Nehovah: A Neologism Creator Nomen Ipsum

Michael R. Smith, Ryan S. Hintze and **Dan Ventura**
Department of Computer Science, Brigham Young University, Provo, UT 84602
msmith@axon.cs.byu.edu, ventura@cs.byu.edu

Abstract

In this paper, we describe a system called *Nehovah* that generates neologisms from a set of base words provided by a user. Nehovah focuses on creating "good" neologisms by evaluating various attributes of a neologism such as how well it communicates the source concepts and how "catchy" it is. Because Nehovah depends on the user to weight the importance of various attributes of the neologism and to choose the source concepts, it is at this point most appropriately considered a collaborative system rather than an autonomous one. To demonstrate the utility of the system, we show several examples of system output and discuss the creativity of Nehovah with respect to several characteristics critical for any computational creative system: appreciation, imagination, skill and accountability.

Introduction

Boden (1994) made one of the first attempts to formalize the notion of creativity. Based on her formalization, computational creativity is often thought of as an exploration of a conceptual space and has been examined in a number of different areas including visual art (Colton, Valstar, and Pantic 2008; Norton, Heath, and Ventura 2011), music (Cope 2005), cooking (Morris et al. 2012), poetry (Rahman and Manurung 2011), metaphor generation (Veale and Hao 2007), and sentence generation (Mendes, Pereira, and Cardoso 2004). In this paper, we describe *Nehovah*, a computational system that generates neologisms.

The generation of neologisms is an important task in many businesses to create a unique brand or company name to distinguish it from its competitors. This often comes in the form of a trademark. Trademarks include words, phrases, symbols and/or designs that identify and distinguish the goods of one party from those of others[1]. According to the United States Patent and Trademark Office, 433,651 trademark applications were filed in 2013; a 4.5% increase from 2012 (The United States Patent and Trademark Office 2014). Thus, developing trademarkable phrases and words is a important step in many businesses.

Additionally, neologisms are often used as a literary device in novels and books to convey meaning more concisely. For example, "cyberspace" was introduced in 1982 by William Gibson to combine the words "cybernetics" and "space" (Gibson 1982). In some cases, neologisms are used to add humor and interest. This technique was used heavily in the many works of Dr. Seuss to help children with limited vocabularies to enjoy reading (Baker 1999).

Neologisms have previously been examined computationally, both from an interpretive standpoint and from a generative one. For example, Cook and Stevenson (2010) propose finding the meaning of neologisms using a statistical model that draws on observed linguistic properties of blends, while Duch and Pilichowski (2007) create neologisms using a neurocognitive model (though, unfortunately, many of the generated neologisms exhibit little to no linguistic/conceptual/cognitive value).

Veale's *Zeitgeist* system rather impressively exhibits both interpretive and generative abilities and is available as a web application. It can be used as a tool for enriching lexical resources such as WordNet (Fellbaum 1998) with modern words that are found in every day speech (Veale 2006) by utilizing Wikipedia[2] to identify neologisms and by reverse engineering their source words using ideas from concept blending (Veale, O'Donoghue, and Keane 2000).

In addition, the Zeitgeist system can be used to generate neologisms by combining prefix and suffix morphemes that overlap by at least one letter (Veale and Butnariu 2006). Morphemes are hand-annotated with their semantic interpretations giving each morpheme a word gloss (such as "astro"="star" and "ology"="study") and a WordNet identifier that indicates where in the WordNet noun taxonomy a neologism with a morphemic suffix should be placed. Given two source words from predefined lists for prefixes and suffixes, the Zeitgeist system creates a set of neologisms that convey the chosen concepts by combining the prefix and suffix morphemes for the source words. The generated neologisms generally have valid word forms and convey the concepts well. On the other hand, Zeitgeist is limited to the morphemes that are annotated. As many of the morphemes are of Greek origin, some of the neologisms are somewhat predictable. For example, if "food" is chosen as a source prefix word, then "gastro" is almost always used. The use of morphemes also requires a knowledge of Greek or Latin word derivatives to understand the neologism. The neolo-

