas. Since the gathering and processing of information is such a private quality, it is likely that this viewpoint is as fluid as a person's psychology wants to make of it. Individuals are not likely to be held directly accountable by social alters for maintaining consistent perceptual disposition; whereas one's friends, family, and lovers do tend to hold an individual accountable for maintaining consistent attitudes. Social milieu is an investment trap and stabilizing force for attitudes. Humor model is not likely to turnover quickly because they are seated in formative life experiences and psychic tensions. Taste buds are likewise slow changing because food preferences are built out of deeply seated life experiences and physiological habituation.

The integrity of one's cultural taste is seemingly stabilized by its reification in one's possessions, and by the stability of one's capital. Everyday possessions, such as symbolic environment of one's habitat, and the habitual activities of everyday life-style, both *echo* and *reinforce* a person's cultural taste identity (Csikszentmihalyi & Rochberg-Halton 1981). Unless an individual were in possession of

unlimited capitals, it would be costly in terms of money and time to constantly replace one's possessions. It would also be disorienting to constantly change our everyday activities. The concreteness and slow turnover of our symbolic environment and habitual lifestyle reinforces into us on a daily basis the image of who we are, our past and recent taste acts—they effect stability in cultural taste. When an individual is fed up and requires a change in taste, it is often a revolutionary change, because the Diderot Effect describes that the aesthetic of each new possession will want to be propagated across the board. Furthermore, the capitals which an individual possesses are fundamentally slow changing, e.g. it takes several years to upgrade one's educational capital, and since lack of sufficient capitals restricts one's access to parts of the field, the relative stability of capital bounds the stability of cultural taste.

§

Monosemization and statistical stabilization. The stability of an individual's viewpoint can be distinguished as being actual or perceived. Just as something can be convincingly aestheticized, a reader can herewith read stability and coherence into a writer's viewpoint. In Greimas's terminology, the tendency for readers to find semantic unification and coherence from a text is called *monosemization*—a reading strategy that is not far from aestheticization since it is rooted in the Romantic hermeneutics of Schleiermacher (1809). As described in Chapter 3, the RATE readers deployed in the model acquisition systems implemented the action of monosemization—the identify ambiguous lexemes in the text, and use the macroscopic textual context to resolve these into semes and classemes. The final pattern of inter-consistent semes and classemes resulting from a reading is the isotopy, and it is stable. Or to be more precise, it is a readerly *stabilization* of the writer's viewpoint.

Stabilization is accomplished in the final, statistical estimation phase of RATE processing, such as by taking the first-order moment of the many instances of a PAD score or significance scores identified in the text. First-order statistics and other information-theoretic computations, almost by definition, means to find underlying stability; or if there is no stability, such approaches can conclude the null hypothesis. For example, Eq. 4.3 implements a reinforcement memory used in reading for attitudes. If a writer expresses a consistent emotive valence regarding topic X across many instances in the text, Eq. 4.3 computes a stable attitude; however, if the emotive valences are inconsistent across many instances in the text, Eq. 4.3 will conclude that no stable attitude exists for topic X.

The results of statistical stabilization also vary greatly depending upon the everyday texts inputted to the RATE reader. For persons constantly becoming, viewpoint constantly evolves, yet as a manifold, viewpoint and the sense of self that it affords can feel quite constant from day to day. A model of an individual's viewpoint

produced from texts spanning years will likely be more diluted than a model produced over the span of months—diluted because there is bound to be more variance and ‘drift’ in the tastes, attitudes, and perceptual dispositions expressed by a writer over the longer span.

Ultimately though, machine readers performing statistical stabilization will not approach the skill of human monosemization—we as well-oiled context maestros are far more capable of determining which of textual particulars are salient and reliable—what in artificial intelligence is called ‘credit assignment’, or what Rosenblatt (1978) has called ‘selective attention’. Human readers will more than machine readers, know when to re-understand passages in order to resolve anomalies and find satisfactory aesthetic closure—what Wolfgang Iser (1979) has called ‘consistency-building’.

§

Cross-context stability. In regards to utility of the produced person models, time-stability should be regarded as secondary to cross-context stability. Rather than falsely hoping to capture the time-invariant essence of Marvin Minsky—a quixotic error owed to those who insist on reading being into becoming, diachronic as synchronic—it would be far more useful to capture how Minsky, at any given period in his life, judges and interacts with the phenomenal circumstances of his life. To see a person as a shadow pattern of involvement in worldly materials may be a rather Heideggerian project. The power of a viewpoint lies in how an economical but amorphous aesthetic hearth of an individual can trickle outward to effect countless, highly specific judgments about materials in the world—judging consumer goods, topics, and interpreting information. To produce a model that can teach others about a viewpoint, the consequences of that viewpoint across many unrelated contexts must be recorded. The essential spirit of a person can only be felt by understanding how seemingly unrelated judgments from many disparate contexts are actually parallel emanations from a common source. The stability of aesthetic judgment across contexts is a central expectation of the aesthetic consistency hypothesis. Indeed, the invocation of semantic relaxation presumes stability and coherence of viewpoint across semantic contexts.

6.5 Dialogics

With human dynamicity, a quality most difficult to capture and simulate is the dialogic imagination. What is dialogics? It is creative, self-conscious, and strategic action taken just-in-context. A point-of-view model records a person’s priors—but to the credit of our kind, a person self-conscious of their priors can willfully overcome and contradict these priors opportunistically; it is through dialogics that persons explore new and unfamiliar positions, drifting and becoming through the process. The person modeling engaged

here has captured a declarative stereotype, and person-model simulation animates the stereotype in the most naïve and straightforward of ways. Dialogics though, is not declarative—it governs how the declarative stereotype is appropriated creatively and opportunistically in procedure. Dialogics can be measured as the deviation of practice from what one’s person model would otherwise predict. Having briefly introduced the idea, the rest of this section 1) discusses some of the ways in which dialogics is integral to, or alters a person model; and 2) reflects upon the account of dialogics in the present modeling, and ponders how it might be addressed in future work.

