COGVIEW & INTELNET: Nuanced Energy-Based Knowledge Representation And Integrated Cognitive-Conceptual Framework For Realistic Culture, Values, and Concept-Affected Systems Simulation

Daniel J. Olsher, Member, IEEE Cognitive Science Program Temasek Laboratories, National University of Singapore dan@intmind.com

Abstract—An increasingly important AI frontier is the ability to represent worldviews, culture, values, and other nuanced structures, and to simulate the effects of these on perception, emotion/affect, judgment, and opinion formation. Such information, however, is notably difficult to model and represent, due to its fine-grained, diffuse nature. Reasoning is also highly challenging in these domains. This paper presents a novel 'Energy-Based' Knowledge Representation formalism (INTELNET) ideal for modeling, fusing, and reasoning about nuanced semantics, cultures, affects, and worldviews. It then introduces the integrated COGVIEW conscious/unconscious psychological simulation framework operating on top of INTELNET and advances a detailed example within the suicide terrorism domain.

Applications include intelligent reasoning systems, humanitarian missions, cultural simulations, knowledge engineering, language processing, anti-discrimination and prejudice reduction, terrorism reduction, and norm change efforts, among others.

Keywords—Culture, Belief, Emotion Simulation, Norm Change, Energy-Based Knowledge Representation, Reasoning, Cognitive Psychology

I. CULTURES, BELIEFS, CONCEPTS, AND MORE: INTELNET ENERGY-BASED KR (EBKR) AND THE COGVIEW FRAMEWORK

The ability to create agents that can reliably integrate and simulate information regarding worldviews, including culture, values, and other conceptual structures represents an increasingly important frontier in artificial intelligence. [1], [2], [3]

Worldviews mediate perception, but are notoriously difficult to model and to represent, in part due to their nuanced, diffuse natures. In computational systems, moreover, it is difficult to define, bound, and reason about whatever 'it' is that culture and worldviews consist of. Beyond this, effects of human cognition add another layer of complexity.

This paper presents a new knowledge representation formalism, *INTELNET*, expressly targeted at nuanced data generated by systems grounded in and arising from the human cognitive capacity. INTELNET is ideally suited to worldview modeling and the nuanced semantics underpinning natural language, enabling simulation of the effects of worldviews on reasoning, persuasion, and understanding.

In order to facilitate enhanced reasoning capabilities, INTELNET seeks to model *nuanced semantics*, representing information with a symbolic opacity intermediate between that of neural networks and typical symbolic systems. INTELNET concepts are represented as *distributed*, *interconnected networks* wherein each part of a network operates in concert with others to define a concept and to model the meaning of a particular semantic domain.

Beyond nuance, computer worldview models are faced with certain 'quirks' of human psychological processing which lead complex conceptual systems to behave differently than they otherwise would if they were not being generated via human cognition. To this end, the paper presents a framework covering specific psychological phenomena that integrate with the INTELNET formalism in order to enable detailed simulations. INTELNET, together with the integrated psychological models discussed in this paper, are referred to as the *COGVIEW* framework.

Potential applications of INTELNET and COGVIEW are extremely wide, including advanced reasoning, perception, persuasion, social systems modeling, and natural language processing. COGVIEW-based cultural networks represent the ways in which specific cultures have concretized and reified meaning components and how these components interact to produce culturally-mediated judgments and behaviors.

COGVIEW/INTELNET is ideal for modeling cultures as it is able to represent the specific conceptual semantics present in particular cultures at a much greater level of specificity and clarity than traditional ontologies, accurately preserving a major source of cultural influence on model outcomes. It is capable of representing essentially infinite nuance by seamlessly combining smaller semantic representations while preserving key semantics.

Moreover, the representation can be made aware which concepts are highly valued and which are stigmatized, and is able to make this information relevant throughout simulation by modulating valences, magnitudes, and other energy aspects. When semantic components interact with previously stigmatized or valorized semantics, the model is able to extend this 'colorization' to connected components.

Taken together, the ability to model: 1) unique interconnection structures between concepts, 2) the nature and semantics of concepts themselves, 3) judgments, 4) expectation/value violations, and 5) the degree to which stimuli generate energy outcomes in accordance with culturally-demanded energy distributions provides a powerful system for cultural simulation.

Previous work ([4], [5], [6], [7]) points towards how graph structures may be used to compute on these representations, identify likely areas of discord and moral questions, design norm-changing communications, and make predictions about how various stimuli are likely to be viewed by members of various cultures.

INTELNET nuance and flexibility also allow for *fusion*, the interconnection of information sources from various domains. In [7], a cultural network is connected to an emotion model, and in ongoing work multiple commonsense knowledge sources are fused and used together to enable reasoning.

II. WHY A NEW KR MECHANISM?

The motivation for Energy-Based Knowledge Representation (INTELNET) and COGVIEW stems from the notion that information originating from psychologically and conceptually-mediated social patterns, principles, and processes (such as *cultures, worldviews and natural language semantics*) possesses unique properties, requiring specific tools in order to be modeled felicitously. In this paper, such information is termed *Cognitively Mediated Process Data*, or CMPD. 'Conceptually mediated' refers to processes that unfold differently in practice depending on the specific concepts present in some knowledge base and on the specific ways in which those concepts are interconnected. An example would be moral perception, in which the specific concepts one has regarding virtue and vice and the implications attached to each of these all play a major role in determining how one will view a particular phenomenon. Similarly, 'psychologically-mediated' evokes phenomena which, in order to be successfully modeled, require reference to the functioning of specific psychological processes. A key example is word association; given the concepts 'sun', 'sand', 'ocean', 'water', and 'waves', a human would immediately reference and activate the concept 'beach', while a computer would *a priori* have no reason to do so. In cases like these, models with no capability to model semantic association would predict incorrect outcomes for incoming stimuli.

Davis, Shrobe and Szolovits [8] suggest that knowledge representation formalisms fulfill a number of roles, including that of a world surrogate enabling reasoning instead of action, ontological commitments suggesting the terms in which we view the world, a theory of intelligent reasoning, a definition of what is 'natural'/easy to say, and a medium of human expression (we can only 'say' what our KR allows us to).

In CMPD-mediated contexts, the goal of INTELNET and COGVIEW is to allow us to say what we need to in a nuanced domain-appropriate manner, to take cognition into account, and to conduct complex reasoning. When modeling cultures and CMPD-heavy contexts, the simplifications required to use symbol- and ontology-based KRs often force us to leave out much of the semantic detail making the system what it is. Standard relations are semantically opaque, generating frame problems and making it difficult to combine multiple pieces of information. Generalizing (repurposing, reconstruing) information during reasoning become difficult as well. Left-out detail makes advanced reasoning difficult indeed, contributing to the perception that deep cultural modeling is an intractable problem.

