

STINFO COPY



AFRL-HE-WP-TR-2007-0046

**Personality and Cultural Modeling
for Agent-Based Representation of a
Terrorist Cell, Phase I**

**CMT Hogan
Robert A. Van Houten
Nini La**

**CACI, Inc. - Federal
1600 Wilson Boulevard, Suite 1300
Arlington VA 22209-2515**

January 2003

Interim Report for March 2002 – January 2003

20071017269

**Approved for public release;
distribution is unlimited.**

**Air Force Research Laboratory
Human Effectiveness Directorate
Warfighter Interface Division
Cognitive Systems Branch
Wright-Patterson AFB OH 45433-7604**

NOTICE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report was cleared for public release by the Air Force Research Laboratory, Det 1, Wright Site, Public Affairs Office and is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (<http://www.dtic.mil>).

AFRL-HE-WP-TR-2007-0046

HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN
ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE DIRECTOR

//SIGNED//

DANIEL G. GODDARD
Chief, Warfighter Interface Division
Human Effectiveness Directorate
Air Force Research Laboratory

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE (DD-MM-YYYY) January 2003		2. REPORT TYPE Interim		3. DATES COVERED (From - To) March 2002 - January 2003	
4. TITLE AND SUBTITLE Personality and Cultural Modeling for Agent-Based Representation of a Terrorist Cell, Phase I				5a. CONTRACT NUMBER N61339-02-C-0084	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER 63832D	
6. AUTHOR(S) CMT Hogan, Robert A. Van Houten, Nini La				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER 0476DM00	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) CACI, Inc. - Federal 1600 Wilson Boulevard, Suite 1300 Arlington VA 22209-2515				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Materiel Command Defense Modeling and Simulation Office Air Force Research Laboratory 1901 N. Beauregard Street, Suite 500 Human Effectiveness Directorate Alexandria VA 22311-1705 Warfighter Interface Division Cognitive Systems Branch Wright-Patterson AFB OH 45433-7604				10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/HECS, DMSO	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-HE-WP-TR-2007-0046	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited. AFRL/PA cleared on 02 July 2007, AFRL-WS-07-1553.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT This report describes the research into the use of personality, cultural and socio-political modeling in order to provide a robust asymmetric opponent for Military Operation in Urban Terrain training. The report defines a plausible framework for modeling psychological and cultural influences and describes the development of rule sets to represent both the personality and cultural domains and demonstrate decision-making influences. The demonstration of collaboration among the rule sets is also shown. Results described in this report include the following: the NEO Personality Inventory Revised (NEO-PI-R) features make it a good basis for describing the behaviors of individuals whom the modeler does not have the luxury of testing; the cultural rule-set adequately supports the goals of exploring how two disparate influences on individual's decision-making might be provided to a simulation-based training exercise but it suffers from an insufficient basis in an excepted sociological model to allow for independent validation; and allowing rule-set collaboration only at the end of the chaining simplifies the analysis and is a therefore a logical first step in exploring collaboration.					
15. SUBJECT TERMS Human Behavior Modeling, Architecture for Behavior and Cognitive Modeling (ABCM), Fuzzy Rule Sets, Rule Set Collaboration, NEO Personality Inventory Revised (NEO-PI-R)					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT SAR	18. NUMBER OF PAGES 44	19a. NAME OF RESPONSIBLE PERSON John L. Camp
a. REPORT UNCLASSIFIED	b. ABSTRACT UNCLASSIFIED	c. THIS PAGE UNCLASSIFIED			19b. TELEPHONE NUMBER (include area code)

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

THIS PAGE LEFT INTENTIONALLY BLANK

TABLE OF CONTENTS

	<u>Page</u>
INTRODUCTION	1
Problem	1
Objectives	1
Background	1
Organization of the Report	2
METHOD	3
The Concept: The Terrorist Cell as a Group of Intelligent Agents	3
Personality Inventory Selection	4
Translation of the NEO-PI-R Into Inputs to a Fuzzy Rule Set	5
Use of the Personality Rule Set	6
Cultural Model Framework	13
Use of the Cultural Rule Set	14
Socio-Political Model Framework	15
Use of the Socio-Political Rule Set	15
RESULTS	16
Fuzzy Rule-Set Inference Engine	16
Single Rule Set Operation	16
Collaborating Rule Sets	16
Design of the Personality Fuzzy Rule Set	17
Design of the Cultural Rule Set	19
CONCLUSIONS AND RECOMMENDATIONS	21
Conclusions	21
REFERENCES	21
APPENDIX A: Personality Trait Recommendation	22
APPENDIX B: Estimation of Facet Scores from Domain Scores	33

LIST OF FIGURES

<u>Figure</u>	<u>Page</u>
1. The fuzzy knowledge bases evaluate sub-goals.....	4
2. The user selects a Big Five domain.....	7
3. Mapping the user inputs.....	18
4. Strong Membership Function.....	19
5. FALSE / LOW Membership Function.....	20

LIST OF TABLES

<u>Table</u>	<u>Page</u>
1. Comparison of Personality Inventories.....	5
2. Domain - Facet - Adjective Pair Mapping.....	8
3. Cultural Domains and Attributes.....	14

EXECUTIVE SUMMARY

PROBLEM

The Defense Modeling and Simulation Office (DMSO) wanted to explore the use of personality, cultural and socio-political modeling in order to provide a robust asymmetric opponent for Military Operation in Urban Terrain training. Increased realism in Human Behavioral Representation is deemed necessary to enhance the training process through exercises against an unpredictable and adaptive enemy.

OBJECTIVE

Define a plausible framework for modeling psychological and cultural influences. Develop rule sets to represent both the personality and cultural domains and demonstrate decision-making influences. Demonstrate collaboration among the rule sets.

APPROACH

Assess three personality inventories (Myers Briggs Type Indicator (MBTI), Minnesota Multiphasic Personality Inventory (MMPI), and NEO Personality Inventory Revised (NEO-PI-R)) against the needs of modeling an unknown enemy, i.e., one who has not taken a personality test. For the selected approach, develop a fuzzy rule set to translate personality descriptions into assessments of likely behaviors of the individual. Develop a similar rule set for cultural aspects of the individual. Develop a mechanism by which the rule sets can collaborate and by which alternate rule sets representing other modeling approaches can be used.

FINDINGS

The NEO PI-R version of the Big Five personality descriptions was found to provide enough descriptors to adequately represent personality, to provide enough granularity in the descriptors to make face-valid assertions more acceptable, to have a high correlation between self-rating and rating by others, and to hold across cultures.

A face-valid framework for cultural descriptions provides a mechanism to explore the interplay between the two influences, but lacks the rigor to support independent validation.

Collaboration should be limited to the end of the rule set chaining. End fact collaboration should be executed through an OR relationship.

CONCLUSIONS

The NEO-PI-R features make it a good basis for describing the behaviors of individuals whom the modeler does not have the luxury of testing. The high inter-rater reliability suggests a rule-set based on NEO-PI-R can be independently validated.

The cultural rule-set adequately supports the goals of exploring how two disparate influences on individual's decision-making might be provided to a simulation-based training exercise. However, it suffers from an insufficient basis in an excepted sociological model to allow for independent validation.

Allowing rule-set collaboration only at the end of the chaining simplifies the analysis and is therefore a logical first step in exploring collaboration. It also supports the use of alternate rule sets whose designers need only know about the end facts in the chaining. The use of an OR relationship is logically consistent with providing a build-up of tendency among the separate influences in an individual's decision process.

RECOMMENDATIONS

Explore an observable-based cultural model to provide the same remote description and validation characteristics of the personality model.

Explore interaction between rule-sets at points other than the end fact.

INTRODUCTION

Problem

A terrorist cell is to be represented by four intelligent agents, providing a leader and three members. The decision-making process of each agent is to be influenced by the individual's personality, the individual's socio-political status, and by its cultural biases.

Objectives

A fuzzy rule set for each of the domains (personality, culture, socio-politics) will be developed. Effects of the rule sets will be demonstrated.

The agent design will specify the entry points for rule-set influences on goal selection and planning.

Background

Training simulations have been seeking to provide more realistic OPFOR behavior. This is especially true in the case of asymmetric warfare where established doctrine may not be available to explain or predict OPFOR decisions. Complaints are often levied that the OPFOR decisions are too scripted and thus less useful in the training process than an unpredictable or adaptive enemy would be.

To address this need, the Defense Modeling and Simulation Office (DMSO) contracted through NAWCTSD with CACI to explore mechanisms for providing personality, cultural and socio-political influences in terrorist cell operations. The operations were to be in the context of urban warfare to support the ultimate goal of providing a Military Operations in Urban Terrain (MOUT) training tool. Early in the program, feedback from the DMSO focused the effort on the personality modeling.