§

Self-reflexive dialogics. Reflexive image of one’s own aesthetics and location in social fields becomes a material and fuel for dialogics. A reflexive image is useful because an individual can see herself as others stereotype her to be. Because she knows that others expect her to make certain judgments and to act in a particular way, she can alter her judgments in order to violate their expectations—thereby releasing potential energy. Self-reflexive dialogics, then, is an exploitation and anticipation of others’ expectations; it is also a demonstration of self-awareness. In Bourdieu’s framework, an individual’s self-acknowledged, self-conscious location in social fields was termed the *doxa*. One’s *doxa* is what is subject to play and exploitation. Bourdieu gives an example of theorists in the literary field (Bourdieu 1993, 184)—theorists are keenly aware of their situation, of what others think of their position, and thus anticipate and deflect criticism by overcoming their own position before other theorists can. It is not unlike firemen using controlled burns to contain a ravaging fire. A person-model simulator would be hard pressed to know if and how a person might contradict or overcome his self-stereotype in response to a particular situation because each situation is fraught of politics, nuance, and opportunities.

Dialogical negotiation in social situations. Goffman wrote that social situations in everyday life are approached as dramatic performances. An individual entering into a new social situation, especially a heavily politicized situation, must negotiate her social position with respect to the other social participants, taking into account which social positions are available for a given situation. Greimas, in the Russian formalist tradition of conceiving narrative plots and situations, termed all the social roles available in a situation the *actants*, and a specific person filling an actant, an actor. Some of life’s social situations have explicit actants—for example, a workplace assigns employees job titles; in a wedding, the participants such as bride, groom, bridesmaid, best man, are explicit. Most social situations, however, only imply actants, and these usually comes in binary oppositions—e.g. hero and villain, insider and outsider, or more abstractly, dominant and submissive. When entering a social situation, an individual will have to negotiate or

fight for a desirable social position, this requiring a self-overcoming, a revision or adaptation of one's prior self-stereotype for the present social context. Ideology mirrors social negotiation, but on a larger scale—just as social participants adapt themselves in order to fulfill an actant, so does political ideology adapt their particular stances on issues in order to oppose and distinguish themselves from competing ideologies. With respect to computation, it would be quite difficult for a machine simulator to anticipate just exactly how a person model should modulate itself when operating in a new social context, especially since negotiation implies back-and-forth repartee between person and other social participants—which is nearly impossible to judiciously simulate between the virtuality of a person and the virtuality of envisaged social participants.

The aesthetic of contrarianism. Although dialogics is usually thought of as a just-in-context strategy, particular to each situation, dialogics can also manifest as a long-term aesthetic—for example, contrarians are those persons who are aesthetically disposed to contradicting and negating. Cultural contrarians, or counter-culturalists, form their viewpoints in polemic relation to cultural thematics—anarchists reject government, anti-fashionistas reject the ideas of fashion, the avant-garde reject the pace of the mainstream. But the power of their voice is dialogic—without government to speak against, anarchists lose their voice; without a mainstream mentality, the avant-garde label loses its sense. The aesthetic of contrarianism is often a long-term aesthetic and a defining quality of a person because contradicting culture produces psychological capital—lambasting culture creates a sense of freedom. Bourdieu (1984, 491) wrote, “The opposition between the tastes of nature and the tastes of freedom introduces a relationship which is that of the body to the soul, between those who are ‘only natural’ and those whose capacity to dominate their own biological nature affirms their legitimate claim to dominate social nature.” Because a contrarian's viewpoint should consistently oppose cultural norms or specific antagonists, it should be possible to recognize contrarians by comparing persons' viewpoints against culturally normative and other viewpoints and detecting negative correspondences.

Sarcasm and irony. Sarcasm and irony are dialogical in so far as their speech-act exploits expectation violation. An employee loathes his boss, but in one rant, sings the boss's praises to a co-worker—thereby committing sarcasm. Sarcasm and irony are forms of meta-speech, because in violating a receiver's expectation of one's viewpoint, it is the idea of violation which is important. By sarcastically praising his boss, the employee is actually communicating the opposite attitude, unequivocally. Sarcasm and irony can be simulated by a machine by simply adopting an opposite attitude, yet for the act to be understood as sarcasm, a perspectival artifact would have to communicate dramatic intention, and develop surrogates for the rich signals typically announcing sarcasm in face-to-face communication such as body language and prosody.

§

Account of dialogics in the present modeling. The built person models are stereotypes, and the person-model simulation cannot anticipate dialogics—at present, or perhaps ever, since so many unanticipatable contextual factors, such as a person’s mood, can cause a revision to the role of dialogics in a particular set of aesthetic judgments. Nonetheless, the omission of dialogics and political strategy which alters how prior a person model manifests just-in-context may actually be a blessing to applications like person-learning and tools for self-reflection. The virtual mentors of *What Would They Think?* which constantly react to the input of a user’s textual context can help the user to learn about each mentor’s perspective. This learning would arguably be impeded, not helped, if virtual mentors implemented dialogics—dialogical mentors would be politicized, would offer overly nuanced and self-contradictory judgments, and this would make it hard for a user, trying to pin down a coherent understanding. The *Identity Mirror* shows a viewer a stereotype of her cultural situation. That the reflection is in fact just a stereotype allows the viewer to contrast her sense-of-self with that stereotype, finding reflexive agency in realizing how they differ.

The simplistic simulation of a person model is rather predictable. Predictability is a positive trait for a computational artifact to possess because it makes it easier for an artifact user to understand and debug the mechanism and limitations of a stereotype. A central concern of interfacing and interacting with something as complex as a person’s perspective is that its rationale could potentially be opaque and unknowable.

Dialogics in future work. It should be possible to estimate the significance of dialogics to an individual in several ways. First, an individual’s dialogical tendency should lie proportional to the variance of his person model as he crosses social milieus. Second, once a stable person model was acquired, other of the individual’s texts occurring in social situations could be analyzed, and instances of sarcasm and irony might be detectable as violations or contradictions of the prior person model.

Ultimately, dialogics should be considered as one strategy out of many, which collectively form an ensemble of strategies for creative, just-in-context appropriation of priors. Another strategy is enrolment—a concept associated with actant networks. That idea is that novel arguments can be made by creatively re-purposing a network of prior Latourian actants²³ (Latour 2005) to give polyvocal support to one’s univocal argument. Together, this ensemble of creative appropriation strategies would constitute an individual’s

²³ being careful to not confuse the actant in Latour versus in Greimas, which are slightly different notions.

everyday poetics, a second-order person model which could operate over the first-order stereotyped person model as a material to be consumed and exploited. A foundation for the study and modeling of everyday poetics, strategies, and tactics was laid by Certeau (1984).