COGVIEW/INTELNET has also demonstrated the simulation of emotional/affective outcomes [7], using knowledge to generate likely appraisals and predict pleasantness, level of emotional engagement, and other key parameters. (see also [9])

In line with the Intrinsic Cognitive Models (ICM) [10] view of mental representation, most cultural and CMPD-grounded contexts involve mostly implicit information. In INTELNET the goal is to store information implicitly and extract it implicitly when you need it, leading to "boundless" inferences ([10]). INTELNET/COGVIEW networks' contextual and dynamic ability to build larger representations from smaller ones is ideal for making implicit knowledge useful during reasoning.

In CMPD contexts, locally and contextually-correct inference often displace the need and desire for complete inference (see, for example, [11]) in that information is widely distributed and conflicting and 'truth' becomes much more relative and contextual in these domains.

A. Nuance and Psychological 'Quirks'

Because it is created by humans operating within the complex human semantic capacity (the capacity for making meaning), CMPD nearly always draws on significantly interconnected concepts within expansive knowledge bases and is defined by *subtle nuances and unlimited meaning gradations* or 'shades of gray'.

CMPD is 'fractal' in the sense that the concepts that comprise it have complex internal structures referencing multiple other concepts, which again possess complex structures, and so on. CMPD is always highly contextualized; intelligent reasoning and social processes take place within and are highly dependent on the nuances of surrounding contexts (which are themselves complex and nuanced, involving implicit reference to knowledge about the world).

With CMPD, the interconnection structure between constituent concepts is a key part of meaning; the *pattern of* *relationships* between a concept and other concepts forms a major part of the original concept's semantics.

CMPD may only fully be modeled by systems capable of handling 'quirky' behavior explicable only with reference to human cognitive functioning, such as priming, associative memory for concepts, emotions, short-term memory limitations, and unconscious processing. Each such 'quirk' exhibits significant influence on the functioning and effectiveness of decision making, perception and persuasion; as an example, in many contexts concepts contribute to perception formation in proportion to the energy they receive. In order to model such perceptions, modelers must pay attention to the (psychological) processes by which various regions of knowledge structures receive energy and contribute to simulation outcomes, and must know specifically how concepts are constructed and interconnected.

A key goal of the COGVIEW architecture is to make various effects of these 'quirks' more explicit by first identifying which psychological processes are most critical and then demonstrating how these processes may be modeled on top of INTELNET.

III. NOVEL KR MECHANISM: ENERGY-BASED KR (INTELNET)

Using purely symbolic tools, it is difficult to represent deeply nuanced, highly interconnected semantics; because symbols are highly granular, with bright-line separations between them, symbolic KR often requires knowledge designers to abandon much of the information otherwise implicit in problem domains because the KR does not offer an easy way to represent it. As a consequence, purely symbolic systems are often unable to perform beyond the original intention and mindset of the knowledge engineer - that is to say, they cannot *reconstrue* the world in new ways based on dynamic task demands. For example, a system which understands a TABLE only as a piece of FURNITURE will not be able to reconstrue it as being capable of serving as SHELTER (something one may hide under) in a context which demands this.

Symbols are opaque, without internal semantics or information about how various aspects could be reused or modified in novel contexts. Neural networks, on the other hand, operate at a level of abstraction too far below concepts to be able to easily replace them in everyday use.

The alternative presented here, 'energy-based' KR, uses *energy flows* between various portions of the larger semantic definition of a concept to connect those components together, permit them to interact and affect one another, and to work together in order to 'build up' a concept definition.

During reasoning, various regions of the concept definition 'space' may be selected according to current needs, allowing contextualized concept reconstrual. INTELNET allows for maximum extraction of the knowledge implicit in any given domain, enabling systems to solve problems unlike those anticipated by the designer and/or that they may have seen before, and allowing for minimal 'pre-cognizing' of problem domains by knowledge engineers. The high 'grain' (capacity for fine detail) of INTELNET representation permits the development of advanced reasoning techniques.

Energy-based KR involves extended spreading combination and recombination of semantic subcomponents within traditional concepts and other normally opaque semantic building blocks.

Concepts as assembled under INTELNET possess many desirable properties, mirroring the real-world functioning of human concepts and cognitive processes and providing sufficient nuance to serve as a base for highly advanced reasoning.

Any KR is only capable of answering certain questions, especially if it demands particular simplifications or is unable to represent a particularly salient part of a problem (such as the psychological processes underlying social models). Systems with finer representational grain are better able to adapt data to new uses as they have access to enough information to make informed decisions about what sorts of new conclusions may be drawn from the original data.

To this end, in INTELNET all information (knowledge, concepts, interaction patterns, and more) is stored in *networks*, defined as augmented graphs. Each element may be 'elaborated' with networks; *even nodes, edges, and so on may all contain networks within them* specifying the semantics of those elements. Opaque symbols are avoided as far as possible, with networks 'all the way down'.

INTELNET network edges channel energy, modify it, and connect disparate concept spaces, concepts, and other elements (as defined below) together. Representations are built up by inserting certain types of energy at start nodes, permitting energy to spread outwards, and collecting results once energy has spread to a sufficiently wide radius or to particularly interesting nodes. The result is a collection of nodes, edges, and a record of the amount and type of energy arriving at each node at each point in time. Analysis is made of energy flows, especially 'clashes' (described below), which occur when energy of one polarity meets energy of another. Clashes are highly meaningful and important, often pointing to semantically significant areas within modeled domains (see examples below).

IV. RELATED WORK AND DISCUSSION: WHAT INTELNET IS... (Hint: It's Not What You Think)

INTELNET and COGVIEW represent fundamentally new approaches to questions of knowledge representation and belief modeling. To support this claim, this section discusses related work, briefly considering key differences vis-à-vis the present approach.

A. Spreading Activation, Marker Passing

Spreading activation, in both its biologically-inspired [12], [13] and semantic network-employed forms, "[serves] the function of quickly spreading an associative relevancy measure over declarative memory". [12].

In COGVIEW/INTELNET, linking one concept to another does not mean that one concept merely *activates* another (raises its relevancy measure), however - it indicates that a concept *interacts* with another to co-create semantic fields and influence other concepts and energy flows. This is a key distinction.

Energy binds smaller semantic components together, driving a dynamic reification process that in essence creates new concepts as it goes. Energy flows carry the action of various information sources (and functions upon these) throughout the network, using network links to distribute information and interaction, binding the components of larger semantic fields together during traversal. Energy quanta include implicit semantic representations of upstream nodes and links, including the means by which they were generated and the judgments attached to the sources of that energy.

Energy quanta do not represent activation; rather, they reflect the summation of activities occurring at upper levels, representing a 'motive force' enabling nodes to participate in semantic simulation. Energy entering a node enables that node to 'do' things like modifying energy flows and sending energy to other network elements. There is no decay; energy flows until it reaches leaf nodes and all loops have been exhausted.

When modeling cultures and beliefs, a key information source is *emotional engagement* - the level of positive or negative emotional 'force' attached to particular concepts (such as family, taboos, and so on.) Emotions drive significant amounts of behavior, suggesting that simulations must be aware of how stimuli generate emotions and 'colorize' other stimuli. The more emotional engagement attached to a particular stimulus, generally the more INTELNET energy it will be able to send to other nodes. Negative energy reaching a node indicates that that node is part of a larger semantic field viewed as unfavorable, impossible, incompatible, or inappropriate, and vice versa for positive. While spreading activation involves ontology traversal, INTELNET energy spreading is a complex process of building up larger wholes from smaller pieces, keeping track of change as interaction and building progresses, identifying the contribution of clashes between energies of differing +/- valences, and calculating the degree to which energy distribution targets have been met once traversal is complete.