[1] http://www.uspto.gov/trademarks/basics/definitions.jsp

[2] www.wikipedia.com

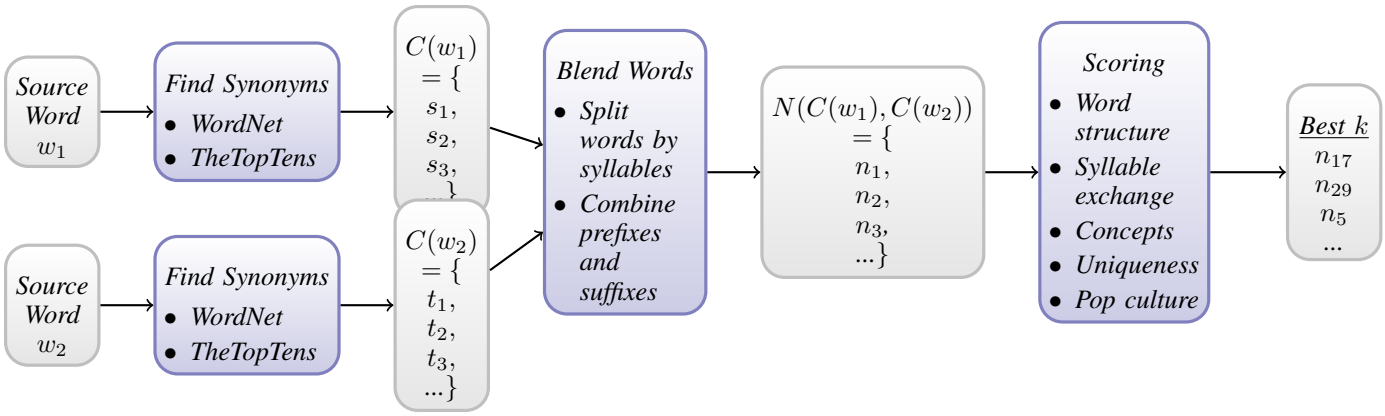

Figure 1: A high-level pipeline view of the process Nehovah uses to generate neologisms through finding synonyms, blending words, and scoring.

gism "ornithoencephalon" is a neologism for "bird-brain" but the meaning is obvious only to the user who knows that the morpheme "ornitho" relates to birds and "encephalon" relates to the brain.

Our system for generating neologisms, Nehovah, is similar to Zeitgeist in that it attempts to preserve the source concepts through blending (as opposed to generating neologisms that represent entirely new ideas by themselves, e.g. "Google"). It differs from Zeitgeist by focusing on blending free-form, user-provided words and their synonyms and by incorporating dynamic web sources of popular cultural information. In addition, the web interface allows a user to weight the importance of several attributes of a neologism, facilitating a creative collaboration between the user and the system.

A Framework for Blending Concepts

The goal of generating neologisms by blending concepts from source words is to convey multiple concepts in a single plausible word, sometimes known as a *portmanteau* (Carroll 1871). We present a framework, containing three major steps, for generating such portmanteau neologisms from two source words:

1. **Finding Synonyms.** Synonyms increase the potential novelty of the neologisms by enriching the set of possible blends that convey the source concept. A greater diversity of synonyms expresses more imagination in the neologism. For example, the word "God" is arguably a more diverse/interesting synonym for "creator" than is the word "maker". We call the set of synonyms for a source word w_i the *concept set* for w_i and denote it as $C(w_i)$. Note that it is always the case that $w_i \in C(w_i)$.
2. **Blending Words.** Once the concept sets for the source words have been generated, the words from each concept set are blended together to create a set of neologisms. Blending the words from the two concept sets consists of three steps. First, each word from the concept sets is split into sets of prefixes and suffixes. Then, each prefix from one concept set is joined with each suffix from the other concept set. Finally, Nehovah checks that the word structure of the neologism is plausible. By plausible, we mean that the letter sequence produced from blending the words is natural compared to other "real" words. Any implausible neologism is discarded. The set of neologisms generated from two concept sets $C(w_1)$ and $C(w_2)$ is denoted $N(C(w_1), C(w_2))$.
3. **Scoring/ranking the Neologisms.** Once a set of neologisms $N(C(w_1), C(w_2))$ is created, they are scored or ranked such that a subset of "best" neologisms can be identified, allowing a potentially large set of neologisms to be quickly filtered. Scoring criteria can be adapted for a particular application and can also potentially incorporate feedback, facilitating online learning and thus dynamic qualification of neologisms.

Nehovah

A functional overview of Nehovah and its implementation of the three steps are shown in Figure 1 and are described in more detail in the following sections. The blue boxes represent each step in the framework for blending concepts and the gray boxes represent sets of words. An on-line version of Nehovah is available at

`http://axon.cs.byu.edu/~nehovah`

from which a screen shot is shown in Figure 2.

Finding Synonyms

In order to populate the set $C(w_i)$, Nehovah searches for synonyms from two different sources: WordNet (Fellbaum 1998) (a lexical database) and TheTopTens[3] (a website of pop culture-inspired "top ten" lists).

Nehovah queries WordNet with each source word w_i (and with its stem) as a noun, verb, adjective, and adverb. If a source word or its stem is defined in WordNet, Nehovah adds to $C(w_i)$ the words contained in the synset for all senses of the word for all parts-of-speech for which it is defined.

[3] `www.thetoptens.com`

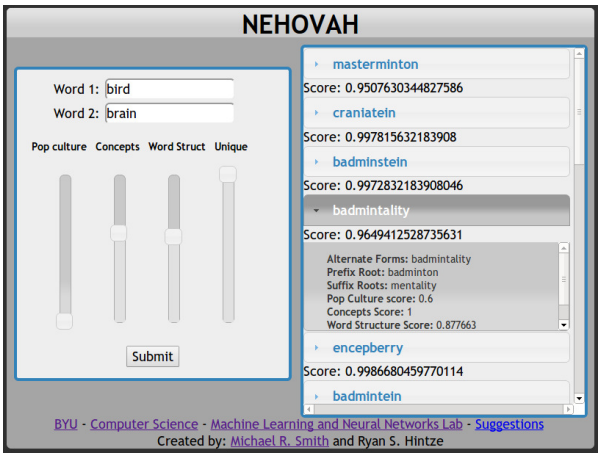

Figure 2: A screenshot of the web interface for Nehovah. Two source words are input in the upper left. The lower left contains sliders that allow relative weighting of the four scoring attributes. On the right is a list of generated neologisms with their scores, in descending order, and these can be expanded to see the base words that Nehovah used to create the neologism and how the neologism is scored for each of the attributes.