The vehicle by which these studies were conducted is CACI's Architecture for Behavior and Cognitive Modeling (ABCM), a reusable architecture designed to support both human and machine behavior representation. Previous versions of the ABCM had successfully incorporated personality into Blue commander decision-making. CACI was to build on the experience gained from the earlier efforts to address the asymmetric environment.

The contracted effort was to be conducted in two Phases. According to the "ABCM Applied to Asymmetric Warfare" Statement of Work (SOW), the first phase centered on the fuzzy rule set development, with enough of the agent design accomplished to understand the use of the rule sets. Complete agent design and development, including integration of the rule sets into the agent goal selection and planning process, was to have taken place during Phase 2.

During the course of this contract, there was a redefinition of DMSO's roles and responsibilities. In due course, government program management responsibility shifted to the Air Force Research Laboratory, which imposed specific goals of its own on the program. The rule-set demonstration, for example, was augmented to include a

comparison between two different personality rule sets, which requires the development of the additional rule set. For that reason, the demonstration has been incorporated as part of another program. Phase 2 of the "ABCM Application to Asymmetric Warfare" effort has been cancelled.

This report presents the Phase 1 effort only.

Organization of the Report

The METHOD section opens with an introduction of the ABCM in order to provide a context for the rest of the discussion. It follows with a description of each rule set in turn. First the rationale for the underlying approach to the rule set design is justified. Because emphasis was placed on the personality rule set, a white paper was prepared earlier and submitted to DMSO for approval. This paper describes the criteria by which three different personality inventories were evaluated and the conclusion reached. That paper is presented in APPENDIX A. The conclusion is presented in the METHOD section.

Following the justification for the approach, the process by which the concept is translated into the fuzzy rules is described. This includes the identification of the fuzzy rule set elements.

The METHOD section concludes with an explanation of how a user might interact with the rule-sets.

The RESULTS section presents the design specifics of each of the rule-sets. Some rule-set statistics are provided. Finally, the collaboration mechanism is described.

The CONCLUSIONS section presents an assessment of the strengths and weaknesses of the system.

Finally, RECOMMENDATIONS are provided for areas that need further study.

All REFERENCES are provided as a part of the APPENDIX A white paper.

METHOD

The Concept: The Terrorist Cell as a Group of Intelligent Agents

The Architecture for Behavior and Cognitive Modeling provides a framework for creating one or more agents according to user specifications. It provides a mechanism for creating a terrorist cell made up of four intelligent agents. Each agent is an instance of the same agent design: each agent is goal-based; each uses a Belief-Desire-Intent planning algorithm; each accesses the same knowledge bases. The difference between the agents is determined by their initial conditions: the personality descriptors, the cultural features, and the socio-political status, plus the role it plays in the cell.

This provides a foundation for a rich variety of experiments. The effects of personality differences can be explored by assessing the individual behavior of agents running different set-up data; the interchange among a small group dedicated to a common purpose can be explored by the interactions of a team of agents. Some interesting questions can be posed, if one of the agents represents a separate culture from the other agents. Would a potential leader's effectiveness be undermined if he acted against the cultural biases of the majority? Would a convert to the belief system, especially one who comes from the culture against which the cell is fighting, be impelled to act more strongly out of a new to establish credibility?

The agent concept, shown in Figure 1, uses a FIPA recognized standard vertical design; i.e., the decision-making occurs in sequential order through a reactive or behavioral layer, a planning layer and a cooperating layer. Initially, the agents were to limit their interaction with each other through environmental influences. That is, each agent would know that the others existed and what they were doing in the same manner it would sense anything else in the environment. Cooperation through negotiation was to be explored in Phase 2.

The behavioral layer makes use of a standard knowledge base that includes the cell's tactics, the tactics of other terrorist groups and a list of resources. This portion of the agent also relies on scripts for quick reaction and to allow the planning layer to use task level sub-goals. If a script fails during execution or if the reaction needed is unscripted, the planning layer is invoked.

Planning is based on a Belief-Desire-Intent algorithm. Belief is determined from the agent's perception of the environment. A goal is established from tactics contained within the tactics database. Multiple sub-goals are generated as the algorithm attempts to find a way to achieve the goal. The utility function assesses risk in the usual fashion. However, the utility is modified by inputs from the fuzzy knowledge bases, which jointly produce an assessment of how likely the given agent is to act in one of nine ways. If the goal is impacted by one of these actions, the utility will be strengthened or weakened accordingly. So final scoring is influenced by the agent's personality, cultural bias, and socio-economic status.

Even though the agent was never implemented, the concept of how the fuzzy rule sets were to interact with the agent's decision-making provides a foundation for the rule set descriptions to follow.




Figure 1. The fuzzy knowledge bases evaluate sub-goals during the planning stage of a vertical agent.

Personality Inventory Selection

The primary purpose of this effort is to provide a capability to a training system. In that context, the most common usage of each of the rule sets will be to describe imaginary people, but to provide them with complex representation. An alternate use could be one in which real people are represented. In the latter case, it is unlikely that the modeler will have the opportunity to give a personality questionnaire to the subject. This implies that the best choice for personality representation will be one in which ratings by others correlate strongly with rating by self. A personality inventory that is strongly based on observable traits will be favored over one that measures subconscious traits.

At some point in the modeling process, the rule set developer will have to produce face-value rules. These rules are more likely to be accurate if the categories are narrow and specific. Broad categories, such as extroversion, can mean too many different things. It is harder to assess face value rules using such large granularity. In addition, a full personality description must be provided, not one or two traits.

Validation of the rule set, particularly the face value rules, is necessary for any faith that the system employing the rule set is providing a useful representation of the enemy. If the personality inventory has been shown to be applicable across cultures, there is a reason to believe it may have a genetic basis. The raters are more likely to recognize such traits--they will feel familiar. In addition, there is the practical consideration that cross-

cultural applicability is more likely to generalize to persons of interest to the system user.

Finally, an inventory using orthogonal, continuous data supports accurate modeling. The availability of normative data also supports this goal. Orthogonality assures the breadth of coverage. Continuousness permits the modeler to capture subtle variations in personality. Norms place trait measures in context—they provide the relative importance.

These requirements formed the basis of an investigation into three personality inventories: the Myers Briggs Type Indicator (MBTI), the Minnesota Multiphasic Personality Inventory (MMPI) and the Big Five (Neuroticism, Extroversion, Openness, Agreeableness, Conscientiousness). The three schools of thought were assessed against the above criteria. Table 1 summarizes the results; the details are provided in APPENDIX A "Personality Trait Recommendation."

Table 1. Comparison of Personality Inventories

Personality Approach	Observable Traits	Cross-Cultural	Orthogonal, Continuous	Personality Variance
MBTI	Reliability for self-report Theory concerns unconscious	yes	Bipolar Non-orthogonal	Not comprehensive: lacks neuroticism Threshold scoring
MMPI	Empirical basis Reliable for self-report Not amenable to observation by others	yes	Strong norms Non-orthogonal	Designed to diagnose psychopathology Factor analysis suggest measures only 2 factors
NEO-PI-R	Linguistic basis universal Self-test correlate highly with "ratings by others"	NEO strongly C mostly A likely to be different among cultures	Factors selected to be orthogonal Measured on a continuous scale Norms available (t scores)	70% personality variance can be accounted for by the 5 factors Wide range of behavioral associations Facet breakdown

Based on the conclusions drawn from the comparison study, the Big Five approach known as the NEO-PI-R was selected.

Translation of NEO-PI-R Into Inputs to a Fuzzy Rule Set

The NEO Personality Inventory Revised breaks the five personality domains of Neuroticism, Extroversion, Openness, Agreeableness and Conscientiousness into six

facets each. The six facets alone provide a level of detail that is useful for our applications. However, each of the six facets can be further broken down into four to eight adjective pairs. These pairs are gleaned by from the 300 questions that comprise the NEO-PI-R test.

As with any personality inventory, the typical manner in which the NEO-PI-R is used is for a subject to answer a questionnaire. The answers are correlated to the related factors and the factors combined to produce a domain score. However, for our application, the user usually will not have enough information on the individual being represented to provide accurate facet scores. To make this approach more widely applicable, it is necessary to convert in the opposite direction. Appendix B provides an explication on how this can be accomplished.

Use of the Personality Rule Set

As indicated in the last section, the user will input the personality profile of the agent by providing scores between 20 and 80 for each of the Big Five domains. Translation tables invisible to the user will provide facet and adjective pair scores. The adjective-pairs will appear together with their calculated value to the user. This is important for debugging because it is the adjective pairs that are the actual input into the rule set. Figure 2 illustrates the Graphical User Interface to the ABCM for using the Big Five Personality Model.