The location in which an energy clash occurs within a network is highly meaningful, often pointing to semantic controversies within the target domain and/or violations of important expectations. The presence of a clash during a cultural simulation suggests that the simulated subject matter will evoke (likely strong) emotions, and, *critically, both the subject matter as well as the incompatibility evoking those emotions can be identified based on the location of the clash within the graph.*

A major component of culture is the arbitrary ways in which concepts are associated with one another. In cultural networks, energy is not distributed over ontological links, as it would be with spreading activation, but instead over links bridging important cultural concepts in ways *described and demanded by the culture itself, not by any pre-existing ontology*. Culturally-specific reification processes generally exert immense effects on the final outcomes of culturally-mediated behaviors and simulations.

INTELNET allows for the modeling of culturally-desired energy distributions - another key cultural component. Specific cultures demand that varying amounts of energy be associated with specific semantic components, and will judge stimuli based on how accurately this takes place. In Asian cultures, for example, 'family' would be expected to receive significant amounts of energy subsequent to a stimulus, much like 'freedom' in American culture.

In COGVIEW, memory retrieval is only touched upon implicitly in that nodes and links are all understood as being ready to influence functioning and outcomes once energy reaches them. Concepts and potential semantic pathways are important to the extent that they are able to participate in and influence energy flows.

This process is intended in part to model the flow of unconscious cognition (discussed in section VI-A below), the way in which concepts interact with and affect one another, the notion that things that are labeled 'good' affect those which are labeled 'bad' and vice versa, and the ways in which societal and cultural approbation reach and interact with specific concepts.

With respect to marker passing [14], instead of traversing ontological links across pre-existing concepts, energy traversal binds together and creates *new* semantic entities. INTELNET/COGVIEW energy quanta are semantically richer than markers, including information about the semantics of prior traversals, magnitude, valence, emotional charge, legitimacy, and approval, and can only be created by energy sources (which impart specific characteristics to the energy they generate).

B. Semantic Networks, Conceptual Graphs, Implicational Networks, MultiNet

Critically, COGVIEW/INTELNET graph elements do not reflect ontological data and do not represent logic or first-order predicate calculus (FOPC) relations. This makes them highly semantically distinct from networks which fundamentally rerepresent FOPC in graph form.

A key goal of INTELNET and COGVIEW is to provide an AI representational system with a capacity for enhanced nuanced and natural representation of cognitivelyand conceptually-mediated systems compared to FOPC.

Semantic network edges typically represent logical predicates and are otherwise generally semantically opaque, taking forms like 'Is-A' or 'Has-A', and involving relations such as meronymy, holonymy, and so on. INTELNET edges, on the other hand, involve the channeling and modification of energy. From time to time, edges may also add particular semantic elements to energy passing through them (cf. [5]); one example would be the marker CHOICE indicating that energy channeled through that edge originally derives from a choice-related context and that downstream nodes and edges will be construed such as to construct that choice.

Semantic network concepts do not have internal structure and are taken in a fundamental sense to be primitive and aligned in some meaningful way with specific, definable 'things' in the wider world. COGVIEW concept nodes, on the other hand, are 'shorthand handles' for the complexity within them and the network locations in which they are embedded. A key function of language is to reify diverse semantic aspects (what would be called subnetworks here) into unified wholes, providing 'handles' (lexical items) which can be used to call these reified wholes into consciousness and allow them to participate in discourse.

In order to model this, in INTELNET/COGVIEW all concepts are viewed as having internal structure and complexity and particular 'concepts' may arise or fall out of use depending on how simulation-driven reification proceeds, making them highly context-dependent. A well-known example (after Whorf) would be the claim that there is a large range of words for various types of snow in certain Inuit languages; this claim suggests that the cultural demands of that environment has caused the reification process in that domain to progress differently than it has in other domains, generating different 'concept handles'.

A final approach, MultiNet [15] is noted here as representing an interesting decompositional approach to some extent intermediate between logic and the type of nuanced representation advocated in this paper.

C. Dependency Graphs, Petri Nets, Causal Networks

While causal information may certainly be represented within the formalism presented here, COGVIEW/INTELNET nodes and links do not represent actions, states, or explicit causal chains that generate states, as in dependency graphs [16], nor do they represent transitions, states or causes (Petri nets and causal networks [17]). Some causal information can be drawn from energy flows in that they reflect the confluence of previously traversed semantics (a form of cause in that these ultimately drive simulation outcomes), but, generally, no causal information need be explicitly represented.

D. System Dynamics

Energy is not 'material' in the system dynamics sense; rather, it is a connective force, binding units together. INTELNET energy does not represent any sort of 'stuff' or tangible physical quantity, and it is not zero-sum. Energy transfer, unlike the transfer of liquid from one tank to another, represents an interaction between two nodes across a particular kind of edge with its own particular characteristics, bringing new nodes into dynamic processes established by earlier network traversals.

Concept nodes do not represent variables or participants in physical processes and causes and effects are not explicitly specified; rather, energy is used to discover those that are implicit in the network, taking full account of culture and concepts. There are no stocks and no rates. If magnitude 20 energy enters a node, the same amount of energy will leave on each outbound link. Energy combination at a particular node is not the combination of various stocks or quantities - it represents the combination of the semantics signaled by participating energy flows and a locus for interaction signaled by these flows' coalescing at a particular location.

COGVIEW/INTELNET is not a process model and system dynamics need not be explicitly specified, arising as they do through the interaction of various semantic components, stimuli, and energy flows. Much of which would be expressed explicitly in a system dynamics approach (dynamic behavior, interaction, and so on) emerges during simulation in the present paradigm. Feedback is generally one-way and not in the form of loops, though some loop support is provided in order to represent key cultural phenomena in which individual concepts support one another, making all constituent concepts more important.

The INTELNET/COGVIEW approach could serve as a semantics representation formalism for social system dynamics models, providing insight into the ways in which culture and knowledge contribute to overall dynamics, and could also benefit from model validation techniques developed under this paradigm. One such area involves modeling of perceptions in humanitarian contexts [6] and development of policy recommendations based on such simulations.

V. TECHNICAL ASPECTS AND DETAILED EXAMPLES

A. Energy

The concept of *Energy* defines the core of the Energy-Based KR (INTELNET) paradigm.

Energy arises from *Energy Sources*, generally created via particular concepts (such as LOVE, FAMILY, AVOIDING PAIN, DESIRE TO OBTAIN MONEY/RESOURCES, and so on, that are instinctual in nature and capable of marshaling great amounts of psychic and physical energy towards their acquisition and defense.

'Energy flows' between components (nodes, concept fields, and so on) are guided by directed edges, with energy always flowing in the indicated direction (except, perhaps, during reasoning, in which INTELNET networks may be traversed in unusual ways).