7 Conclusion

This thesis reported an end-to-end exploration of *person modeling*—from acquisition of persons’ textual traces from everyday texts; to generalization of those textual traces using semantic resources and mined cultural patterns; to application of person modeling to simulate a person’s perspective in tasks such as self-reflection, person learning, and deep customization.

This work contributes a methodology for extracting person models from personal texts to the user modeling literature. In category-based modeling, user profiles are typically acquired via solicitation and generalized using manually crafted stereotypes. In behavior-based modeling, users’ histories and *behavioral traces* are typically acquired from applications and generalized using statistical inference. In the modeling technique introduced here, text processing and mining enable the acquisition of *textual traces* from first-person everyday texts; these traces are then generalized using mined cultural patterns and compiled semantic resources. By focusing on a person’s everyday personal texts—such as their blogs, homepages, etc.—we hope to capture persons as they are in the social everyday, rather than as users of specific applications.

The person models thus acquired were modestly successful in several evaluations, and thusly they show promise as tools for simulating a person’s reactions to new experiences. In the cultural taste realm, the generalized model demonstrated an accuracy of 0.86 in a complete recommendation task, outperforming a baseline accuracy of 0.73 for an item-item collaborative filtering algorithm. When attitude modeling was applied to evaluate a political corpus, the results corresponded with those from a well-known study that used different metrics. Another evaluation for attitude modeling measured the deviation of predicted reactions from actual reactions to news stories; its results showed that prediction of arousal was statistically significant over both baselines, and predictions of pleasure and dominance were statistically significant over the randomized baseline. In the perception realm, an evaluation involving 3800 weblog diaries showed that bloggers’ disposition for thinking versus feeling could be distinguished with 0.62 accuracy, which was statistically significant over chance, and at about a midpoint between chance and an estimated upper bound of 0.73; the perception model was less successful at discerning blogger’s disposition for sensing versus feeling, and that result was not statistically significant over chance.

To the computational reading and information extraction literatures, this thesis contributes a reading method for excavating models of taste etc. from personal texts and a method for mining cultural corpora. The ‘reading for affective themes’ (RATE) technique, introduced in this work, can excavate summaries of a person’s tastes

and attitudes from their texts. By focusing on everyday texts (first-person and self-expressive) like weblog diaries and social network profiles, the difficult task of viewpoint attribution was obviated, and the text's emergent affective themes could be attributed to the writer. Building upon both knowledge-based and statistical approaches, RATE is an associative reader that overlays a commonsense-based topic skimmer with an affect skimmer which gists textual affect using a pleasure-arousal-dominance model. The affect skimmer builds upon a combination of previous word-based and event-based techniques in textual affect analysis. Using different reading schemas, RATE can be repurposed to read for the various realms' models. Evaluation of attitude modeling on a political corpus produced results consistent with intuition, but did contain some attribution errors that associative methods have a known vulnerability for, e.g. Democrats were attributed highly displeasurable attitudes toward topics such as 'god' and 'elderly', when really their displeasure was directed at the invocation of 'god' in the public sphere. The technique of 'culture mining' can acquire a cultural topology from a cultural text corpus, and was applied to 100,000 social network profiles to produce a taste fabric with emergent features such as identity hubs and taste cliques. The contribution of these topological features to improving generalization accuracy was validated in the cultural taste model evaluation—without identity nodes, accuracy deteriorated from 0.86 to 0.81, and further declined to 0.79 when taste cliques were also weakened.

Finally the ensemble of six implemented applications contributes a methodology for embedding perspective in the interface to the intelligent user interfaces literature. The paradigm of just-in-time information systems was applied to create mirrors and simulators that observed and reacted in real-time to user context. By designing desktop companions that afford just-in-time and just-in-context reactions, users are able to understand someone else's (or their own) perspective deeply by "walking through it." The virtual mentors application was evaluated in a 36-person user study—results showed with statistical significance that participants using virtual mentors learned more effectively about someone else's personality than participants using a text search interface; virtual mentor users were also shown to learn more effectively about someone else's explicitly stated attitudes than text search users and users who were allowed to read a weblog.

Personal style, point-of-view, and taste are such nuanced and ineffable matters that they will always evade monolithic representation. Without trying to represent (and reduce) the whole, the computational investigations reported in this thesis instead captured some of the shadows that a person casts against life's materials—cultural tastes, attitudes, ways of perceiving, taste for food, and sense of humor. The portrait of a person modeled under the ensemble of these realms preserves richness, and amounts to what Geertz called a "thick description."

8 References

- Aberer, K. *et al.*: 2004, Emergent semantics, *Proc. of 9th International Conference on Database Systems for Advanced Applications (DASFAA 2004)*, LNCS 2973, 25-38 Heidelberg.
- Adorno, T.: 1970/1998, *Aesthetic Theory*, University of Minnesota Press.
- Alm, C.O., D. Roth & R. Sproat: 2005, Emotions from text: machine learning for text-based emotion prediction, *Proceedings of HLT/EMNLP 2005*.
- Badue, C., W. Vaz & E. Albuquerque: 1998, Using an Automatic Retrieval System in the Web to Assist Cooperative Learning, *Proceedings of The 1998 World Conference of the WWW, Internet and Intranet, (Orlando, FL, USA)*, AACE Press.
- Bakhtin, M.: 1935, *The Dialogic Imagination*, Austin, TX: University of Texas Press.
- Barthes, R.: 1964/1967, *Elements of Semiology*. (Translated by Annette Lavers & Colin Smith). London: Jonathan Cape.
- Barzilay, R. & M. Elhadad: 1997, Using lexical chains for text summarization, *ACL Workshop on Intelligent Scalable Text Summarization*.
- Baumann, S. & O. Hummel: 2005, Enhancing Music Recommendation Algorithms Using Cultural Metadata, *Journal of New Music Research*, 34(2), Routledge, part of the Taylor & Francis Group.