For example, the word "school" as a noun has the following senses:

- "school→educational institution"
- "school, schoolhouse→building, edifice"
- "school, schooling→education"
- "school→body"
- "school, schooltime, school day→time period, period of time, period"
- "school, shoal→animal group"

and, additionally as a verb has the following senses:

- "school→educate"
- "educate, school, train, cultivate, civilize, civilise→polish, refine, fine-tune, down"
- "school→swim"

(it has no senses as either an adjective or adverb). Therefore, the set of WordNet-derived synonyms for the word "school", $C(\text{"school"}) = \{$school, educational institution, schoolhouse, building, edifice, schooling, education, body, school time, school day, time period, period of time, period, shoal, animal group, educate, train, cultivate, civilize, polish, refine, fine-tune, down, swim$\}$.

Because a source word is specified without context, neither its part-of-speech nor its intended sense can be inferred, and, as a result the space of possible synonyms is increased, providing greater creative potential in the generated neologisms at the risk of potentially conveying an awkward or unintended conceptual blend.

Nehovah queries TheTopTens with each source word w_i using a custom API that returns lists of words from a set of "top ten" lists that match the query. For example, a query to TheTopTens using the source word "car" would return lists with titles such as "Top Ten Best Car Companies," "Best Car Brands," "Greatest Songs by the Cars" and "Best Car Insurance Companies."

Of course, some lists will be much more relevant than others. To minimize the number of included irrelevant words, Nehovah determines which of the returned lists are relevant based on their titles, by the identifying descriptive and plural words in the title. Descriptive words are identified as words that end with "-est" – as is common practice on TheTopTens. If a descriptive word in a list title directly precedes the source word, then the list is deemed relevant. For example, the list "Top Ten Best Car Companies" would be accepted since the descriptive word "best" is describing the source word "car". Also, if there are multiple plural words in a list title, Nehovah assumes the first plural word in the title identifies the subject of the list. For example, in the list "Greatest Songs by the Cars," there are two plural words: "Songs" and "Cars." The list is determined to be about songs rather than cars since "Songs" appears before "Cars" and because the descriptive word "greatest" proceeds "Songs" rather than "Cars." Nehovah also includes lists that have the source word directly before the first plural word such as "Top Ten Car Movies", inferring that the source word is being used as a descriptor for the plural word.

Once a list is determined to be relevant, the list items also need to be processed. Because TheTopTens is composed of user-defined free-form lists, some list items are more descriptive than others. For example, the "Best Muscle Cars" list may contain items such as "1961 Ford GT Mustang From Gone in 60 Seconds." While this information is beneficial for determining why an item made the list, it is difficult to use to generate neologisms. To compensate, Nehovah parses the list items so that any words or symbols that indicate descriptive information ("from", "in", "–", ",", etc) and any words that follow are not included. Another issue with user-defined lists is the lack of quality control. To filter out obscure (and/or misspelled) words and references, Nehovah only keeps list items that are also found in Wikipedia. Any list entries that survive this level of parsing and filtering are also included in $C(w_i)$. Note that using the words from TheTopTens adds hyponyms (e.g. "Ford Mustang" for "car") rather than synonyms in some cases. We allow the use of hyponyms as the pop culture reference adds to the creativity and uniqueness of Nehovah and because it is difficult to distinguish between hyponyms and synonyms.

Blending Words

Given two concepts sets $C(w_1)$ and $C(w_2)$, Nehovah blends the words from the two concept sets to create a set of neologisms $N(C(w_1), C(w_2))$. Each word $u \in C(w_i)$ is into into a set of prefixes $P(u)$ and a set of suffixes $S(u)$. The words are split between syllables to maintain conceptual coherence and to reduce the likelihood of introducing invalid letter combinations during blending.

Unfortunately, for English it is a non-trivial task to algorithmically identify syllable boundaries because pronunciation information is not (consistently) encoded in the spelling

computational		method	
Prefixes:	Suffixes:	Prefixes:	Suffixes:
-	computational	-	method
co	mputational	me	thod
com	putational	meth	od
compu	tational	method	-
computa	tional		
computati	onal		
computatio	nal		
computation	al		
computational	-		

Table 1: Examples of how Nehovah splits words into prefix/suffix pairs by attempting to split on syllable boundaries.

of the word. For example, "io" could create two separate vowel sounds as in "lion" or be a diphthong as in "motion". To account for this, Nehovah conservatively splits each word u after every vowel (except the last) and between any two consecutive consonants (with exception of "sh," "th," and "ch") after the first vowel and before the last vowel. Each such split yields one prefix to be added to the set $P(u)$ and one suffix to be added to the set $S(u)$. In addition, u is also added to both $P(u)$ and $S(u)$. For example, the word "track" would be split up into the prefixes "track" and "tra" and the suffixes "ack" and "track". See Table 1 for additional examples.