Each of the tabs in the upper right corner of the adjective list corresponds to a Big Five domain. By selecting a tab, the user can see the list of adjectives that are active for this rule set, together with the calculated value. Other uses of the Big Five personality profiles may concentrate on different adjectives, depending on the application. To facilitate reuse of the system, all of the adjectives will be retained in the translation tables, regardless of whether they are visualized in the particular application. The full list of adjective pairs associated with each of the thirty facets is listed in Table 2.

The user provides the original domain scores through the section of the GUI entitled Big Five Descriptors. The snapshot in Figure 2 was taken during development where placeholder numbers were used for formatting purposes. This section should only have values between 20 and 80.




Figure 2. The user selects a Big Five domain to visualize the adjectives

Table 2 includes a column labeled "Reflects Facet." The adjectives pairs were derived by association with the 300 questions McCrae and Costa provide for the NEO-PI-R assessment. In some cases, the relationship between question and facet is weak. For rule set development considerations, we want to limit the use of an adjective pair for more than one facet, especially across domains. Where an adjective pair is weak in one facet association, but strong in another, the pair will be dropped from the weak one. Where the adjective pair is strong in more than one facet, an alternative expression will be tried. The system is limiting the number of adjective pairs per facet to four; in some cases, value judgments must be made. Again, all adjective pairs will be represented in the translation tables; a subset will be active in any given application.

The remaining rule sets developed for this effort were subjected to fewer rigors in their theoretical foundation. They are fundamentally notional in nature.

Table 2. Domain - Facet - Adjective Pair Mapping

Domain	Facet	Descriptors from Study	Reflects Facet	Possible Alternatives
Neuroticism	Anxiety	anxious---calm	yes	
		fearful---not fearful	yes	
		worrying---not worrying	yes	
		tense---relaxed	yes	
		nervous---composed	yes	
		doubting---confident	yes	
		pessimistic---optimistic	yes	
	Angry Hostility	anxious---calm	no	quick tempered---not quick tempered
		irritable---easy going	yes	
		impatient---patient	yes	
		excitable---imperturbable	yes	
		moody---even tempered	yes	
		callous---gentle	no	resentful---not resentful
		tense---relaxed	no	
	Depression	worrying---not worrying	OK	self blaming---other blaming
		troubled---contented	yes	
		doubting---confident	OK	self-deprecating---immodest
		not self-confident---self-confident	yes	
		pessimistic---optimistic	yes	
		moody---even tempered	OK	sad---cheerful
		anxious---calm	no	
	Self-Consciousness	shy---outgoing	yes	
		not self-confident---self-confident	OK	not confident in social situations---confident
		timid---bold	no	easily embarrassed---not easily embarrassed
		doubting---confident	OK	
		defensive---open	OK	
		inhibited---expressive	OK	feels inferior---feels superior
		anxious---calm	no	sensitive to criticism---thick skinned
	Impulsiveness	moody---even tempered	no	impulsive---slow to act
		irritable---easy going	no	self indulgent---self denying
		sarcastic---sincere	no	given to excess---exercising constraint
		self-centered---not self-centered	no	lacking self-control---exhibiting self-control
		loud---soft spoken	no	
		hasty---studied	yes	
		excitable---imperturbable	OK	
	Vulnerability	muddled---clear thinking	yes	
		not self-confident---self-confident	OK	
		doubting---confident	OK	

Domain	Facet	Descriptors from Study	Reflects Facet	Possible Alternatives
		anxious---calm	no	dependent---independent
		inefficient---efficient	no	easily stressed---cool under pressure
		distractible---alert	no	indecisive---decisive
		careless---thorough	no	emotionally fragile---emotionally stable
Extraversion	Warmth	friendly---unfriendly	yes	
		warm---distant	yes	
		sociable---not sociable	yes	
		cheerful---gloomy	OK	
		involved---aloof	yes	
		affectionate---impassive	yes	
		outgoing---reserved	OK	
	Gregariousness	sociable---not sociable	yes	
		outgoing---reserved	OK	prefers groups---prefers working alone
		pleasure-seeking---not pleasure-seeking	no	
		involved---aloof	OK	seeks out others---seeks isolation
		talkative---reticent	OK	
		spontaneous---not spontaneous	no	
		interactive---withdrawn	OK	
	Assertiveness	aggressive---retiring	OK	seeks leadership---prefers to follow
		bold---shy	OK	dominant---submissive
		assertive---compliant	yes	
		self-confident---not self-confident	no	takes initiative---waits for others to act
		forceful---not forceful	yes	
		enthusiastic---apathetic	no	
		confident---doubting	no	
	Activity	energetic---not energetic	yes	
		hurried---deliberate	yes	
		quick---hesitant	yes	
		determined---quitting	no	works vigorously---works sporadically
		enthusiastic---apathetic	no	
		aggressive---retiring	no	
		active---not active	yes	
	Excitement Seeking	pleasure-seeking---not pleasure-seeking	OK	
		daring---cowardly	OK	excited by danger---frightened by danger
		adventurous---timid	yes	
		charming---not charming	no	seeks novelty---content with the familiar
		handsome---ordinary	not under any	seeks stimulation---does not seek stimulation

Domain	Facet	Descriptors from Study	Reflects Facet	Possible Alternatives
			circumstances	
		spunky---cautious	OK	
		clever---straightforward	no	prefers action oriented activities---does not
	Positive Emotions	enthusiastic---apathetic	yes	
		humorous---humorless	yes	
		praising---not praising	no	ebullient---placid
		spontaneous---not spontaneous	no	cheerful---sorrowful
		pleasure-seeking---not pleasure-seeking	no	
		optimistic---pessimistic	yes	
		jolly---gloomy	yes	
Openness	Fantasy	dreamy---concrete	yes	
		imaginative---practical	yes	
		humorous---humorless	no	given to fantasy---grounded in reality
		mischievous---not mischievous	no	capable of fantastic plans---rejection of fantastic plans
		idealistic---realistic	no	
		artistic---structured	OK	
		complicated---not complicated	no	
	Aesthetics	imaginative---prosaic	no	values art---does not value art
		artistic---structured	yes	
		original in thought---common in thought	no	experiences in feelings and images---does not
		enthusiastic---not enthusiastic	no	derives meaning from form as well as words--does not
		inventive---unoriginal	no	
		idealistic---practical	no	
		versatile---limited	no	
	Feelings	excitable---placid	no	relies on feelings---relies on thoughts
		spontaneous---planned	no	has strong emotional associations--does not
		insightful---imperceptive	no	experiences deep feelings---does not
		imaginative---prosaic	no	aware of emotional reactions---is not
		affectionate---unaffectionate	OK	
		talkative---quiet	no	empathizes easily---does not
		outgoing---reserved	no	
	Actions	wide interests---narrow interests	yes	
		imaginative---prosaic	no	open to experience---set in ways
		adventurous---timid	OK	
		optimistic---not optimistic	no	tries different approaches---sticks to one approach

Domain	Facet	Descriptors from Study	Reflects Facet	Possible Alternatives
		strong mannered---mild mannered	no	
		talkative---quiet	no	enjoys experimenting---content with the status quo
		versatile---limited	yes	
	Ideas	idealistic---practical	no	enjoys abstract ideas---concrete
		wide interests---narrow interests	yes	
		inventive---mundane	OK	
		curious---not curious	yes	
		original in thought---common in thought	yes	
		imaginative---prosaic	OK	
		insightful---imperceptive	OK	
	Values	liberal---conservative	OK	believes values can change---values should not change
		unconventional---conventional	OK	
		willing to try new things---not willing	yes	
		flirtatious---reserved	no	willing to restructure beliefs---relies on enduring rules
Agreeableness	Trust	forgiving---unforgiving	Ok	
		trusting---distrustful	yes	
		accepting---suspicious	yes	
		quick to trust---wary	yes	
		optimistic---pessimistic	no	believes people are honest---skeptical of people
		peaceable---combative	no	takes things at face value---suspicious of motives
		soft-hearted---hard-hearted	Ok	
	Straightforwardness	straightforward---crafty	yes	
		accepting---demanding	no	above board---manipulative
		conventional---clever	Ok	
		direct---flirtatious	Ok	
		unpolished---charming	no	values honesty---willing to be unscrupulous
		simple and direct---shrewd	yes	
		democratic---autocratic	no	
	Altruism	warm---unfeeling	yes	
		soft-hearted---hard-hearted	Ok	
		gentle---harsh	Ok	thoughtful---inconsiderate
		generous---stingy	yes	
		kind---cruel	Ok	
		tolerant---intolerant	no	courteous---rude
		selfless---selfish	yes	
	Compliance	flexible---stubborn	yes	
		accepting---demanding	no	sincere---sarcastic