Energy flows between disparate conceptual fields, beyond transferring energy, have the function of connecting those conceptual fields together. Information flows from one part of the concept field to another, semantically integrating all concepts involved in a manner based on the dynamic needs of the domain being modeled and the problem currently being examined. Context is taken into account in that relevant qualities of energy found upstream (at previous timesteps) are transmitted downstream and affect downstream nodes in a highly dynamic way as simulation progresses. Context changes upstream may lead to non-obvious effects on other nodes.

The energy reaching a given component is understood as a measure of both its desirability (positive energy means that a concept is 'approved of', 'should' happen, is 'useful' to the goal, and other 'positive' evaluations) as well as its current level of spreading activation.

During simulation, subjects are expected to seek to maximize the positive energy achievable in any particular network, and are *a priori* expected to exhibit behavior grounded in those concepts. For example, if PARENT receives significant positive energy in the belief network of a particular subject, we would expect that subject to act in accordance with the roles and expectations of a parent and to avoid contrary actions.

As an initial example, in Fig.1 the role of PARENT and the act of RAISING CHILDREN both send energy to the concept of IDENTITY, which then sends that energy to PERSONAL SIGNIFICANCE, WORTH, AND MORALITY. This suggests that parenting and raising children contributes to identity strength and positivity, which then in turn contributes to personal significance and worth. In addition, COMMUNITY is a source of identity and personal worth enhancement; the more positive energy (representing strength and importance of community) associated with this concept, the more identity and personal worth will be enhanced.

Fig.3 presents an example within the suicide terrorism domain, demonstrating how concepts as disparate as the desires to preserve life, to be moral, to work, to earn a living for one's family, and to obtain personal significance in one's life may all be brought together, interacting with and affecting one another. (Note: 'propel' as an edge label indicates a node sending (that is, 'propelling') energy to other nodes.) Drawn from the more comprehensive example in Figure 5, the figure represents two different sets of beliefs, denoted with solid and dotted lines.



Fig. 1. Conceptual Participation in Identity



Fig. 2. Distributing Energy Related to Personal Significance

Solid lines denote beliefs held before one is persuaded to adopt suicide terrorism, and dotted lines those expected to hold post-persuasion.

Pre-persuasion, the desire to preserve life and to be moral both send highly negative energy to the concept SUICIDE TERRORISM (suggesting it is highly undesirable and unlikely to be adopted behaviorally) and terrorism is not connected to the desire to obtain financial support for one's family. Postpersuasion, PEACEFUL SIGNIFICANCE ACHIEVEMENT begins to receive negative energy and death is seen as a more desirable option, as links sending negative energy to the concept DIE are broken and SUICIDE TERRORISM begins sending more positive energy to these alternative concepts.

In Fig.2, energy arising from the need for personal significance is parceled out to various concepts, crossing edges which gate the total energy reaching each node to 1/5 of the total sent.

1) Structure of Energy Quanta: As alluded to above, INTELNET energy has internal structure; a given quantum of energy contains Descriptive Quantities and Semantic Aspects. Each quantum may be described as follows : Descriptive Magnitude (on the interval [0,1000]), Descriptive Polarity ([-1,1]), Descriptive Approval ([-1,1]), Aspect Approval ([-1, 1]), Aspect Legitimacy ([-1, 1]), and Aspect Focus ([-1, 1]). Magnitudes are relative, with meaning only in relation to other magnitudes in the same graph.

Descriptive Magnitude and Descriptive Polarity, respectively, describe how much energy is present and whether that energy is positive or negative. Positive energy entering a node suggests a positive promotion of the associated concept. For example, pushing (in INTELNET terms, 'propelling') +50 energy into the concept FAMILY suggests that FAMILY is important and good and that both it and concepts associated with it should be promoted. It also means that FAMILY is important in that it is able to transfer its energy to other connected concepts, endowing it with the capability to generate significant effects during simulation.

Approval is present both as *Descriptive Quantity* and *Aspect*. In general, approval represents the general 'emotional quality' of particular energy, ranging from highly positive to highly negative, conveying information about the levels of social approval attached to nodes traversed thus far. Approval can be understood in some sense as an approach/avoidance measure (high approval=high desire for oneself and/or others to approach the concept at hand).

The Approval *Aspect* is a diffuse measure of general judgment, often generated at the energy source. As an example, a highly negative concept like MURDER would generate negative Approval aspect, which would be propagated downstream to represent the notion that any phenomenon (even if it is otherwise positive) resulting from MURDER should be viewed negatively (that is, as 'tainted') as well.

The Approval *Descriptive Quantity*, on the other hand, is tied to specific energy flows, marking the quality of a specific quantum of energy as it traverses the network and encounters nodes with varying semantics. An example would include a node which receives energy after passing through a number of nodes with high Approval, which would then spread.

Aspect Legitimacy finds use with respect to particular broad social concepts or competency/social acceptance judgments made about powerful actors. This can be contrasted with Approval, which is allocated based on the outcome of *moral* judgments.

Finally, *Aspect Focus* is used to indicate the insight, validated in priming studies, that concepts that have been recently referred to retain a greater role in information processing (that is, some of the energy applied to them 'lingers'). Work on focus magnitude decay schedules is ongoing, grounded in previous work on priming. Focus represents energy transferred from upstream nodes that should be expected to decay over time, as opposed to regular energy, which does not have a time component.

Edges exert effects on the energy flowing across them. Examples include 'gate' edges which limit energy and the NEG and FAIL edges which reverse the valence of the energy flowing across them while maintaining the same magnitude of flow.

A very wide range of models of social phenomena, including judgment and prejudice, can be generated through the paradigm put forth here. A spreading-energy-based view of *disapproval*, for example, leads to several corollaries. One is the notion that if a certain part of a concept or concept field is 'disapproved of' (receives negative energy) then the spreading of that negative energy will generally mean that other parts of that concept *and*, *critically*, *interconnected and/or related concepts*, will receive some disapproval energy as well. During argumentation, an oft-heard claim is that if a particular phenomenon is 'approved of', then this will lead to an effectively infinite number of potential negative consequences – the 'slippery slope'. In the INTELNET + COGVIEW framework, this can be understood as effectively equivalent to a claim that preventing the application of negative legitimacy/disapproval energy to a particular portion of the concept network will stop disapproval from flowing to related concepts that previously received negative legitimacy (stigma).

2) Concepts and Concept Fields: Concepts and Concept Fields (or Spaces) are generally created through decomposition of the semantics of a particular natural language word or concept, such as 'Father' or 'Community Values'. Concepts include information about what other concepts a particular concept tends to be associated or to reoccur with, connections to emotions and emotion-laden concepts, detailed information about the semantic extension of the concept, and so on.

A concept node represents the reification of detailed concept semantics as a whole entity and serves as the starting point for a concept definition.

