Modeling Knowledge, Expression, and Aesthetics via Sensory Grounding

Brad Spendlove and Dan Ventura
Computer Science Department
Brigham Young University
Provo, UT 84602 USA
brad.spendlove@byu.edu, ventura@cs.byu.edu

Abstract

In many instances, computationally creative systems need to meaningful interact with human agents—by existing in the environment they live in and/or interfacing with one or more forms of humans perception and/or understanding the way that humans structure their perceptions of the environment. In this preliminary work, we present a model of a mind's knowledge base consisting of concepts that are grounded in sensory perception of environmental phenomena. Although the individual components of this model are complex and difficult to fully define, our model's structure describes how knowledge is acquired and used for expression—including creative expression—and aesthetic assessment. Our formalization aims to serve as an abstract modularization that effectively factorizes these complex operations.

Introduction

The field of computational creativity includes the goal of developing increasingly powerful autonomous systems that can be considered truly creative in their own right. Achieving that goal is challenging in part due to the deep, inherent differences between computers and humans. Wiggins offers the following definition of creativity: "The performance of tasks which, if performed by a human, would be deemed creative" (2006). Implicit in this definition is the requirement that any agent, human or non-human, must engage with creativity on human terms[1].

This means that creative endeavors must exist within the scope of not only the environment in which humans exist but also the particular ways that environment is perceived by humans. Furthermore, creative performances presented for comparison to human creativity must be compatible with the way human thought is organized.

These factors indicate that the following related mechanisms may be useful inclusions in a model of computational creativity: a model of an environment, a model of how minds perceive that environment, and a model of how minds structure those perceptions. We will refer to the quantization of environment as *phenomena*, the quantization of structured perceptions of phenomena as *concepts*, and the structure of all concepts in a mind as *knowledge*.

Modeling knowledge is challenging, but the lack of such knowledge in a computational model leaves a vacuum which is easy to underestimate. Identifying that vacuum and filling it with a useful knowledge model appears to be a fruitful avenue toward improving computational models of creativity.

Current computational knowledge bases such as ConceptNet (Speer and Havasi 2012), WordNet (Miller 1995), Datamuse[2], and word2vec (Mikolov et al. 2013) have proven useful for certain language-based tasks but fall short of a complete knowledge model due, at least in part, to a lack of grounding of the concepts contained therein.

Similar to how the words in a dictionary are defined exclusively in terms of other words, existing knowledge bases contain concepts and relations between them but are missing a crucial element that grounds those concepts' meanings and relationships. Their concepts exist as if in a cloud, floating above the ground and impossible for a computer to align with the real world (at least as perceived by humans) without additional information.

Grounding provides intrinsic meaning that is independent from any observer or other knowledge. While not all concepts must be directly grounded, ensuring that all concepts are at least indirectly grounded (i.e., related to or built up from grounded concepts) will aid in attaining these properties of intrinsic, independent meaning.

We propose that perception—sensory experiences that arise from a mind observing its environment—is a useful basis upon which to ground concepts. Because the environment is consistent and exists independently from any observer, it can be relied upon to serve as a foundation for inherently meaningful concepts and knowledge.

We present a model that formalizes the acquisition of a knowledge base of concepts which are grounded in the sensory experiences that arise when a mind observes phenomena in its environment. This knowledge base may be employed by many faculties of that mind including expression of concepts as generated phenomena and aesthetic assess-

[1]Although later work, such as that by Colton and Wiggins, has defined creativity in a way that omits direct deference to humanity (2012), we argue that the now somewhat unfashionable human-centric definition is still relevant to an important subset of unbounded trans-human creativity, namely that subset that relates to humans in any way. As such, we are content with limiting our scope to that definition in this work.

[2]http://www.datamuse.com/api/

ment of observed phenomena. While the ultimate goal of this model is to provide a framework for improved models of computational creativity, we expect that it may have application in other disciplines as well.

Our model could be useful to the field of computational creativity by serving either as a framework for the creation of computational minds or to approximately model human minds as part of a computational system, for example in order to predict how humans would react to a given creative expression. Although each element of our model is in itself complex and difficult to implement, we hope that our formalization allows an abstract modularization that effectively factorizes the problem.

Finally, while this model is inspired by the human mind and we draw illustrative analogies to it in this work, we make no claims that this is an accurate model of the human mind itself. We do posit that it is one lens through which to view a mind, whether natural or artificial, and that it could serve as a useful blueprint for designing computational minds.

Modeling Knowledge Acquisition via Sensory Grounding

In this section we present a formal model of a mind's knowledge base that is grounded in atomic sensory experiences of natural phenomena. Grouping, labeling, and categorizing sensory experiences gives rise to knowledge concepts, which can be further combined into higher-level concepts. Ultimately grounding these concepts in environmental phenomena provides concrete points of reference for manipulating and sharing concepts between minds, which is a chief concern of creativity.