Domain	Facet	Descriptors from Study	Reflects Facet	Possible Alternatives
		docile---headstrong	yes	
		patient---impatient	no	prefers cooperation---prefers competition
		tolerant---intolerant	no	reluctant to express anger---quick to express anger
		reserved---outspoken	Ok	
		soft-hearted---hard-hearted	Ok	
	Modesty	unassuming---show-off	yes	
		straightforward---clever	Ok	
		compliant---assertive	no	self effacing---self promoting
		agreeable---argumentative	no	sees self as equal--sees self as superior
		not self confident---self confident	no	arrogant---humble
		docile---aggressive	no	egalitarian---not egalitarian
		concrete---idealistic	no	
	Tender-mindedness	friendly---not friendly	no	meeting human needs is imperative---not an imperative
		warm---cold	no	merciful judge---rule oriented judge
		sympathetic---indifferent	yes	
		soft-hearted---hard-hearted	yes	
		gentle---harsh	Ok	
		stable---unstable	no	compassionate---insensitive to others
		kind---cruel	Ok	
Conscientiousness	Competence	efficient---inefficient	yes	
		self confident---not self-confident	no	common sensical---not well grounded
		thorough---sloppy	yes	
		resourceful---not resourceful	no	responsible---irresponsible
		confident---hesitant	no	prepared---extemporaneous
		clear thinking---confused	Ok	good judgment---poor judgment
		intelligent---not intelligent	no	accustomed to success---never certain of success
	Order	organized---disorderly	yes	
		thorough---sloppy	no	detailed planner---plans only broadly
		efficient---inefficient	Ok	
		precise---not precise	Ok	
		methodical---not methodical	yes	
		mentally focused---absent minded	no	exacting---not exacting
		careful---careless	no	
	Dutifulness	open---defensive	no	conscientious---not conscientious
		focused---distractible	no	reliable---unreliable
		careful---careless	yes	
		industrious---lazy	Ok	

Domain	Facet	Descriptors from Study	Reflects Facet	Possible Alternatives
		thorough---sloppy	Ok	
		mentally focused---absent minded	no	bound by ethics---not bound by ethics
		confirming---fault finding	no	follows rules faithfully---cuts corners
	Achievement Striving	thorough---sloppy	no	stern task master---easy going
		ambitious---unambitious	yes	
		industrious---lazy	yes	
		enterprising---not enterprising	no	
		determined---irresolute	yes	
		confident---not confident	no	
		persistent---quitting	yes	
	Self Discipline	organized---disorderly	no	punctual---not punctual
		industrious---lazy	Ok	self starter---procrastinator
		efficient---inefficient	yes	
		mentally focused---absent minded	Ok	completes by prioritizing---tends to be overwhelmed
		energetic---not energetic	no	maintains focus on goals---easily distracted
		thorough---sloppy	no	productive---wasteful
		even tempered---moody	no	finishes tasks---leaves tasks unfinished
	Deliberation	slow to act---hasty	yes	
		deliberate---impulsive	yes	
		careful---careless	yes	
		patient---impatient	Ok	
		mature---immature	no	considers consequences before acting---does not
		thorough---sloppy	no	
		even tempered---moody	no	

Cultural Model Framework

Four aspects of culture are used to develop a set of face-valid fuzzy rules for determining cultural influences on decision-making. These aspects are relationships with authority, pluralism and tolerance, institutionalized violence, and time and work ethic. Similar to the approach with the psychology rule-set, these broad categories are divided into finer areas. Table 3 lists the cultural rule-set considerations.

Table 3. Cultural Domains and Attributes

Domain	Facet	Specifics
Authority Relationships	Social Organization	Autocratic / Democratic Government
		Caste or Implied Caste
	Class System	Economic, Aristocracy
	Beholdenness	Obligations Through Family Relations
		Obligations Through Favors
	Family Position	Head of household / clan Birth position
Tolerance / Pluralism	Cultural Sub-groups	Women
		Secondary cultures
		Race
	Belief Systems	Religions
		Ideologies
	Outsiders	Openness to other ways Acceptance of strangers
Institutionalized Violence	Legal System	Individualism
		Harsh sentences Death penalty
	Suicide	Expected to save face or honor
		Accepted path to martyrdom
Time and Work Ethic	Time	Specific (5:17) vs. Approximate (at dawn)
		Promptness
	Work Ethic	Priorities

An examination of the third column in Table 3 shows several points that look like they might be natural bridges to the psychology rule set. "Openness to other ways" and "Acceptance of strangers" under Tolerance reminds one of the features of both the Openness and Agreeableness domains. However, the cultural rule-set considers these and the other features, not from the individual's point of view, but from the predilections provided to the individual due to his cultural exposure.

Use of the Cultural Rule Set

The cultural rule-set is initiated by establishing facts related to the attributes listed in Column 3 of Table 3. In most cases, the attributes are established by Yes or No answers

to questions such as "Is the Death Penalty recognized?" in other cases selections must be made, such as a particular religion. These become the inputs to the rule-set.

A number of ways were considered to allow the user to describe an agent's cultural bias without falling back on the user's own biases. No good way was established. Either the user has to set a number of input parameters, which is the option we chose, or the user could describe a culture by equating it with one of a few stereotypical cultures. The user interface for the cultural rule-set is not yet finalized and either option will still establish the same inputs to the rule-set.

Socio-Political Model Framework

The Socio-Political rule set is little more than a stub. It considers only two aspects of the individual's social position: economic level and education level.

Use of the Socio-Political Rule Set

The user is presented with two drop down menus. The first provides the following options for educational level:

- No formal education
- Religious Indoctrination
- Vocational training
- Primary
- Secondary
- College

The user is to select the highest level achieved. If the individual has attended a level, whether graduation has occurred or not, that is to be selected.

The second drop down menu provides the following options for economic level:

- Poor
- Lower Medium Income
- Medium Income
- Upper Medium Income
- Wealthy

The options were so stated to avoid any sense of class in these selections. Class is an integral part of the cultural rule set.

RESULTS

Fuzzy Rule-Set Inference Engine

The existing ABCM fuzzy rule set inference engine was used to establish the design parameters for the rule sets. The following sections present some basics about its use.

Single Rule Set Operation

Each of the rule sets consists of a set of fuzzy facts. Associated with each fact is a list of fuzzy rules that help to establish that fact. The rules associate the fact with a particular membership function. Rules that associate the fact with a common membership function are executed through an algebraic OR. The algebraic AND is only employed when two phrases of a rule are ANDed. If two conditions are to equally contribute to a facts value, they are represented by separate rules. In such a way, the inference engine will perform the OR. Normal logic relationships hold within the fuzzy system.

A fact's membership in a function is determined by the y value associated with the x value derived from the above calculation. A fact always has some membership value, even if it is 0. That is, all of the rules are always active. Unlike a standard rule set, where a rule can be said to "not fire," each rule in the fuzzy rule set fires. The chaining is in the forward direction from input to a set of concluding facts. At any given time, only some of the rules associated with a fact may produce non-Nil values. Each time a rule achieves a value, the fuzzy fact is recalculated entirely. Eventually, all the data that is presented to the rule set will be taken into account.

Collaborating Rule Sets

Rules that establish the same fact in the same membership function are executed through an algebraic OR. As long as more than one rule set does not employ the same membership functions or the same fact names, their rules will never collide. In this case, common membership functions are used with the two rule sets. Collaboration between the two rule sets is controlled through the choice of fact name. For this initial exploration of cultural and psychological interplay, that collaboration was limited to the final facts. Each rule sets produces the same nine final facts. Thus each rule set is able to consider the set-up data and the situational data provided by the simulation test bed within which it is hosted, to independently assess the agent's likelihood to behave in a particular manner.

The results from the rule sets pool in a manner consistent with accumulating evidence. Remembering that the application for this work is the modeling of a terrorist cell, the behavior under question is in most cases extreme. Certainly, the nine final facts represent extreme behavior. The incremental accumulation of input towards this behavior is consistent with an Or-ing function. An AND-ing function would allow one rule set to override the consensus of any other rule sets in play. Unless a person has a deviant psychology, it is unlikely that psychological make-up alone would drive him to kill a child,

for example. Unless a person is a fanatic, it is unlikely that cultural or ideological influences alone would drive him to kill a child. However, a personality with callous characteristics born of a culture that predisposes the individual to accept thinking of ones enemies as less than human could indeed be willing to take this action.

Design of the Psychology Rule Set

Figure 3 depicts the flow of information from the user into the rule set. The adjective pairs map into five areas that deal with attitudes and relationships. The first layer of the rule-set associates the adjectives that represent observable behavior with the five rule sub-sets. These groupings provide rules about relationships to authority, rules about attitudes towards success, rules about judgment by others, rules about tolerance and sympathy and rules about unpredictability. The next layer focuses on establishing the willingness to perform certain actions. Finally an assessment is made of nine "likelihoods". These are:

- LikelihoodToCommitViolence,
- LikelihoodToHarmSelf,
- LikelihoodToHarmOthers,
- LikelihoodToHarmAChild,
- LikelihoodToKillSelf,
- LikelihoodToKillOthers,
- LikelihoodToKillAChild,
- LikelihoodToDestroyProperty,
- LikelihoodToBlindlyObey.