Groupings of related, interconnected (and 'inter-defining') concepts generate *concept fields*, defined as delineated subregions with semantic coherence sufficient to define them as 'meaningfully' connected. As an example, the concept field BACHELOR, far from being a simple conjunction of the predicates UNMARRIED and MALE, is embedded within a cloud containing concepts related to common expectations of bachelors, marriage as a social phenomenon more generally, courting, and so on.

Meaning is stored in a 'cloud' of interconnected information components, which assists in modeling the fractal nature of concepts. Fields are not separable; should any portion be removed, some important aspect of that concept's meaning must also necessarily disappear. Figure 5 provides examples of two concept fields: COMMUNITY and IDENTITY.



Fig. 3. Suicide Terrorism Pre- and Post-Persuasion

Generally, when multiple semantic components are reified into (viewed as) larger groupings of concepts with a single name, this exerts critical effects on cognition, enabling complex wholes to be evoked via a single word retaining while internal complexity (complexity which affects the behavior of models in critical ways) within larger processes. INTELNET decomposes and represents the semantics underneath such reification while maintaining the general conceptual identity of the concepts within which meaning is stored.

Concepts build meaning in part through the collectivity of the other nodes that link to them and that they link to. The ability to define concepts in this manner is the source of much of the present approach's power and of its ability to represent fine-grained semantics.

a) Importing Concepts From Other Fields: Concepts and/or concept subfields (or even complete fields) are often 'borrowed' or 'imported' from other concept fields. In this way, concept fields are built out of others, creating a rich semantic 'tapestry'. Many interesting and complex effects arise when this interrelated tapestry structure interacts with activation flows, enabling models of complex social process effects not easily explained with traditional techniques. Effects such as these offer a rich path towards determining model correctness, offering predictions of 'telltales' that should be observable if a domain has been modeled correctly.

B. Loops

Normally, INTELNET networks are intended to be loopfree in order to ease the development of reasoning algorithms. In cases of highly entrenched conceptual networks, however, loops may be generated, within which concentrated energy reinforces and strengthens links between member concepts. Energy emanating from one concept reinforces others, and so on.

Attempted changes to looped links are more disruptive to overall energy maximization and thus such links are more resistant to change. In Figure 5, a loop is present between the concepts FAMILY, COMMUNITY, RAISE CHILDREN, and other nodes denoted by a long-dashed line. The presence of a loop suggests that 'family' is especially important in this model, that this concept is capable of marshaling significant energy, and that this state of affairs cannot easily be changed.

C. Clashes

In COGVIEW, a *clash* occurs when energy flowing in one direction meets energy flowing in the other. Clashes are especially important when the valences of the clashing energies are opposed to one another. Clashes are a key mechanism by which the network-based architecture underlying COGVIEW is capable of creating specific 'results' during processing activities. Clashes represent the point at which conceptual incompatibilities have become manifest and when the ability and/or need to take some sort of meaningful action has been identified, comprising one of the most important types of 'conclusion' that can be drawn under the INTELNET paradigm.

In Figure 5, clashes are identified with asterisks (*) placed next to the labels of concepts that serve as clash sites; the concept SA [SIGNIFICANCE ACHIEVEMENT]: SUICIDE TER-RORISM is the locus of three clashes, as denoted by the multiple asterisks (see below for further analysis).

Interestingly, the sites of clashes often coincide with the most relevant and important (moral) issues that a human would identify within particular conceptual fields. This points to the cognitive reality and importance of clashes and offers a method for validating specific conceptual networks (and the COGVIEW framework) as accurately reflecting underlying semantics and corresponding accurately to the real world. Clash sites as predicted can be compared to the work of human raters asked to decide where controversies exist within specific conceptual domains.

The subnetwork within which a clash occurs will be activated most when a result propagates to consciousness, and thus is more likely to be the subnetwork from which the conscious knower will consider the 'insight' of the clash to have originated. Critically, this is true even (as is the common case) when the energy which caused the clash came from other concept fields. This phenomenon allows for persuasive processes in which the semantic/conceptual content of a clash is removed from its original context (usually one in which conscious processing would have caused the desired conceptual understanding to be rejected) and shifted to another context (i.e. another subportion of the conceptual cloud) wherein which the attribution of the clash output to that context is 'safer' (or more desired).

Tentative support for a neurological basis of clashes can be found in the psychological literature (see for example [18]).

During COGVIEW network design, clashes should not be expressly 'engineered for'; rather, they should be expected to occur naturally (and in reasonable places) in a well-designed network.

Clashes represent a rare instance where the results of otherwise inaccessible processing are made conscious, providing a promising mechanism by which such processing could be made amenable to analysis by experiment. Experiments can be designed to elicit clashes at certain points in hopes of validating conceptual analyses generated by various candidate models.

D. Stored energy

Within an INTELNET/COGVIEW network, certain concepts (such as FAMILY or FREEDOM) tend to have a significant baseline level of energy associated with them at all times as a reflection of their positive or negative status within society. Concepts with significant stored energy may act as persistent energy sources and/or dominate the energy flows they participate in. In order to model this phenomenon INTELNET supports 'stored energy', drawing in part on the concept of resting activation.

Stored energy is well-attested in the persuasion and neurosciences literatures. Burdein, Lodge, and Taber [19] suggest that, in persuasion processes, "much of the influence of group identifications comes from the affective coloration that they automatically invoke, even before cognitive appraisal brings semantic associations to mind." The Hot Cognition hypothesis from psychology suggests that concepts have long-term affective information attached to them in memory, and that this information exerts differential, generally significant, effects upon cognition in various contexts. INTELNET stored energy assists in modeling the mechanism by which this occurs.

Stored energy may be applied to concept fields as well as single concepts; in the former case this energy may be envisioned as being distributed across the concept field much in the way that applying a voltage to an electrical circuit brings each point of the circuit to the same electrical potential.

Persistent values of stored energy do not require replenishment from active energy flow. They tend to act more as 'intensifiers' or 'deintensifiers' of energy passing through the concepts they are attached to, especially depending on whether or not the valence of a concept's persistent energy is the same as the energy flowing through the concept itself.

VI. COGVIEW: INTEGRATING CONCEPTS AND MODELS OF PSYCHOLOGICAL 'QUIRKS'

As suggested above, fully modeling CMPD requires the concomitant ability to model various psychological 'quirks' affecting the functioning of CMPD-mediated systems. In this section we identify specific psychological phenomena that should be considered and suggest how they can be modeled in practice. Westen [20] provides a convenient guide to the types of phenomena that should be taken into account (quotes in this section refer to Westen's article).

The first is *associative memory*, involving the "unconscious activation of networks of association –thoughts, feelings, wishes, fears, and perceptions that are connected, so that activation of one node in the network leads to activation of the others." In INTELNET/COGVIEW, this phenomenon is realized via energy flows between the various components of concept definition fields.

Next is the understanding that information is encoded in "distributed networks of neurons whose coactivation constitutes a representation" instead of in isolated "places" or 'pieces'. Thus, "memory is not a 'thing' that is stored somewhere in a mental warehouse and can be pulled out and brought to the fore. Rather, it is a *potential for reactivation* of a set of neurons [in INTELNET, a set of nodes] that together constitute a particular meaning." (emphasis in original) In line with this, INTELNET stores information about concept meaning within extended and distributed 'fields', each having internal structure.