- Baumgarten, A.G.: 1735/1954, *Reflections on Poetry*, trans. K. Aschenbrenner and W.B. Holther, Berkeley: University of California Press.
- Berwick, R.C.: 1989, Learning Word Meanings from Examples, *Semantic Structures: Advances in Natural Language Processing*, Lawrence Erlbaum Associates
- Bhabha, H.K.: 1994, *The Location of Culture*, New York: Routledge.
- Binsted, K., G. Ritchie: 1994, An implemented model of punning riddles. *Proceedings of the American Association of Artificial Intelligence*, AAAI Press.
- Bloom, H.: 1982, *Agon: Towards a Theory of Revisionism*, New York : Oxford University Press.
- Bourdieu, P.: 1984, *Distinction: A Social Critique of the Judgment of Taste* (Trans. Richard Nice), Cambridge: Harvard University Press.
- Bourdieu, P.: 1993, *The Field of Cultural Production*, Columbia University Press.
- Boyd, D.: 2004, Friendster and publicly articulated social networks, *Conference on Human Factors and Computing Systems (CHI 2004)*, ACM Press.
- Bradley, M.M., & P.J. Lang: 1999, Affective norms for English words (ANEW): Instruction manual and affective ratings. *Technical Report C-1*, The Center for Research in Psychophysiology, University of Florida.
- Brajnik, G. & C. Tasso: 1994, A shell for developing non-monotonic user modeling systems, *International Journal of Human-Computer Studies* 40:31-62.
- Breese, J., D. Heckerman & C. Kadie: 1998, Empirical analysis of predictive algorithms for collaborative filtering, *Proc. of the Fourteenth Annual Conference on Uncertainty in Artificial Intelligence*: 43-52.
- Briggs, K. C., & I. B. Myers: 1976, *Myers-Briggs Type Indicator: Form F*, Palo Alto: Consulting Psychologists Press.
- Brill, E.: 1992, A simple rule-based part of speech tagger, *Proceedings of the Third Conference on Applied Natural Processing*.
- Brin, S., R. Motwani, J.D. Ullman, S. Tsur: 1997, *Proceedings ACM SIGMOD International Conference on Management of Data*, May 13-15, 1997, Tucson, Arizona, USA.
- Brusilovsky, P.: 2001, Adaptive Hypermedia, *User Modeling and User-Adapted Interaction* 11: 87-110, Kluwer Academic Publisher.
- Budzik, J. & K. Hammond: 1999, Watson: Anticipating and Contextualizing Information Needs, *Proceedings of the Sixty-second Annual Meeting of the American Society for Information Science*.
- Budzik, J. & K.J. Hammond, User Interactions with Everyday Applications as Context for Just-in-time Information Access, *Proceedings of Intelligent User Interfaces 2000*, ACM Press.
- Budzik, J., K.J. Hammond, L. Birnbaum, and M. Krema: 2000, Beyond similarity, *Working Notes of the AAAI 2000 Workshop on AI for Web Search*.
- Carbonell, J.: 1979, *Subjective understanding: Computer models of belief systems* (Technical Report YALE/DCS/tr150), New Haven, CT: Computer Science Department, Yale University.
- Cardie, C.: 1993, Using Decision Trees to Improve Case-Based Learning, *Proceedings of the 10th International Conference on Machine Learning*: 25-32, Amherst, MA, Morgan Kaufmann.
- Cardie, C., J. Wiebe, T. Wilson, & D. Litman: 2003, Combining Low-Level and Summary Representations of Opinions for Multi-Perspective Question Answering, *Proceedings of the AAAI Spring Symposium on New Directions in Question Answering*: 20-27.

- Cassell, J. & T. Bickmore: 2003, Negotiated Collusion: Modeling Social Language and its Relationship Effects in Intelligent Agents, *User Modeling and User-Adaptive Interaction* 13(1): 89-132.
- Certeau, M.de: 1984, *The Practice of Everyday Life* (Trans. S. Rendall), University of California Press.
- Charniak, E.: 1972, *Toward a Model of Children Story Comprehension*. MIT School of EECS PhD Thesis.
- Charniak, E.: 1977, Ms. Malaprop, a language comprehension program, *Proceedings of the Fifth International Conference on Artificial Intelligence*.
- Charniak, E.: 1986, A neat theory of marker passing, *Proceedings of the Fifth National Conference on Artificial Intelligence*: 584-588, Menlo Park, CA: AAAI Press.
- Church, K.W., P. Hanks: 1990, Word association norms, mutual information, and lexicography, *Computational Linguistics*, 16, 22-29.
- Collins, A. M., and E. F. Loftus: 1975, A spreading-activation theory of semantic processing, *Psychological Review* 82: 407-428.
- Collins, A.M. & M.R. Quillian: 1969, Retrieval from semantic memory, *Journal of Verbal Learning and Verbal Behavior*.
- Csikszentmihalyi, M., E. Rochberg-Halton: 1981, *The Meaning of Things: Domestic Symbols and the Self*, Cambridge University Press.
- Cullingford, R.E.: 1978, *Script application: Computer understanding of newspaper stories* (Technical Report YALE/DCS/tr116), New Haven, CT: Computer Science Department, Yale University.
- Das, S.R., & M.Y. Chen: 2001, Yahoo! for Amazon: Sentiment Parsing from Small Talk on the Web, *Proceedings of EFA 2001, European Finance Association Annual Conference*.
- Dave, K., S. Lawrence, D.M. Pennock: 2003, Mining the peanut gallery: Opinion extraction and semantic classification of product reviews, *Proceedings of WWW-03, 12th International Conference on the World Wide Web*: 519-528, Budapest, HU: ACM Press.
- Deerwester, S., S. T. Dumais, G. W. Furnas, T. K. Landauer, R. Harshman: 1990, Indexing by Latent Semantic Analysis, *Journal of the Society for Information Science*, 41(6), 391-407.
- DeJong, G.F.: 1979, *Skimming stories in real time: An experiment in integrated understanding* (Technical Report YALE/DCS/tr158), New Haven, CT: Computer Science Department, Yale University.
- Dewey, J.: 1934, *Art as Experience*, New York: Perigee Books.
- Dolan, C.: 1989, *Tensor manipulation networks: Connectionist and symbolic approaches to comprehension, learning, and planning* (Technical Report 890030), Los Angeles, CA: Computer Science Department, University of California, Los Angeles.
- Don, A.: 1989, *We Make Memories*, Master's Thesis, Interactive Telecommunications Program, New York University.
- Doyle, J.: 1980/1987, A truth maintenance system, in *Reading in nonmonotonic reasoning*: 259-279, Morgan Kaufmann.
- Drummond, C.: 1992, *Automatic Goal Extraction From User Actions to Accelerate the Browsing of Software Libraries*, M. A. Sc. Thesis, University of Ottawa.
- Dyer, M.G.: 1983, *In-depth understanding*. Cambridge, Mass.: MIT Press.