The nine "likelihoods" are the only common points with other rule-sets in the system.

In the case of the Psychology Rule Set, , the membership functions Strong and Weak associate the adjective pairs in Column 3 of Table 2 with the right and left hand side of a bell curve distribution. Each adjective can range from 20 to 80, centered at 50. Thus the binary adjective pairs, such as "Worrying / Not Worrying" have numbers that add to 100. The non-binary adjective pairs, such as "Imaginative / Practical" are not mirror images. A person can be both imaginative and practical. Figure 4 shows the fuzzy membership function for Strong.




Figure 3. Mapping the user inputs into the rule set sub-sets




Figure 4. Strong Membership Function

Design of the Cultural Rule Set

The first layer of the rule-set associates these observable features of the society and underlying culture with rule-sets representing the four areas listed in Column 1. The output of this layer is mostly a set of attitudes. The rules in this layer are mostly generic and are unlikely to change much when a different culture is represented. However, the subsequent layers must necessarily get into the specifics of the particular cultures. For example, it is suitable in layer 1 to associate the Beholdenness features with general attitudes on obligations. However, it takes details of the specific culture to understand that a favor done one by an influential patron is such an obligating act in Filipino culture that it has been deemed one of the root causes of much of the corruption in past Filipino administrations. It also takes details of the specific culture to understand what activities are likely to create this sense of obligation.

As with the Psychology Rule Set, the next two layers focus on establishing the willingness to perform certain actions and the assessment of the nine "likelihoods". The nine "likelihoods" are the only common facts between the two rule-sets. The rule-sets do have portions that discuss similar things; however they do so with different fact names. This design constraint facilitates the ease with which the interplay between the two can be established.

The membership functions utilized throughout the cultural rule set are right and left hand trapezoids. Figure 5 illustrates the Cultural Rule Set membership functions False and Low. True and High are right-hand versions of Figure 5.




Figure 5. FALSE / LOW Membership Function

Some parts of the Psychology rule set also use the trapezoid membership function. As long as rules do not try to mix the x-axis ranges of facts, this is allowable.

CONCLUSIONS AND RECOMMENDATIONS

Conclusions

The Big Five provides the depth and breadth for rich personality modeling on the basis of observable behavior, making it an excellent choice when the modeler must describe individuals for whom test results are unknown.

The collaboration of rule sets at the end points facilitates the ability to plug in alternate rule sets. More complex dependencies between psychological and sociological foundations require further study before mid point cross overs are recommended.

REFERENCES

1. Costa, P. T., Jr., & McCrae, R. R. (1992). *Professional Manual, The Revised NEO Personality Inventory*. Psychological Assessment Resources: Lutz, FL.

**APPENDIX A:
PERSONALITY TRAIT RECOMMENDATION**

Personality Trait Recommendation
Robert Van Houten, Ph D

INTRODUCTION

Phase one of the statement of work requires conduct of an initial proof of principle demonstration of Architecture for Behavior and Cognitive Modeling (ABCM) technology to represent individual terrorist behavior. Terrorist behavior will be modeled through the use of system agents. The personality trait profiles of these agents will be used within a psychology rule set to provide one of several bases for agent moves. It is expected that the system will be used in either of two ways. First, when used as a training system the user will specify personality trait profiles of hypothetical terrorists by referring to trait descriptions. The user will thus be able to discern within an interactive environment how differences in personality profiles affect differences in terrorist behaviors. Second, the mature system may be used to model the behavior of known terrorists. For this use the user must be capable of assessing the personality traits of an individual by examining available intelligence and biographical information. In either case the personality traits that are ultimately input to the system must describe individual behavioral tendencies in sufficient detail to permit agents to discriminate between alternative moves.

Although the system to be built is not a psychological test, parallels between developing this system and developing a structured personality test suggest that such tests are a good place to look for a trait set. A personality test infers the existence of traits based on a limited behavior sample (such as a paper and pencil test) and then generalizes the expression of these traits within a different behavioral domain. Thus, implicit in the process of undertaking test construction are the assumptions that likely behavior patterns can be predicted from a set of trait constructs and that these traits can be accurately assessed by a test. The constructs that personality tests purport to measure may be labeled personality types, traits, scale scores, factors or dimensions. Regardless of label, these measures perform the same function; they are constructs that mediate a relationship between a sample of known behaviors and a set of forecast behaviors.

The system to be developed will perform an analogous function. Further, similar to test construction, developing a system that accurately transforms behavioral samples into intermediate trait constructs and then translates those into agent moves requires adherence to many of the same psychometric principles. The extensive effort involved in obtaining psychometric data for the system is clearly beyond the scope of the current effort. As an alternative, however, traits assessed by tests having good psychometric properties can be selected. A comprehensive set of traits that can be reliably assessed and then related to valid moves are needed. This paper discusses desirable characteristics of the trait set to be selected, reviews several alternative trait sets obtainable through existing instruments, and recommends a set of personality traits to be used.

DESIRABLE CHARACTERISTICS OF TRAITS

The following paragraphs discuss the characteristics of the trait constructs in order of desirability with the most desirable discussed first.

Trait constructs should be reliably measured. The system will use a set of trait vectors to characterize individual personality variables. There must be confidence that these trait vectors can be reliably measured using the tools available to the user of the system. Typically, personality tests rely on self-reports to infer the existence of traits. The user of the system is not likely to have that luxury. Instead, traits must be inferred on the basis of a sample of observed behavior (biography, intelligence summaries, writings, speeches, etc.). The user of the system must be capable of assessing traits by rating another's behavior. Thus, the self-report behavior statements used to measure a trait must be capable of being used for 'ratings by others' as well. The most predictive personality trait would be of little use in the system if it could not be assessed reliably by observing the individual's behavior.

A distinct advantage of using traits defined by a structured test is that reliability data are usually published for the items that measure the traits. Reliability for personality tests is most often reported as either test-retest, or internal consistency coefficients. All of the instruments reviewed have adequate published reliability for self-report data. The reliability of 'ratings by others' for personality tests is not particularly important to the psychological community and is hard to find. Nevertheless, extending the use of reliable self-report items to 'ratings by others' is preferable to writing new items. It does, however, require a judgment as to how well these items can be used by raters.

The set of traits selected should account for a high percentage of the variance in personality. The more fully the individual's personality is described the greater the potential for the system to accurately predict behavior. There are obviously practical limits on the number of traits that can be used. Care should be taken, however, not to omit an important trait. The traits selected should be reviewed for completeness. There are statistical methods that are able to provide some evidence that a set of traits accounts for significant variance in personality. Further, reviewers are quick to point out traits that are not well represented by specific tests. Their critiques can be used to ensure that the set of traits selected provide a reasonably complete measure of personality variance.

Each trait selected should have a rich array of behavioral associations. The fuzzy rule sets for personality will connect traits to behaviors through a network of propositions that begins with traits and terminates in attitudes that relate to moves. The power of the system to discriminate among possible moves will depend in large part on the validity of these statements. To the extent that one behavior can accurately be inferred from the existence of another behavior or characteristic, a proposition is valid. Ideally, the validity of each proposition within the rule set should rest on an empirical foundation. Inevitably though, at some point in the network of the rule set, there will be a need to insert propositions that are simply face-valid in order to extend the path in the direction of the possible moves. A trait that is measured by a variety of concrete behaviors and described

by an array of adjectives that are demonstrated to relate to the trait is more valuable in building the fuzzy rule set because more valid associations are available. Further, the availability of more associations reduces the risk that a face-valid association will prove to be unrelated to actual behavioral propositions.

The set of traits selected should exhibit evidence of applicability across multiple cultures. Two advantages accrue from using personality traits that have been demonstrated in multiple cultures. First, research results are more likely to generalize to populations of interest to the system user. Personality traits that show up in multiple cultures are more likely to have a genetic component. Traits with a genetic component are in turn, more likely to appear universally. Extending the applicability of the traits to an untested population thereby entails less risk. Second, universal traits are likely to be more easily recognized by raters as familiar behavioral tendencies. Evidence of universality facilitates making good hypotheses concerning the manner in which differences in levels of the trait will be manifested in daily activities in other cultures. Raters can then look for these markers.

The traits selected should be orthogonal; they should be measured on a continuous, interval scale; and normative data should be available. Measurement of multiple orthogonal traits along continua permits the system to capture subtle variations in personality profiles that should ultimately increase the system's ability to discriminate between target behaviors. Quantities of a trait are not particularly useful without norms that place the trait measure in context. Normative data is also important in understanding the relative importance of traits within an individual's personality profile. Traits exhibited to an extreme degree may dominate a personality. Conversely, traits present in normal amounts may not differentiate a personality in a useful way for modeling. Thus, data measured on an interval scale may be most useful if first transformed to normative data.