Wiggins' formalization (Wiggins 2006) of Boden's model of creativity (Boden 2004) introduces a universe $\mathscr{U}$ consisting of all possible concepts. Such concepts serve as the atomic unit of Wiggins' framework, and we expand upon his definition to model how concepts are represented in minds. Our concept model is thus compatible with Wiggins' model of conceptual systems as they pertain to exploratory and transformational creativity, while also expanding to incorporate expression and aesthetic evaluation which we will explore in a later section. We begin with an abstraction of natural phenomena and model how a mind can convert such phenomena into concepts.

Sensory Grounding

Let $\mathscr{S}$ represent the set of all senses available to a given mind by which it may perceive its environment. For example, for the human mind this set could include the five basic senses commonly taught to children as well as the many more nuanced senses humans possess such as temperature, depth perception, and emotional sensations. Regardless of how the methods by which a mind experiences the universe are partitioned into senses, it necessary that $\mathscr{S}$ include every one of those experience-enabling mechanisms.

Let a given sense $S \in \mathscr{S}$ represent a multi-dimensional *sense space* in which each distinct *sense point* $s \in S$ corresponds to a distinct instance of that sense. Again using human senses as an example, a sense point corresponds to a particular smell, tone, hue, emotion, etc. Furthermore, this concept of sense spaces includes comparative differences in otherwise similar sensory information such as intensity. Gärdenfors presents a model of representing information in terms of quality dimensions based on sensory perception which could be viewed as one method for modeling sense points more granularly (2004).

Let I represent the set of all possible sense points regardless of which sense space they occur in, such that

$$I = \bigcup_{S \in \mathscr{S}} S.$$

Let the set $\mathscr{P}$ represent all possible natural phenomena, which we define as quantized elements of the environment in which the mind exists. For example, in the real world a flower exists of itself and may be experienced in different way by different human senses such as sight, touch, and smell.

Such phenomena exist independently from the minds that observe them, and minds experience phenomena by way of their senses. Therefore, let

$$\psi : \mathcal{P}(\mathscr{P}) \to \mathcal{P}(I)$$

where $\mathcal{P}$ is the power set, represent an observation function by which the phenomena that a mind encounters are experienced sensorily (i.e., translated into sense points). Thus, a given phenomenon $p \in \mathscr{P}$ is interpreted by ψ to yield a set of zero or more sense points, which may occur in disparate sense spaces.

This function ψ is immutable and mechanical, not consciously directed by the mind. Recognizing that senses are inherent and out of the mind's control is an important consideration for modeling and reasoning about minds. However, ψ may differ between distinct minds, even those of the same structure, resulting in different perceptions of identical phenomena. This corresponds, for example, to how a person with colorblindness perceives color differently than one without.

In a later section we will discuss phenomena that are generated by a mind. We note here that such phenomena necessarily exist in the same environment as "natural" phenomena, and as such we draw no distinction between them in this model.

Knowledge Acquisition

With a model for sensory experience of phenomena in hand, we turn to modeling the process by which a mind organizes sense points into concepts. Such concepts may be very complex and draw from many disparate sense points. We acknowledge the difficulty of building a robust model of knowledge acquisition and present this model as a means to abstract and compartmentalize the various mechanisms at play so that they may be addressed independently.

Let K represent the set of all concepts in a given mind's knowledge base, with $K \subseteq \mathscr{U}$, where $\mathscr{U}$ is Wiggins' universe. We inductively define a concept $c \in K$ as follows:

Say that c is a ***concept*** if c is

1. $\{\} = \top$, the empty concept,

2. $\{s\}$ for some $s \in I$,
3. $c_1 \cup c_2$, where c_1 and c_2 are concepts, or
4. $\{c\}$, where c is a concept.

Note that a concept may simply consist of a set of sense points. Such low-level concepts directly describe natural phenomena (as filtered through the mind's senses), such as the concept of the color blue in the human mind. Higher-level concepts may contain a mix of sense points and other concepts, with the highest-level concepts being exclusively comprised of other high-level concepts.

We now to turn to modeling the method by which a mind's knowledge base is assembled from the sensory input it experiences. Although minds may develop very different knowledge bases depending on the phenomena each mind experiences and the senses by which they perceive such, our model assumes that minds of the same structure acquire knowledge via the same underlying process.

Modeling this process is a daunting task due to the complexities of both the inputs into a mind and the network of knowledge that results from performing the knowledge-assembling process on those inputs. Our model abstracts those complex inputs as sense points which are used to assemble a knowledge base consisting of abstract concepts, facilitating the modeling of the the knowledge-assembling process itself.