Slightly abusing notation, we define the set of neologisms formed by blending two words u and v using the sets $P(u)$, $S(u)$, $P(v)$ and $S(v)$ as

$$\begin{aligned} N(u,v) =& \{yz | y \in P(u) \wedge z \in S(v) \wedge K(yz)\} \cup \\ & \{yz | y \in P(v) \wedge z \in S(u) \wedge K(yz)\} \end{aligned}$$

where $K()$ is a predicate that returns FALSE if its argument contains a letter combination not found in WordNet and TRUE otherwise.

Then, the full set of neologisms for the synonym sets $C(w_1)$ and $C(w_2)$ is generated by iterating over all pairs of words from these synonym sets:

$$N(C(w_1), C(w_2)) = \bigcup_{u \in C(w_1), v \in C(w_2)} N(u,v)$$

Scoring

Nehovah scores each neologism $n \in N(C(w_1), C(w_2))$ using four scoring criteria: word structure, concepts, uniqueness, and pop culture. Each scoring criterion can be assigned a relative weight, allowing the creation of different types of neologism.

Word Structure. The word structure score $\mathcal{W}(n)$ measures how well a neologism retains aspects of the word structure of one or both source words, as maintaining source word structure tends to produce catchier neologisms that better convey the meaning of the base words. For example, "ginormous" is a combination of "giant" and "enormous" created by replacing the first syllable from enormous with the first syllable from giant. Enough of enormous is left that the meaning is still apparent. Another example is "Linsanity," which replaces the first syllable in insanity with the single syllable word "Lin" (the last name of a professional basketball player). In this case, the overlap of "Lin" and "insanity" makes it easy to recognize the source words.

To attempt to capture this kind of desirable structure, given base words $u = y_1 z_1$ and $v = y_2 z_2$, Nehovah calculates a raw structure score for a candidate neologism $n = y_1 z_2$ as

$$S(n) = \sigma(y_1, y_2) + \pi(z_1, z_2) + B(n, u, v)$$

where $\sigma(y_1, y_2)$ is the length of the suffix common to y_1 and y_2, $\pi(z_1, z_2)$ is the length of prefix common to z_1 and z_2 and

$$B(n,u,v) = \max\{\delta(\#(n), \#(u)), \delta(\#(n), \#(v))\}$$

where $\#(x)$ returns the number of syllables in x and δ is the Kronecker delta function [$B(n,u,v)$ equals 1 if neologism n maintains the same syllable count as either base word and 0 otherwise]. $S(n)$ therefore quantifies "catchiness" by measuring base word overlap and syllable count conservation.

Given this, the word structure score $\mathcal{W}(n)$ of neologism n is the normalized raw score, with normalization taken over the set of all candidate neologisms.

$$\mathcal{W}(n) = \frac{S(n)}{\max_{\tilde{n} \in N(C(w_1), C(w_2))} S(\tilde{n})}$$

Concepts. One of the primary goals of Nehovah is to convey the concepts of the source words in the neologism. While word structure *can* aid in conveying a concept, Nehovah also explicitly measures concept clarity for a neologism by scoring how well the base concepts are communicated in its prefix and suffix.

How clearly a concept is conveyed by the prefix or suffix of a base word obtained from WordNet is measured using MoreWords[4], a tool for crossword puzzles and other word games. MoreWords uses the words from the Enable2k North American word list that is used in well-known word games. It contains 173,528 words and does not include any hyphenated words, abbreviations, acronyms, or proper nouns. Querying MoreWords with a prefix/suffix x returns the set of words W_x that have x as a prefix/suffix in MoreWords and the approximate number of times each word $\tilde{u} \in W_x$ occurs per million words ($FPM(\tilde{u})$). $FPM(\tilde{u})$ is estimated from studies on the British National Corpus[5].

Nehovah determines how apparent the concept is in a prefix/suffix by comparing the frequency of the word that the prefix/suffix is derived from with the frequencies of other words that begin/end with the same prefix/suffix. A distinctiveness score for a prefix/suffix x of base word u is calculated by first calculating a distinctiveness ratio:

$$\phi(x,u) = \frac{FPM(u)}{\sum_{\tilde{u} \in W_x} FPM(\tilde{u})}$$

[4] www.morewords.com

[5] http://www.natcorp.ox.ac.uk/

The distinctiveness score is then calculated using (an empirically determined) piecewise linear interpolation on the value of the distinctiveness ratio:

$$\chi(x,u) = \begin{cases} 1, & \text{if } \phi(x,u) \geq 0.1 \\ 0.8 + 2\phi(x,u), & \text{if } 0.01 < \phi(x,u) < 0.1 \\ 80\phi(x,u), & \text{if } 0 \leq \phi(x,u) \leq 0.01 \end{cases}$$

This score differentiates between prefixes/suffixes that do not convey the concept, that partially convey the concept, and that completely convey the concept.