REVIEW OF TESTS

The following tests were reviewed as candidates for the trait set.

Myers Briggs Type Indicator (MBTI).

The MBTI is a pencil and paper, forced choice, self-report instrument, having 94 scoreable items assessing eight traits measured on four bi-polar scales (Extravert/Introvert, Sensing/Intuiting, Thinking/Feeling, Perceiving/Judging). The test was created by Isabel Myers and Katherine Briggs and is based on a personality theory of psychological types advanced by Carl Jung. Test items provide logical content in conformance with the theory that inspired the test. The MBTI returns 16 – four letter personality designators formed by forcing the selection of one of two descriptors from each scale. General behavioral tendencies of each type are described.

Strengths: The MBTI may be the most popular personality test in use today. It is used extensively for career counseling and management/team building. Although the writer

did not view any predictive validity data in these areas, the continued use of the MBTI for more than 50 years provides a de facto measure of validity. Simply stated, the instrument could not have endured if it were not found to be useful in categorizing personalities and typing individuals with similar traits and interests.

Concurrent validity was evidenced by reported significant correlation with predictions of the Strong Campbell Interest Inventory. Reliability of the measures is adequate to good with only the T/F scale reliability coefficient reported below .8. Moreover, twin studies have revealed strong evidence of heritability of E/I and T/F scales and weaker but significant heritability of S/N and J/P. Further, the types have been demonstrated across cultures.

Weaknesses: Reliability of measurement of the scales for potential use in the system is questionable. I found no data to support reliable measurement using MBTI scales by other than self-report. Although reliability of the self-report items is usually reported as good, individuals taking the test often complain of difficulty in making forced choices on the items. If it is difficult for individuals to rate themselves on these measures, it is likely to be even more difficult for others to rate them reliably using these items. Further, much of Jung's theoretical position concerned the unconscious, by definition inaccessible to self-report and difficult for others to assess.

There is evidence that the MBTI is not a comprehensive measure of personality. Nor is it clear that it was ever intended to be. One personality component consistently reported by other researchers, neuroticism, has not been recovered by the MBTI. Further, even though the MBTI is theory driven, it has a questionable theoretical foundation. Jung proposed two attitudes (Extravert/Introvert) that can be paired with any one of four functions: sensing, intuiting, thinking, and feeling. This 2 by 4 representation gave rise to eight two-letter personality types formed differently from the 16 proposed by Myers and Briggs (after adding a fourth scale). Thus, the theory may have been extended beyond its founder's original intentions.

The forced choice format and non-continuous scales of the MBTI provide for categorical outcomes that may mask useful individual differences. For example, an individual is typed as either extravert or introvert based upon a threshold score. Individuals with widely differing responses to questions could nevertheless be typed in the same way. Further, using the quantitative score on a scale as a representation of the strength of the underlying trait is a misuse of the MBTI. It has been argued that such dichotomous, non-continuous scales will not be able to account for differences in behavior of individuals with the same personality type (Barbutto, 1997). This criticism is particularly trenchant when applied to using the MBTI traits as the basis for the system. Finally, there is no evidence that the scales represent orthogonal dimensions.

Conclusion: The MBTI provides a good profile of individual temperament types. Its enduring success appears predicated on the assumption that people with the same temperament will gravitate toward the same interests, occupations, and roles within organizations. Much of its success may also derive from the extraversion/introversion

scale, a factor that shows up consistently in other studies. The other scales appear to be unique to the MBTI and could well be derived from other, more basic personality factors. The scales of the MBTI are a means to an end, that end being personality type. Description and norms exist at that level. To equate scales to traits is a questionable. Thus, the scales of the MBTI do not necessarily constitute the best choice for the traits of the system.

Minnesota Multiphasic Personality Inventory (MMPI).

The MMPI and its successor, the MMPI-2, contain over 550 true-false, self-report, largely self referent items selected to measure ten clinical scales (Hypochondriasis, Depression, Hysteria, Psychopathic Deviate, Masculinity-Femininity, Paranoia, Psychasthemia, Schizophrenia, Hypomania, Social Introversion) and three validity scales. When the MMPI was developed, the use of face valid, self-report items was being questioned on the grounds that respondents could determine how they wished to portray themselves. As a consequence, a more empirical approach was followed. A large pool of items without obvious scale associations was administered to both a control and criterion groups (psychiatric inpatients). The items that were able to differentiate the criterion groups from the control groups were subsequently included in the MMPI. Validity scales were also added to specifically detect respondents who attempted to portray themselves either favorably or unfavorably. As a result the MMPI does a good job of what it was expressly constructed to do, help diagnose certain types of psychopathology.

Strengths: The strength of the MMPI lies in its norms. Raw scores on each scale are converted to standard t scores. Results are interpreted in terms of the scales having t scores at least one standard deviation above the norm. Typically, personality types are expressed in terms of the scales that are high, such as 2-4 (Depression-Psychopathic Deviate). This process of developing scale norms and typing personalities based on salient traits recognizes that certain traits can dominate an individual's behavior. A clinician skilled in interpreting the MMPI is able to describe abnormal personalities by reading the scale elevations. The MMPI has been demonstrated across cultures as well.

Weaknesses: Although there is a wealth of items to rate, the self-referent nature of these items make it unclear that even an intimate friend would be able to rate another. Further, the scales were not designed to be orthogonal. Inter-correlations among scales are high (nor surprising since some items contribute to two scales). Factor analyses have suggested that the MMPI measures only two factors, suggesting that it fails to capture personality fully.

The weakness that renders the MMPI inappropriate as a basis for the traits in the system is its failure to account for a high percentage of variance in the normal personality. The MMPI was developed to detect psychopathology. The terrorist leaders that will be modeled are likely high functioning individuals who would score in the normal range on the MMPI. Scores in the normal range, however, are thrown away. Thus, the items do not differentiate in a meaningful way among the population of interest.

Conclusion: The MMPI scales would not provide a useful basis for trait selection for the system in that it fails to characterize normal personalities.

NEO Personality Inventory (NEO-PI-R).

The NEO-PI-R is a theory based test designed to measure what has come to be known as the Big Five personality traits. The instrument, developed by Robert McCrae and Peter Costa, follows in the tradition of factor analytic approaches to personality definition embodying the work of J. P. Guilford (Guilford-Zimmerman Temperament Survey) and R. B. Cattell (16 PF). The original justification for exploratory factor analytic approaches to the study of personality is rooted in linguistics. The Lexical Hypothesis holds that adjectives that have been used to describe facets of personality are likely to refer to real traits. Moreover, more important traits will be described by a greater number of words. By determining the adjectives that tend to be used to refer to the same characteristic, i.e., that load on the same underlying factor, the existence of a trait can be inferred. The adjectives themselves then provide an indication as to what the factor should be named. Cattell found 16 factors accounting for the variance in personality descriptors. Later, these factors were reduced to four super factors when the sixteen factors were factor analyzed. Eysenck maintains that the Big Five can be collapsed to three that are measured by the EPQ. Zuckerman et al have proposed an alternative five factors measured by the ZKPQ-III. Interpreting factor analytic solution ultimately requires judgment concerning the number of factors that are meaningful and how they should be named. This debate over basic personality factors is likely to go on for some time. It appears to this reviewer, however, that presently the factors of the NEO-PI-R enjoy slightly wider acceptance.

The NEO-PI-R provides a description of personality along the dimensions of Neuroticism, Extraversion, Openness, Agreeableness, and Conscientiousness. The NEO-PI contains 280 items measuring the 18 facets of the N, E, and O scales as well as measuring the A and C scales. Items employ a 5-point Likert format. In order to avoid biasing the instrument, statements on each scale are evenly divided between stating an attitude positively and negatively. NEO-PI scale construction has been guided by both theory and confirmatory factor analysis. The broad, ambitious intent is to use the instrument to predict interests, health and illness behavior, psychological well-being, and characteristic coping styles. Studies have also confirmed the incremental validity of the NEO-PI over the MMPI in predicting certain deviant behaviors.

Strengths: Reported reliability is excellent. More importantly, research indicates that self-test scores have correlated highly with ratings by others, thereby increasing the likelihood that portions of the NEO-PI can be used to obtain personality ratings for individuals to be modeled. This result is not surprising because the Lexical Hypothesis virtually guarantees it by starting with words that have been used to describe others.