We model the relationship between sense points and a mind's knowledge base via two functions: an *add* function, and a *lookup* function. The add function modifies the knowledge base to incorporate new sense points, creating new concepts or modifying existing concepts as necessary, and the lookup function determines what concepts a set of sense points evokes in the mind, respectively. Both of these functions are nontrivial to model and compute, but we present an abstraction for the latter.

Let $\lambda : \mathcal{P}(I) \rightarrow \mathcal{P}(K)$ represent the lookup function by which a mind relates sense points to existing concepts in its knowledge base. Given a set of sense points $Y \subseteq I$, the set of concepts returned by $\lambda(Y)$ may be very simple (i.e., a subset of I) or may consist of potentially many combined high-level concepts. Modeling the means by which a knowledge base is traversed from low-level concept representations of sense points to high-level concepts that represent a synthesis of those points appears to be a challenging avenue for future work.

Our abstraction of the "add" function does not lend itself to further subdivision as it directly mutates the knowledge base to incorporate new sense points. As such, we leave exploration of this function as future work. Describing the process by which a knowledge base is constructed and augmented as the mind encounters new sense points is central to implementing this knowledge model. As stated previously, our intent in presenting this model is to compartmentalize and isolate difficult aspects of modeling minds to reduce their complexity as much as possible.

Expression, Perception, & Aesthetics

A mind does not acquire knowledge simply to hoard it; its knowledge base is a deep well of resources that informs and powers other useful mental faculties. One of the most direct utilizations of knowledge is the act of expression, which we define as the production of phenomena meant to evoke certain concepts.

Creativity, which is the primary focus of this work, is encompassed within the larger umbrella of expression. Although creatively combining concepts or artfully expressing concepts in novel ways may differ substantially from merely reciting simple or well-known concepts, we draw no distinction between creative and non-creative expression in this model. Both can be modeled as the conversion of knowledge concepts into phenomena.

Expression

Expression can be modeled as a function that maps concepts to phenomena. Let $\xi : \mathcal{P}(K) \rightarrow \mathcal{P}(\mathscr{P})$ represent an expression function by which a mind generates phenomena from concepts. For example, the human mind has various modes of expression that may result in different phenomena representing the same concepts, such as drawing a picture of a bird versus flapping one's arms to imitate a bird in flight.

Let Ξ represent the set of all expression functions available to a given mind. This model encapsulates all differences in type of expression between the various functions whether they be in the type of phenomena they produce (i.e., the medium) or the quality of those phenomena. For example, two humans' functions for expressing concepts via paintings could differ due to their different experience levels with working in that medium.

Expressing concepts as phenomena is a challenging task to which there exist a large number of potential approaches and solutions. The human mind accomplishes expression via a physiological linkage between the brain and other body parts and requires practice to attain proficiency. Identifying useful ways to computationally generate phenomena from concepts is an ongoing quest in computational creativity and other artificial intelligence disciplines, and in fact the invention of novel expression functions could itself be an interesting creative task. By abstracting the other complexities of a mind's knowledge base, our model seeks to isolate and, to the extent possible, simplify this act of expression for further investigation.

Perception

The way other minds perceive phenomena generated through expression is an important consideration for expressive and creative endeavors. In particular, it is useful to compare the set of concepts $C \subseteq K$ from which the expresser generates phenomena and the set of concepts that those phenomena elicit in a perceiver.

Given two minds a and b, a set of concepts C_a that a wishes to express, a's expression function ξ_a, and b's observation and lookup functions ψ_b and λ_b, let

$$C_b = \lambda_b(\psi_b(\xi_a(C_a))).$$

Thus, C_b is the set of concepts that b infers from the phenomena a generates when attempting to express C_a via ξ_a.

By comparing C_a to C_b, we can model how closely b understood the concepts that a intended to express. We hypothesize that λ lookup functions will commonly be more successful at linking phenomena to the intended concept when a lower-level concept is being expressed. Higher-level concepts require the traversal of more connections from sense points to final concept and therefore seem more likely to be misunderstood.

In the special case that $a = b$, this process results in

$$C'_a = \lambda_a(\psi_a(\xi_a(C_a)))$$

which represents a mind evaluating its own expression, perhaps to compare how well the phenomena it generated reflect its intended concepts. This models a common and useful operation in creativity processes.

Aesthetics

The creative process often includes elements of aesthetic evaluation to complement generation. Such evaluation can be applied to one's own creative works as well as to those of others. Indeed, aesthetic appreciation is a significant factor in what gives meaning to some important creative endeavors.