Because many pop culture words are not contained in MoreWords, Nehovah measures how clearly a concept is conveyed by a pop culture base word obtained from TheTopTens as the normalized count of the number of times that a pop culture word u appears in the set of lists $L(w)$ returned from TheTopTens for a given source word w:

$$\psi(u,w) = \frac{\lambda(u, L(w))}{\max_{\tilde{u} \in T(w)} \lambda(\tilde{u}, L(w))}$$

where $\lambda(u, L(w))$ represents the number of times a base word u appears in $L(w)$, and $T(w)$ represents the set of unique pop culture words in $L(w)$.

Note that this distinctiveness score indicates the "popularity" of the concept for a pop culture reference in the neologism by comparing the prevalence of other pop culture words to the prevalence of the entire base word (rather than by considering just some prefix or suffix of the base word).

Under the assumption that these distinctiveness scores correlate with conceptual content, given a source word w, a base word $u \in C(w)$ and a prefix/suffix x of u, a concept score for the base word is computed as

$$c(x,u,w) = \begin{cases} \chi(x,u), & \text{if } u \text{ appears in WordNet} \\ \psi(u,w), & \text{otherwise} \end{cases}$$

Finally, given a concept score for both a prefix y of base word u and a suffix z of base word v, the concept score $\mathcal{C}(n)$ of the created neologism $n = yz$ is simply the average of the concept scores of the base words and their prefix/suffix:

$$\mathcal{C}(n) = \frac{c(y,u,w_1) + c(z,v,w_2)}{2}$$

Uniqueness. A score for uniqueness should place greater value on words that are not commonly used (but still convey the source concept). For example, for the source word "pants," the base word "trousers" is more common than the base word "bloomers," although both convey the same concept. Uniqueness for a base word $u \in C(w)$ is calculated using the frequency per million words score from MoreWords ($FPM(u)$) relative to all of the other synonymous words in the concept set:

$$v(u,w) = 1 - \frac{FPM(u)}{\max_{\tilde{u} \in C(w)} FPM(\tilde{u})}$$

The uniqueness score $\mathcal{U}(n)$ for a neologism n formed from the base words u and v is simply the average of their uniqueness scores:

$$\mathcal{U}(n) = \frac{v(u,w_1) + v(v,w_2)}{2}$$

Neologism	Base Words		Source Words	
Nehovah	neologism	Jehovah	neologism	creator
divinage	divine	coinage	neologism	creator
machinative	machine	creative	machine	creative
Spritependency	Sprite	dependency	soda	addiction
Pepsidiction	Pepsi	addiction	soda	addiction
pisome	pizza	awesome	awesome	pizza
pimazing	pie	amazing	awesome	pizza
iniquitivate	iniquity	cultivate	evil	school
immoralize	immorality	civilize	evil	school
coalesception	coalesce	conception	concept	blend
portmanception	portmanteau	conception	concept	blend

Table 2: A set of example neologisms generated by Nehovah with their base words and the source words that were provided to Nehovah.

Pop Culture. The pop culture score indicates if one or both of the base words are pop culture words, allowing the emphasis of pop culture references. The pop culture score $\mathcal{P}(n)$ for a neologism n created from base words u and v is given by

$$\mathcal{P}(n) = \begin{cases} 1 & \text{if } u \text{ and } v \text{ are pop culture words} \\ 0.5 & \text{if } u \text{ or } v \text{ is a pop culture word} \\ 0 & \text{otherwise} \end{cases}$$

Combining Scores

The final score for a neologism is computed as a linear combination of the four attribute scores, weighted by user-selected coefficients (cf. the sliders in Figure 2):

$$\mathcal{S}(n) = \alpha_{\mathcal{W}}\mathcal{W}(n) + \alpha_{\mathcal{C}}\mathcal{C}(n) + \alpha_{\mathcal{U}}\mathcal{U}(n) + \alpha_{\mathcal{P}}\mathcal{P}(n)$$

Evaluation of Nehovah

We now examine Nehovah in the context of the creative tripod, which consists of skill, imagination, and appreciation (Colton 2008). Skill is the ability of a system to produce something useful. Imagination is the ability of the system to search the space of possibilities and produce something novel. Appreciation is the ability of the machine to self-assess and produce something of worth. We also evaluate Nehovah with respect to its accountability–the ability of the system to explain why it generated the artifact it generated.

Skill

Nehovah demonstrates skill by generating neologisms that convey the concepts in the base words and have proper word structure. First, proposed neologisms with invalid word structure are discarded. Next, Nehovah determines if a pop culture word is valid based on its presence in Wikipedia. Wikipedia is a dynamic source that does contain neologisms (Veale 2006) and consulting Wikipedia provides a safe-guard against low quality user-supplied content in TheTopTens. Finally, only splitting the words on their syllable boundaries aids in creating word fragments that convey

meaning and are able to be blended in a way that forms a plausible word.

The skill of any system is most easily demonstrated in the artifacts that it produces. Exhibit A for the Nehovah system is its own name, which is the direct result of providing the (originally anonymous) system with the source words "neologism" and "creator." The name Nehovah is a mix of the words "neologism" and "Jehovah", and it is readily apparent that Nehovah incorporates the word "Jehovah"; another candidate neologism was "Neohovah," which conveys a bit more of the meaning of "neologism" but is not as structurally pleasing since an additional syllable is added.