The next component, *parallel processing architecture*, shifts from a view of information as processed in a step-bystep, linear manner towards one in which multiple areas of concept structures may be receiving and transmitting energy at any given moment. This suggests that INTELNET systems should be capable of modeling energies arriving at multiple nodes and simulating multiple energy paths and the interactions between them.

X-System	C-System
Parallel processing	Serial processing
Fast operating	Slow operating
Slow learning	Fast learning
Nonreflective consciousness	Reflective consciousness
Sensitive to subliminal presentations	Insensitive to subliminal presentations
Spontaneous processes	Intentional processes
Prepotent responses	Regulation of prepotent responses
Typically sensory	Typically linguistic
Outputs experienced as reality	Outputs experienced as self-generated
Relation to behavior unaffected by cognitive load	Relation to behavior altered by cognitive load
Facilitated by high arousal	Impaired by high arousal
Phylogenetically older	Phylogenetically newer
Representation of symmetric relations	Representation of asymmetric relations
Representation of common cases	Representation of special cases (e.g., exceptions)
	Representation of abstract concepts (e.g., negation, time)

Fig. 4. Satpute & Lieberman Two-Level Model [23]

In INTELNET, "relevance" is understood as being relative to the semantic or cognitive content of a particular social process or to the conditions required to generate the set of emotional or psychological states necessary for the functioning of some process. Where Anderson's declarative memory posits 'chunks' of explicit information, INTELNET expands this to include anything within the INTELNET/COGVIEW universe, including other concept fields and, critically, relevant subportions of the networks used to define other concepts.

A. Two-level Structure (U(nconscious) and C(onscious))

Symbolic and sequential information tend to be easily consciously available to human modelers, and therefore form the core of most Artificial Intelligence models. The cognitive perspective, however, points to the critical importance of processes that occur subconsciously and without awareness. Such unconscious processes exert significant effects on the functioning of human actors, especially in the case of typical INTELNET/COGVIEW target phenomena such as perception and persuasion, phenomena which operate quite differently in practice than they would if only conscious processes were involved.

Specifically, in the psychological literature there is extensive evidence (see, for example, [21], [22], [23]), of the existence of multiple (usually two) processing systems within the human brain, one that involves automatic, fast, parallel, unconscious processing, and one involving slow, serial, more conscious processing. Indeed, Chaiken and Trope (as cited in [23]) suggest that "[d]ual-process models of automatic and controlled social cognition have been proposed in nearly every domain of social psychology."

In order to model the differential effects of each of these two levels, COGVIEW proposes two software layers on top of the INTELNET core, the U(nconscious)-Level and the C(onscious)-Level. Figure 4 enumerates typical properties of each of these processing systems.

Diverse aspects of simulation problems lend themselves to processing at each level. The C-Level handles phenomena amenable to conscious awareness, concept selection, energy introduction into INTELNET networks, and other tasks related to 'interfacing' with the outside world.

The C-Level also handles those aspects of problems amenable to logic. When modeling complex social processes, however, while logic may be used to model belief or disbelief in particular ideas or arguments at the C-Level, particularly difficult arguments, fatigue, or other impediments to the cognitively taxing process of resisting argumentation are in practice likely to cause a cognitive shift to U-Level processing, suggesting that models should shift techniques as well.

The U-Level is ideal for modeling problem aspects involving connections between concepts, accessing energy sources, and the shifting/transferring of energy between sources. Emotions are also triggered here, all without conscious or rational intervention. U-Level processing is not accessible to conscious awareness and is dominated by the effects of associational memory in that the primary unconscious activity is the spreading of energy between concepts and concept fields and the associative detection of congruence/similarity between various portions of the extended concept universe. Clash detection, the focusing of energy on specific nodes and the decay of those energies also take place here.

The C-Level is best for modeling skills that must be explicitly developed, and the U-Level those that are more autonomous and 'endogenously endowed'.

In general, U-Level processing is claimed to be far more important than that of the C-Level for understanding worldview- and cognition-mediated phenomena (the opposite of the order observed in traditional systems).

Software implementations of U-Level processes should be compatible with Bargh's four criteria for *automaticity*; they should be "spontaneous; that is, the process ... [is] triggered even if the individual is not consciously engaged." They should be "unconscious; the processes themselves ... occur outside of awareness." U-Level implementations should be "uncontrollable; once triggered, [processing] runs its course without conscious guidance." Importantly, U-Level processes should also be modeled as being "invoked and carried through while expending little or no cognitive resources." (Bargh, as cited in [19]).

The U-Level is also tailor-made for modeling *priming*, a memory-based psychological phenomenon in which early stimuli influence later ones. COGVIEW and INTELNET are ideal for modeling the knowledge-based, associative spread, and memory-based aspects of the phenomenon. Priming is important to cognitive modelers; Bargh [24] notes it as capable of affecting phenomena as diverse as:

Social norms (Aarts & Dijksterhuis, 2003; Hertel & Kerr, 2001) to guide or channel behavior within the situation; goals to achieve high performance, to cooperate with an opponent, or to be fair minded and egalitarian (Bargh, Gollwitzer, Lee-Chai, Barndollar, & Troetschel, 2001; Moskowitz, Gollwitzer, Wasel, & Schaal, 1999); emotions that shape our reactions and responses to subsequent, unrelated stimuli (Lerner, Small, & Loewenstein, 2004); and of course, knowledge structures such as stereotypes and trait constructs for use in the comprehension and encoding of often ambiguous social behavior (see Bargh, 1989; Higgins, 1996, for reviews). And social behavior itself can be produced nonconsciously in the same fashion (Bargh, Chen, & Burrows, 1996; Chartrand & Bargh, 1999; Dijksterhuis & van Knippenberg, 1998). (emphasis in original)

Bargh also notes that "priming effects ... have been discovered ... in the nonconscious activation of deep cultural ideologies ... and other interpersonal relations (e.g., power differentials ...) that seem to alter even basic (i.e., non-social) information processing; as well as representations of close relationship partners [and] self-regulatory effects."

Clearly, U-Level processing is especially powerful, suggesting why communicators so diligently seek to design appeals capable of bypassing the C-Level in favor of the U-Level. Some evidence (cf. [25]) even suggests that, at the U-Level, mere understanding may be cognitively tantamount to acceptance, requiring an active act of 'disbelief' in order to be overcome.

Even though U-Level processing may not be amenable to conscious analysis, this does not mean that it is impossible to gain any knowledge about its functioning. In particular, hypotheses about the structure or contents of various portions of the concept universe may be tested, providing useful information along the way. The presence of clashes at various points may be verified, generating information about the correctness of COGVIEW/INTELNET graphs and energy flow patterns.

VII. COGVIEW: WORLDVIEW NETWORKS

Having now presented the general concept of INTELNET and the conscious and unconscious mechanisms that operate on top of it, we introduce the concept of COGVIEW Worldview Networks, employing INTELNET and COGVIEW to represent the concepts, concept semantics and interconnections that comprise cultures and worldviews.