- Eagle, N., P. Singh, A. Pentland: 2003, Common sense conversations: understanding casual conversation using a common sense database, *Proceedings of the Artificial Intelligence, Information Access, and Mobile Computing Workshop (IJCAI 2003)*.
- Ekman, P.: 1993, Facial expression and emotion, *American Psychologist*, 48, 384-392.
- Elliott, C.: 1992, *The Affective Reasoner: A Process model of emotions in a multi-agent system*. Institute for the Learning Sciences, Evanston, IL: Northwestern University, Ph.D. Thesis.
- Elliott, C., & A. Ortony: 1992, Point of view: Reasoning about the concerns of others, *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, 809-814, Bloomington, IN: Cognitive Science Society.
- Ellis, D., B. Whitman, A. Berenzweig & S. Lawrence: 2002, The Quest For Ground Truth in Musical Artist Similarity, *Proceedings of the 3rd International Conference on Music Information Retrieval: 13-17 October 2002, Paris, France*.
- Fauconnier, G. and M. Turner: 2002, *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities*, Basic Books.
- Fellbaum, C.: 1998, *WordNet: an electronic lexical database*. Cambridge, Mass.: MIT Press.
- Festinger, L.: 1957, *A theory of cognitive dissonance*, Stanford University Press.
- Finin, T. et al.: 2005, Social Networking on the Semantic Web, *The Learning Organization*.
- Finin, T.W. & D. Drager: 1986, A general user modeling system, *Proc. of the 6th Canadian Conference on Artificial Intelligence: 24-29, Montreal, Canada*.
- Francis, M.E., & R.J. Booth: 2001, *Linguistic Inquiry and Word Count: LIWC 2001*. Mahwah, NJ: Lawrence Erlbaum.
- Freud, S.: 1905/1963, *Jokes and their relation to the unconscious*, WW Norton.
- Freud, S.: 1915/1959, Instincts and their Vicissitudes, in *Collected Papers of Sigmund Freud, vol 4, 70-89*, New York, Basic Books
- Freud, S.: 1919/1990, The 'uncanny', in Dickson, A. (Ed.), *The Penguin Freud Library, Vol. 14: Art and Literature: 335-6*, Penguin, Harmondsworth.
- Gärdenfors, P. and K. Holmqvist: 1994, Concept formation in dimensional spaces, *Lund University Cognitive Studies* 26.
- Gaver, W.W.: 1991, Technology affordances, *Proceedings of CHI 91*, 79-84. New York, NY: ACM Press.
- Geertz, C.: 1973, *The interpretation of cultures*. New York: Basic.
- Gentner, D.: 1983, Structure-mapping: A theoretical framework for analogy, *Cognitive Science* 7: 155-170.
- Gibson, J.J.: 1979, *The ecological approach to visual perception*, Boston, MA: Houghton Mifflin.
- Goffman, E.: 1959, *The Presentation of Self in Everyday Life*, Garden City, NY: Doubleday.
- Goldstone, R.L., & B.J. Rogosky: 2002, Using relations within conceptual systems to translate across conceptual systems, *Cognition*, 84, 295-320.
- Grefenstette, G., Y. Qu, J.G. Shanahan, D.A. Evans: 2004a, Coupling Niche Browsers and Affect Analysis for an Opinion Mining Application, *Proceeding of RIAO-04*.

- Greimas, A.J.: 1966/1983, *Structural semantics: an attempt at a method*, Lincoln: University of Nebraska Press.
- Greimas, A.J.: 1987, *On Meaning: Selected Writings in Semiotic Theory* (trans. Paul J Perron & Frank H Collins), London: Frances Pinter.
- Groseclose, T. & J. Milyo, 2005, A Measure of Media Bias, *The Quarterly Journal of Economics*, vol. 120(4), 1191-1237, MIT Press.
- Haberlandt, K., C. Berian & J. Sandson: 1980, The episode schema in story processing, *Journal of Verbal Learning and Verbal Behavior* 19: 635-650
- Halpin, H.R.: 2003, *The plots of children and machines: The statistical and symbolic semantic analysis of narratives* (MSc. thesis), School of Informatics, University of Edinburgh.
- Hatzivassiloglou, V., & K.R. McKeown: 1997, Predicting the semantic orientation of adjectives, *Proceedings of ACL-97, 35th Annual Meeting of the Association for Computational Linguistics*: 174-181, Madrid, ES: Association for Computational Linguistics.
- Hatzivassiloglou, V., & J.M. Wiebe: 2000, Effects of adjective orientation and gradability on sentence subjectivity, *Proceedings of COLING-00, 18th International Conference on Computational Linguistics*: 299-305, Saarbrücken, GE: Morgan Kaufmann.
- Haug, W.F.: 1986, *Critique of Commodity Aesthetics*, Minneapolis: University of Minnesota Press.
- Hegel, G. W. F.: 1807, *Phänomenologie des Geistes*.
- Herlocker, J., J. Konstan, A. Borchers & J. Riedl: 1999, An algorithmic framework for performing collaborative filtering, *M. Hearts, F. Gey and R. Tong (eds.), Proc. of the 22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*: 230-237.
- Hofstadter, D.R. and M. Mitchell: 1995, The copycat project: A model of mental fluidity and analogy-making, in *D. Hofstadter and the Fluid Analogies Research group: Fluid Concepts and Creative Analogies*, 205-267, Basic Books.
- Horvitz, E. et al.: 1998, The Lumiere Project: Bayesian User Modeling for Inferring the Goals and Needs of Software Users, *Proceedings of the Conference on Uncertainty in AI*.
- Hovy, E.H. and C.H. Lin: 1997, Automated text summarisation in SUMMARIST, *Proceedings of the ACL97/EACL97 Workshop on Intelligent Scalable Text Summarisation, Madrid, Spain*.
- Iser, W: 1979, *The Act of Reading: A Theory of Aesthetic Response*. Baltimore: Johns Hopkins UP.
- Itten, J.: 1961, *The art of color: the subjective experience and objective rationale of color*, New York: Van Nostrand Reinhold.
- Jakobson, R.: 1960, Closing Statement: Linguistics and Poetics, in *Style in Language*, ed. T.A. Sebeok: 350-77, New York: Wiley.
- Jensen, D. & J. Neville: 2002, Data mining in social networks, *National Academy of Sciences Symposium on Dynamic Social Network Analysis*.
- Joachims, T., D. Freitag, T.M. Mitchell: 1997, Web Watcher: a tour guide for the World Wide Web. *Proceedings of the 15th International Joint Conference on Artificial Intelligence (IJCAI'97) vol.1*: 770-777.
- Jung, C. G.: 1921/1971, *Psychological Types*, trans. by H. G. Baynes, Princeton, NJ: Princeton University Press.
- Kant, I.: 1790/1929, Critique of judgment, in *T. M. Greene (Ed.), Kant selections*: 375-432, New York: Scribner.