We model the aesthetic sensibilities of a mind as partitions of its sense spaces, with each partition being considered of different aesthetic quality by the mind. For a simplified example, consider a human mind that finds certain smells appealing and others off-putting. This corresponds in our model to a partition of that mind's olfactory sense space into two subsets.

Let Θ represent a mind's set of aesthetic partitions of its sense spaces, such that

$$\forall \theta \in \Theta . \exists S \in \mathscr{S} . \theta \subseteq S$$

and

$$\forall S \in \mathscr{S} . \exists T \subseteq \Theta . \bigcup_{\theta \in T} \theta = S.$$

Note that this model requires complete coverage of each sense space by aesthetic partitions. Thus, a mind will always have an aesthetic opinion of any newly encountered sense point. Descriptions of the extent to which a mind is aware of its aesthetic partitions, the degree to which those partitions may change over time, and the degree to which such changes can be consciously enacted by the mind are left to future work.

We assume that the mind maintains an aesthetic opinion for each θ, perhaps by means of a function $q : \theta \to \mathbb{R}$ that maps a partition to an abstract, real-number representation of the mind's opinion of the sense points in that subset. We concede that aesthetic opinion may be more nuanced than can be represented by a single real number and leave further exploration into useful models of such opinion as future work.

Let $\pi : I \to \Theta$ represent a function that identifies to which aesthetic partition a sense point belongs. Then, to model a mind's aesthetic opinion of a set of phenomena $P \subseteq \mathscr{P}$, construct

$$A = \{q(\pi(s)) \mid s \in \psi(P)\}$$

which contains the set of aesthetic opinions that correspond to the sense points experienced from those phenomena. Constructing A for a set of phenomena generated by an expresser's ξ reflects the perceiver's aesthetic assessment of that expression.

Similarly to calculating C'_a to evaluate the concepts evoked by a mind's own expressed phenomena, constructing A for those same phenomena allows the mind to aesthetically evaluate that expression. This "self-criticism" operation is useful in creative processes.

Our model also allows for the aesthetic assessment of (grounded) concepts. Let $c \rightsquigarrow s$ represent the *grounded in* relation between a concept c and a sense point $s \in I$. Say $c \rightsquigarrow s$ if and only if

1. $s \in c$ for some $s \in I$ or
2. $b \in c$ and $b \rightsquigarrow s$

Then, given a concept c, constructing O such that

$$O = \{q(\pi(s)) \mid c \rightsquigarrow s\}$$

allows us to model the mind's aesthetic opinion of the concept in question via the sense points that ground the concept.

Discussion

Our model of knowledge is by no means exhaustive. A mind's knowledge base influences many other important mental functions such as language, reasoning, and imagination.

Language in particular is interesting to consider under our model. Sensory experiences of language phenomena seem to circumvent a mind's add and lookup functions by serving as "machine code" that executes on the knowledge base directly. We view language and other facets of minds as interesting avenues for future work to be explored using our knowledge model.

This model seeks to be environment- and mind-agnostic so that it allows for many interpretations. As such, it may be useful to explicitly model a given mind's *environment* as $E \subseteq \mathscr{P}$ in order to reason about the types of phenomena that that mind can and cannot observe.

If minds a and b exist in environments E_a and E_b, respectively, such that $E_a \neq E_b$, then there exist some phenomena which each mind can experience that the other cannot. As phenomena inform sense points which in turn inform concepts, this means that there could exist concepts that the minds cannot share. Identifying the differences between human and computational environments could give better insight into the limitations of direct human-to-computer communication, and vice versa.

We believe that further exploration of this model of knowledge that is grounded in sensory perception of environmental phenomena will be useful for designing computationally creative systems that must ultimately operate with in the environment of the natural world as perceived by human senses.

References

Boden, M. A. 2004. *The creative mind: Myths and mechanisms*. Routledge.

Colton, S., and Wiggins, G. A. 2012. Computational creativity: The final frontier? In *Proceedings of the European Conference on Artificial Intelligence*, 21–26.

Gärdenfors, P. 2004. Conceptual spaces as a framework for knowledge representation. *Mind and Matter* 2(2):9–27.

Mikolov, T.; Chen, K.; Corrado, G.; and Dean, J. 2013. Efficient estimation of word representations in vector space. *arXiv* abs/1301.3781.

Miller, G. A. 1995. WordNet: A lexical database for English. *Communications of the Association for Computing Machinery* 38(11):39–41.

Speer, R., and Havasi, C. 2012. Representing general relational knowledge in ConceptNet 5. In *Proceedings of the International Conference on Language Resources and Evaluation*, 3679–3686.

Wiggins, G. A. 2006. A preliminary framework for description, analysis and comparison of creative systems. *Knowledge-Based Systems* 19(7):449–458.