COGVIEW networks reflect beliefs - information about how people view the construction and content of the universe. Worldview Networks include answers to questions about which concepts are important/relevant in specific domains, how those concepts are interconnected, how various phenomena affect and interact with others, and which important energy sources are operating in a particular domain.

A Worldview Network can be understood as a 'snapshot' of beliefs, relevant to some particular set of concept fields, held by a group of people at a particular time. These networks act as substrates for the operation of U-Level processing, providing pathways for energy flow between concepts and a base for unconscious processes such as clash and similarity detection between various portions of extended networks. Worldview Networks support C-Level processing by specifying sets of concepts that should be considered relevant with respect to the contexts and goals held within a particular set of beliefs.

Worldview Network concepts provide clearly delineated points where COGVIEW graphs may 'interface' with the wider world and with natural language, and act in essence as 'filters' between what a person is capable of perceiving and expects to find in the world and the world itself, thus acting as a key determinant of behavior.

Worldview Networks embody significant domain knowledge, reflecting deep correlational and practical information drawn from cultural and domain experts. The development and testing of Worldview Networks is an evolving area of practice, but the design process roughly begins with the identification of relevant concepts and energy sources. These concepts/energy sources are placed in relation to one another, with test simulations used to demonstrate correctness. Such simulations involve introducing energy into particular nodes, tracing the changes occurring to that energy as it traverses the network, and determining if the semantic outcomes of network traversal are accurate within the problem domain. Path length counts; as longer paths involve more network elements, the longer a path is, the more evidence it provides of overall network correctness.

VIII. COGVIEW WORLDVIEW NETWORK EXAMPLE

Fig. 5 provides a detailed example of a COGVIEW Worldview Network, modeling the 'before and after' belief system of a person persuaded that suicide terrorism is a viable option. Following Kruglanski et al. [26], suicide terrorism is framed as a quest for personal significance.

In Fig. 5, each graph node represents a concept. In general, although INTELNET and COGVIEW advocate for nuanced representation, and thus for concepts with internal structure, in this example a level of reification has been chosen that is sufficient to illustrate the concepts at hand, but that still provides examples of both reified concepts and concept fields with internal structure.

Solid lines represent pre-persuasion links; edges with short dashes represent post-persuasion beliefs. There is a small set of edges with long dashes; these edges form an 'energy loop' (see section V-B) between FAMILY, COMMUNITY, RAISE CHILDREN, PARENT ROLE, and PERSONAL SIGNIFICANCE, WORTH, MORALITY.

The ability of COGVIEW to model energy transmission across concept field boundaries via inter-concept connection pathways is critical in modeling unconscious conceptuallymediated processes which contribute significantly to social complexity.

There are two concept fields in the diagram: COMMU-NITY and IDENTITY, each denoted via rectangles. In keeping



Solid Lines: Pre-Persuasion Beliefs Dotted Lines: Post-Persuasion Beliefs Long-Dashed Lines: Loop (Reinforced Energy Flow) 'SA' Denotes 'Significance Achievement' * Denotes Clashes

Fig. 5. COGVIEW Network Example

with the INTELNET paradigm, each concept field represents multifaceted semantics through the combination of semantic components.

A. Discussion

COGVIEW diagrams may be best understood through a process of identifying energy sources of interest, following the flow of energy across various paths, and then determining the sub-beliefs and prediction outcomes indicated by each concept node/edge/energy result. Key is noting the energy that ends up in important nodes and the edges crossed while reaching these nodes. An example would be the observation that, pre-persuasion, energy arising from the source NEED-FOR-PERSONAL-SIGNIFICANCE flows through subordinate nodes to MEANS OF ACHIEVING PERSONAL SIGNIFICANCE to CON-TRIBUTION: CHARITY and CONTRIBUTION: COMMUNITY EF-FORTS, and then, finally, into COMMUNITY. Should persuasion be achieved, however, energy from the nodes under NEED-FOR-PERSONAL-SIGNIFICANCE would be redirected to SUICIDE-TERRORISM. From this we may infer that, indirectly, charity, community efforts, and the community as a whole would be expected to suffer if successful persuasion towards the course of suicide terrorism were achieved.

It may also be noted that the energy received by SUICIDE TERRORISM crosses a NEG link on exit from the node (reversing its polarity), reaching CAUSED-SUFFERING, VICTIMS and COMMUNITY. This suggests that the more positively suicide terrorism is viewed, the more suffering and victims will be likely to arise and the more likely that the community will be negatively affected.

As suggested above, when interpreting COGVIEW graphs a general principle is that those concepts that best maximize positive energy equilibrium are most likely to drive behavior. It is therefore instructive to observe locations wherein certain energy flows, if overwhelmed by other flows, would be expected to significantly shift overall energy balance. As an example, the energy sources DESIRE-TO-BE-MORAL and DESIRE-TO-PRESERVE-LIFE propel highly significant negative energy into SUICIDE TERRORISM, suggesting that the latter concept is very unlikely to drive behavior in the initial (pre-persuasion) configuration. Concepts such as REVENGE and SHAME can contribute positive energy to this node, but unless the energy they generate is quite strong, no shift will occur. It is much more likely, however, that a shift will occur if one or more of the subnodes under NEED-FOR-PERSONAL-SIGNIFICANCE begins to send positive energy to SUICIDE TERRORISM, and the more nodes activated, the more likely a shift is to occur. This points both to potential strategies for suicide terrorism reduction (that is, tailoring efforts towards those nodes capable of sending energy to SUICIDE TERROR-ISM) and to means of understanding which concepts are most likely to generate shifts under various scenarios.

1) Clashes: As indicated above, clashes point to critical 'dilemmas' and moral issues obtaining within problem domains.

In Fig. 5, at node DIE a clash occurs between positive energy arriving from SUICIDE TERRORISM and negative energy from DESIRE-TO-LIVE. This points to the inherent tension between suicide terrorism as a practice and the general strong desire to maintain life. At SUICIDE TERRORISM, positive energy from SHAME \rightarrow REDEMPTION and REVENGE \rightarrow PUNISHMENT \rightarrow PER-PETRATORS clashes with negative energy from DESIRE-TO-LIVE, DESIRE-TO-PRESERVE-LIFE, and DESIRE-TO-BE-MORAL. This clash site illustrates the immense psychological tension between the negative experiences often stated in the literature as motivators during interviews with those who have failed to activate their suicide devices and the core need to protect and preserve life.

The clash at PEACEFUL SIGNIFICANCE ACHIEVEMENT between positive energies from DESIRE-TO-BE-MORAL and DESIRE-TO-PRESERVE-LIFE and negative energy from SUI-CIDE TERRORISM offers a precise demonstration of the terms of the conflict: namely, questions of how one should achieve significance in one's life vs. the human costs of suicide terrorism.