- Kobsa, A.: 2001, Generic User Modeling Systems, *User Modeling and User-Adapted Interaction 11*: 49-63, Kluwer Academic Publisher.
- Kolodner, J.L.: 1984, *Retrieval and organizational strategies in conceptual memory: A computer model*, Hillsdale, NJ: Erlbaum.
- Kristeva, J.: 1986, Word, dialogue, and the novel, in T. Moi (Ed.), *The Kristeva reader*, 35-61, New York: Columbia University Press.
- Krueger, M.W.: 1983, *Artificial Reality*, Addison-Wesley.
- Kulyukin, V.: 1999, Application-Embedded Retrieval From Distributed Free-Text Collections, *Proceedings of AAAI-99*.
- Lakoff, G. and M. Johnson: 1980, *Metaphors We Live by*, University of Chicago Press.
- Landauer, T.K. & S.T. Dumais: 1997, A solution to Plato's problem: The Latent Semantic Analysis theory of the acquisition, induction, and representation of knowledge, *Psychological Review 104*: 211-140
- Lange, T.E. & M.G. Dyer: 1989, High-level inferencing in a connectionist network, *Connection Science 1*, 181-217.
- Langston, M.C., T. Trabasso & J.P. Magliano: 1999, A connectionist model of narrative comprehension. *Ashwin Ram, & Kenneth Moorman (Eds.), Understanding Language Understanding*: 181-226, Cambridge, MA: MIT Press.
- Latour, B.: 2005, *Reassembling the Social: An Introduction to Actor-Network-Theory*, Oxford University Press.
- Lebowitz, M.: 1980, *Generalization and memory in an integrated understanding system* (Technical Report YALE/DCS/tr186), New Haven, CT: Computer Science Department, Yale University.
- Lehnert, W.G.: 1982, Plot units: A narrative summarization strategy. *W.G. Lehnert & Martin H. Ringle (Eds.), Strategies for natural language processing*: 375-412, Hillsdale, NJ: Erlbaum.
- Lenat, D.: 1995, Cyc: a large-scale investment in knowledge infrastructure. *Communications of the ACM 38(11)*. ACM Press.
- Levenshtein, V.: 1965/1966, Binary codes capable of correcting deletions, insertions, and reversals, *Doklady Akademii Nauk SSSR*, 163(4):845-848, 1965 (Russian). English translation in *Soviet Physics Doklady*, 10(8):707-710.
- Lieberman, H.: 1995, An agent that assists web browsing, *Proceedings of IJCAI, 1995*.
- Lin, D.: An Information-Theoretic Definition of Similarity, *15th International Conference on Machine Learning*: 296--304, Morgan Kaufmann, 1998.
- Liu, H.: 2002, MontyLingua: Commonsense-Enriched NLP, Toolkit and API. Accessed at: <http://web.media.mit.edu/hugo/montylingua/>
- Liu, H.: 2005, Experimental System for Character Affect Dynamics Analysis. MIT Media Lab Technical Note.
- Liu, H., G. Davenport: 2005, Self-reflexive performance: dancing with the computed audience of culture , *International Journal of Performance Arts and Digital Media 1(3)*, Intellect Ltd.
- Liu, H., M. Hockenberry & T. Selker: 2005, Synesthetic Recipes: foraging for food with the family, in taste-space, *Proceedings of the 32nd Annual Conference on Computer Graphics and Interactive Techniques*, Los Angeles, CA.
- Liu, H., H. Lieberman and T. Selker: 2003, A model of textual affect sensing using real-world knowledge, *Proceedings of the 7th International Conference on Intelligent User Interfaces, IUI 2003*, 125-132, ACM Press.

- Liu, H. and P. Maes: 2004, What Would They Think? A Computational Model of Attitudes. *Proc. of the 2004 ACM Conference on Intelligent User Interfaces*, 38-45. ACM Press.
- Liu, H. and P. Maes: 2005a, InterestMap: harvesting social network profiles for recommendations, *Proceedings of ACM Beyond Personalization 2005: A Workshop on the Next Stage of Recommender Systems Research*, 54-59, ACM Press.
- Liu, H. and P. Maes: 2006, Rendering aesthetic impressions of text in color space, *International Journal on Artificial Intelligence Tools* 15. World Scientific Press.
- Liu, H., P. Maes, G. Davenport: 2006, Unraveling the taste fabric of social networks, *International Journal on Semantic Web and Information Systems* 2(1): 46-78, Hershey, PA: Idea Academic Publishers.
- Liu, H. and P. Singh: 2004a, Commonsense reasoning in and over natural language, M Negoita, RJ Howlett, LC Jain (Eds.): *Knowledge-Based Intelligent Information and Engineering Systems*, LNCS 3215:293-306, Springer.
- Liu, H. and P. Singh: 2004b, ConceptNet: A Practical Commonsense Reasoning Toolkit. *BT Technology Journal* 22(4), 211-226. Kluwer Academic Publishers.
- Liu, H., T. Selker, H. Lieberman: 2003, Visualizing the affective structure of a text document, *Proceedings of the Conference on Human Factors in Computing Systems, CHI 2003*, April 5-10, 2003, Ft. Lauderdale, FL, USA, 740-741, ACM Press.
- Luhmann, N.: 1984, *Social Systems*, Stanford University Press.
- Maglio, P.P., R. Barrett, C.S. Campbell, T. Selker: SUITOR: an attentive information system, *Proceedings of the 5th international conference on Intelligent user interfaces*: 169-176, ACM Press.
- Mauldin, M.L.: 1991, Performance in FERRET: A Conceptual Information Retrieval System, *SIGIR 1991*: 347-355.
- McCallum, A., A. Corrada-Emmanuel, X. Wang: 2005, Topic and Role Discovery in Social Networks, *IJCAI, 2005*.
- McCloud, S.: 1993, *Understanding Comics: The Invisible Art*, New York, NY: Harper Collins Inc.
- McCracken, G.: 1988, *Culture & Consumption*, Indiana University Press.
- McCracken, G.: 2006, *Plenitude*, Bloomington: Indiana University Press.
- Mehrabian, A.: 1995b, Framework for a comprehensive description and measurement of emotional states. *Genetic, Social, and General Psychology Monographs*, 121, 339-361.
- Mihalcea, R. & H. Liu: 2006, A corpus-based approach to finding happiness, *Proceedings of AAAI-CAAW-06, the Spring Symposia on Computational Approaches to Analyzing Weblogs*.
- Miikkulainen, R. & M.G. Dyer: 1991, Natural language processing with modular PDP networks and distributed lexicon, *Cognitive Science*, 15, 343-400.
- Minsky, M.: 1980, K-lines: A Theory of Memory, *Cognitive Science* 4, 117-133.
- Minsky, M.: 1992, Future of AI Technology, *Toshiba Review* 47(7).
- Minsky, M.: forthcoming, *The Emotion Machine*, Simon & Schuster.
- Mishne, G.: 2005, Experiments with Mood Classification in Blog Posts, *Style2005 - the 1st Workshop on Stylistic Analysis Of Text For Information Access*, at *SIGIR 2005*.
- Mitkov, R.: 1998, Robust Pronoun Resolution with Limited Knowledge, *Proceedings of ACL*.