Other examples of neologisms created by Nehovah are shown in Table 2. As a further demonstration, consider the following arguably coherent sentence constructed from some of the neologisms from Table 2:

> Spritependency *is a* machinative *neologism created through* portmanception *to describe someone who is addicted to Sprite.*

We also point out that the neologism "immoralize" is an actual word found in some dictionaries (it is not found in WordNet). According to the Merriam-Webster on-line dictionary, it means "to make immoral"[6] which is what is conveyed by the neologism. In other words, the system (re)invented a real word, a nice demonstration of Boden's P-creativity.

Accountability

In addition to producing a set of neologisms, Nehovah also includes the base words that were blended together to produce the neologism (see the expansion of the third neologism in the righthand pane of Figure 2). Therefore, at some level Nehovah can explain how it created a neologism. The perceived creativity of the neologisms in Table 2 is likely increased with the available explanation of which base words were blended together as well as what the source words are. For example, "portmanception" is created from the source words "concept" and "blend" using "portmanteau" and "conception" as base words. Using "portmanteau" in the place of "blend" and "conception" in the place of "concept" conveys similar meaning; revealing the connection between the base words and source words helps justify the quality and creativity of the neologism.

Imagination

A Google search for most of the generated neologisms will show that Nehovah provides novel artifacts. The hits for "Nehovah" contain references to this project and an individual's name. Most of the neologisms have no hits when searched for in Google or the hits returned are names or screen names ("divinage" is a *World of War Craft* user name).

Nehovah explores all possible combinations of prefixes and suffixes derived from the base words. Further, Nehovah also considers the synonyms for all possible senses of each base word for each possible part of speech. Using all of the possible senses for all of the parts of speech for a source word along with an ever-expanding set of free-form, user-defined (pop culture) lists can create a potentially very large search space and produce unpredictable results. For example, if "evil" and "school" are used as the source words with the intended sense of school being an "educational institution", then seeing a neologism such as "Darth_swim" would likely be somewhat unexpected (the base words of the neologism are "Darth_Vader" from the TheTopTens list "The 10 Most Evil Villains in Video Games" and "swim", a hypernym of one of the senses of the verb "school"). This, however, demonstrates the imagination of Nehovah, since it takes into consideration other and unintended senses of a source word to produce more creative neologisms. Of course, the flip side of such imaginative creations is that unintended senses can cause problems, if the main goal is to create a neologism that captures a specific sense of a source word. Thus, there is a tension between creating a rich concept set that includes all of the possible senses for a source word and generating neologisms that convey the concept of the intended sense.

Best Dog Breeds	Best Hot Dog Toppings
Pitbull	Coney Sauce
Rottweiler	Mustard
Chihuahua	Stadium Mustard
Great Dane	Relish
Miniature Pinscher	Ketchup

Table 3: The top five words returned from two lists from TheTopTens for the source word "dog", demonstrating the range of synonyms that Nehovah uses as base words.

Using the pop culture references allows Nehovah to demonstrate imagination in an unusual and contemporary fashion by using social/popular connections between words to convey meaning. Most people who are familiar with the Star Wars series would recognize the word "Darth" as having an evil connotation. As with using all the senses for a base word, some of the words from TheTopTens do not capture the intended concept of the base word. For example, consider the top five entries from two of the TheTopTens lists returned for the word dog shown in Table 3. The "Best Dog Breeds" list conveys the concept of dog to most users better than the "Best Hot Dog Toppings" list. An example set of neologisms is shown in Table 4 that shows the unintended use of the "Best Hot Dog Toppings" versus using "Best Dog Breeds" when blending the source words "robot" and "dog". Despite being irrelevant for the animal dog, these examples demonstrate the imagination of Nehovah in generating neologisms. And, in fact, the neologism "Terminaise" could be a serendipitous discovery for an exciting new condiment if the intended sense of the word"dog" was "hot dog".

Appreciation

Nehovah's appreciation is demonstrated by determining which neologisms are the "best" given a set of base words and which scoring criteria are weighted the highest. Ta-