The diagram also demonstrates clashes taking place at the following sites: FAMILY, COMMUNITY, and IDENTITY, as follows:

- **Family** This clash site highlights various influences on the family made by those who do not engage in suicide terrorism: the positive inputs of direct personal contributions (especially time spent with the family), the negative general impact of deaths and pro-suicide philosophies on the family as a whole, and the psychological importance of the concept of family.
- **Community** This clash site represents effects of peaceful action and suicide terrorism on the community, taking into account general contributions and specific actions which are likely to affect welfare in the long run. Any positive energy that enters the node SUICIDE TERRORISM will be transmitted across the NEG link (flipping its polarity) to CAUSED-SUFFERING and COMMUNITY. This negative energy clashes with positive energy at COMMUNITY resulting from CHARITY and COMMUNITY EFFORTS, representing a clash of methods of *serving the community* either through peaceful means or through violence. This clash suggests the negative effects of suicide terrorism on the community as a whole and highlights its direct interactions with positive social phenomena.
- **Identity** This clash site receives positive energy from the roles PARENT and RAISE CHILDREN. It receives negative energy through the sequence CAUSED-SUFFERING \rightarrow VICTIMS \rightarrow COMMUNITY \rightarrow IDENTITY. This demonstrates the often unobvious effects of energy transmission across multiple concepts, suggesting that, notwithstanding family-sanctioned 'martyr' identities, suicide terrorism may in fact result in *loss* of specific components of community- and family-related identity. This may be contrasted with strongly identity-based theories of suicide terrorism (see, for example, [27]).

These clash sites identify both the sites of ideological conflict within specific social processes as well as the specific concepts and connections contributing to them, enabling strategy development.

IX. CONCLUSION

Energy-Based Knowledge Representation (INTELNET) and the COGVIEW framework together represent a powerful paradigm for producing nuanced, highly functional models of data arising from the human cognitive capacity.

Such models enhance intelligent agents' abilities to model how beliefs and unconscious processes work together to generate judgments and, ultimately, behavior.

INTELNET/COGVIEW models of conceptually-mediated belief systems have already been successfully applied within intelligent reasoning systems, humanitarian mission modeling, cultural simulations, knowledge engineering, language processing, anti-discrimination and prejudice reduction, terrorism reduction, and norm change efforts, among others.

REFERENCES

- R. A. Grier, B. Skarin, A. Lubyansky, and L. Wolpert, "SCIPR: A computational model to simulate cultural identities for predicting reactions to events," in *Proceedings of the Second International Conference on Computational Cultural Dynamics*, V. S. Subrahmanian and A. Kruglanski, Eds. AAAI Press, 2008, pp. 32–38.
- [2] K. Carley, "Computational modeling for reasoning about the social behavior of humans," *Computational & Mathematical Organization Theory*, vol. 15, pp. 47–59, 2009.
- [3] R. Conte and M. Paolucci. (2011) On agent based modelling and computational social science. [Online]. Available: http://dx.doi.org/10. 2139/ssrn.1876517
- [4] D. Olsher, "Changing discriminatory norms using models of conceptually-mediated cognition and cultural worldviews," in *Proceed*ings of the 34th Annual Meeting of the Cognitive Science Society. Cognitive Science Society, 2012.
- [5] —, "Full spectrum opinion mining: Integrating domain, syntactic and lexical knowledge," in *Proceedings, 2012 IEEE 12th International Conference on Data Mining (ICDM)*, 2012, pp. 693–700.
- [6] —, "Cognitive-cultural simulation of local and host government perceptions in international emergencies," 2013, submitted.
- [7] , "Novel methods for energy-based cultural modeling and simulation: Why eight is great in Chinese culture," in *Proceedings, 2013 IEEE Symposium Series on Computational Intelligence*, 2013.
- [8] R. Davis, H. Shrobe, and P. Szolovits, "What is a knowledge representation?" AI Magazine, vol. 14, no. 1, 1993.
- [9] E. Cambria, T. Mazzocco, A. Hussain, and C. Eckl, "Sentic medoids: Organizing affective common sense knowledge in a multi-dimensional vector space," in *Advances in Neural Networks ISNN 2011*. Springer Berlin Heidelberg, 2011, vol. 6677, pp. 601–610.
- [10] J. A. Waskan, "Intrinsic cognitive models," Cognitive Science, vol. 27, no. 2, pp. 259–283, 2003.
- [11] S. E. Fahlman. (2008) In defense of incomplete inference. [Online]. Available: http://scone1.scone.cs.cmu.edu/nuggets/?p=34
- [12] J. R. Anderson and P. L. Pirolli, "Spread of activation," Journal of Experimental Psychology: Learning, Memory, & Cognition, vol. 10, no. 4, pp. 791–799, 1984.
- [13] A. M. Collins and E. F. Loftus, "A spreading-activation theory of semantic processing," *Psychological Review*, vol. 82, no. 6, pp. 407– 428, 1975.
- [14] S. E. Fahlman, NETL: A System for Representing and Using Real-World Knowledge. MIT Press, 1979.
- [15] H. Helbig, Knowledge Representation and the Semantics of Natural Language. Springer, 2006.
- [16] R. C. Schank, "Conceptual dependency: A theory of natural language understanding," *Cognitive Psychology*, vol. 3, no. 4, pp. 552 – 631, 1972.
- [17] C. Rieger, "An organization of knowledge for problem solving and language comprehension," *Artificial Intelligence*, vol. 7, no. 2, pp. 89 – 127, 1976.
- [18] M. D. Lieberman, D. Schreiber, and K. N. Ochsner, "Is political cognition like riding a bicycle? How cognitive neuroscience can inform research on political thinking," *Political Psychology*, vol. 24, 2003.
- [19] I. Burdein, M. Lodge, and C. Taber, "Experiments on the automaticity of political beliefs and attitudes," *Political Psychology*, vol. 27, no. 3, pp. 359–371, 2006.
- [20] D. Westen, "Developments in cognitive neuroscience for psychoanalytic psychotherapy," *Harvard Review of Psychiatry*, vol. 10, pp. 369–373, 2002.
- [21] E. R. Smith and J. DeCoster, "Dual process models in social and cognitive psychology: Conceptual integration and links to underlying memory systems." *Personality and Social Psychology Review*, vol. 4, pp. 108–131, 2000.
- [22] S. Chaiken and Y. Trope, *Dual-process theories in social psychology*. Guildford Press, 1999.
- [23] M. D. Lieberman, "Social cognitive neuroscience: A review of core processes," *Annual Review of Psychology*, vol. 58, no. 1, pp. 259–289, 2007.
- [24] J. A. Bargh, "What have we been priming all these years? On the development, mechanisms, and ecology of nonconscious social behavior," *European Journal of Social Psychology*, vol. 36, no. 2, pp. 147–168, 2006.
- [25] M. Shermer, Why People Believe Weird Things, 2nd ed. Holt, 2002.
- [26] A. W. Kruglanski, X. Chen, M. Dechesne, S. Fishman, and E. Orehek, "Fully committed: Suicide bombers' motivation and the quest for personal significance," *Political Psychology*, vol. 30, no. 3, pp. 331– 357, 2009.
- [27] M. Harrison, "An economist looks at suicide terrorism," World Economics, vol. 7, no. 3, pp. 1–15, 2006.