There is reason to believe that the five-factor solution is comprehensive. One criticism of factor analytic solutions is that a factor cannot be recovered if the items that would load on it have not been included. In this case, however, the Lexical Hypothesis and the

number of personality adjectives used in studies (around 500) provide some assurance that most of the important factors have emerged. There have been attempts to determine if words not tested, e.g. clever, would load on different factors. The results have not been clear. Overall, it appears that about 70 percent of personality variance can be accounted for by the five factors.

The NEO-PI, by virtue of the way it came into being, offers a wide range of behavioral associations. Each factor is derived from an array of adjectives and behaviors with varying levels of applicability. These adjective correlations provide more than a list of synonyms. The data can be extrapolated to more complex behaviors. For example, the factor labeled agreeableness shows a heavy negative loading for aggressiveness. The NEO-PI has also been demonstrated across different language groups (universality). The factors were selected to be orthogonal, are measured on a continuous scale, and have norms available. The NEO-PI uses t scores.

Weaknesses: The five-factor approach is missing an integrating theory that relates traits to behaviors. Research has focused on identifying the factors rather than understanding their interplay. McCrae and Costa believe that the five-factor approach will provide the basis of a new generation of personality theories. That sounds very plausible. For the last 70 years that has been an attempt to measure the traits that theories require. Perhaps it is time to develop theories around traits that have been measured.

The factor analytic approach results in traits that are somewhat general and diluted. This outcome is an artifact of the factor analysis process. The need to reduce the number of factors to a few necessarily results in clustering entities that are somewhat dissimilar. As a result, the names of the factors alone do not provide pure traits that suggest how the individual is likely to behave. Three of the factors are composed of six facets each. Depending on the composition of these facets, the interpretation of the factor can be different. Thus, to use the data one must be familiar with more than five factors, not unlike the need to understand the subscales of the MMPI.

Conclusion: The use of the NEO-PI is recommended. It is amenable to ratings by others, provides multiple associations that will be useful in developing the rule set, and is comprehensive and appears to hold across cultures. Further, it is on the path of much current research. If the system is to be viable, there will be a need to update it, as more learned.

SELECTION OF TRAITS

There is by no means agreement on the five factors. Eysenck maintains that there are three factors, Extraversion, Neuroticism, and Psychoticism. Zuckerman has a stronger behavioral-ethological orientation, maintaining that basic traits should be recognizable in animal behaviors. His alternative five includes sociability, neuroticism, impulsive unsocialized sensation seeking, aggression-hostility, and activity. Attempts have been made to reconcile the differences with mixed results. There is general agreement on extraversion and neuroticism, but then the numbers and names diverge. These attempts

at reconciliation have, however, produced further associations that should be useful both in obtaining personality ratings and in building the rule sets. For example, sociability and activity are highly related to extraversion. Psychoticism correlates positively with conscientiousness, impulsive unsocialized sensation seeking, and aggressiveness and correlates negatively with agreeableness. For this reason, the choice of construct names does not appear to be critical. What is critical is that the constructs are used to reflect what was measured. In order to remain consistent, however, it is recommended that Neuroticism, Extraversion, Openness, Agreeableness, and Conscientiousness from the NEO-PI be used.

One last point in justification of the use of traits derived from factor analysis is offered. It can be conjectured that at the genotypic level there exist heritable sequences related to variability in basic behavioral tendencies. Zuckerman would suggest that these sequences should have been highly conserved and therefore apparent in other animals as well. How then do these sequences manifest in specific behaviors observable at the phenotypic level? If the answer were as simple as relating different alleles to behavioral tendencies, these alleles would have been identified long ago. Instead, there are likely a multitude of epistatic interactions that mediate their expression and produce a dazzling array of phenotypic behavioral variations even without the considerable contribution of learning and experience. According to the Lexical Hypothesis the members of this array of behavioral tendencies have been named in all languages because they are important in communal living. Clustering these words describing phenotypic variation is tantamount to tracing the leaves of a tree back to its branches, limbs and trunk. At some point the structures that give shape to the tree can be identified. That's what the Big Five appear to do. That's why it seems the best place to start.

REFERENCES

- Barbuto, John E. (1997). A critique of the Myers Briggs Type Indicator and its operationalization of Carl Jung's psychological types. *Psychological Reports*, Vol. 80(2), pp 611-625.
- Bouchard, T. J. & Hur, Y. (1998). Genetic and environmental influences on the continuous scales of the Myers Briggs Type Indicator: An analysis based on twins reared apart. *Journal of Personality*, Vol. 66(2), pp135-149.
- Buss, D. M. (1996). Social adaptation and the five major factors of personality. In J. S. Wiggins (Ed.), *The five-factor model of personality: Theoretical perspectives*, pp 180-207.
- Caparo, Robert M. & Caparo, Mary Margaret (2002). Myers Briggs Type Indicator score reliability across studies: A meta-analytic reliability and generalization study. *Educational & Psychological Measurement: Special Issue: Reliability generalization as a measurement meta analytic method*, Vol. 62(4), pp 590-602.

- Costa, P. T., Jr., & McCrae, R. R. (1992). Four ways five factors are basic. *Personality and Individual Differences*, Vol.13, pp 653-665.
- Digman, J. M. (1997). Higher order factors of the big five. *Journal of Personality and Social Psychology*, Vol. 73, pp 1246-1256.
- Eysenck, H. (1992). Four ways five factors are not basic. *Personality and Individual Differences*, Vol. 13, pp 667-674.
- Furman, Adrian (1996). The Big Five versus the Big Four: The relationship between the MBTI and NEO-PI five factor model. *Personality and Individual Differences*, Vol. 21(2), pp. 303-307.
- Greene, R. L. (1991). *The MMPI-2/MMPI: An Interpretive Manual*. Allyn and Bacon: Needham Heights, MA.
- Hall, R. P., Ball, J. D., Stutts, M. L., Barth, J. T., Ingrisawant, L. & Turf, E. (2002). Correlations for scores on the 180 item version of the MMPI-2 and the Neuroticism scale of the NEO-Personality Inventory. *Psychological Reports*, Vol. 90(1), pp 227-230.
- Jang, K. L., McCrae, R. R., Angleitner, A., Riemann, R. & Livesley, W. J. (1998). Heritability of facet-level traits in a cross-cultural twin sample: Support for a hierarchical model of personality. *Journal of Personality and Social Psychology*, Vol. 74, pp 1556-1565.
- John, O. P. (1990). The "Big Five" factor taxonomy: Dimensions of personality in the natural language and in questionnaires. In L. A. Pervin (Ed.), *Handbook of personality: Theory and research*, pp 66-100. New York: Guilford.
- Kaplan, R. M. & Saccuzzo, D. P. (1993). *Psychological Testing*. Brooks/Cole: Pacific Groove, CA.
- MacDonald, D. A., Anderson, Peter E., Tsagarkis, C. I. & Holland, C. J. (1995). Correlation between MBTI and NEO-PI facets. *Psychological Reports*, Vol. 76(2), pp 449-450.
- McCrae, R. & Costa, P. (1989). Re-interpreting the Myers Briggs Type Indicators from the perspective of the five factor model of personality. *Journal of Personality*, Vol. 57, pp 17-40.
- McCrae, R. R. & Costa, P. T., Jr. (1996). Toward a new generation of personality theories: Theoretical contexts for the five-factor model. In J. S. Wiggins (Ed.), *The five-factor model of personality: Theoretical perspectives*, pp 51-87. New York: Guilford Paunonen.

- Mitchell, S. V. & Jackson, D. N. (2000). What is beyond the Big Five? Plenty! *Journal of Personality*, Vol. 68, pp 821-835.
- Saucier, G. & Goldberg, L. R. (1998). What is beyond the Big Five? *Journal of Personality*, Vol. 66, pp 495-524.
- Sharpe, J. P. & Desai, S. (2001). The revised NEO-Personality Inventory and the MMPI-2 Psychopathology Five in the prediction of aggression. *Personality and Individual Differences*, Vol. 31(4), pp 505-518.
- Zuckerman, M. (1992). What is a basic factor and which factors are basic? . *Personality and Individual Differences*, Vol. 13, pp 675-681.
- Zuckerman, M. (1994). Impulsive unsocialized sensation seeking: The biological foundations of a basic dimension of personality. In J.E. Bates & T.D. Wachs (Eds.), *Temperament: Individual differences at the interface of biology and behavior*, pp 219-255. Washington, D.C.: American Psychological Association.

APPENDIX B:
ESTIMATION OF FACET SCORES FROM DOMAIN SCORES

Estimation of Facet Scores from Domain Scores

Recall that Costa and McCrae subdivided each of the big five domains into six facets and then wrote items to measure each facet. The NEO-PI-R produces standardized T scores for each facet. The domain scores are then computed as an average of the facet scores within the domain. Although the six facets relate to the same domain, they are not completely correlated. An individual can be above average on one facet and below on another. Thus, many distinctly different facet score profiles can provide the same domain score. For example, consider the facets of Neuroticism, N1 through N6. T scores of (50, 44, 48, 52, 59, 61) and (61, 50, 52, 59, 48, 44) would yield the same domain score of 52.33. Yet the nature of these personalities would be different. The first is high in vulnerability and impulsiveness while being low in angry hostility. The second is high in anxiety and self-consciousness, but low in vulnerability.