[6] http://www.merriam-webster.com/dictionary/immoralize

	Neologism	Base Words		Score
Best 10	rottweilers: Revenge of the Fallen	rottweiler Top Ten Best Dog Breeds	Transformers: Revenge of the Fallen Top Ten Best Robot Movies of All Time	0.786
	rottweilerminator 3	rottweiler Top Ten Best Dog Breeds	Terminator 3 Top Ten Best Robot Movies of All Time	0.786
	automaton terrier	automaton	boston terrier Top Ten Best Dog Breeds	0.762
	automatian	automaton	dalmatian Top Ten Best Dog Breeds	0.755
	chihuahuaton	chihuahua Top Ten Best Dog Breeds	automaton	0.754
	automestic	automaton	domestic	0.752
	golden retrievers: Revenge of the Fallen	golden retriever Top Ten Best Dog Breeds	Transformers: Revenge of the Fallen Top Ten Best Robot Movies of All Time	0.750
	dobermansformers: Revenge of the Fallen	doberman Top Ten Worst Dog Breeds	Transformers: Revenge of the Fallen Top Ten Best Robot Movies of All Time	0.714
	doberminator 3	doberman Top Ten Worst Dog Breeds	Terminator 3 Top Ten Best Robot Movies of All Time Rise	0.714
	chihuahuanic attack	chihuahua Top Ten Best Dog Breeds	panic attack Greatest Robot Wars Robots Of All Time	0.714
Worst 10	panicpoodle	panic attack Greatest Robot Wars Robots Of All Time	poodle Top Ten Best Dog Breeds	0.143
	bulroadblock	bull terrier Top 10 Guard Dog Breeds	roadblock Greatest Robot Wars Robots Of All Time	0.143
	cheeatomic	cheese Top Ten Best Hot Dog Toppings	atomic Greatest Robot Wars Robots Of All Time	0.143
	labradorroadblock	labrador retriever Top Ten Best Dog Breeds	roadblock Greatest Robot Wars Robots Of All Time	0.143
	borderrobots	border collie Top Ten Best Dog Breeds	robots Top Ten Best Robot Movies of All Time	0.143
	bulrobots	bull terrier Top 10 Guard Dog Breeds	robots Top Ten Best Robot Movies of All Time	0.143
	borderroadblock	border collie Top Ten Best Dog Breeds	roadblock Greatest Robot Wars Robots Of All Time	0.143
	labradorrobots	labrador retriever Top Ten Best Dog Breeds	robots Top Ten Best Robot Movies of All Time	0.143
	atomustard	atomic Greatest Robot Wars Robots Of All Time	mustard Top Ten Best Hot Dog Toppings	0.143
	shetlandtornado	shetland sheepdog Top 10 Smartest Dogs	tornado Greatest Robot Wars Robots Of All Time	0.143

Table 5: Highest rated 10 and lowest rated 10 neologisms generated by Nehovah using the source words "dog" and "robot" with all scoring attributes equally weighted. The higher rated neologisms tend to flow better and convey the concepts of the base words better than the lower rated neologisms.

ble 5 shows the highest rated 10 and lowest rated 10 neologisms created using the source words "dog" and "robot" as scored with all attributes equally weighted. The source words "dog" and "robot" were chosen for this example because both source words have pop culture references and clearly demonstrate the effects of the different scoring attributes. Comparing the two sets of neologisms in Table 5, the highest rated 10 neologisms flow better and better capture the source concepts. The bottom 10 do not flow as well and this often contributes to (further) obfuscation of the source concepts. For example compare "rottweilerminator" and "cheeatomic"—the former better follows the word structure of both base words and the concepts are more clearly conveyed.

Each of Nehovah's scoring attributes can be weighted by a user to increase or decrease its relative importance. Table 6 shows a sampling of neologisms derived from blending the source words "robot" and "dog", when weighting is skewed completely to one of the four scoring factors. Each sub-table gives a set of neologisms weighted exclusively for the factor titled above it. For example, looking at the first sub-table (titled Pop Culture), for all neologisms, both source words are from the TheTopTens, although the word structures may be awkward and the concepts may not

Best Dog Breeds

Neologism	Base Words	
dobermaton	doberman	automaton
rottweilerminator_3	rottweiler	Terminator_3
dobermansformers	doberman	transformers

Best Hot Dog Toppings

Neologism	Base Words	
sauerminator_3	sauerkraut	Terminator_3
Terminaise	Terminator_3	mayonnaise
mustardmaton	mustard	automaton

Table 4: A set of sample neologisms for the source words "dog" and "robot" using two different lists from TheTopTens for the source word "dog".

be apparent e.g. "alasdo" from the source words "alaskan malamute" and "tornado". Neologisms in the list weighting only the Concept score tend to have prefixes and suffixes that are evocative of distinct base words, such as "bot" from the base word "robot". When Word Structure is the sole factor, the created neologisms look the most like real words, e.g., "Terman shepherd", strongly overlaps "Terminator" with "German shepherd" and preserves the number of syllables in "German shepherd." In the case of weighting solely for Uniqueness, the resulting neologisms and their base words are often quite unusual, sometime at the expense of understandability, e.g. "godiron" from "golem" and "andiron". As expected, weighting according to a single factor filters the neologisms, presenting only those that have a particular attribute, often at the expense of other factors.

Overall, we tend to favor the word structure and concepts factors for creating the best neologisms. These help to convey the concepts contained in the base words and also produce more realistic appearing words as they have valid letter sequences and are similar to the base words. While favoring the concept and word structure factors, the pop culture and unique factors can be used as a secondary bias towards certain types of base words to be blended together.

Conclusions and Future Work

In this paper, we have presented Nehovah, a system that generates neologisms from a set of user-provided source words by searching the space of synonyms and then blending two base words. We have argued for Nehovah's ability to demonstrate some necessary characteristics for creativity, including skill, imagination, appreciation and accountability.