- Montgomery, H.: 1994, Towards a perspective theory of decision making and judgment, *Acta Psychol (Amst)* 87(2-3):155-78.
- Moorman, K. and A. Ram: 1994, Integrating creativity and reading: A functional approach. In *Sixteenth Annual Conference of the Cognitive Science Society*.
- Mueller, E.T.: 1998, *Natural language processing with ThoughtTreasure*, New York: Signiform, <http://www.signiform.com/tt/book/>
- Myers, I. B., & M.H. McCaulley: 1985, *Manual: A guide to the development and use of the Myers-Briggs Type Indicator*, Palo Alto, CA: Consulting Psychologists Press.
- Narayanan, S.S.: 1997, Knowledge-based action representations for metaphor and aspect (KARMA) (Unpublished doctoral dissertation). University of California, Berkeley
- Newman, M.: 2001, Who is the best connected scientist? A study of scientific coauthorship networks, *Phys. Rev. E* 64.
- Ogden T: 1979, On projective identification, *International Journal of Psychoanalysis* 60, 371-394.
- Oren, T. et al.: 1990, Guides: characterizing the interface, Laurel, B. (Eds.) *The art of human-computer interface design*, Addison-Wesley.
- Ortony, A., G.L. Clore, M.A. Foss: 1987, The Referential Structure of the Affective Lexicon, *Cognitive Science* 11(3): 341-364.
- Orwant, J.: 1995, Heretogeneous Learning in the Doppelganger User Modeling System. *User Modeling and User Adapted Interaction*, 4(2):107-130.
- Pang, B. & L. Lee: 2004, A Sentimental Education: Sentiment Analysis using Subjectivity Summarization based on Minimum Cuts, *Proceedings of ACL-04, 42nd Meeting of the Association for Computational Linguistics: 271-278*, Barcelona, ES: Association for Computational Linguistics.
- Pang, B., L. Lee & S. Vaithyanathan: 2002, Thumbs up? Sentiment Classification using Machine Learning Techniques, *Proceedings of EMNLP-02, the Conference on Empirical Methods in Natural Language Processing*, 79-86, Philadelphia, US: Association for Computational Linguistics.
- Polanyi, I. & A. Zaenen: 2004, Contextual valence shifters, *Computing Attitude and Affect in Text: Theory and Applications: 1-9*, The Information Retrieval Series 20, Springer.
- Ram, A.: 1994, AQUA: Questions that drive the explanation process. In Roger C. Schank, Alex Kass, & Christopher K. Riesbeck (Eds.), *Inside case-based explanation: 207-261*, Hillsdale, NJ: Erlbaum.
- Resnick, P. & H. Varian: 1997, Introduction to Special Section on Recommender Systems, *Communications of the ACM* 40(3).
- Rhodes, B. and P. Maes: 2000, Just-in-time information retrieval agents, *IBM Systems Journal* 39(4): 685-704.
- Rhodes, B. & T. Starner: 1996, Remembrance Agent: A continuously running automated information retrieval system, *The Proceedings of The First International Conference on The Practical Application Of Intelligent Agents and Multi Agent Technology: 487-495*.
- Rich, E.: 1979, User modeling via stereotypes, *Cognitive Science*.
- Riloff, E., & J.M. Wiebe: 2003, Learning extraction patterns for subjective expressions, *Proceedings of EMNLP-03, 8th Conference on Empirical Methods in Natural Language Processing: 105-112*.
- Roget, P.: 1911, *Roget's Thesaurus of English Words and Phrases*. Retrieved from gutenberg.net/etext/10681