The problem we have is that we would like to take advantage of the fine distinctions made by facet scores even though we only have estimates of the domain scores. How do we go about it? First, let me say that ultimately we will want to measure the facets in order to be accurate. There is a way, however, that we can estimate facet scores based on the domain scores. Costa and McCrae have provided information in the manual [NEO-PI-R testing material] that will allow us to develop a method to estimate facet scores. They factor analyzed data from all 30 facets to determine which ones tended to load on the same domain. They called the constructs factors as opposed to the domain scores derived from averaging the facets. The result was that there were some fairly significant loadings of facets on factors other than the domain in which they were defined. Let's look at an example. If we wanted to estimate a factor score for the domain N, we would combine the scores from the 30 domains in the following fashion.

$$\begin{aligned} N = & (.26N1 + .18N2 + .23N3 + .22N4 + .11N5 + .18N6 + \\ & .01E1 - .06E2 - .07E3 + .08E4 - .02E5 + .02E6 + \\ & .02O1 + .09O2 + .16O3 - .06O4 - .02O5 - .06O6 - \\ & .09A1 + .05A2 + .05A3 - .02A4 + .07A5 + .05A6 - \\ & .03C1 + .10C2 + .05C3 + .09C4 + .01C5 + .02C6) - 31.00 \end{aligned}$$

This tells us that the major determinants of factor N are the facet scores N1 through N6. From the small coefficients on the other terms, it is apparent that most of these other facets do not contribute significantly to N. There are some exceptions, however. O3 provides a strong contribution to N. Does this make sense? O3 measures depth of feeling. A person with strong feelings is not necessarily neurotic, but my sense is that it can help. There is a distinct emotional component to many neuroses. Behavior patterns are adopted early in life in order to regulate emotional experience. These patterns become the basis of personality traits that can be maladaptive later in life. Thus, it appears feasible to me that high O3 would be associated with high N. What about the facets with negative numbers? How do they contribute? Let's look at A1 (note the minus sign is on the preceding line). A1 is trust. The negative sign tells us that being low in trust is related to high N. That also makes sense. Costa and McCrae point out that the domain scores obtained from simply averaging the six facet scales and the factor

scores obtained by using all 30 facets are highly correlated, but that the factor scores provide slightly better markers of the big five because they benefit from the secondary loadings of other facets. I think we can see from this example that both aspects of this statement are true.

We're not there yet, however. We still need to get from domain scores to facet scores. If the factor scores are related to the facet scores, then the facet scores relate to the factor scores in the same way. Thus, N1 is related to N, but N1 is also related to each of the other domains as well. If we look at the formulas in the material, we find that N1 is related to N by .26 and to C by .09. We also see that the relationship between N1 and the other three factors is negligible (.03 or below). Let's assume that any contribution of .05 or below is too small to be considered. This cutoff is arbitrary. We can consider all of the numbers if it is desired. However, given that the factor scores are going to be behavioral estimates such precision is not merited. We now estimate N1 by taking roughly 3 parts and N and one part C by using the formula,

$$N1 = (.26N + .09C) / .35$$

If N and C are measured by T scores, then N1 will also be a T score. We are able to average them in this way because the variances are equal. Let's look an example.

Suppose an individual has a score of N=60, E=40, O=40, A=60, C=40. Let's say we want to estimate facet N1, anxiety, N2, angry hostility, and N5, impulsiveness.

The formula for N2 is:

$$N2 = (.18N + .09C + .12(100 - A)) / .39$$

The formula for N5 is:

$$N5 = (.11N + .16E + .06(100 - O)) / .33$$

Note that the term related to the O domain is negative in Costa and McCrae's table. In order to include its influence on N5, I added its negative. Because we are using a normal distribution with mean of 50, a 40 on openness is the equivalent of (100-40) on "not openness".

If we substitute values in the equations, we obtain N1=54.86, N2=49.23, N5=50.3. Thus, we are able to obtain facet scores that differ from each other. Further, if we had simply let the facet scores equal the domain scores, then all would have been 60, substantially different from what we obtained. The question is, do these scores make sense?

N1: The 60 score on N was moderated by a low score on C. Does an individual high in anxiety tend to be high in conscientiousness? Probably. Obsessive compulsive people tend to both conscientious and anxious. An extreme conscientiousness serves to reduce

anxiety. The low score in C indicated that the score on N1 was probably lower than the overall domain score on N. Thus, we can estimate N1 as lower.

N2: Part of this formula is the same. It is telling us that high C tends to moderate angry hostility. This result is not as obvious until we read through the facets of C. Is an angry-hostile person orderly, dutiful, self-disciplined, achievement striving? Maybe. High scores in these areas could apply to a person who is rigid, unyielding, and intolerant. It would seem reasonable for high C to add to N2. What about the third term? The 60 on this domain tells us that the person is agreeable. Agreeable people are hardly angry and hostile. In this case we are adding “not agreeable” by taking the complement of a high agreeable score of 60. Thus, we are averaging in a low “not agreeable” and bringing the angry hostility score down. Even though the individual scored 60 on the N scale, he is only average in angry hostility because of his high agreeableness score and low conscientiousness score.

N5: This formula says that impulsiveness tends to be increased by high scores on E. Again, the facets of E include activity, and excitement seeking. This relationship is also feasible. The relationship between N5 and openness is negative. Thus, a small amount of “not openness” increases N5. This one is more difficult to explain. The only explanation I can provide is that a person not open to fantasy, ideas, feelings, etc. is more likely to seek emotional and intellectual satisfaction through impulsive acts. Again, this effect is very small.

In summary, by using the scores on the other domains, we can correct scores within specific facets. I reiterate that this does not yield accurate facet scores. However, by adjusting the facet scores in this way, we obtain better estimates than we would have by using the domain score as an estimate for every one of its facets.

The last page contains all of the formulas.

How can the formulas be used?

I would expect that the user would be queried for a personality profile. Placing the cursor over N,E,O,A,or C would reveal a description of the domain and request entry of a score from 20 to 80 with a description of what specific ranges of scores mean. A domain description will be provided to help users rate. The scores are as follows:

20-35 = Very Low

35-45 = Low

45-55 = Average

55-65 = High

65-80 = Very High

In the NEO-PI-R raw scores on the various domains have been normalized and scaled to T scores to provide a means of comparing scales. The user scale above reflects T scores.

Based on the domain entries we can then compute the facet scores by using the formulas as described above

FORMULAS: (Be careful of the O's)

$$N1 = (.26N + .09C) / .35$$

$$N2 = (.18N + .12(100-A) + .09C) / .39$$

$$N3 = (.23N) / .23 \text{****for some we must simply use the domain score as an estimator***}$$

$$N4 = (.22N + .07C) / .29$$

$$N5 = (.11N + .16E + .06(100-O)) / .33$$

$$N6 = (.18N) / .18$$

$$E1 = (.21E + .12A) / .33$$

$$E2 = (.06(100-N) + .24E + .09(100-O) + .09(100-C)) / .48$$

$$E3 = (.07(100-N) + .10E + .12(100-A)) / .29$$

$$E4 = (.08N + .15E + .09(100-A) + .13C) / .45$$

$$E5 = (.21E + .06(100-O) + .11(100-A)) / .38$$

$$E6 = (.24E) / .24$$

$$O1 = (.23O + .08(100-C)) / .31$$

$$O2 = (.09N + .12(100-E) + .34O + .08A + .08C) / .71$$

$$O3 = (.16N + .07E + .17O + .08C) / .48$$

$$O4 = (.06(100-N) + .22O) / .28$$

$$O5 = (.14(100-E) + .35O) / .49$$

$$O6 = (.06(100-N) + .21O + .07(100-C)) / .34$$

$$A1 = (.09(100-N) + .16A + .08(100-C)) / .33$$

$$A2 = (.20A + .07C) / .27$$

$$A3 = (.19E + .09(100-O) + .16A) / .44$$

$$A4 = (.23A) / .23$$

$$A5 = (.07N + .19A) / .26$$

$$A6 = (.08E + .20A) / .28$$

$$C1 = (.16C) / .16$$

$$C2 = (.10N + .09(100-O) + .24C) / .43$$

$$C3 = (.07(100-E) + .06A + .21C) / .34$$

$$C4 = (.09N + .06(100-A) + .25C) / .40$$

$$C5 = (.21C) / .21$$

$$C6 = (.14(100-E) + .18C) / .32$$