Future work includes incorporating a learning mechanism so that users can indicate which neologisms they prefer. Nehovah could then use this information to better score the neologisms. An interesting line of future work includes generating a definition for a neologism using the base words. This would involve solving at least two difficult problems. The first problem is generating the definitions. Candidate definition components could be found by searching Wikipedia, an on-line dictionary, and/or another source for definitions for each source word. A potential definition would then be formed by blending candidate components in a way that both conveys the concept from each source word and is readable (i.e. correct grammar). The second problem is validation of the potential definition, which may be accomplished, for example, through a user study/game where Nehovah could *learn* to match definitions to neologisms based on users' votes.

Pop Culture

Neologism	Base Words	
1 labrador retrogates	labrador retriever	surrogates
1 alasdo	alaskan malamute	tornado
1 lharestorm	lhasa apso	firestorm
1 ketchupsycat	ketchup	pussycat
1 iroadblock	ibizan hound	roadblock

Concepts

Neologism	Base Words	
1 supnism	support	mechanism
1 scountomaton	scoundrel	automaton
1 domesrobot	domestic	robot
1 supbot	support	robot
1 scounrobot	scoundrel	robot

Word Structure

Neologism	Base Words	
1 pomers	pomeranian	transformers
1 automatian	automaton	dalmatian
1 Terman shepherd	Terminator 3	german shepherd
1 firestic	firestorm	domestic
1 Terman pinscher	Terminator 3	doberman pinscher

Uniqueness

Neologism	Base Words	
1 wiegolem	wiener	golem
1 gomiliaris	golem	familiaris
1 bliglem	blighter	golem
1 godiron	golem	andiron
1 gofiredog	golem	firedog

Table 6: Sample of neologisms created from the base words "dog" and "robot" using weighting schemes skewed completely toward a single factor, demonstrating Nehovah's appreciation for each scoring measure. Each set of neologisms possesses the desired attribute, often at the expense of others, e.g., the neologisms weighted for uniqueness are difficult to interpret and those weighted for pop culture have poor structure.

Acknowledgements

We would like to thank Dylan Mills from TheTopTens for providing an API for Nehovah.

References

Baker, K. 1999. Seussisms and violations to universal language constraints. In Hisagi, M., and Bradinova, M., eds., *Working Papers in Linguistics*, volume 6. George Mason University.

Boden, M. 1994. Creativity: A framework for research. *Behavioral and Brain Sciences* 17(3):558–568.

Carroll, L. 1871. *Through the Looking-Glass*. Macmillan.

Colton, S.; Valstar, M. F.; and Pantic, M. 2008. Emotionally aware automated portrait painting. In *Proceedings of the Third International Conference on Digital Interactive Media in Entertainment and Arts*, 304–311.

Colton, S. 2008. Creativity versus the perception of creativity in computational systems. In *AAAI Spring Symposium: Creative Intelligent Systems*, 14–20. AAAI.

Cook, P., and Stevenson, S. 2010. Automatically identifying the source words of lexical blends in English. *Computational Linguistics* 36:129–149.

Cope, D. 2005. *Computer Models of Musical Creativity*. The MIT Press.

Duch, W., and Pilichowski, M. 2007. Experiments with computational creativity. *Neural Information Processing – Letters and Reviews* 11(4-6):123–133.

Fellbaum, C., ed. 1998. *WordNet: an electronic lexical database*. MIT Press.

Gibson, W. 1982. *Neuromancer*. Ace Books.

Mendes, M.; Pereira, F. C.; and Cardoso, A. 2004. Creativity in natural language: Studying lexical relations. In *Proceedings of the Workshop on Language Resources for Linguistic Creativity, 4th International Conference on Language Resources and Evaluation (LREC)*.

Morris, R.; Burton, S.; Bodily, P.; and Ventura, D. 2012. Soup over bean of pure joy: Culinary ruminations of an artificial chef. In *Proceedings of the Third International Conference on Computational Creativity*, 119–125.

Norton, D.; Heath, D.; and Ventura, D. 2011. Autonomously creating quality images. In *Proceedings of the Second International Conference on Computational Creativity*, 10–15.

Rahman, F., and Manurung, R. 2011. Multiobjective optimization for meaningful metrical poetry. In *Proceedings of the Second International Conference on Computational Creativity*, 4–9.

The United States Patent and Trademark Office. 2014. *Performance and Accountability Report, Fiscal Year 2013*. Government Printing Office.

Veale, T., and Butnariu, C. 2006. Exploring linguistic creativity via predictive lexicology. In *The Third Joint Workshop on Computational Creativity*, ECAI 2006.

Veale, T., and Hao, Y. 2007. Comprehending and generating apt metaphors: A web-driven, case-based approach to figurative language. In *Proceedings of the Association for the Advancement of Aritificial Intelligence*, 1471–1476. AAAI Press.

Veale, T.; O'Donoghue, D.; and Keane, M. T. 2000. Computation and blending. *Computational Linguistics* 11(3/4):253–281.

Veale, T. 2006. Tracking the lexical zeitgeist with WordNet and Wikipedia. In *Proceedings of the 17th European Conference on Artificial Intelligence*, 56–60.