- Rokeby, D.: 1995, Transforming Mirrors: Subjectivity and Control in Interactive Media, in Penny, S. (Ed.), *Critical Issues in Electronic Media*: 133-158, Series in Series in Film History and Theory, Albany: SUNY Press.
- Rosenblatt, L.: 1978, Efferent and Aesthetic Reading, *The Reader, The Text, The Poem: A Transactional Theory of the Literary Work*: 22-47, Carbondale: Southern Illinois UP.
- Rubin, V.L., J.M. Stanton & E.D. Liddy: 2004, Discerning Emotions in Texts, *The AAAI Symposium on Exploring Attitude and Affect in Text, AAAI-EAAT 2004*, Stanford, CA.
- Sack, W.: 1994, On the Computation of Point of View, *Proceedings of the National Conference of Artificial Intelligence (AAAI 94)*, Seattle, WA., July 31-August 4.
- Sack, W.: 2001, Actor-Role Analysis: Ideology, Point of View and the News, *New Perspectives on Narrative Perspective (Will Van Peer and Seymour Chatman, Editors)* New York: SUNY Press.
- Salton, G. & C. Buckley: 1988, Term weighting approaches in automatic text retrieval, *Information Processing and Management*, 24(5): 513-- 523.
- Salton, G. & M.J. McGill: 1983, *Introduction to Modern Retrieval*, McGraw-Hill Book Company.
- Sarwar, B.M., G. Karypis, J.A. Konstan, J.T. Riedl: 2001, Item-based collaborative filtering recommendation algorithms. In *Proceedings of the Tenth International World Wide Web Conference*, 285--295.
- Schapire, R.E. & Y. Singer: 2000, BoosTexter: A boosting-based system for text categorization, *Machine Learning* 39(2/3), 135-168. Springer.
- Schleiermacher, F.: 1809/1998, General Hermeneutics, in A. Bowie (Ed.) *Schleiermacher: Hermeneutics and Criticism*: 227-268. Cambridge University Press.
- Selker, T.: 1994, COACH: a teaching agent that learns, *Communications of the ACM* 37(7): 92-99.
- Shardanand, U. and P. Maes: 1995, Social information filtering: Algorithms for automating 'word of mouth'. *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*: 210-217.
- Silber, G., K.F. McCoy: 2002, Efficiently computed lexical chains as an intermediate representation for automatic text summarization, *Computational Linguistics* 28(4): 487-496.
- Silverman, K.: 1983, *The subject in semiotics*, Oxford University Press.
- Simmel, G.: 1971, D.N. Levine (Ed.): *On Individuality and Social Forms: Selected Writings of Georg Simmel*, University of Chicago Press.
- Simon, H.: 1967, Motivational and emotional controls of cognition, *Psychological Review*, 74(1), 29-39.
- Singh, P., T. Lin, E. T. Mueller, G. Lim, T. Perkins, W. L. Zhu: 2002, Open Mind Common Sense: knowledge acquisition from the general public. *Proceedings of ODBASE'2002*.
- Slovan, A. & M. Croucher: 1981, Why Robots Will Have Emotions, *Proceedings of the Seventh International Conference on Artificial Intelligence*.
- Solomon, M.R., H. Assael: 1987, The forest or the trees - a gestalt approach to symbolic consumption, in Umiker-Sebeok (Eds), *Marketing and Semiotics - New Directions in the Study of Signs for Sale*, 189-218, Mouton-Gruyter.
- Subasic, P. & A. Huettner: 2001, Affect analysis of text using fuzzy semantic typing, *IEEE Transaction on Fuzzy Systems* 9(4): 483-496.

- Staab, S., S. Santini, F. Nack, L. Steels, and A. Maedche: 2002, Emergent semantics, *IEEE Intelligent Systems*, 17(1):78--86.
- St. John, M.F.: 1992, The story gestalt: A model of knowledge-intensive processes in text comprehension, *Cognitive Science*, 16, 271-306.
- Talmy, L.: 1988a, Force Dynamics in language and cognition, *Cognitive Science*, 12(1), 49-100.
- Tan, E.S.: 1994, Story processing as an emotion episode. *Naturalistic Text Comprehension*, H. van. Oostendorp and RA Zwaan (eds.), 165-188
- Todorov, T.: 1968/1981, *Introduction to Poetics*, Theory and History of Literature Vol. 1, University of Minnesota Press.
- Tulving, E.: 1983, *Elements of episodic memory*. Oxford: New York
- Turney, P.D.: 2001, Mining the Web for Synonyms: PMI-IR versus LSA on TOEFL, *Proceedings of the Twelfth European Conference on Machine Learning*.
- Turney, P.: 2002, Thumbs Up or Thumbs Down? Semantic Orientation Applied to Unsupervised Classification of Reviews, *Proceedings of ACL-02, 40th Annual Meeting of the Association for Computational Linguistics*: 417-424, Association for Computational Linguistics.
- Turney, P.D. & M.L. Littman: 2003, Measuring praise and criticism: Inference of semantic orientation from association. *ACM Transactions on Information Systems (TOIS)*, 21(4): 315-346.
- Tversky, A.: 1977, Features of similarity, *Psychological Review* 84:327--352.
- Wahlster, W.: 1991, User and Discourse Models for Multimodal Communication, J. Sullivan and S. Tyler (eds.), *Intelligent User Interfaces*: 45-67, ACM Press.
- Watson, D., A. Tellegen: 1985, Toward a consensual model of mood, *Psychol Bull* 98(2): 219-35
- Whissell, C.: 1989, The Dictionary of Affect in Language. In R. Plutchnik and H. Kellerman, editors, *Emotion: Theory, Research and Experience*, 113--131. NY: Academic Press.
- Whitman, B. & S. Lawrence: 2002, Inferring Descriptions and Similarity for Music from Community Metadata, *Voices of Nature, Proceedings of the 2002 International Computer Music Conference*: 591-598.
- Whitman, B. & P. Smaragdis: 2002, Combining Musical and Cultural Features for Intelligent Style Detection, *Proceedings of the 3rd International Conference on Music Information Retrieval*.
- Wiebe, J., R.F. Bruce & T.P. O'Hara: 1999, Development and use of a gold-standard data set for subjectivity classifications, *Proceedings of the 37th annual meeting of the Association for Computational Linguistics on Computational Linguistics*: 246-253, College Park,US: Association for Computational Linguistics.
- Wiebe, J.M. & W.J. Rapaport: 1988, A computational theory of perspective and reference in narrative, *Proc. 26th Annual Meeting of the Assoc. for Computational Linguistics (ACL-88)*: 131-138.
- Wiebe, J., T. Wilson, R. Bruce, M. Bell, M. Martin: 2004, Learning Subjective Language, *Computational Linguistics* 30(3), 207-308.
- Wilensky, R.: 1978, *Understanding goal-based stories* (Technical Report YALE/DCS/tr140), New Haven, CT: Computer Science Department, Yale University.
- Wilson, T., J. Wiebe & P. Hoffman: 2005, Recognizing Contextual Polarity in Phrase-Level Sentiment Analysis, *Proceedings of HLT/EMNLP 2005*

- Winograd, T.: 1971, Procedures as a Representation for Data in a Computer Program for Understanding Natural Language. *MIT AI Technical Report 235*.
- Yarowsky, D.: 1992, Word-sense disambiguation using statistical models of Roget's categories trained on large corpora, *Proc. of the 33th Annual Meeting of the Association for Computational Linguistics*: 189--196.
- Zholkovsky, A.K.: 1970b: Porozhdayuschaya poetika v rabotakh S. M. Eizenshteina (Generative poetics in the writings of Eisenstein), in *Greimas (ed.) Sign-Language-Culture*: 451-68, The Hague, Mouton.
- Zukerman, I. & D.J. Litman: 2001, Natural Language Processing and User Modeling: Synergies and Limitations, *User Modeling and User-Adapted Interaction 11*: 49-63, Kluwer Academic Publisher.
- Zwaan, R.A. and G.A. Radvansky: 1998, Situation models in language comprehension and memory. *Psychological Bulletin*, 123(2): 162-185